

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays on the Price of Risk

Permalink

<https://escholarship.org/uc/item/6rg7g5rx>

ISBN

9798189275975

Author

Perry, Andrew

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Essays on the Price of Risk

By

Andrew Maurice Perry

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Eben Lazarus, Chair

Professor Christine Parlour

Professor Martin Lettau

Professor David Sraer

Spring 2026

Essays on the Price of Risk

Copyright 2026

by

Andrew Maurice Perry

Abstract

Essays on the Price of Risk

By

Andrew Maurice Perry

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Eben Lazarus, Chair

This dissertation explores how risk is priced across different sectors of the economy. In particular, Chapter 1 focuses on equity markets while Chapters 2 and 3 look at the commercial banking sector.

In Chapter 1 I develop an information-geometric framework for selecting priced risk factors from a high-dimensional candidate set. The stochastic discount factor is cast as a point on a statistical manifold of exponential-affine pricing kernels, and the manifold's dual flatness yields an exact Pythagorean decomposition of Kullback-Leibler divergence across nested SDFs. The decomposition drives a greedy forward-selection algorithm that is isomorphic to the portfolio problem of an investor with multiplier preferences. The investor's ranking of candidates by worst-case welfare loss coincides with the statistical ranking by KL divergence, and the Hansen-Jagannathan distance emerges as a leading-order approximation. A Wald test on estimated risk prices and an entropy bound calibrated to the investor's detection-error probability provide distinct stopping criteria. Applied to a rich set of asset returns, the algorithm selects a parsimonious SDF spanning roughly nine thematic categories. Debt issuance, quality, and profit growth form a stable core; accruals is absent from every specification. Selected factors derive a substantial share of their pricing content from non-Gaussian features of the return distribution.

Chapter 2 introduces a novel dataset of bank resolution terms constructed from over forty years of FDIC press releases to analyze the determinants of franchise value in failed bank

auctions. Using a structural model of the auction process that encompasses both the liability and asset sides of the failed bank’s balance sheet, I characterize the deposit franchise as a wasting asset and endogenize the acquirer’s optimal asset take-up decision as a function of comparative advantage relative to the FDIC’s own resolution technology. The “Toxic Ratio”, the proportion of assets retained by the FDIC in receivership, emerges as the dominant driver of resolution outcomes across the pricing, cost, and selection channels. Because the retained asset share and the deposit premium are jointly determined at the auction, I instrument the realized toxic ratio with the pre-failure loan-loss provision (LLP) ratio. The headline cost result is that a one percentage point increase in the loan-loss provision ratio raises the cost to the Deposit Insurance Fund by approximately 86 basis points of total assets, making asset quality the dominant quantitative determinant of fiscal exposure. Accelerated deposit runoff prior to failure further reduces bid premiums and increases resolution costs, consistent with the franchise-value mechanism. This trade-off highlights a shift in regulatory technology from liquidation to financial engineering, with significant implications for moral hazard and the “Too-Big-To-Fail” doctrine.

Finally, Chapter 3 investigates the competitive distortions introduced by “Too-Big-To-Fail” (TBTF) guarantees within the U.S. banking sector, specifically examining the flight-to-safety dynamic that redistributes funding during periods of systemic distress. To formalize the mechanism through which implicit government backstops alter deposit market competition, I develop a continuous-time structural model that integrates spatial differentiation with endogenous default risk. The model yields a “scissors” dynamic, predicting that while distressed small banks must compress margins to retain uninsured depositors, TBTF incumbents strategically widen spreads to monetize their perceived safety. To empirically test these predictions across the entire size spectrum of the banking system, I construct a novel, continuous measure of default probability for the universe of regional and local banks using a Weighted Ensemble machine learning framework. Utilizing a difference-in-differences identification strategy, I find robust evidence of a post-GFC regime shift: uninsured deposit demand has become highly elastic to bank solvency, but with a stark size hierarchy. A one-percentage-point increase in default probability is associated with an approximately

48 percent increase in uninsured deposit share for Global Systemically Important Banks (GSIBs), while local banks experience an essentially flat or slightly negative relationship. These findings suggest that regulatory reforms have not eliminated the TBTF subsidy but have instead repriced it, allowing systemic institutions to act as liquidity sinks that absorb franchise value from smaller competitors during economic contractions.

“The true mind can weather all the lies and illusions without being lost.

The true heart can tough the poison of hatred without being harmed.”

- Lion Turtle

Contents

List of Figures	vi
List of Tables	vii
1 Geometry of the SDF: The Factor Zoo as Structured Ambiguity	1
1.1 Introduction	1
1.2 Related Literature	4
1.3 Theory	10
1.3.1 Statistical Geometry	10
1.3.2 Investor's Portfolio Choice Problem	19
1.4 Methodology	29
1.4.1 Forward Selection Algorithm	29
1.4.2 Estimating the Exponential-Affine SDF	33
1.5 Empirics	39
1.5.1 Discussion of Parameters	40
1.5.2 Illustrative Example	41
1.5.3 The Full Factor Zoo	45
1.5.4 Path Dependence	48
1.5.5 Stopping Criteria in Action	54
1.5.6 Wald v.s. ESS	58
1.5.7 Higher-Order-Moments	62
1.5.8 Summary	65
1.6 Conclusion	66

2	Selling Distress: The Economics of Failed Bank Resolutions	70
2.1	Introduction	70
2.2	Institutional Details	75
2.2.1	S&L Crisis and the FDICIA	76
2.2.2	FDIC Resolution Process	78
2.2.3	Related Literature	86
2.3	Press Release Data	88
2.3.1	Construction and Coverage	89
2.3.2	Failure Landscape	90
2.3.3	Balance Sheet Primitives	92
2.3.4	Resolution Outcomes	96
2.4	Theory	101
2.4.1	Environment	102
2.4.2	Valuing the Deposit Franchise	103
2.4.3	Asset Valuation	104
2.4.4	Auction Mechanism	106
2.4.5	Retained Share of Assets	110
2.4.6	Comparative Statics	112
2.4.7	Probability of P&A	114
2.4.8	Testable Predictions	116
2.5	Empirics	117
2.5.1	Data and Sample	118
2.5.2	Deposit Pricing Channel	119
2.5.3	FDIC Cost Channel	121
2.5.4	P&A Selection Channel	123
2.5.5	Heckman Corrected Deposit Pricing Channel	125
2.5.6	Summary	127
2.6	Conclusion	127

3	Depositing Trust: Banking Demand and Too-Big-to-Fail	130
3.1	Introduction	130
3.2	Related Literature	136
3.3	Model	138
3.3.1	Model Setup	138
3.3.2	Valuing Default Risk	139
3.3.3	Deposit Market	140
3.3.4	Bank Competition	141
3.3.5	Equilibrium	143
3.4	Empirics	154
3.4.1	Data	155
3.4.2	Default Model	158
3.4.3	Causal Tests	164
3.4.4	Results	166
3.5	Conclusion	171
	Bibliography	173
A	Appendix to Chapter 1	190
A.1	Proofs	190
A.1.1	Proposition 1	190
A.1.2	Proposition 2	193
A.1.3	Lemma 1	194
A.1.4	Proposition 3	195
A.1.5	Lemma 2	201
A.1.6	Lemma 3	202
A.2	Derivations	203
A.2.1	Closed Form Kullback-Leibler Divergence	203
A.2.2	Minimum Cross Entropy Estimator	205
A.2.3	Asymptotic Standard Errors	207

A.2.4	Evaluating the Significance of a Candidate Factor	210
A.3	Numerical Implementation	211
A.3.1	Bounding the Concentration Ratio	211
A.3.2	Numerical Implementation of the MCE Estimator	212
B	Appendix to Chapter 2	222
B.1	Model Derivations	222
B.1.1	Deposit Franchise Value	222
B.1.2	Asset Valuation	224
B.1.3	Scoring Rule and Reduction to a Scalar Auction	224
B.1.4	Optimal Asset Take-Up	226
B.1.5	Maximized Private Value	226
B.1.6	Symmetric Equilibrium Bidding	227
B.1.7	Deposit-Side Comparative Statics	228
B.1.8	Asset-Side Comparative Statics	229
B.1.9	Premium–Retention Relationship	230
B.1.10	Probability of P&A	231
B.2	Robustness	232
B.2.1	IV First-Stage Regression	232
B.2.2	Heckman Selection-Corrected Pricing	232
B.2.3	Bias Decomposition	234
B.3	Data Sources	237
B.4	FDIC Failed Bank Press Release	238
C	Appendix to Chapter 3	240
C.1	Model Derivations	240
C.1.1	First Passage Time ODE	240
C.1.2	First Passage Time Security	242
C.1.3	Equity Value ODE	243
C.1.4	Equity Value and Optimal Default Bound	244

C.1.5	Proposition 1	246
C.1.6	Corollary 1	248

List of Figures

1.1	Kullback-Leibler Scree Plot	49
1.2	KL Divergence by Theme	55
1.3	Wald Statistic by Theme	57
1.4	Higher-Order Moment Structure	63
2.1	Annual Number of Bank Failures	72
2.2	Location of Domestic Bank Failures	91
2.3	Deposit Runoff Distribution	93
2.4	Deposit Premium v.s. Deposit Beta	94
2.5	Loan Loss Provision Ratio	95
2.6	Retained Asset Share v.s. LLP Ratio	96
2.7	Retained Asset Share Distribution	100
2.8	Uninsured Deposit Recovery Rate	101
2.9	Motivating Bivariates	102
3.1	GSIB Deposit Rate Spread	131
3.2	Equilibrium Deposit Rates	150
3.3	Equilibrium Deposit Market Shares	151
3.4	Simulation of Bank Fundamentals	152
3.5	Simulation of Bank Deposit Spreads and Shares	153
3.6	Simulation of Bank Profits	154
3.7	Distribution of Estimated Default Probabilities	162
3.8	Wells Fargo Time Series of Estimated Default Probabilities	163

List of Tables

1.1	Factor Zoo Taxonomy	28
1.2	Illustrative Example Results	43
1.3	Full Factor Zoo Results	46
1.4	Selected Factor Themes	51
1.5	Results of the ESS Line Search	59
1.6	ESS Violations by Theme	60
1.7	Final Selected Factor Roster	67
2.1	Coverage of FDIC Press Release Data	90
2.2	In-Sample Bank Failure Summary Statistics	97
2.3	Resolution Outcomes by Period	98
2.4	Resolution Outcome Summary Statistics	99
2.5	Deposit Premium Regressions	120
2.6	FDIC Cost Regressions	122
2.7	Probability of P&A Regressions	124
2.8	Heckman Second Stage	125
3.1	In-Sample Banking Data Summary Statistics	157
3.2	Estimated Default Probabilities Summary Statistics	161
3.3	Bank Deposit Share Regressions	167
3.4	Bank Deposit Spread Regressions	170
A1	LLP IV First-Stage	232

A2	Heckman First Stage	233
A3	Heckman Decomposition	234
A4	FDIC Press Release Data Description	237

Acknowledgements

Dad, it is almost surreal that this fifteen-year journey is coming to an end. The magnitude of this accomplishment is difficult to fully grasp, in part because you made it so seamless. When someone as dedicated as you moves the barriers out of the way, it is easy to overlook just how many there were. You handled everything, what class to take next, how to pay for the following semester, and a hundred other logistical concerns I never had to think about, so that I could focus entirely on learning. You always ask me where I picked up a piece of knowledge when I share it, and thinking about that now, I can say with confidence that this dissertation represents the sum of everything I have absorbed: the linear algebra, computer science, and probability from Stanford; the economics, finance, and econometrics from Berkeley; and even the calculus and the Pythagorean theorem from high school. I have all of that knowledge, and much more, because of you. Thank you for this extraordinary gift.

To Neal and Bill: even from across the country, you have always treated me like one of your own. You checked in on me consistently and ensured I was genuinely happy. The perspective and life advice you have both shared have been central to the confidence I carry today. Neal, your motto— “Never give up”—is the Perry way, and it has served me well. I know that mindset will continue to serve me in the future. Bill, your advice to always enjoy life is incredibly important to me; it has shaped who I am just as much as anything else I have learned. Thank you both, and I look forward to continuing to learn from you.

My advisors—Eben, Christine, David, and Martin—have been the finest group of mentors I could have hoped for. The knowledge you shared, the frameworks for tackling hard problems, the patience, and the life advice have shaped me in ways I am still discovering. Learning how to do research, and more fundamentally how to ask a good question, has been the most difficult thing I have ever had to learn, and you guided me through every step. Eben, thank you for entertaining my half-formed thoughts on factor models and helping me turn them into something coherent. Christine, you did not have to sit with me every Thursday for two years working through ideas: that kind of dedication is rare in any setting, let alone academia. I would not have made it through the program without you. Thank you

for not giving up on me. To Mrs. Grano, Mr. Hoag, Hung Le, Douglas Bernheim, Monica Piazzesi, Guenther Walther, Noé Pablo Lozano, Peter Maxted, Richard Stanton, Matteo Benetton, and Johan Walden: thank you for the advice and guidance.

Dr. Professor Leegee, in six short (long) years, you have become one of my dearest friends. This would have been impossible without you. I know San Diego State is gaining an exceptional researcher and educator, because you taught me to be both. I was truly honored to be one of your groomsmen, and I remain constantly struck by your patience and kindness. Thank you—oh, and the Green Line takes only twenty-three minutes to get from the airport to campus. To Karin and Soobin, thank you for being wonderful friends and collaborators. I cannot wait to see what you both produce. To Jun, Yao, Sam, Clint, Gene, Jake, Minghao, Anand, Tim, Jae, Ethan, Seungjoo, Pietro, and many others, thank you for the help and the laughs.

Darrith, for over a decade you have supported me and made me feel appreciated. I am so glad we both stayed in the dorm over Thanksgiving break freshman year. You have always given me a space to work and to be comfortable: some of my biggest breakthroughs came in Castaño, East, and even at Everlaw. Thank you for being my best friend. To Kavita, my true accomplice: some of my most joyful hours over these past few years have been spent on your couch, baking, and wandering grocery store aisles in search of charcuterie. The second chapter of this dissertation does not exist without you two. To Grace, Chris, Perris, Kelsey, Lauren, Autumn, Blake, and Jeffrey: high school and college would have been miserable without you. The camping trips and 2018 Yost are the memories I hold closest.

Dew, I hope you found your forever home, where you get all the love and springs you deserve.

Chapter 1

Geometry of the SDF: The Factor Zoo as Structured Ambiguity

1.1 Introduction

The cross-section of expected stock returns remains unsettled. Cochrane (2011) characterized the state of empirical asset pricing as having a “zoo of [...] factors” and the zoo has only continued to expand. Harvey et al. (2015) documents over 300 factors published in top journals. The proliferation is not merely troubling; it is epistemically disabling. Every test of a candidate factor is conducted against a benchmark, and the benchmark itself is contested. Statistical significance is therefore conditional on which fraction of the zoo one chooses to control for, which makes the evaluation problem circular. The zoo grows because there is no benchmark model. There is no procedure for factor selection that simultaneously respects economic theory and provides criteria for which factors meaningfully contribute to pricing. This paper closes that gap.

I develop an information-geometric framework that recasts factor selection as the traversal of a statistical manifold of admissible pricing kernels. The manifold $\mathbb{M}$ consists of exponential-affine stochastic discount factors (SDF), which form an exponential family whose dual flatness equips the space with a canonical divergence measure, the Kullback-Leibler (KL) divergence. I show that the traversal of $\mathbb{M}$ is isomorphic to the portfolio problem of

an investor with multiplier preferences in the sense of Hansen and Sargent (2001), and that an exact Pythagorean decomposition of KL divergence across nested SDFs drives a greedy forward-selection algorithm. The algorithm is not a heuristic approximation. It is an exact manifold traversal at each step.

The mechanism rests on the dual flatness of $\mathbb{M}$. The exponential family of SDFs generates an e -flat manifold, and the moment-restricted space is m -flat under the mixture connection. This dual flatness yields an exact Pythagorean decomposition of KL divergence. This allows for a decomposition of residual pricing error to the true (unobservable) SDF. Therefore, minimizing the residual error after adding candidate factor $\mathbf{g}$ is mathematically equivalent to maximizing the divergence of the new factor relative to the baseline set of factors $\mathbb{S}$.

The economic isomorphism follows from the same exponential-tilt structure. A multiplier-preference investor's worst-case distortion h^* is itself an exponential tilt and lives on $\mathbb{M}$. The investor's ranking of candidate factors by worst-case welfare loss coincides with the purely statistical ranking by KL divergence. The two are not aligned by assumption. They coincide as a consequence of the e -flat geometry interacting with the exponential form of the worst-case distortion. The Hansen and Jagannathan (1997) (HJ) distance is recovered from this framework as a second order approximation. Unlike the HJ metric, KL divergence respects the higher order moments of the return distribution. The geometric criterion is therefore crash-sensitive in a way that mean-variance pricing is not, prioritizing factors whose pricing content concentrates in the left tail of the return distribution where marginal utility is high.

The framework yields two distinct stopping criteria. A Wald test on the estimated risk price $\hat{\theta}_{\mathbf{g}}$ screens for statistical detectability and an entropy bound $\bar{\eta}$, calibrated from the investor's detection-error probability ε as in Hansen and Sargent (2010), screens for economic materiality. Each guards a distinct failure mode. The Wald test prevents false inclusion driven by sampling noise; the entropy bound prevents kitchen-sink overfitting that admits factors whose genuine pricing signal falls below the threshold the adversary's distortion can rationalize. Neither criterion can substitute for the other.

I apply the procedure to a candidate universe of 157 factors, comprising the 151 characteristic-based long-short factors of Jensen et al. (2023) augmented with the six standard Fama-French

factors, and price them against 73 principal components of 4,686 characteristic-sorted decile portfolios spanning 666 months from Chen and Zimmermann (2022). At the baseline calibration, the algorithm selects 14 factors spanning 9 thematic categories.

Across a sweep of parameter values a stable core of themes emerges. Debt issuance and quality each enter in 6 of 7 trials, profit growth in 5 of 7. Seasonality, value, the equity premium, and size form a middle tier represented in 3 to 4 trials. Accruals never enters the selected model, despite the candidate universe containing 6 accruals factors, a pattern consistent with substantial overlap between the accruals signal and the quality theme that does enter. Path dependence operates primarily within themes through rotation among near-substitutes. Between-theme orderings are comparatively stable. The algorithm therefore recovers a consistent ranking of the cross-sectional pricing structure rather than an arbitrary subset of the candidate universe.

The selected factors derive a substantial share of their pricing content from non-Gaussian features of the return distribution. The KL contribution of each factor decomposes additively into a Gaussian component, which captures the mean-variance pricing content available to a standard SDF, and a higher-order component, which captures skewness, kurtosis, and tail structure. This pattern is the structural counterpart of the robust investor’s sensitivity to model misspecification. The adversarial tilt concentrates probability mass on tail states, and factors whose pricing information loads on those states are correspondingly informative.

The framework relates closely to several strands of prior work. Feng et al. (2020) develop a double-selection LASSO procedure that evaluates the marginal contribution of new factors against a high-dimensional control set. Their criterion is statistical, namely significance after controlling for redundancy, whereas the criterion here is economic, derived from the welfare loss of an investor under ambiguity, and it coincides with the statistical criterion. Lettau and Pelger (2020) modify the construction of factors through risk-premium PCA to incorporate cross-sectional pricing information. This paper takes the candidate factors as given and modifies the selection rule.

Ghosh et al. (2016) introduce KL divergence to asset pricing as a tool for nonparametric estimation of an unobservable SDF component in the consumption CAPM framework. Their

KL is a fitting criterion for a single SDF; here KL is a ranking criterion across nested SDFs. ichi Amari (2020) supplies the dually flat manifold structure and the generalized Pythagorean theorem that make the decomposition exact. Kitamura and Stutzer (1997) develop the Minimum Cross-Entropy (MCE) estimator that I use to estimate the exponential-affine SDF. Prior implementations in asset pricing solve the inner problem for fixed risk prices or restrict the parameter space. This paper is, to my knowledge, the first to estimate the full nested exponential-affine SDF over both Lagrange multipliers and risk prices under the pricing restrictions.

The paper is organized as follows: Section 1.2 reviews the literature this paper contributes to, Section 1.3 develops the geometric and economic theories, Section 1.4 describes the forward selection algorithm and the architecture of the MCE optimizer, Section 1.5 gives the result of the method applied to progressively richer universe of assets, and Section 1.6 concludes.

1.2 Related Literature

This paper connects several distinct literatures: the development of equilibrium and no-arbitrage factor pricing models, the subsequent proliferation and critique of proposed risk factors, statistical and machine-learning approaches to dimension reduction in factor space, information-theoretic and entropy-based methods in econometrics, the geometry of exponential families, robust decision theory under model ambiguity, and greedy approximation algorithms. I review each strand in turn, emphasizing how the present paper draws on and departs from existing work.

The theoretical foundations of factor pricing trace to the Capital Asset Pricing Model (CAPM) of Sharpe (1964), Litner (1965), and Mossin (1966), which identifies the market portfolio as the sole source of systematic risk and implies that expected excess returns are proportional to market betas. Black (1972) extended the model to accommodate restricted borrowing, replacing the risk-free rate with a zero-beta portfolio return. Fama and MacBeth (1973) developed the two-pass cross-sectional regression methodology that became the standard empirical test of the CAPM and its successors. Early empirical tests, including those

of Black et al. (1972), revealed departures from the single-factor prediction: small firms, high book-to-market firms, and high-earnings-yield firms earned returns inconsistent with their market betas alone Banz (1981); Basu (1983); Fama and French (1992).

Ross (1976) generalized the theoretical framework with the Arbitrage Pricing Theory (APT), demonstrating that in large economies, the absence of arbitrage implies expected returns are approximately linear in exposures to a few common factors. The APT permits multiple sources of systematic risk but does not identify which factors are priced. Chamberlain and Rothschild (1983) formalized this insight within the approximate factor model, showing that when the eigenvalues of the idiosyncratic covariance matrix remain bounded as the cross-section grows, the return covariance structure is well approximated by a few pervasive factors. This theoretical architecture—equilibrium pricing through the stochastic discount factor (SDF), no-arbitrage restrictions, and approximate factor structure—provides the scaffolding on which the subsequent empirical literature was built.

The empirical program of identifying observable risk factors accelerated rapidly after the early CAPM anomalies. Fama and French (1992, 1993) documented that a three-factor model incorporating market, size, and value factors explains a substantial share of the cross-section of average stock returns. Jegadeesh and Titman (1993) uncovered momentum as a powerful return predictor, subsequently incorporated as a fourth factor by Carhart (1997). This initial set of factors was soon joined by many others. Sloan (1996) identified the accruals anomaly. Novy-Marx (2013) documented a gross profitability premium. Lakonishok et al. (1994) linked value effects to extrapolation and contrarian investment strategies.

By the mid-2000s, the rate of factor discovery had increased markedly. Cochrane (2011), in his Presidential Address, characterized the state of the literature as a “zoo of new factors” and called for organizing principles. Efforts to consolidate the zoo into parsimonious multifactor models followed. Fama and French (2015) extended the three-factor model to five factors by adding profitability and investment. Hou et al. (2014) proposed a four-factor q -model grounded in investment theory, arguing that investment and profitability factors subsume many anomalies. Stambaugh and Yuan (2017) developed mispricing factors based on a composite of anomaly clusters. Daniel et al. (2020) constructed behavioral factors cap-

turing short- and long-horizon mispricing. Barillas and Shanken (2018) proposed a Bayesian model comparison framework for evaluating competing factor specifications, finding that the choice of factors matters more than the choice of test assets.

Despite these organizational efforts, the proliferation continued. Harvey et al. (2015) documented over 300 factors published in top journals and argued that standard significance thresholds are inadequate in the face of extensive multiple testing, proposing that a t -statistic exceeding 3.0 is a more appropriate hurdle. McLean and Pontiff (2015) provided out-of-sample evidence that many anomalies decay significantly post-publication, consistent with data mining or with investors learning to exploit and thereby erode anomalous returns. Chen and Zimmermann (2022) constructed a comprehensive open-source replication of the factor zoo, documenting that roughly half of published anomalies fail to replicate under uniform methodology. The diagnosis is now widely shared: the cross-section of expected returns is high-dimensional, many proposed factors are redundant or spurious, and the discipline requires principled selection methods to identify the factors that genuinely drive the SDF.

One strand of the literature seeks to extract the pricing-relevant factors from the data without imposing the identity of the factors *a priori*. Connor and Korajczyk (1986, 1988) pioneered the use of asymptotic principal component analysis (PCA) to estimate latent factor realizations from large cross-sections of asset returns. PCA provides consistent estimates of the factor space as both the cross-sectional and time-series dimensions grow, under the approximate factor model assumptions of Chamberlain and Rothschild (1983). Bai and Ng (2002) developed information criteria for consistently determining the number of latent factors in large panels, introducing penalized criteria that balance fit and parsimony without imposing strong restrictions on the relative magnitudes of N and T .

Standard PCA targets the covariance structure of returns rather than the pricing kernel, and the two objectives can diverge when the factors carrying risk premia are not the factors with the largest variance. Lettau and Pelger (2020) address this gap with Risk-Premium PCA (RP-PCA), a generalization that incorporates a penalty for the first moment of returns, ensuring that the extracted components explain both the cross-section of expected returns and the time-series variation. Kelly et al. (2019) develop Instrumented PCA (IPCA), which

models factor loadings as functions of observable characteristics, allowing the latent factor structure to vary with firm-level attributes. Lettau (2023) examines the factor zoo through the lens of high-dimensional factor models, emphasizing the need to separate genuine risk premia from idiosyncratic noise in settings where the number of candidate factors is large relative to the time series.

A parallel literature imposes structure directly on the SDF rather than extracting latent factors. Kozak et al. (2020) formulate the SDF as a linear combination of observable factors and show that a ridge-regression prior, which shrinks the SDF toward the directions of maximum return variance, produces out-of-sample pricing gains. Their approach is grounded in the insight that, under an approximate factor structure, the SDF weights must align with the principal components of the return covariance matrix, and factors explaining little variance should receive small weights. Feng et al. (2020) develop a double-selection LASSO procedure that evaluates the marginal contribution of new factors while controlling for the factors already in the model, providing a testing framework for individual factor significance.

The application of flexible nonparametric methods to asset pricing has expanded rapidly. Gu et al. (2020) conduct a comprehensive comparison of machine-learning techniques including: neural networks, random forests, gradient-boosted trees, and elastic net, for predicting the cross-section of returns. They find that nonlinear methods achieve substantial improvements in out-of-sample R^2 relative to linear benchmarks, suggesting that the true SDF may involve complex interactions among characteristics. Gu et al. (2021) extend this approach to autoencoder-based latent factor models that simultaneously learn both the factors and their loadings from large characteristic-sorted portfolios. These methods are powerful but largely atheoretical: they optimize prediction accuracy without imposing the economic structure. This paper takes a different approach, exploiting the geometric structure of the exponential-affine SDF family to provide both economic foundations and a tractable selection criterion.

The econometric evaluation of SDF models rests on the framework developed by Hansen (1982), who introduced the Generalized Method of Moments (GMM) and established the conditions under which moment restrictions implied by asset pricing models can be consistently estimated and tested. Hansen and Jagannathan (1997) proposed the Hansen-Jagannathan

(HJ) distance as a measure of model misspecification, defined as the minimum Euclidean distance between a candidate SDF and the set of admissible pricing kernels in L^2 . The HJ distance is closely linked to the Sharpe ratio of the maximum-correlation portfolio and to the GRS test of Gibbons et al. (1989), which evaluates whether a set of test assets has jointly zero alphas relative to a benchmark model.

The HJ distance operates in the space of square-integrable random variables and inherits Euclidean geometry: it treats positive and negative pricing errors symmetrically and imposes no constraint on the sign of the SDF. In practice, the HJ-optimal SDF is frequently negative in extreme states, precisely the crash states where the true pricing kernel should be largest. This limitation motivates the information-theoretic alternative developed in this paper, which replaces the Euclidean distance with the Kullback-Leibler (KL) divergence and thereby enforces SDF positivity by construction. As I show in Section 1.3, the HJ distance is recovered as the leading term of a Taylor expansion of the KL divergence, so the two criteria agree locally but diverge precisely in the tails where positivity and crash sensitivity matter most.

The application of entropy and information theory to asset pricing has a substantial history. Stutzer (1995) proposed a variational characterization of the set of state-price densities based on minimizing the Kullback-Leibler divergence, deriving an information-theoretic bound on the SDF that automatically satisfies a positivity constraint absent from the Hansen and Jagannathan (1997) variance bound. Kitamura and Stutzer (1997) developed the formal econometric theory, showing that the Minimum Cross-Entropy (MCE) estimator—which minimizes the KL divergence between the model-implied and empirical probability measures subject to moment constraints—is asymptotically equivalent to two-step efficient GMM. Kitamura and Stutzer (2002) further explored the connections between entropic and linear projections, clarifying when and how the two estimation paradigms diverge.

Ghosh et al. (2016) applied the information-theoretic framework to the consumption CAPM, nonparametrically estimating the SDF via relative entropy minimization and deriving entropy bounds that reveal the failure of standard consumption-based models to capture crash risk. Their work demonstrates that the gap between the entropy-implied SDF and the

consumption-based SDF is concentrated in states of market stress, a finding that motivates the crash sensitivity built into the KL-based selection criterion I develop here. Almeida and Garcia (2012) and Backus et al. (2014) study entropy bounds on the SDF in the context of option pricing and the term structure, showing that the entropy of the pricing kernel is linked to the variance risk premium and the shape of the volatility smile.

This paper extends this literature in two directions. First, I move from estimation of a single SDF to *selection* among candidate factors, using the KL divergence not as a fitting criterion but as a ranking criterion for the marginal contribution of each factor to the pricing kernel. Second, I estimate the full nested exponential-affine SDF under pricing restrictions via a two-stage MCE procedure. To my knowledge, prior implementations of the Kitamura and Stutzer (1997) estimator in asset pricing have solved only a portion of the problem for fixed risk prices or imposed additional restrictions that reduce the parameter space. In contrast, the implementation developed in Section 1.4 solves the full MCE optimization subject to the pricing restrictions.

The geometric structure exploited in this paper originates in the field of information geometry, which studies the differential geometry of probability distributions. Efron (1975) introduced statistical curvature and showed that the geometric properties of parametric families govern the second-order efficiency of estimators. ichi Amari (2020) developed the comprehensive framework of dually flat manifolds that describes the space of admissible pricing kernels. The canonical divergence on such manifolds is the Bregman divergence induced by the cumulant generating function Bregman (1967), which for exponential families coincides with the KL divergence. The generalized Pythagorean theorem for dually flat manifolds, a central result in ichi Amari (2020), guarantees an exact orthogonal decomposition of the divergence. Nielsen (2024) provides a modern treatment of divergences induced by the cumulant and partition functions of exponential families.

The economic grounding in this paper draw on the robust control literature initiated by Hansen and Sargent (2001) and developed extensively in Anderson et al. (2003) and Hansen and Sargent (2020). Maccheroni et al. (2006) provided the axiomatic foundation for this class of preferences, showing that multiplier preferences are a special case of the varia-

tional preferences that generalize the max-min expected utility of Gilboa and Schmeidler (1989). Maenhout (2004) applied robust portfolio rules to asset pricing, demonstrating that robustness decreases equity demand and is observationally equivalent to increased effective risk aversion under recursive preferences.

The forward-selection procedure I employ belongs to the broader class of greedy approximation algorithms, which build solutions incrementally by making locally optimal choices at each step. Greedy algorithms have a long history in approximation theory and signal processing. Mallat and Zhang (1993) introduced Matching Pursuit, a greedy algorithm for decomposing signals into sparse linear combinations of dictionary elements. Pati et al. (1993) developed Orthogonal Matching Pursuit (OMP), which re-orthogonalizes at each step to ensure that previously selected elements do not contaminate the greedy criterion. Natarajan (1995) established that the general sparse approximation problem is NP-hard, motivating the use of greedy heuristics as computationally tractable alternatives to exhaustive search. Davis et al. (1997) provided early convergence guarantees for adaptive greedy algorithms under restricted isometry conditions, and Tropp (2004) established that greedy algorithms recover the optimal sparse solution.

1.3 Theory

The central question of empirical asset pricing is which risks the market prices, and how severely. I approach this question geometrically. The search for the true stochastic discount factor is equivalent to a search through a space of probability measures, and that space has a structure defined by the distribution of asset returns. My goal in this section is to make that structure precise, exploit it, and show that it delivers both a tractable estimation algorithm and a complete economic theory of why investors hold the factors they do, and no more. The factor zoo is not tamed by an arbitrary cutoff; it is tamed by a criterion that a rational agent under uncertainty would independently arrive at.

The section proceeds in two parts. Section 1.3.1 develops the statistical geometry of the SDF manifold. I define the manifold of admissible pricing kernels, equip it with a divergence measure, derive the projection geometry that drives the selection algorithm, and extract a

statistical stopping rule that determines when a candidate factor's risk price is detectable. This machinery delivers a complete algorithm and a significance filter, but it does not provide an economic criterion for when residual mispricing is tolerable. Section 1.3.2 supplies that criterion. I introduce a representative investor with multiplier preferences in the sense of Hansen and Sargent (2001, 2020), derive the robust SDF and the misspecification premium that governs pricing anomalies, and show that the investor's preference-based ranking of candidate factors by welfare loss is identical to the purely statistical ranking by KL divergence. The investor's welfare analysis produces an economic stopping rule that determines when the pricing kernel should cease expanding. The section closes by observing that the geometric algorithm and the investor's equilibrium selection describe the same phenomenon from two vantage points.

1.3.1 Statistical Geometry

The search for the true stochastic discount factor can be cast as a projection of the true data-generating process onto a statistical manifold of admissible pricing models. By embedding asset pricing into the framework of information geometry, the heuristic practice of hunting for factors is replaced by a principled traversal of a statistical manifold. This subsection formalizes the space of admissible SDFs, derives the projection geometry that drives the selection algorithm, and establishes a statistical stopping rule for the greedy procedure.

Let $(\Omega, \mathfrak{F}, \mathcal{P})$ be a probability space representing the physical economy, supporting a vector of asset excess returns $\mathbf{r}_t^e \in \mathbb{R}^N$ and a vector of candidate risk factors $\mathbf{f}_t \in \mathbb{R}^K$. Standard asset pricing theory posits the existence of a strictly positive stochastic discount factor m_t such that

$$\mathbf{E}_{\mathcal{P}} [m_t \mathbf{r}_t^e] = \mathbf{0}_N. \quad (1.1)$$

The family of admissible pricing kernels must satisfy several requirements simultaneously. The SDF must be strictly positive everywhere, both to enforce no-arbitrage and to ensure that it defines a valid probability measure via its Radon-Nikodym interpretation. It must be closed under projection, so that restricting the factor set produces a valid SDF within

the same family. It must admit a canonical divergence, enabling a well-defined measure of distance between nested models. And it must nest cleanly, so that an augmented model embeds the restricted model as a special case. These requirements are not independent. The unique family of strictly positive distributions satisfying all four is the exponential family.

I restrict attention to the family of exponential-affine pricing kernels

$$m_t(\boldsymbol{\theta}) := \exp \left(-\boldsymbol{\theta}^\top \mathbf{f}_t - \psi(\boldsymbol{\theta}) \right), \quad (1.2)$$

where $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^K$ is the vector of factor risk prices and $\psi(\boldsymbol{\theta})$ is the cumulant generating function

$$\psi(\boldsymbol{\theta}) := \ln \left(\mathbf{E}_{\mathcal{P}} \left[\exp \left(-\boldsymbol{\theta}^\top \mathbf{f}_t \right) \right] \right). \quad (1.3)$$

The normalization in (1.3) ensures that $\mathbf{E}_{\mathcal{P}} [m_t(\boldsymbol{\theta})] = 1$, so that $m_t(\boldsymbol{\theta})$ can be interpreted as the Radon-Nikodym derivative inducing a risk-neutral probability measure $\mathcal{Q}(\boldsymbol{\theta})$. The construction guarantees $m_t(\boldsymbol{\theta}) > 0$ for all valid parameters, placing every candidate model within the open simplex of probability distributions where information geometry is valid. Economically, this positivity enforces the no-arbitrage condition by construction, precluding the negative state prices that plague linear factor models during market stress.

Assumption 1 (Regularity). *The joint distribution of $(\mathbf{f}, \mathbf{r}^e)$ under the physical measure $\mathcal{P}$ has a well-defined moment generating function for $\mathbf{f}$ in a neighborhood of the origin, and the support of $\mathbf{f}$ under $\mathcal{P}$ is not contained in any proper affine subspace of $\mathbb{R}^K$.*

Under the finite-moment and full-rank conditions of Assumption 1, the set

$$\mathbb{M} = \{ \mathcal{Q}(\boldsymbol{\theta}) : \boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^K \}$$

forms a K -dimensional statistical manifold, which I term the *SDF manifold*. In the language of information geometry, $\mathbb{M}$ constitutes an exponential family of distributions. This structure is not merely a statistical convenience; it endows the space of pricing models with a rigorous form of geometric flatness.

A statistical manifold is flat if there exists a coordinate system in which geodesics, the shortest paths between distributions, are straight lines. The manifold $\mathbb{M}$ is specifically *e-flat* (*exponentially flat*) because its geodesics are linear in the natural parameters $\boldsymbol{\theta}$ (Ichi Amari, 2020). This linearity in log-probability space is the property that enables orthogonal decomposition of pricing errors, analogous to projecting vectors in Euclidean space, but on a surface whose geometry is dictated by the distribution of asset returns.

To quantify the distance between the true pricing kernel and a candidate model, I employ the Kullback-Leibler (KL) divergence, also known as relative entropy. This measure generalizes Shannon entropy Shannon (1948) by measuring the expected informational loss incurred when the unobservable true risk-neutral measure $\mathcal{Q}^\dagger$ is approximated by a candidate model $\mathcal{Q}(\boldsymbol{\theta})$.¹

Let $\mathcal{Q}^\dagger$ denote the true, unobservable risk-neutral measure that correctly prices all assets. The divergence from the truth to a candidate model $\mathcal{Q}(\boldsymbol{\theta}) \in \mathbb{M}$ is

$$D_{KL}(\mathcal{Q}^\dagger || \mathcal{Q}(\boldsymbol{\theta})) := \int_{\Omega} \ln \frac{d\mathcal{Q}^\dagger}{d\mathcal{Q}(\boldsymbol{\theta})} d\mathcal{Q}^\dagger = \mathbf{E}_{\mathcal{Q}^\dagger} \left[\ln \left(\frac{d\mathcal{Q}^\dagger}{d\mathcal{Q}(\boldsymbol{\theta})} \right) \right]. \quad (1.4)$$

Minimizing the KL divergence minimizes the expected informational loss incurred when approximating the true pricing distribution $\mathcal{Q}^\dagger$ with the model $\mathcal{Q}(\boldsymbol{\theta})$. Although statistically equivalent to maximizing the expected log-likelihood, this criterion serves a distinct economic purpose: it selects the pricing kernel that least distorts state probabilities relative to investors' true marginal utilities. The estimator naturally aligns risk premia with observed returns while penalizing unnecessary complexity in the factor representation.

The KL divergence is not a true metric in the topological sense, satisfying neither symmetry nor the triangle inequality. It is, however, the canonical divergence for exponential families. This status is not accidental. The variational preference axioms of Maccheroni et al. (2006) establish that KL divergence is the unique measure of model distance, up to affine transformations, consistent with coherent decision-making under ambiguity. The

¹Shannon entropy quantifies the absolute uncertainty of a single probability distribution. The KL divergence generalizes this to the relative uncertainty between two distributions, measuring the excess information required to describe the true distribution using an approximating model. The KL divergence is foundational across disciplines, from information theory and coding theory to machine learning and statistical physics.

exponential-entropy structure is dictated by axioms, not assumed for tractability.

The geometry of the SDF manifold $\mathbb{M}$ is governed by a dual structure that is fundamental to the estimation strategy. The manifold $\mathbb{M}$ is e -flat under the exponential connection, while the space of mixture distributions is m -flat under the mixture connection. A central result in information geometry (ichi Amari, 2020) establishes that this dual flatness guarantees the projection of the true distribution onto the model manifold is unique and satisfies a generalized Pythagorean theorem.

The econometric goal is to identify the model $\mathcal{Q}^{\mathbb{S}}$ within a restricted factor span $\mathbb{S} \subseteq \mathbb{F}$ that best approximates the true pricing kernel $\mathcal{Q}^{\dagger}$. This optimization defines $\mathcal{Q}^{\mathbb{S}}$ as the m -projection of the true measure $\mathcal{Q}^{\dagger}$ onto the e -flat submanifold $\mathbb{M}^{\mathbb{S}}$.

Consider a nested factor structure where we augment a base set of factors $\mathbb{S}$ with a candidate factor $\mathbf{g}$, so that $\mathbb{M}^{\mathbb{S}} \subset \mathbb{M}^{(\mathbb{S}, \mathbf{g})}$. Since both the restricted measure $\mathcal{Q}^{\mathbb{S}}$ and the augmented measure $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ are absolutely continuous with respect to $\mathcal{P}$, the Radon-Nikodym derivative of the augmented measure with respect to the restricted measure exists and equals the ratio of their pricing kernels. Because $\mathbb{M}^{\mathbb{S}}$ is a submanifold of $\mathbb{M}^{(\mathbb{S}, \mathbf{g})}$ and both are e -flat, the generalized Pythagorean theorem for dually flat manifolds (ichi Amari, 2020) applies. For the chain of distributions $\mathcal{Q}^{\dagger}$ (true), $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ (augmented model), and $\mathcal{Q}^{\mathbb{S}}$ (restricted model), the divergence decomposes exactly as

$$D_{KL}(\mathcal{Q}^{\dagger} || \mathcal{Q}^{\mathbb{S}}) = D_{KL}(\mathcal{Q}^{\dagger} || \mathcal{Q}^{(\mathbb{S}, \mathbf{g})}) + D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}). \quad (1.5)$$

This identity is the theoretical cornerstone of the approach. It admits a transparent geometric interpretation. The restricted model $\mathcal{Q}^{\mathbb{S}}$ is the projection of the true measure $\mathcal{Q}^{\dagger}$ onto the submanifold $\mathbb{M}^{\mathbb{S}}$. When the submanifold is expanded by one dimension through the addition of factor $\mathbf{g}$, the new projection $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ lies closer to the truth. The three points, $\mathcal{Q}^{\dagger}$, $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$, and $\mathcal{Q}^{\mathbb{S}}$, form a right triangle on the manifold in the information-geometric sense, with the two legs orthogonal. The decomposition is exact due to the e -flat structure of $\mathbb{M}$.

The economic content is equally direct. The left-hand side is the total mispricing under the restricted model, fixed once $\mathbb{S}$ is chosen. The first term on the right is the residual mispricing after adding $\mathbf{g}$, the pricing error that remains unexplained. The second term is

the information gain from $\mathbf{g}$, the mispricing that the new factor resolves. Since the total mispricing is fixed *a priori*, minimizing the residual mispricing is mathematically equivalent to maximizing the information gain

$$\arg \min_{\mathbf{g} \in \mathbb{F} \setminus \mathbb{S}} D_{KL}(\mathcal{Q}^\dagger \| \mathcal{Q}^{(\mathbb{S}, \mathbf{g})}) \equiv \arg \max_{\mathbf{g} \in \mathbb{F} \setminus \mathbb{S}} D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^\mathbb{S}). \quad (1.6)$$

This equivalence motivates greedy forward selection: at each step, add the factor that maximizes the divergence between the augmented and baseline models. The equivalence in (1.6) holds at each selection step for the current restricted model $\mathcal{Q}^\mathbb{S}$. Across steps, as $\mathbb{S}$ expands and the model is re-estimated, the total error $D_{KL}(\mathcal{Q}^\dagger \| \mathcal{Q}^\mathbb{S})$ itself changes, so the decomposition does not telescope into a cumulative sum.

On the e -flat manifold, the KL divergence between two exponential-affine SDFs admits a representation as the Bregman divergence generated by the cumulant generating function $\psi(\boldsymbol{\theta})$

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^\mathbb{S}) = \psi(\boldsymbol{\theta}^\mathbb{S}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) + \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})^\top (\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^\mathbb{S}). \quad (1.7)$$

The Bregman form makes the pricing content of the divergence transparent. The first two terms, $\psi(\boldsymbol{\theta}^\mathbb{S}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})$, measure the change in the log-normalizing constant when the new factor is added to the SDF. Since $\psi(\boldsymbol{\theta})$ is precisely the adjustment that ensures the pricing kernel integrates to unity, this difference reflects how much the model must redistribute probability mass across states of the world to maintain no-arbitrage once the new risk dimension is priced. A large change indicates that the new factor substantially reshapes the state-price density; a small change indicates that the existing factors already account for most of the relevant variation in state prices. The third term is an inner product between the gradient $\nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})$, which is the vector of expected factor payoffs under the augmented model's risk-neutral measure, and the parameter displacement $\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^\mathbb{S}$. Together, these terms decompose the information gain into a normalization effect and a repricing effect. A large divergence occurs when the inclusion of factor $\mathbf{g}$ forces a substantial revision of the state-price density, either by shifting the level of the normalizing constant or by altering

the risk-neutral expectations of factor payoffs. If $\mathbf{g}$ is redundant with the factors already in $\mathbb{S}$, the augmented parameters will differ negligibly from $\boldsymbol{\theta}^{\mathbb{S}}$, the normalization will barely change, and the divergence will be near zero.

The gradient of the Bregman divergence links this representation directly to pricing errors. For a candidate factor $\mathbf{g}$ outside the current span $\mathbb{S}$, the derivative of the KL divergence with respect to its loading $\theta_{\mathbf{g}}$ is proportional to the factor's alpha under $\mathcal{P}$:

$$\frac{\partial}{\partial \theta_{\mathbf{g}}} D_{KL} \propto \mathbf{E}_{\mathcal{P}} [m_t(\boldsymbol{\theta}^{\mathbb{S}}) g_t] \neq 0. \quad (1.8)$$

A nonzero alpha is a displacement in the tangent space of the manifold at $\mathcal{Q}^{\mathbb{S}}$, and its magnitude is measured by the KL divergence in that direction. The Bregman representation therefore provides a direct link between the statistical criterion used for factor selection and the economic magnitude of the pricing error that the omitted factor generates. Maximizing the divergence is equivalent to searching the zoo for the source of risk that carries the largest unexplained risk premium in information units.

The local geometry of $\mathbb{M}$ provides the connection to the Hansen and Jagannathan (1997) (HJ) distance that is prevalent in empirical finance. For infinitesimally close models parameterized by $\boldsymbol{\theta}$ and $\boldsymbol{\theta} + d\boldsymbol{\theta}$, the KL divergence is locally quadratic and governed by the Fisher information metric

$$D_{KL}(\mathcal{Q}(\boldsymbol{\theta}) || \mathcal{Q}(\boldsymbol{\theta} + d\boldsymbol{\theta})) \approx \frac{1}{2} d\boldsymbol{\theta}^{\top} I(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (1.9)$$

where $I(\boldsymbol{\theta}) = [i_{jk}(\boldsymbol{\theta})]$ is the Fisher information matrix, given by the Hessian of the cumulant generating function

$$i_{jk}(\boldsymbol{\theta}) = \frac{\partial^2 \psi}{\partial \theta_j \partial \theta_k}(\boldsymbol{\theta}) = \text{Cov}_{\mathcal{Q}(\boldsymbol{\theta})}(f_{jt}, f_{kt}).$$

This metric acts as a volatility-adjusted ruler on the manifold. A displacement in parameter space represents a meaningful change in the pricing model only when it occurs in a direction of high factor variance under the risk-neutral measure. The intrinsic nature of the metric ensures that the selection criterion is invariant to the rescaling or rotation of the factor space.

The HJ distance D_{HJ} measures the minimum Euclidean distance between a candidate SDF and the set of admissible pricing kernels in the space of square-integrable random variables. It is a second-moment metric, fundamentally linked to mean-variance efficiency and linear beta pricing, and it inherits the geometry of L^2 . Deviations from the admissible set are measured symmetrically: positive and negative pricing errors receive equal weight, and the strict positivity of the SDF plays no formal role in the distance calculation. This last point is consequential. The HJ projection frequently implies an SDF that is negative in extreme states, precisely the crash states where the true SDF should be largest, because the Euclidean geometry has no mechanism to penalize violations of positivity.

The information-geometric approach subsumes the HJ framework as a special case. A Taylor expansion of D_{KL} around the null hypothesis $\theta_{\mathbf{g}} = 0$ yields

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}}) = \frac{1}{2} D_{HJ}^2(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}}) + \mathcal{O}_{\mathcal{P}} \left(\|\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}\|_2^3 \right). \quad (1.10)$$

For small deviations from the baseline model, or equivalently in the Gaussian limit where all cumulants of order three and above vanish, maximizing the KL divergence is identical to maximizing the reduction in HJ distance. The two criteria agree precisely where the linear, quadratic world is an adequate approximation to the true return distribution. The HJ distance is therefore not superseded by the geometric framework; it is recovered from it as the leading term of a curvature expansion on $\mathbb{M}$.

The KL divergence departs from HJ through the higher-order terms, and it does so in a directionally asymmetric way that is economically meaningful. The divergence penalizes models that assign insufficient probability mass to states where the true SDF is high, which are generically the crash states where marginal utility is elevated. A candidate factor that hedges downside tail risk generates a strictly larger KL divergence than a factor with symmetric volatility and an identical unconditional Sharpe ratio. The geometric selection criterion is therefore inherently crash-sensitive, prioritizing factors that capture the wedge between the physical and risk-neutral measures in the left tail even when standard mean-variance analysis would treat them as equivalent to their symmetric counterparts. These higher-order terms encode the skewness, kurtosis, and co-kurtosis of the return distribution, the precise features

that drive the premium demanded for bearing asymmetric downside exposure and fat-tailed losses in states of high marginal utility.

The geometric machinery delivers a complete ranking algorithm: project, measure divergence, add the factor that moves furthest along the manifold. It does not, however, say when to stop. Two questions must be addressed. First, *can we detect a nonzero risk price at all?* Second, *is the resulting mispricing economically meaningful?* This subsection addresses the first question; Section 1.3.2 addresses the second.

The natural approach to determining whether adding $\mathbf{g}$ significantly changes the model is to compare the nested models via their KL divergences. However, as Vuong (1989) established, nested model comparison via KL divergence cannot distinguish the models when the restricted model lies on the boundary of the augmented model's parameter space, as it does here when $\theta_{\mathbf{g}} = 0$. The appropriate test statistic is instead the Wald statistic on the estimated risk price $\hat{\theta}_{\mathbf{g}}$.

Proposition 1 (Statistical Stopping Rule). *Under the exponential-affine SDF with factor $\mathbf{g}^\perp$ as the orthogonal component of $\mathbf{g}$ with respect to $\mathbb{S}$, the Wald statistic $W_T := \hat{\theta}_{\mathbf{g}}^2 / \hat{\sigma}_{\hat{\theta}_{\mathbf{g}}}^2$ satisfies:*

1. Under $H_0 : \theta_{\mathbf{g}} = 0$, $W_T \rightarrow \chi_1^2$.
2. Under $H_1 : \theta_{\mathbf{g}} \neq 0$, $W_T \rightarrow \chi_1^2(\delta_{\mathbf{g}})$, with noncentrality

$$\delta_{\mathbf{g}} = \frac{T(\theta_{\mathbf{g}}^*)^2}{\sigma_{\hat{\theta}_{\mathbf{g}}, \infty}^2}$$

where $\sigma_{\hat{\theta}_{\mathbf{g}}, \infty}^2$ denotes the long-run asymptotic variance of $\sqrt{T}\hat{\theta}_{\mathbf{g}}$.

3. To second order in $\theta_{\mathbf{g}}^*$, the noncentrality parameter coincides with twice the KL divergence between the augmented and restricted models scaled by the long run variance,

$$\delta_{\mathbf{g}} = \frac{T}{\Omega_{\mathbf{gg}}} \frac{2D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})}{(\nabla_{\theta_{\mathbf{g}}} \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}))^2}.$$

Proposition 1 admits a dual interpretation. Statistically, the Wald test asks whether the estimated risk price $\hat{\theta}_{\mathbf{g}}$ is large relative to its standard error. Economically, the noncentrality $\delta_{\mathbf{g}}$ links the test's power directly to the KL divergence. A factor with a large KL divergence but a poorly estimated risk price, perhaps because of heavy tails or serial dependence in its returns, can produce a small $\delta_{\mathbf{g}}$ and fail the test despite its economic relevance.

The Wald test creates a feasible set of factors

$$\mathbb{J} := \{\mathbf{g} : W_T(\mathbf{g}) > \chi_1^2(1 - \alpha)\}$$

for which a true pricing signal is detected at α significance level. Within this set, factors can be ranked by their KL divergence D_{KL} . What the test does not provide is an economic criterion for when a factor's omission creates tolerable mispricing. A factor might pass the Wald test with a highly significant but economically negligible risk price. This gap, the absence of an economic threshold for acceptable mispricing, motivates the investor's problem in the following subsection.

The geometric machinery of this subsection delivers a selection algorithm and a statistical significance filter. The Pythagorean decomposition provides an exact, step-by-step criterion for ranking candidate factors. The KL divergence supplies a canonical objective directly related to model implied pricing errors. The Taylor expansion around the HJ distance clarifies precisely where the geometric approach departs from conventional methods and what it adds: crash sensitivity, positivity enforcement, and higher-order risk premia. The Wald test provides a statistical filter. What remains missing is an economic criterion for when residual mispricing is tolerable. For that, I need an investor.

1.3.2 Investor's Portfolio Choice Problem

The purpose of this subsection is twofold. First, I provide economic foundations for the geometric procedure developed above, showing that a rational investor facing model uncertainty would follow precisely the ranking criterion that the manifold geometry dictates. Second, I derive an economic bound on tolerable mispricing that, combined with the statistical stopping rule of Proposition 1, determines the equilibrium composition of the pricing kernel.

The economic model is not novel; I adopt the multiplier preferences of Hansen and Sargent (2001, 2020). What is new is the connection to the manifold geometry, the application to the factor zoo, and the resulting characterization of the stopping threshold.

Consider a representative investor who lives for one period, is endowed with initial wealth W_0 , and must allocate across N assets with excess returns $\mathbf{r}^e \in \mathbb{R}^N$. The economy is driven by K candidate risk factors $\mathbf{f} \in \mathbb{R}^K$ that jointly determine the distribution of returns. The investor holds a reference probability model $\mathcal{P}$ that they believe is approximately correct, but they acknowledge that $\mathcal{P}$ may be misspecified. To formalize this concern, the investor evaluates their portfolio under a family of distorted probability measures, guarding against the worst-case alternative. Following Hansen and Sargent (2001), the investor has multiplier preferences

$$V(\mathbf{w}) := \min_{h \geq 0, \mathbf{E}_{\mathcal{P}}[h]=1} \{ \mathbf{E}_{\mathcal{P}}[hu(W_1)] + \vartheta \mathbf{E}_{\mathcal{P}}[h \ln(h)] \} \quad (1.11)$$

subject to the budget constraint $W_1 = W_0(R_f + \mathbf{w}^\top \mathbf{r}^e)$, where $u : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing and concave, $\vartheta \in \mathbb{R}_{++}$ is the *welfare sensitivity to misspecification*, and R_f is the gross risk-free rate.

The distortion h is a Radon-Nikodym derivative satisfying $h \geq 0$ and $\mathbf{E}_{\mathcal{P}}[h] = 1$, so that the reweighted measure $\mathcal{Q} = h\mathcal{P}$ is a valid probability distribution. When $h(\omega) > 1$, the distortion assigns state ω more probability mass than the reference model suggests; when $h(\omega) < 1$, the state is downweighted. The distortion represents the “adversary’s” hypothesis about how the world differs from the investor’s baseline.

The first term in (1.11) evaluates expected utility under the distorted measure $h d\mathcal{P}$. The second term penalizes deviations from the reference model by the relative entropy $D_{KL}(h\mathcal{P}||\mathcal{P}) = \mathbf{E}_{\mathcal{P}}[h \ln(h)]$, which measures the statistical cost of the distortion. The parameter ϑ governs the tradeoff between severity and plausibility. A large ϑ makes large distortions prohibitively expensive, confining the adversary to a small entropy neighborhood of the reference; the investor is nearly confident in the baseline. A small ϑ imposes little cost on large distortions, allowing the adversary to consider distortions far from the reference; the investor is highly ambiguity-averse. In the limit $\vartheta \rightarrow \infty$, the penalty dominates, forcing

$h \equiv 1$, and the investor reduces to a standard expected-utility maximizer.²

In the context of the factor zoo $\mathbb{F}$, the investor anchors to a baseline pricing kernel that loads on a subset $\mathbb{S} \subseteq \mathbb{F}$ of established factors, with the baseline SDF given by (1.2) and the induced risk-neutral measure $\mathcal{Q}^{\mathbb{S}}$. The investor has confidence in $\mathcal{Q}^{\mathbb{S}}$; it embodies the factors that have survived rigorous empirical testing. However, the investor acknowledges that the factor zoo contains $K - |\mathbb{S}|$ additional candidates, each representing a direction along which the baseline might be misspecified. This is structured ambiguity: the set of plausible alternative models is parameterized by the factor zoo, and the investor's ambiguity set is the SDF manifold $\mathbb{M}$.

The investor's full problem is

$$\max_{\mathbf{w} \in \mathbb{R}^N} V(\mathbf{w}) \equiv \max_{\mathbf{w} \in \mathbb{R}^N} \min_{h \geq 0, \mathbf{E}_{\mathcal{P}}[h]=1} \{ \mathbf{E}_{\mathcal{P}}[hu(W_1)] + \vartheta \mathbf{E}_{\mathcal{P}}[h \ln(h)] \}. \quad (1.12)$$

The inner problem in (1.12) is a constrained minimization of a functional that is linear in h plus a strictly convex entropy penalty. The solution is the standard Gibbs variational result.

Proposition 2 (Worst-Case Distortion). *For any fixed portfolio $\mathbf{w}$, the worst-case density distortion that solves the inner problem in (1.12) is*

$$h^*(\mathbf{w}) = \frac{\exp(-\vartheta^{-1}u(W_1))}{\mathbf{E}_{\mathcal{P}}[\exp(-\vartheta^{-1}u(W_1))]} \quad (1.13)$$

The worst-case measure $\mathcal{Q}_{\mathbf{w}}^ := h^*(\mathbf{w})d\mathcal{P}$ is an exponential tilt of the reference model in the direction of low utility.*

The worst-case model is not selected by the econometrician; it emerges endogenously from the investor's preferences. The exponential tilt in (1.13) upweights states where utility is low and downweights states where utility is high. States in which the investor's portfolio suffers

²Multiplier preferences admit an axiomatic foundation through the variational preferences of Maccheroni et al. (2006), which generalize the max-min expected utility of Gilboa and Schmeidler (1989). The entropy penalty $\vartheta D_{KL}(\mathcal{Q}||\mathcal{P})$ is the unique penalty, up to affine transformations, that satisfies translation invariance and is generated by a convex, lower-semicontinuous function on the space of probability measures. The KL divergence therefore emerges as the unique model distance consistent with the preference axioms.

severe losses, the crash states where marginal utility is elevated, receive the most probability mass under the distortion. The severity of the tilt is governed by ϑ : investors with smaller welfare sensitivity to misspecification produce more extreme distortions, concentrating the worst-case measure more sharply on adverse outcomes.

Each omitted factor represents a direction along which the adversary can attack the baseline model. If a candidate factor $\mathbf{g} \notin \mathbb{S}$ carries a nonzero risk price $\theta_{\mathbf{g}}^* \neq 0$, the baseline model $\mathcal{Q}^{\mathbb{S}}$ misprices assets exposed to $\mathbf{g}$, generating an anomaly $\alpha_{\mathbf{g}}^{\mathbb{S}}$ in the pricing errors. The adversary exploits precisely these unhedged directions: the worst-case distortion h^* tilts probability mass toward states where the omitted factor takes extreme values, concentrating on the crash states that the baseline ignores. Including factor $\mathbf{g}$ in the SDF removes it from the adversary's toolkit by pricing the risk that was previously left as an anomaly. An investor with small ϑ faces a powerful adversary who can create severe distortions along every unhedged direction, making even modest unexplained premia costly. This investor hedges aggressively, including more factors to shrink the adversary's action set. An investor with large ϑ faces a constrained adversary who can inflict only mild distortions; unexplained premia of moderate size are tolerable, and the investor includes fewer factors.

The distortion h^* is central to the stopping rule derived below. Its entropy, $\eta := \mathbf{E}_{\mathcal{P}}[h^* \ln(h^*)]$, bounds the amount of mispricing that the investor can rationalize as misspecification rather than as evidence of a genuinely omitted factor.

The worst-case distortion $h^*(\mathbf{w})$ is an exponential function of $u(W_1)$, which in turn depends on the factor realizations through the portfolio return $(\mathbf{w}^*)^\top \mathbf{r}^e$. When utility is affine in the factors, as it is to first order for any smooth utility function in a neighborhood of the optimum, the distortion $h^*(\mathbf{w})$ is an exponential-affine function of $\mathbf{f}$. An exponential-affine function of $\mathbf{f}$ is precisely a point on the SDF manifold $\mathbb{M}$. The worst-case model therefore lives on the same manifold that the geometric algorithm traverses. This is the structural reason why the investor's problem and the manifold projection problem produce the same answer: the adversary's optimal distortion, derived from preference axioms, lands on the very surface that the information-geometric framework uses to organize pricing models.

Substituting the worst-case distortion back into the outer problem and taking first-order

conditions yields the Euler equation

$$\mathbf{E}_{\mathcal{P}} [h^*(\mathbf{w}^*)u'(W_1)\mathbf{r}^e] = \mathbf{0}_N,$$

from which the equilibrium SDF is identified as

$$m^{robust} \propto \underbrace{h^*(\mathbf{w}^*)}_{\text{misspecification}} \times \underbrace{m^{ccapm}}_{\text{consumption risk}} \quad (1.14)$$

where the proportionality constant ensures that m^{robust} correctly prices the risk-free asset, and $m^{ccapm} \propto u'(W_1)$ is the baseline CCAPM SDF. The equilibrium SDF is the product of two components. The marginal-utility component prices risk exactly as in the standard consumption CAPM. The ambiguity distortion h^* tilts the pricing kernel toward worst-case states, adding a premium for risks that the reference model may be ignoring. The standard SDF is nested as the special case $\vartheta \rightarrow \infty$. The model-implied pricing restriction is $\mathbf{E}_{\mathcal{P}} [m^{robust}\mathbf{r}^e] = \mathbf{0}_N$.

The decomposition of the SDF in (1.14) generates a corresponding decomposition of expected excess returns that separates compensation for known risk from compensation for model uncertainty.

Lemma 1 (Risk Premia Decomposition). *Under the SDF (1.14), the expected excess return on asset i decomposes as*

$$\mathbf{E}_{\mathcal{P}} [r_i^e] = - \left(\frac{\text{Cov}_{\mathcal{P}}(m^{ccapm}, r_i^e)}{\mathbf{E}_{\mathcal{P}} [m^{ccapm}]} + \frac{\text{Cov}_{\mathcal{P}}(h^*, m^{ccapm} r_i^e)}{\mathbf{E}_{\mathcal{P}} [m^{ccapm}]} \right). \quad (1.15)$$

The first term is the standard risk premium under the baseline model, compensation for covariance with consumption risk. The second term is the misspecification premium, the additional compensation that the investor demands for bearing exposure to the worst-case distortion h^* . The relationship between the misspecification premium and the baseline pricing error is made explicit by the identity

$$\mathbf{E}_{\mathcal{P}} [m^{ccapm} r_i^e] = - \text{Cov}_{\mathcal{P}}(h^*, m^{ccapm} r_i^e).$$

The left-hand side is the pricing error of the baseline, the component of the expected return that consumption risk fails to explain. The right-hand side states that this anomaly is governed by the covariance between the worst-case distortion and the SDF-weighted return, exactly as the standard risk premium is governed by the covariance between the SDF and excess returns. The distortion h^* therefore functions as a second pricing kernel, operating on the directions the baseline misses.

The sign of the misspecification premium has a transparent interpretation. When h^* and $m^{ccapm}r_i^e$ move in opposite directions, the asset delivers low SDF-weighted returns in the states that the distortion makes more probable. The alpha is positive, and investors demand compensation for bearing exposure to a risk the baseline ignores. When they comove, the asset provides insurance against model uncertainty and commands a lower return. Since h^* is an exponential tilt toward low-utility states, this covariance is largest for assets exposed to omitted factors that are active in the left tail of the return distribution. When no factors are omitted, $h^* \equiv 1$, the covariance vanishes, and the standard Euler equation is recovered.

The misspecification premium is the economic content of the nonzero tangent displacement identified in the geometry of Section 1.3.1. The displacement in the manifold's tangent space at $\mathcal{Q}^{\mathbb{S}}$, measured by the gradient $\partial_{\theta_g} D_{KL}$ in (1.8), is the investor's misspecification premium.

The misspecification premium establishes that each omitted factor generates a welfare cost for the investor. The natural question is how to rank candidate factors by the severity of this cost. The following result shows that the investor's preference-based ranking reduces to a purely statistical quantity.

Proposition 3 (Factor Ranking). *Under the multiplier preferences (1.11) and the exponential-affine SDF specification, the worst-case welfare loss from misspecification in the direction of factor $\mathbf{g}$, to leading order, satisfies*

$$|\Delta V_{\mathbf{g}}| \leq \sqrt{2\vartheta \mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(u(W_1^*))} \sqrt{D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})}. \quad (1.16)$$

The ranking of candidates by worst-case welfare loss is identical to the ranking by KL diver-

gence:

$$|\Delta V_{\mathbf{g}}| > |\Delta V_{\mathbf{g}'}| \iff D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) > D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g}')} || \mathcal{Q}^{\mathbb{S}}).$$

Proposition 3 is the result that connects the economic model to the statistical geometry. The investor solves a portfolio problem, evaluates the worst-case welfare loss from each direction of misspecification, and finds that the KL divergence between the augmented and restricted models is the quantity governing how much indirect utility can deteriorate under the worst-case perturbation. The KL divergence appears as a consequence of this economic analysis, not as an assumed objective.

The content of the proposition is that the preference-based ranking of factors by welfare loss is identical to the purely statistical ranking by divergence. This is not guaranteed *a priori*; it relies on the specific interaction between the exponential-tilt structure of the worst-case distortion and the e -flat geometry of the SDF manifold. The prefactor $\sqrt{2\vartheta \mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(u(W_1^*))}$ is common to all candidates. An investor with smaller ϑ suffers a smaller absolute welfare loss from each omission, but the relative ranking of candidates is invariant to ϑ . All investors with multiplier preferences, regardless of their individual sensitivity to misspecification, agree on the ordering.

The ranking in Proposition 3 tells the investor which factor to add next, but it does not tell them when to stop. Without a stopping criterion, the investor would include the entire factor zoo. The key observation is that not every anomaly reflects a genuinely omitted factor. The investor knows their pricing model is wrong, and some of the apparent mispricing can be attributed to misspecification rather than to a truly missing risk source.

The entropy of the worst-case distortion h^* provides the criterion. The investor holds the baseline $\mathcal{Q}^{\mathbb{S}}$ and considers extending to $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$. Both are members of the exponential family on $\mathbb{M}$. The marginal divergence $D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})$ measures how far the augmented model is from the baseline on $\mathbb{M}$. The investor's decision reduces to the following question: *Is this difference evidence of a genuine missing factor, or could it be an artifact of the adversary's distortion?*

If $D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) \leq \bar{\eta}$, where $\bar{\eta}$ is the adversary's entropy budget, then there exists a distortion h within the adversary's feasible set that could generate the same divergence. The

investor cannot rule out that the apparent improvement from adding $\mathbf{g}$ was manufactured by the adversary. Including $\mathbf{g}$ would attribute to systematic risk what could be misspecification, and there is no basis for preferring $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ over $\mathcal{Q}^{\mathbb{S}}$.

If $D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) > \bar{\eta}$, the divergence exceeds what any distortion within the entropy budget can produce. The investor must expand the model to account for the missing factor.

Corollary 1 (Economic Stopping Criterion). *The investor includes factor $\mathbf{g}$ only if*

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) > \bar{\eta}. \quad (1.17)$$

The entropy budget $\bar{\eta}$ bounds the total divergence the adversary can produce in any direction on $\mathbb{M}$, so it is an upper bound on the divergence the adversary could produce specifically in the direction of $\mathbf{g}$. The condition is therefore conservative.

The stopping criterion admits a connection to the Hansen-Jagannathan (HJ) bound and the GRS test statistic. Bounding h^* by the HJ inequality yields $SR_{h^*} \leq \sqrt{e^{2\eta} - 1}$, where $\eta = \mathbf{E}_{\mathcal{P}}[h^* \ln(h^*)]$ is the entropy of the worst-case distortion. The set of pricing errors that can be rationalized by misspecification is

$$\mathbb{A}(\eta) := \{\boldsymbol{\alpha} \in \mathbb{R}^N : \boldsymbol{\alpha}^\top \Sigma^{-1} \boldsymbol{\alpha} \leq e^{2\eta} - 1\},$$

where Σ is the covariance matrix of the test assets.

An investor includes factor $\mathbf{g}$ if the pricing errors created by its omission lie outside $\mathbb{A}(\eta)$, that is, $\boldsymbol{\alpha}_{\mathbf{g}} \notin \mathbb{A}(\eta)$. Geometrically, the entropy bound $\bar{\eta}$ defines a ball on the manifold centered at $\mathcal{Q}^{\mathbb{S}}$. Factors whose marginal divergence falls within this ball are indistinguishable from adversarial distortion. Factors outside the ball represent genuine manifold curvature that must be captured by expanding the pricing kernel.

The entropy budget $\bar{\eta}$ is in principle endogenous to the selected factor set, since the entropy of h^* depends on which factors are already included in the SDF. I resolve this endogeneity by pre-calibrating $\bar{\eta}$ before starting the forward selection, analogous to fixing the significance level before conducting a sequence of hypothesis tests. The pre-calibration is accomplished through the detection error probability ε as in Hansen and Sargent (2010).

The detection error probability is defined as the average of Type I and Type II error when testing $\mathcal{Q}^{\mathbb{S}}$ against $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$

$$\varepsilon := \frac{1}{2} \left(\mathbf{P}_{\mathcal{Q}^{\mathbb{S}}} \left\{ \text{Reject } \mathcal{Q}^{\mathbb{S}} \text{ for } \mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \right\} + \mathbf{P}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}} \left\{ \text{Fail to reject } \mathcal{Q}^{\mathbb{S}} \right\} \right).$$

A high ε (close to $1/2$) means the two models are statistically indistinguishable; a low ε (close to 0) means the investor is confident in their ability to discriminate between the baseline and the alternative. Under a Gaussian approximation, the closed-form mapping from ε to $\bar{\eta}$ is

$$\bar{\eta} = \frac{2 (\Phi^{-1}(1 - \varepsilon))^2}{T}, \quad (1.18)$$

where Φ^{-1} is the inverse standard normal CDF and T is the sample size. The investor's choice of ε and the sample size T are the only inputs to the stopping condition.³

The relationship between ε and factor inclusion runs counterintuitively and warrants careful discussion. Low ε means the investor can clearly attribute observed pricing errors to the adversary's distortion rather than to an omitted priced factor. This implies a large entropy budget $\bar{\eta}$. The investor tolerates more apparent mispricing before concluding that a genuine factor is missing, so fewer factors pass the inclusion threshold. High ε means the investor cannot reliably distinguish adversarial distortion from genuine factor omission. This implies a small entropy budget $\bar{\eta}$, so more factors must be included to account for residual pricing errors that cannot be confidently assigned to misspecification. The mapping is counterintuitive because one might expect an ambiguity-averse investor (high ε) to include fewer factors, having already hedged against model uncertainty. The resolution is that when the adversary's power is large and indistinguishable from a missing factor, the investor is compelled to add factors precisely because they cannot rule out that the mispricing is genuine.

The economic stopping criterion of Corollary 1 and the statistical stopping rule of Proposition 1 address distinct failure modes, and neither can perform the other's role.

The KL divergence is bounded below by zero, so sampling noise will produce a nonzero

³This mapping has been validated by simulation to be an excellent approximation under non-Gaussian distributions (Student- t , Gaussian mixtures) at empirically relevant sample sizes.

$\hat{D}_{KL}$ in any finite sample. Even when the economic criterion is satisfied, the investor must verify that the signal is detectable and not merely sampling noise. This is the role of the Wald test. The statistical criterion prevents false inclusion, the admission of factors that carry no true pricing signal. The economic criterion prevents overfitting, the inclusion of factors whose genuine signal is swamped by the amount of mispricing the adversary's entropy budget can rationalize.

Consider the two failure modes separately. A factor with $\theta_{\mathbf{g}}^* = 0$ will produce a nonzero estimated divergence in finite samples through sampling variation alone. If this spurious divergence exceeds $\bar{\eta}$, the economic criterion would erroneously admit the factor. The Wald test guards against this by requiring that the estimated risk price be large relative to its standard error. Conversely, a factor with a small but highly significant risk price might pass the Wald test easily while producing a divergence below $\bar{\eta}$. Including it would add a parameter for minimal pricing improvement, the classic kitchen-sink problem. The economic criterion guards against this by requiring that the divergence exceed the entropy budget.

Combining the two criteria yields the investor's complete inclusion decision.

Definition 1 (Equilibrium Stopping Rule). *The robust investor includes the candidate $\mathbf{g}$ with the highest $\hat{D}_{KL}$ if and only if both conditions hold:*

$$\hat{D}_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}}) > \bar{\eta}, \quad (1.19)$$

$$W_T(\mathbf{g}) > c_{\alpha}, \quad (1.20)$$

where $\bar{\eta}$ is calibrated via (1.18) and c_{α} is the χ_1^2 critical value at significance level α .

The four-quadrant structure of the $(\xi_{\mathbf{g}}, \delta_{\mathbf{g}})$ space, where $\xi_{\mathbf{g}} := D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}})$ is the population KL divergence and $\delta_{\mathbf{g}}$ is the Wald noncentrality, generates a taxonomy of pricing anomalies

Table 1.1: Factor Zoo Taxonomy		
	Large δ_g	Small δ_g
Large ξ_g	Both pass $\rightarrow$ Include	Econ pass, stat fail $\rightarrow$ Exclude
Small ξ_g	Econ fail, stat pass $\rightarrow$ Exclude	Both fail $\rightarrow$ Exclude

Table 1.1: Taxonomy of the factor zoo. ξ_g is the population KL divergence and δ_g is the Wald noncentrality.

The upper-left quadrant contains well-behaved factors with detectable, meaningful risk prices. A market-based factor may have a large and precisely estimated risk price because market returns are relatively well-behaved statistically. Both criteria are satisfied, and the factor enters the SDF.

The upper-right quadrant contains factors whose mispricing is genuine but whose data cannot yet support inclusion. An options-based tail-risk factor may carry substantial misspecification content, reflected in a large ξ_g , but the extreme kurtosis of option returns inflates long-run variance Ω_{gg} relative to information I_{gg} , simultaneously raising the economic threshold and depressing the Wald noncentrality δ_g . These factors may migrate to the upper-left as sample sizes grow or estimation techniques improve.

The lower-left quadrant contains factors with small but tightly estimated risk prices, the kitchen-sink problem. A low-volatility factor might pass the Wald test with ease because its returns are smooth and Ω_{gg} is small, but the divergence itself is negligible, falling below $\bar{\eta}$. Including it would add a parameter that buys almost no pricing improvement in the SDF.

The lower-right quadrant contains noise or data mined factors.

The geometric algorithm of Section 1.3.1 and the robust investor's equilibrium selection of Section 1.3.2 describe the same phenomenon from two vantage points. The geometry provides the ranking criterion and the computational machinery. The economic model provides the stopping threshold and the welfare interpretation. The two converge: the investor's preference-derived ranking by welfare loss coincides with the geometry's ranking by KL divergence, and both stopping conditions, statistical detectability and economic materiality, emerge from each framework independently. As $T \rightarrow \infty$, the Wald noncentrality grows without bound for every factor with a nonzero risk price, $\bar{\eta} \rightarrow 0$ through the calibration in

(1.18), and both criteria are eventually satisfied for every priced factor. The selected set converges to $\mathbb{S}^* := \{\mathbf{g} \in \mathbb{R}^K : \theta_{\mathbf{g}}^* \neq 0\}$, the population truth. The following Methodology section operationalizes this unified framework for empirical implementation.

1.4 Methodology

This section details the construction of the Kullback-Leibler Stochastic Discount Factor (KL-SDF) and the estimation techniques required to implement it. The methodology proceeds in two distinct stages. First, I describe a greedy forward-selection algorithm to identify the set of risk factors that maximally expands the span of priced risk, systematically enlarging the effective dimension of the stochastic discount factor. Second, I estimate the parameters of the resulting exponential-affine SDF using a Minimum Cross-Entropy (MCE) approach that is robust to the heavy tails and non-linearities inherent in financial data. This paper is, to my knowledge, the first to estimate the full nested exponential-affine SDF under pricing restrictions.

1.4.1 Forward Selection Algorithm

A primary challenge in empirical asset pricing is the identification of the true pricing kernel from a factor zoo populated by hundreds of potential risk candidates. Since the number of possible factor combinations grows exponentially with the size of the candidate pool, an exhaustive search is computationally infeasible. The factor zoo consists of approximately 151 factors, corresponding to 2^{151} model specifications, more than the number of nanoseconds elapsed since the Big Bang ($\approx 4 \times 10^{26}$), meaning no algorithm running since the formation of the universe could have evaluated even a negligible fraction of the candidate space. Instead, I construct the SDF using a stepwise forward-selection procedure rooted in information geometry.

The algorithm builds the pricing kernel dimension by dimension. It begins with a null model and sequentially adds the single factor that yields the largest marginal increase in the manifold's curvature. Formally, let $\mathbb{S}$ denote the set of currently selected factors and $\mathcal{Q}^{\mathbb{S}}$

denote the risk-neutral probability measure induced by those factors. In each iteration, the algorithm considers every remaining candidate factor $\mathbf{g}$ from the pool $\mathbb{F}$. It computes the KL divergence between the augmented model $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$, which includes the new factor, and the current restricted model $\mathcal{Q}^{\mathbb{S}}$. The candidate that maximizes this divergence is permanently added to the basis set $\mathbb{S}$. Econometrically, the procedure echoes forward stepwise regression, but with an intrinsic curvature-based criterion rather than the R^2 or F -statistic used in classical model selection. The analogy to Akaike (1974) is useful: both approaches use KL-oriented scoring, but the selection criterion here operates directly on the statistical manifold of exponential-affine SDFs rather than on a linear likelihood.

The theoretical justification for this greedy approach rests on the dualistic geometry of the exponential family. The manifold of admissible SDFs is e -flat, while the space of distributions satisfying the moment conditions is m -flat. This dual flatness admits a generalized Pythagorean decomposition (1.5). At each step of the algorithm, the factor maximizing the KL divergence is the one whose orthogonal component under $\mathcal{Q}^{\mathbb{S}}$ creates the largest independent displacement in the tangent space of the SDF manifold. The economic intuition is direct. Consider the population analogue: the omission of a priced factor from the SDF induces a non-zero alpha, a systematic pricing error, on portfolios loaded on that factor. The KL divergence translates this pricing error into information units, so the factor with the largest divergence is the one whose omission would generate the largest alpha under the current model. By targeting this factor at each step, the algorithm expands the effective dimension of the manifold rather than traversing directions already spanned by $\mathbb{S}$.

The complete procedure is presented in Algorithm 1. I now walk through its structure, explaining the rationale for each step.

The algorithm begins with an empty factor set $\mathbb{S}$ and a zero parameter vector. Under the null model, the SDF is a constant $m_t = 1$ for all t , implying that the risk-neutral measure coincides with the physical measure and no risk is priced. Each subsequent iteration perturbs this baseline by adding the single factor whose inclusion forces the largest revision of the state-price density.

At each iteration, the algorithm computes the implied probability weights $\pi_t(\hat{\boldsymbol{\theta}}^{\mathbb{S}})$ induced

Algorithm 1 Greedy D_{KL} Forward Selection Algorithm**Require:** Factor zoo $\mathbb{F}$, asset excess returns $R^e \in \mathbb{R}^{T \times N}$, entropy threshold $\bar{\eta}$

```

1:  $\mathbb{S} \leftarrow \emptyset$ 
2:  $\hat{\boldsymbol{\theta}}^{\mathbb{S}} \leftarrow \mathbf{0}_K$ 
3: while true do
4:    $\mathcal{Q}^{\mathbb{S}} \leftarrow m_t(\boldsymbol{\theta}^{\mathbb{S}}) \left( \sum_{s=1}^T m_s(\boldsymbol{\theta}^{\mathbb{S}}) \right)^{-1}$ 
5:    $\tilde{\mathbf{f}}_t^{\mathbb{S}} \leftarrow \text{ZCA}_{\mathcal{Q}^{\mathbb{S}}}(\mathbf{f}_t^{\mathbb{S}})$ 
6:   for all  $\mathbf{g}_t \in \mathbb{F} \setminus \mathbb{S}$  do
7:      $\tilde{\mathbf{g}}_t \leftarrow \text{ZCA}_{\mathcal{P}}(\mathbf{g}_t)$ 
8:      $\tilde{\mathbf{g}}_t^{\perp} \leftarrow \tilde{\mathbf{g}}_t - \text{Proj}_{\text{span}(\mathbb{S})}^{\mathcal{Q}^{\mathbb{S}}}(\tilde{\mathbf{g}}_t)$ 
9:      $X_t \leftarrow [\tilde{\mathbf{f}}_t^{\mathbb{S}}, \tilde{\mathbf{g}}_t^{\perp}]$ 
10:     $\hat{\boldsymbol{\theta}}^{\mathbb{S}, \mathbf{g}} \leftarrow \text{MCEOPTIMIZATION}(X_t, R^e)$ 
11:     $D_{KL}^{\mathbf{g}} \leftarrow \psi(\hat{\boldsymbol{\theta}}^{\mathbb{S}}) - \psi(\hat{\boldsymbol{\theta}}^{\mathbb{S}, \mathbf{g}}) + \nabla_{\boldsymbol{\theta}} \psi(\hat{\boldsymbol{\theta}}^{\mathbb{S}, \mathbf{g}})(\hat{\boldsymbol{\theta}}^{\mathbb{S}, \mathbf{g}} - \hat{\boldsymbol{\theta}}^{\mathbb{S}})$ 
12:     $\mathbb{J} \leftarrow \{\mathbf{g} : W_T(\mathbf{g}_t) > \chi_1^2(1 - \alpha)\}$ 
13:     $\mathbf{g}^* \leftarrow \arg \max_{\mathbf{g} \in \mathbb{J}} D_{KL}^{\mathbf{g}}$ 
14:    if  $\mathbb{J} = \emptyset$  or  $D_{KL}^{\mathbf{g}^*} \leq \bar{\eta}$  then
15:      break ▷ no candidate can explain remaining pricing errors
16:     $\mathbb{S} \leftarrow \mathbb{S} \cup \{\mathbf{g}^*\}$ 
17:     $\hat{\boldsymbol{\theta}}^{\mathbb{S}} \leftarrow \hat{\boldsymbol{\theta}}^{\mathbb{S}, \mathbf{g}^*}$ 
18: return  $\mathbb{S}$ 

```

by the current SDF. These weights define the risk-neutral measure $\mathcal{Q}^{\mathbb{S}}$ under which all subsequent operations within the iteration are performed. As factors are added and $\hat{\boldsymbol{\theta}}^{\mathbb{S}}$ changes, the tilted measure evolves, so that orthogonalization, whitening, and divergence evaluation all reflect the geometry of the current model rather than the geometry of the raw data.

Before evaluating a candidate, I apply zero-phase component analysis (ZCA) whitening to the factor panel under the current tilted probabilities. ZCA whitening sphericalizes the factor innovations, converting elliptical contours in the optimization surface to circular ones, and improves the condition number of the Hessian for the MCE estimation that follows.

Lemma 2 (Affine Reparameterization Invariance). *The forward selection algorithm is invariant to affine reparameterization of the exponential family.*

Since whitening is an affine transformation of the factor space via $\hat{\Sigma}_{\mathcal{Q}^{\mathbb{S}}}^{-1/2}$, Lemma 2 guarantees that the divergence improvement $D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})$ and the ordering of candidates are unchanged. ZCA therefore serves purely as a numerical conditioning device, with no effect on the economic content of the selection criterion.

Candidate factors in the zoo are highly correlated. In the raw data, multicollinearity inflates the variance of parameter estimates and flattens the optimization landscape, making it difficult to isolate the marginal contribution of any single factor. Line 7 projects each whitened candidate $\tilde{\mathbf{g}}_t$ onto the orthogonal complement of $\text{span}(\mathbb{S})$ under $\mathcal{Q}^{\mathbb{S}}$, retaining only the residual $\mathbf{g}_t^\perp$.

Lemma 3 (Orthogonal Projection Invariance). *Replacing a candidate factor $\mathbf{g}$ with its orthogonal residual $\mathbf{g}^\perp := \mathbf{g} - \text{Proj}_{\text{span}(\mathbb{S})}^{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g})$ under the current risk-neutral measure $\mathcal{Q}^{\mathbb{S}}$ leaves the augmented SDF manifold and the KL divergence unchanged.*

Lemma 3 states that the information gain of a factor is determined solely by its component outside the existing span. The proof, provided in Appendix A.1.6, follows from the fact that the component of $\mathbf{g}$ within $\text{span}(\mathbb{S})$ is already absorbed by the existing parameters $\hat{\boldsymbol{\theta}}^{\mathbb{S}}$, contributing zero marginal divergence. This Gram-Schmidt (Gram, 1883; Schmidt, 1907) orthogonalization step delivers two operational benefits. First, it stabilizes the numerical conditioning of the design matrix $X_t = [\tilde{\mathbf{f}}_t^{\mathbb{S}}, \tilde{\mathbf{g}}_t^\perp]$, since the columns are now uncorrelated under $\mathcal{Q}^{\mathbb{S}}$. Second, it acts as a natural filter against redundant factors: if $\mathbf{g}^\perp \approx \mathbf{0}$, the candidate is nearly spanned by $\mathbb{S}$, and its KL contribution will be negligible regardless of its raw correlation with returns.

For each candidate, the algorithm constructs the augmented design matrix X_t , estimates the full parameter vector $\hat{\boldsymbol{\theta}}^{(\mathbb{S}, \mathbf{g})}$ via MCE optimization, detailed in Section 1.4.2, and computes the KL divergence of the augmented model relative to the current model using Equation (1.7).

Not every candidate with positive KL divergence warrants inclusion. Line 12 applies a Wald test to all remaining candidates, retaining in the feasible set $\mathbb{J}$ only those factors whose

risk price is statistically distinguishable from zero at the chosen significance level. Line 13 selects from this feasible set the factor with the highest KL divergence. This two-stage screening separates the statistical question, *Is the risk price significantly different from zero?*, from the economic question *Which surviving candidate contributes the most to the span of priced risk?*.

The algorithm terminates when either the feasible set is empty, meaning no remaining candidate has a statistically significant risk price, or the maximum KL divergence within the feasible set falls below the entropy bound $\bar{\eta}$.

The output of the algorithm is an ordered set of factors $\mathbb{S}$ and a corresponding parameter vector $\hat{\boldsymbol{\theta}}^{\mathbb{S}}$ that together define a parsimonious exponential-affine pricing kernel. The ordering is economically meaningful: each factor is selected because, conditional on those already in the model, it resolves the largest remaining mispricing in the cross-section. The stopping rule ensures that the final model includes every factor whose contribution to the span of priced risk exceeds the threshold implied by the investor's tolerance for model misspecification, and excludes those that do not.

1.4.2 Estimating the Exponential-Affine SDF

Once a candidate factor set $\mathbb{S}$ is selected, the second stage of the methodology requires a robust technique to estimate the risk prices. The natural candidate is the Generalized Method of Moments (GMM), which minimizes a quadratic form of the sample pricing errors $\mathbf{h}_T(\boldsymbol{\theta})^\top W \mathbf{h}_T(\boldsymbol{\theta})$. However, the exponential structure of the pricing kernel creates a specific pathology. The exponential function $m_t(\boldsymbol{\theta})$ allows the SDF to take values arbitrarily close to zero. In finite samples, a GMM optimizer can exploit this property by driving $\boldsymbol{\theta}$ to extreme magnitudes, attenuating $m_t(\boldsymbol{\theta})$ in crash states and artificially minimizing the quadratic objective. The result is a boundary solution where the estimated SDF concentrates probability mass on a subset of observations while ignoring the rest. Economically, such a solution implies that the marginal utility of wealth is negligible during market crashes, the precise opposite of standard equilibrium theory.

I instead frame estimation as an information projection problem. The goal is to find a

risk-neutral probability measure $\mathcal{Q}$ that satisfies the no-arbitrage conditions while minimizing the informational distance to the physical measure $\mathcal{P}$. This is the Minimum Cross-Entropy (MCE) estimator.

Let $\mathbf{f}_t \in \mathbb{R}^K$ denote the vector of selected risk factors and $\mathbf{r}_t^e \in \mathbb{R}^N$ denote the vector of asset excess returns, with $N > K$. The objective is to estimate the parameter vector $\boldsymbol{\theta}$ for the exponential-affine pricing kernel (1.2) subject to the fundamental asset pricing restrictions $\mathbf{E}_{\mathcal{P}} [m_t \mathbf{r}_t^e] = \mathbf{0}_N$.

Fix a risk-price vector $\boldsymbol{\theta}$ and define the set of all probability distributions satisfying the pricing restrictions as

$$\mathbb{P}(\boldsymbol{\theta}) := \{\mathcal{P} : \mathbf{E}_{\mathcal{P}} [m_t(\boldsymbol{\theta}) \mathbf{r}_t^e] = \mathbf{0}_N\}.$$

The primal problem minimizes the Kullback-Leibler divergence to the true distribution $\mathcal{P}^\dagger$

$$\min_{\mathcal{P} \in \mathbb{P}(\boldsymbol{\theta})} D_{KL}(\mathcal{P}^\dagger || \mathcal{P}) \equiv \min_{\mathcal{P} \in \mathbb{P}(\boldsymbol{\theta})} \int \ln \frac{\mathcal{P}^\dagger}{\mathcal{P}} d\mathcal{P}. \quad (1.21)$$

The formulation asks for the minimal reweighting of the pricing kernel required to fairly price the assets. If the factor model is correctly specified, there exists a $\boldsymbol{\theta}$ such that the pricing errors are zero under the empirical measure, and the divergence vanishes. If the model is misspecified, the estimator finds the closest valid pricing model in the information-geometric sense. A key advantage over GMM is that the MCE approach enforces non-negativity of the SDF by construction: exponential tilting cannot produce negative probability weights, so the implied pricing kernel remains strictly positive even under misspecification, eliminating the boundary solutions that arise when GMM drives parameters to extreme magnitudes.

Solving the primal problem directly is intractable because it requires optimization over the space of all probability distributions. Kitamura and Stutzer (1997) show that this constrained optimization can be transformed into an unconstrained dual problem via Lagrange multipliers. Introducing Lagrange multipliers $\boldsymbol{\gamma} \in \mathbb{R}^N$ for the moment constraints, the dual objective function is the cumulant generating function of the pricing errors

$$\phi_T(\boldsymbol{\gamma}; \boldsymbol{\theta}) := \ln \left(\frac{1}{T} \sum_{t=1}^T \exp(\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e) \right). \quad (1.22)$$

The estimation proceeds as a nested optimization

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^K} \max_{\boldsymbol{\gamma} \in \mathbb{R}^N} \phi_T(\boldsymbol{\gamma}; \boldsymbol{\theta}), \quad (1.23)$$

which can equivalently be written as a nested minimization problem via Fenchel duality by introducing the convex conjugate of the cumulant generating function of the pricing restrictions. In practice, I solve the nested min-min formulation because the inner problem is strictly convex with a unique global solution, and the envelope formula delivers analytic outer gradients. As shown by Kitamura and Stutzer (1997), the MCE estimator is asymptotically equivalent to two-step efficient GMM, inheriting its consistency and asymptotic normality while avoiding the finite-sample pathologies described above.

The estimation procedure separates the optimization into two stages that exploit the different curvature properties of the objective. Although the MCE framework has been developed theoretically by Kitamura and Stutzer (1997) and applied in simplified settings (Ghosh et al., 2016), that implementation solves only the inner problem for fixed risk prices or impose additional restrictions that reduce the parameter space. The full nested optimization over both the Lagrange multipliers $\boldsymbol{\gamma}$ and the exponential-affine risk prices $\boldsymbol{\theta}$ has not, to my knowledge, been implemented before. Because the outer problem is non-convex and the exponential structure amplifies small parameter perturbations into large changes in the state-price density, the practical details of the optimization matter for reproducibility. I therefore describe the two loops in turn.

The estimation procedure separates the optimization into two distinct stages that exploit the different curvature properties of the objective.

The inner- $\boldsymbol{\gamma}$ loop, detailed in Algorithm 2, solves for the optimal Lagrange multipliers $\boldsymbol{\gamma}^*(\boldsymbol{\theta})$ holding the factor parameters $\boldsymbol{\theta}$ fixed. This sub-problem minimizes the dual objective $\phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta})$, which is strictly convex in $\boldsymbol{\gamma}$, guaranteeing a unique global solution. The resulting vector $\boldsymbol{\gamma}^*$ has a clear economic interpretation: it represents the shadow prices of the moment

Algorithm 2 Inner γ Loop
$$\begin{aligned}
&1: \text{Precompute } m_t(\boldsymbol{\theta}) \leftarrow \exp(-\boldsymbol{\theta}^\top \mathbf{f}_t - \psi(\boldsymbol{\theta})) \text{ for } t = 1, \dots, T \\
&2: \quad \text{where } \psi(\boldsymbol{\theta}) = \ln \frac{1}{T} \sum_{t=1}^T \exp(-\boldsymbol{\theta}^\top \mathbf{f}_t) \\
&3: \text{Set } \mathbf{h}_t \leftarrow m_t(\boldsymbol{\theta}) \mathbf{r}_t^e \text{ for } t = 1, \dots, T \\
&4: \text{Initialize } \boldsymbol{\gamma}^{(0)} \leftarrow \mathbf{0}; \ell \leftarrow 0 \\
&5: \textbf{while } \ell \leq L_{\max} \textbf{ do} \\
&\quad 6: \quad s_t^{(\ell)} \leftarrow \boldsymbol{\gamma}^{(\ell)\top} \mathbf{h}_t \text{ for } t = 1, \dots, T \\
&\quad 7: \quad \tilde{s}_t \leftarrow s_t^{(\ell)} - \max_t s_t^{(\ell)} \text{ for } t = 1, \dots, T \quad \triangleright \text{Used for numerical stability in log-sum-exp.} \\
&\quad 8: \quad w_t^{(\ell)} \leftarrow \frac{\exp(\tilde{s}_t)}{\sum_{s=1}^T \exp(\tilde{s}_s)} \text{ for } t = 1, \dots, T \\
&\quad 9: \quad \mathbf{h}^{(\ell)} \leftarrow \sum_{t=1}^T w_t^{(\ell)} \mathbf{h}_t \\
&\quad 10: \quad \textbf{if } \|\mathbf{h}^{(\ell)}\|_\infty < \varepsilon_\gamma \textbf{ then} \quad \triangleright \nabla_\gamma \phi_T(\boldsymbol{\gamma}^{(\ell)}) = \mathbf{0} \\
&\quad \quad \textbf{break} \\
&\quad 12: \quad d^{(\ell)} \leftarrow \text{L-BFGS} \\
&\quad 13: \quad \alpha^{(\ell)} \leftarrow \text{Guarded Strong Wolfe} \\
&\quad 14: \quad \boldsymbol{\gamma}^{(\ell+1)} \leftarrow \boldsymbol{\gamma}^{(\ell)} + \alpha^{(\ell)} d^{(\ell)} \\
&\quad 15: \quad \ell \leftarrow \ell + 1 \\
&\textbf{return } \boldsymbol{\gamma}^* \leftarrow \boldsymbol{\gamma}^{(\ell)}, \quad w_t^* \leftarrow w_t(\boldsymbol{\gamma}^*; \boldsymbol{\theta}) \text{ for } t = 1, \dots, T
\end{aligned}$$

The outer- $\boldsymbol{\theta}$ loop, detailed in Algorithm 3, minimizes the profiled objective function $J_T(\boldsymbol{\theta}) = \phi_T(\boldsymbol{\gamma}^*(\boldsymbol{\theta}), \boldsymbol{\theta})$ with respect to the factor loadings $\boldsymbol{\theta}$. This step identifies the SDF parameters that minimize the distributional distortion found in the inner loop. While the inner problem is convex, the outer problem is generally non-convex due to the nonlinearity of the exponential SDF. I exploit the Envelope Theorem to compute analytic gradients for the outer optimization: because $\boldsymbol{\gamma}^*$ minimizes the inner objective, the gradient of the profiled objective depends only on the partial derivative with respect to $\boldsymbol{\theta}$, avoiding the computational expense and instability of numerical differentiation.

Algorithm 3 Outer $\boldsymbol{\theta}$ Loop

Require: $\{\mathbf{r}_t^e \in \mathbb{R}^N\}_{t=1}^T$, $\{\mathbf{f}_t \in \mathbb{R}^K\}_{t=1}^T$, tolerances $\varepsilon_\theta, \varepsilon_\gamma, \varepsilon_{\text{rel}}$, max iterations L_{max}

- 1: Initialize $\boldsymbol{\theta}^{(0)}$ (e.g., $\mathbf{0}$ /multi-start)
 - 2: $\ell \leftarrow 0$
 - 3: $J^{(-1)} \leftarrow \infty$
 - 4: **while** $\ell < L_{\text{max}}$ **do**
 - 5: $\psi^{(\ell)} \leftarrow \ln T^{-1} \sum_{s=1}^T e^{-\boldsymbol{\theta}^{(\ell)\top} \mathbf{f}_s}$
 - 6: $m_t \leftarrow e^{-\boldsymbol{\theta}^{(\ell)\top} \mathbf{f}_t - \psi^{(\ell)}}$ for $t = 1, \dots, T$
 - 7: $g_t \leftarrow m_t \mathbf{r}_t^e$ for $t = 1, \dots, T$
 - 8: $(\boldsymbol{\gamma}^*, \{w_t^*\}) \leftarrow \text{INNERGAMMALOOP}(\{g_t\}, \varepsilon_\gamma)$
 - 9: $\nabla^{(\ell)} \leftarrow - \sum_{t=1}^T w_t^* m_t(\boldsymbol{\theta}^{(\ell)}) (\boldsymbol{\gamma}^*)^\top \mathbf{r}_t^e \left(\mathbf{f}_t - T^{-1} \sum_{\tau=1}^T m_\tau(\boldsymbol{\theta}^{(\ell)}) \mathbf{f}_\tau \right)$
 - 10: $J^{(\ell)} \leftarrow \ln T^{-1} \sum_{t=1}^T e^{\boldsymbol{\gamma}^{*\top} g_t}$
 - 11: **if** $\|\nabla^{(\ell)}\|_\infty \leq \varepsilon_\theta$ **and** $|J^{(\ell)} - J^{(\ell-1)}| \leq \varepsilon_{\text{rel}}(1 + |J^{(\ell)}|)$ **then**
 - 12: **break**
 - 13: $d^{(\ell)} \leftarrow \text{L-BFGS}$
 - 14: $\alpha^{(\ell)} \leftarrow \text{Guarded Strong Wolfe}$
 - 15: $\boldsymbol{\theta}^{(\ell+1)} \leftarrow \boldsymbol{\theta}^{(\ell)} + \alpha^{(\ell)} d^{(\ell)}$
 - 16: $\ell \leftarrow \ell + 1$
 - return** $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}^{(\ell)}$
-

The dual architecture offers several advantages over standard GMM in the exponential-affine setting. The objective function penalizes higher moments of the pricing errors, providing inherent robustness against the outliers that often derail quadratic GMM objectives. The log-sum-exp structure of ϕ_T acts as a smooth approximation to the maximum function, preventing the optimizer from fixating on single extreme observations or crash states. The separation of the linear moment constraints (handled in the inner loop) from the nonlinear parameter search (handled in the outer loop) ensures that the Euler equations are satisfied to high precision at every step of the estimation, in contrast to GMM, which often treats constraints as soft targets and may violate the fundamental pricing restrictions in finite samples. The full derivation of the dual objective and its gradient is provided in Appendix A.2.2.

The low signal-to-noise ratio characteristic of financial returns creates three interrelated computational challenges for the exponential-affine SDF. First, the objective surface is asymptotically flat over broad regions of the parameter space, depriving gradient-based methods of a reliable descent direction. Second, the exponential map amplifies small perturbations in $\boldsymbol{\theta}$ into large changes in $m_t(\boldsymbol{\theta})$, producing degenerate probability distributions in which a handful of observations absorb nearly all the implied probability mass. Third, these two features interact to produce numerical cliffs: abrupt changes in the objective from parameter perturbations that are negligible in Euclidean norm but large in the natural geometry of the exponential family.

To address these challenges, I developed a bespoke optimization architecture whose full details are provided in Appendix A.3. The key techniques are as follows. ZCA whitening of factors and test assets before estimation sphericalizes the Hessian, eliminating the ill-conditioning induced by correlated factors. Guard-safe line searching monitors the effective sample size and weight concentration at every candidate step, rejecting steps that produce pathological probability distributions. Backtracking and trust-region strategies handle the non-convexity of the outer problem, preventing stagnation in shallow local minima. Multi-start initialization from both structured starting points (linear risk premia estimates $\boldsymbol{\lambda} = \Sigma_F \boldsymbol{\theta}$) and random perturbations provides robustness against convergence to local optima.

A separate constraint arises from random matrix theory. The concentration ratio N/T governs the eigenvalue distribution of the sample covariance matrix of test asset returns via the Marchenko-Pastur law. When N/T is large, the smallest sample eigenvalues are contaminated by noise, falling below the Marchenko-Pastur edge and destabilizing both the ZCA whitening operator and the inner-loop Hessian. I impose $N/T < 0.1$ throughout the estimation to ensure that the sample eigenvalues remain well above the noise floor. The practical consequence is that the method requires long time series. This constraint limits the number of test assets that can be used at any given estimation step and motivates the specific test asset specifications employed in the empirical section.

The MCE estimator translates the economic requirement of no-arbitrage into a well-posed information-geometric optimization. By working in the dual space, the estimator

avoids the pathologies that afflict quadratic objectives in exponential models and recovers a pricing kernel that is positive, internally consistent, and driven by the full distribution of asset returns rather than a handful of influential observations. The nested structure is more demanding computationally than standard GMM, but the payoff is an estimator whose finite-sample behavior reflects the theoretical properties that motivate the exponential-affine specification in the first place.

1.5 Empirics

This section applies the forward selection algorithm to two progressively richer empirical settings. The first is an illustrative exercise on the 49 Fama-French industry portfolios with the six standard Fama-French factors, which makes the algorithm’s mechanics transparent in a well-understood cross-section. The second applies the algorithm at full scope, using 157 candidate factors drawn from the Jensen et al. (2023) characteristic universe and the six Fama-French factors, priced against 73 principal components of 4,686 characteristic-sorted decile portfolios. The central finding is that the algorithm selects a parsimonious set of thematic clusters rather than individual factors. Profit growth, profitability, quality, and value account for the majority of the cumulative KL budget; momentum, investment, low leverage, and low risk are marginal entrants to the selected model. Accruals is never selected by the Algorithm across all parameter configurations examined. The remainder of the section documents this result, traces its origins to the path dependence of the greedy procedure, and assesses its stability across the space of tuning parameters.

1.5.1 Discussion of Parameters

The forward selection algorithm is governed by three tuning parameters, each controlling a distinct aspect of factor inclusion. Before presenting the main results, I outline each parameter and its economic interpretation. These three criteria form a hierarchy, not merely a set of interchangeable robustness checks. First, the Wald test at significance level α acts as a necessary condition: a candidate factor must show statistically significant incremental

pricing content before entering the feasible set at a given iteration. Second, the entropy distortion budget $\bar{\eta}$ imposes an economic significance criterion, conditional on statistical significance. Third, the minimum effective sample size (ESS) threshold serves as a diagnostic guard against degenerate importance weights, rather than as a model-quality criterion.

A loose Wald significance level admits factors whose incremental KL contribution may be driven by sampling noise, inflating the apparent entropy budget absorbed by the selected set. A tight significance level may exclude genuinely priced factors in finite samples, particularly those whose risk price is economically meaningful but imprecisely estimated. The costs of these two errors are asymmetric. A false inclusion overstates robustness by attributing noise to priced risk; a false exclusion understates the entropy budget required to span the worst-case distortion. Because the multiplier-preference framework penalizes worst-case outcomes, the welfare consequences need not offset each other. I report the main results at $\alpha = 0.005$ and examine sensitivity to the alternatives in the robustness section.

Recall from Section 1.3 that $\bar{\eta}$ represents the maximum admissible KL divergence of the worst-case measure from the data-generating measure. Instead of calibrating $\bar{\eta}$ directly, I use the detection error probability ε for calibration. When ε is high, $\bar{\eta}$ is small: the worst-case measure is nearly statistically indistinguishable from the baseline, so the investor attributes less to model misspecification and perceives more factors as priced risk. Conversely, when ε is low, $\bar{\eta}$ is large, and more pricing error is attributed to the adversary. Throughout, I set $\varepsilon = 0.10$ as the baseline⁴.

Finally, consider the ESS threshold. The exponential tilt that defines the MCE measure Q^S reweights each observation by its compatibility with the pricing kernel. Factors that achieve large Sharpe ratios by exploiting only a few extreme returns concentrate probability mass on those observations. This produces a distribution that is tilted far from uniform. The ESS quantifies this concentration. It measures the effective number of observations that receive non-negligible weight under Q^S , and is inversely related to the squared Euclidean distance between the uniform distribution and the vector of model-implied proba-

⁴A detection error probability of $\varepsilon = 0.10$ means the Bayesian agent misclassifies the two measures 10% of the time. As ε rises toward 0.50, the two measures become indistinguishable. Equation (1.18) relates the entropy budget of the adversary to the error detection probability.

bility weights (Martino et al., 2017). When the tilt is mild and the pricing content is broadly supported, ESS remains close to T . In contrast, when a factor’s contribution is driven by only a few influential data points, ESS falls sharply. This signals in-sample overfitting rather than genuine risk pricing. To guard against this, I impose a threshold of $ESS/T \geq 0.30$, requiring that at least thirty percent of the sample retain meaningful probability weight under the tilted measure.⁵

1.5.2 Illustrative Example

I apply the forward selection algorithm to a setting designed to make the procedure’s mechanics transparent before turning to richer specifications. The test assets are the 49 Fama-French Industry value-weighted portfolios, observed monthly from July 1969 to August 2025. The candidate factor set consists of the six standard Fama-French factors (EQP⁶, SMB, HML, RMW, CMA) augmented with momentum (MOM). This configuration serves as a useful baseline for several reasons. The 49 industry portfolios are among the most widely studied test asset cross-sections in empirical asset pricing, and their limited dispersion in expected returns is well documented. The candidate factors have clear economic interpretations grounded in decades of research, which permits direct evaluation of the algorithm’s selections against prior knowledge. The combination of a low-dimensional factor zoo and a well-understood test asset universe also allows the results to be cross-checked against the large existing literature on HJ-distance bounds and mean-variance efficiency tests applied to these same portfolios. Table 1.2 reports the full selection path, displaying the top-ranked feasible and infeasible candidates at each iteration.

Column (1) reports the estimated risk price $\hat{\theta}_{\mathbf{g}}$ for each candidate factor, with heteroskedasticity-robust standard errors in parentheses. Column (2) reports the incremental KL divergence $\Delta D_{KL}^{\mathbf{g}}$, which measures the information gain from adding factor $\mathbf{g}$ to the current SDF. The standard error is reported below. Column (3) translates this information gain into an an-

⁵In unreported results, ESS/T can fall below 0.05 before the exponential tilt encounters numerical convergence failures. However, specifications in this range yield implausible implied Sharpe ratios exceeding 10, far beyond any economically defensible bound.

⁶“Equity Premium” (Mkt-RF)

nualized Sharpe ratio. This represents the maximum reward-to-risk ratio an investor could extract from the factor’s marginal contribution to the pricing kernel. Column (4) reports the entropy bound $\bar{\eta}$, calibrated from the detection error probability ε . A factor whose incremental KL divergence exceeds this threshold contributes pricing information that an ambiguity-averse investor cannot attribute to model misspecification. Column (5) reports the Wald statistic W_T testing $\hat{\theta}_{\mathbf{g}} = 0$. Column (6) records the reason for exclusion. “Wald” indicates failure of statistical significance at conventional levels; “Econ” indicates the factor’s incremental KL falls below the entropy bound. A factor must pass both criteria to enter the model.

The algorithm selects two factors: the EQP in iteration 1 and SMB in iteration 2. The 49 industry portfolios exhibit limited cross-sectional dispersion in expected returns. Therefore, only a few systematic risk sources can be reliably identified from these test assets. Both selected factors enter with highly significant risk prices, and Wald statistics exceed 10. The annualized Sharpe ratios implied by their incremental contributions (0.503 for EQP, 0.588 for SMB) fall within plausible ranges for priced risk factors estimated over this sample. For each selected factor, the incremental KL divergence exceeds the entropy bound of 0.0032. This result implies that an ambiguity-averse investor calibrated to a detection error probability of $\varepsilon = 0.15$ would recognize these factors as genuine sources of systematic risk rather than artifacts of model misspecification.

The feasible set contains exactly one factor in each of the first two iterations. The algorithm terminates at iteration 3 with an empty feasible set. This pattern illustrates a property of the sequential procedure. As additional factors enter the SDF, the conditional risk price estimates for remaining candidates become less contaminated by omitted-variable bias. At the same time, the bar for incremental contribution rises because the unexplained pricing variation shrinks. The evolution of SMB across iterations is instructive. In iteration 1, SMB does not appear among the top-ranked infeasible candidates. This implies an incremental KL divergence below the 0.002 recorded for HML. In iteration 2, with the market factor included, SMB’s incremental KL rises to 0.014. This order-of-magnitude increase shows that size-related cross-sectional variation in industry returns is partially confounded with broad

Table 1.2: Illustrative Example Results

Iter.	Rank	Factor	$\hat{\theta}_{\mathbf{g}}$ (1)	$\Delta D_{KL}^{\mathbf{g}}$ (2)	SR (3)	$\bar{\eta}$ (4)	W_T (5)	Exc. (6)
1	<i>Feasible</i>							
	1	EQP	3.081*** (0.862)	0.011* (0.006)	0.503 —	0.0032 —	12.78*** —	— —
	<i>Infeasible</i>							
	1	CMA	-4.446 (2.864)	0.004 (0.006)	0.321 —	0.0032 —	2.41 —	Wald —
	2	HML	-2.215 (1.618)	0.002 (0.003)	0.237 —	0.0032 —	1.87 —	Wald, Econ —
	⋮							
2	<i>Feasible</i>							
	1	SMB	-5.674*** (1.776)	0.014 (0.009)	0.588 —	0.0032 —	10.21*** —	— —
	<i>Infeasible</i>							
	1	MOM	2.367 (1.533)	0.005 (0.007)	0.356 —	0.0032 —	2.38 —	Wald —
	2	RMW	4.221 (2.684)	0.005 (0.006)	0.335 —	0.0032 —	2.47 —	Wald —
	⋮							
3	<i>Feasible</i>							
	1.	—	—	—	—	—	—	—
	<i>Infeasible</i>							
	1	RMW	3.232 (2.898)	0.003 (0.005)	0.278 —	0.0032 —	1.24 —	Wald —
	2	MOM	1.593 (1.620)	0.003 (0.005)	0.246 —	0.0032 —	0.97 —	Wald, Econ —
	⋮							
$T = 674 \quad N = 49 \quad N/T = 0.073$								
$\sum \hat{D}_{KL}^{\mathbf{g}} = 0.0250$								
HJ-Distance = 0.2835								
$\hat{D}_{KL}(\hat{\mathcal{Q}}^{\mathbf{S}} \mathcal{P}) = 0.0251$								
ESS = 640.3 ESS/T = 0.950								

Table 1.2: The table reports the results from Algorithm 1 applied to the Fama-French 49 industry portfolios. Each iteration is partitioned into the feasible and infeasible sets, showing the top 2 factors in each set. The Algorithm was run with the error detection probability $\varepsilon = 0.15$, the Wald significance $\alpha = 0.005$, and minimum ESS at 90%. Standard errors in parentheses are HAC with a Bartlett kernel and 6 lags.

Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

market exposure. Once market risk is absorbed, the independent pricing content of SMB becomes identifiable. The changing ranks of the infeasible set reflect this as well. For example, MOM and RMW reverse order between iterations 2 and 3. The remaining candidates in iterations 2 and 3 (MOM, RMW, CMA) fail the Wald criteria.

The algorithm terminates at iteration 3 with an empty feasible set. At this point, the two highest-ranked infeasible candidates are RMW and MOM, each demonstrating how different constraints can lead to exclusion. Specifically, RMW meets the entropy bound but fails the Wald test, indicating that while an ambiguity-averse investor might view its incremental pricing content as economically meaningful, the associated risk price lacks estimation precision, precluding inclusion. In contrast, MOM fails both the entropy bound and the Wald test, justifying its exclusion on both economic and statistical grounds.

These exclusion patterns reflect the taxonomy from Table 1.1. A factor that passes the entropy bound but fails the Wald test provides economically relevant pricing information, but it is too imprecisely estimated to act on. CMA in iteration 1 and MOM and RMW in iteration 2 show this: their incremental KL divergences exceed $\bar{\eta}$, but risk prices remain statistically indistinguishable from zero. A factor that fails both criteria, such as HML in iteration 1 and MOM in iteration 3, makes neither a statistically reliable nor an economically significant contribution to the SDF, conditional on the selected factors. The table shows no factor passing the Wald test but failing the entropy bound. In such cases, a precisely estimated but economically negligible risk price emerges, which is more likely in a larger cross-section of asset returns with weak but well-measured signals.

The panel at the bottom of Table 1.2 reports summary diagnostics for the selected two-factor SDF. The cumulative KL divergence $\sum \hat{D}_{KL}^g = 0.0250$ quantifies the information content of the pricing kernel relative to the risk-neutral benchmark and closely matches the direct estimate $\hat{D}_{KL}(\hat{\mathcal{Q}}^S || \mathcal{P}) = 0.0251$, indicating minimal distortion of the physical measure to price the 49 industry portfolios.

No factor triggered an ESS warning at any iteration. The final two-factor model yields an ESS of 640.3, or 95.0% of the 674-month sample. The pricing kernel fits cross-sectionally by modest, diffuse changes to state probabilities, rather than concentrating mass on a few

extreme returns. Economically, market and size factors price the 49 industries by pervasive exposure differences, not rare-event mechanisms. In later specifications with broader test-asset sets and more factors, this property sometimes fails, making the ESS diagnostic crucial for guarding against overfitting due to extreme probability distortion.

1.5.3 The Full Factor Zoo

I now turn to the full candidate universe. The factor set consists of the 151 characteristic-based long-short factors constructed by Jensen et al. (2023), together with the six factors from the illustrative example. This deliberate augmentation introduces intentional redundancy, as several JKP factors overlap with or replicate the Fama-French and momentum factors already included. For example, the JKP factor `ret_12_1` and the standard momentum factor `MOM` represent the same underlying return predictor, and the algorithm ranks these two factors identically at every iteration of the forward selection. By including near-duplicate factors, I directly test whether the algorithm is robust to redundancy in the candidate set.

The test assets are principal components extracted from 4,686 characteristic-sorted decile portfolios spanning 623 anomaly dimensions. The anomalies are collected from Jensen et al. (2023) and Chen and Zimmermann (2022). I drop firms that have less than 5 years of characteristic data. The full sample contains 17,296 firms spanning from July 1969 to December 2024. To respect the concentration bound of $N/T < 0.1$, I retain the first 73 principal components of the test assets, which collectively explain 95% of the cross-sectional return variation.

The columns of Table 1.3 match those of Table 1.2, with two additions. Column (6) reports the projected effective sample size (ESS) that each candidate factor would produce if added to the SDF. Factors with projected ESS below the feasibility threshold are excluded, regardless of their KL divergence or Wald statistic. Each factor name includes its thematic classification as defined by Jensen et al. (2023).

The algorithm selects 14 factors spanning 9 thematic categories. The largest themes are profit growth (3 factors), profitability (2), quality (2), and value (2). Together these four themes account for approximately 62% of the cumulative KL divergence $\sum \hat{D}_{KL}^g = 2.27$.

Table 1.3: Full Factor Zoo Results

Iter.	Factor	$\hat{\theta}_{\mathbf{g}}$ (1)	$\Delta D_{KL}^{\mathbf{g}}$ (2)	SR (3)	$\bar{\eta}$ (4)	W_T (5)	Proj. ESS (6)	Exc. (7)
1	dsale_dinv (Profit Growth)	49.888*** (6.654)	0.256** (0.100)	2.477	0.005	56.21***	0.401	
2	nfna_gr1a (Debt Issuance)	51.586*** (7.481)	0.163*** (0.045)	1.979	0.005	47.55***	0.562	
3	seas_11_15an (Seasonality)	-26.671*** (6.269)	0.132** (0.054)	1.778	0.005	18.10***	0.399	
4	saleq_su (Profit Growth)	-22.788*** (5.074)	0.082** (0.033)	1.401	0.005	20.17***	0.330	
5	cop_at (Quality)	13.595*** (3.369)	0.091*** (0.016)	1.474	0.005	16.28***	0.568	
6	sale_me (Value)	17.663*** (2.897)	0.258** (0.103)	2.487	0.005	37.18***	0.431	
7	tax_gr1a (Profit Growth)	-16.384*** (5.321)	0.070*** (0.027)	1.298	0.005	9.48***	0.322	
8	EQP (Market)	6.251*** (1.191)	0.119*** (0.016)	1.691	0.005	27.53***	0.503	
9	ret_1_0 (ST Reversal)	-12.693*** (3.200)	0.145* (0.078)	1.865	0.005	15.74***	0.453	
10	ocf_at (Profitability)	23.707*** (6.450)	0.207*** (0.030)	2.231	0.005	13.51***	0.403	
11	ni_me (Value)	10.085*** (3.131)	0.288*** (0.012)	2.628	0.005	10.37***	0.383	
12	at_turnover (Quality)	13.004*** (4.627)	0.200*** (0.015)	2.193	0.005	7.90***	0.315	
13	ami_126d (Size)	8.176*** (2.753)	0.214*** (0.009)	2.264	0.005	8.82***	0.386	
14	o_score (Profitability)	-15.304*** (5.235)	0.044*** (0.013)	1.025	0.005	8.55***	0.466	
15 (rej.)	op_atl1 (Quality)	18.348* (10.136)	0.055*** (0.016)	1.152		3.28*	0.317	Wald
	⋮							
$T = 666 \quad N = 73 \quad N/T = 0.110$ $\sum \hat{D}_{KL}^{\mathbf{g}} = 2.2678$ HJ-Distance = 0.5549 $\hat{D}_{KL}(\hat{Q}^{\mathbf{S}} \mathcal{P}) = 0.3537$ ESS = 310.7 ESS/T = 0.466								

Table 1.3: The table reports the results from Algorithm 1 applied to the first 73 principal components of the 4,686 decile sorted portfolios and the 151 JKP + 6 FF factors. The table shows the final selected factor at each iteration. The Algorithm was run with the error detection probability $\varepsilon = 0.10$, the Wald significance $\alpha = 0.005$, and minimum ESS at 30%. Standard errors in parentheses are HAC with a Bartlett kernel and 6 lags.

Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

The remaining 5 selections span debt issuance, seasonality, short-term reversal, size, and the equity premium, each contributing a single factor. No momentum, accruals, investment, low leverage, or low risk factor enters the model.

The shift from the six-factor illustrative setting to the full 157-factor and test asset universe changes not only the size of the selected model but the identity of the factors that represent each economic theme. In the illustrative example, EQP enters at iteration 1 with an incremental KL divergence of 0.011, accounting for approximately 44% of the total KL budget against the 49-industry cross-section. In the full specification, EQP enters at iteration 8, after three profit-growth factors, one debt-issuance factor, one seasonality factor, one quality factor, and one value factor have sequentially reduced the unexplained cross-sectional KL. Despite entering later, EQP's incremental contribution in the full model is an order of magnitude larger than in the illustrative setting. The PCA test assets span substantially richer cross-sectional variation than the 49 industry portfolios, and the market factor faces stronger competition from characteristic-based factors that price aggregate-return exposure through firm-level sorts. The conditional risk price 6.25 is roughly twice its illustrative value, and the Wald statistic more than doubles. The market premium is therefore robustly priced across both test-asset universes, with its conditional pricing content strengthening once the characteristic factors that compete with it on the PCA test assets are controlled for.

The size factor behaves differently. SMB enters the illustrative model at iteration 2 with 0.014 and an annualized Sharpe ratio of 0.588. In the full specification, SMB does not enter at any iteration. The size theme is represented instead by the Amihud illiquidity measure at iteration 13, which enters with an incremental KL of 0.214 and a Sharpe ratio of 2.264. The Amihud measure loads on the same small-cap premium that drives SMB but captures it through illiquidity rather than through a market-cap-sorted spread (Amihud, 2002). Its conditional pricing content against the PCA test assets substantially exceeds what SMB contributes at any comparable point in the selection path. The displacement reflects the richer information set of the full test-asset cross-section. Against the PCA test assets, size-related pricing information is summarized more efficiently by a liquidity-based measure than by the market-cap-sorted spread that dominates against the industry portfolios.

The contrast between these two factors points to a general feature of greedy sequential selection. The identity of the representative that captures a given economic theme is determined partly by the ordering of earlier selections, not solely by the unconditional pricing content of the factor. In the illustrative model, SMB is the best available incremental source of size information given EQP. In the full model, twelve other factors enter before the size theme is reached, and a different representative wins the conditional race. The next section formalizes this argument and shows that while the representative standing for each theme may rotate across specifications, the set of themes itself is broadly stable.

1.5.4 Path Dependence

The patterns documented in the previous section are not specific to EQP and SMB. They reflect a general property of the greedy sequential algorithm, whereby within-theme path dependence rotates the factor representative but leaves the set of selected themes broadly intact.

At each iteration, the algorithm selects the candidate with the highest incremental KL divergence from the current conditional distribution. Two factors from the same thematic cluster share a large fraction of their pricing information, so whichever factor enters first substantially reduces the conditional KL available to the other. The second factor may then fail the Wald test, not because its unconditional pricing content is weak, but because its marginal contribution conditional on the first representative falls below the significance threshold. The algorithm therefore selects one or two representatives per theme rather than a dense within-theme bundle. Path dependence manifests primarily as rotation among near-substitutes within a theme rather than as substitution across themes.

The ordering of themes in the selection path has more consequence than within-theme rotation. Themes that enter early absorb a disproportionate share of the KL budget, reducing the conditional KL available to later themes. When two themes are partially correlated, the order in which their representatives enter can determine whether the second theme clears the entropy bound. For example, profitability and quality, which Novy-Marx (2013) and Asness et al. (2019) show share substantial overlap in their pricing content. Between-theme ordering

is itself path-dependent, depending on which factor from which theme attains the highest unconditional KL in the first iteration.

Despite within-theme rotation and between-theme ordering effects, the set of themes represented in the selected model is broadly stable across parameter configurations. To document this formally, I pool all ESS trials from $\text{ESS}/T \in \{0.3, 0.4, \dots, 0.9\}$ at the baseline $\alpha = 0.005$ and $\varepsilon = 0.10$. This yields seven model estimates. For each, I record which themes are represented in the selected set. Pooling across trials averages out within-theme rotation and surfaces the themes that reliably attract a representative regardless of the overfitting threshold.




Figure 1.1: This plot shows the cumulative D_{KL} gain across algorithm steps for the 7 aggregated scenarios. The scenarios differ in the minimum ESS threshold as a percentage of T .

Figure 1.1 plots the cumulative KL divergence $\sum \hat{D}_{KL}$ as a function of iteration k for each ESS trial. Through the first several iterations the paths are tightly bunched across thresholds, accumulating comparable KL at comparable rates. Trials with ESS between 40% and 80% subsequently decelerate and terminate between iterations 6 and 10, a shape consistent with a diminishing-returns greedy algorithm in which the dominant pricing information is absorbed early, and successive admissible candidates contribute progressively less. The ESS = 30% trial does not exhibit this pattern. Its slope is approximately constant across the full selection path, with the marginal KL contribution in the later iterations comparable to or larger than

in the early iterations, and the curve terminates abruptly at iteration 14 with a terminal cumulative KL more than twice that of the next-highest trial.

A possible explanation is that at the loosest thresholds the algorithm admits factors whose effective sample sizes are small enough that their realized KL contributions substantially exceed their population values. The sampling distribution of $\hat{D}_{KL}$ at small T can produce inflated point estimates, and once the algorithm exhausts the Wald-significant candidates in the common admissible set, the ESS = 30% filter permits short-sample factors whose apparent marginal contribution reflects a mixture of true pricing content and small-sample noise. The abrupt termination of the curve, rather than a smooth asymptotic approach, is consistent with this interpretation: the algorithm continues to find nominally significant additions until the residual set of short-sample candidates is exhausted, at which point selection halts. The tight bunching of the early iterations therefore indicates that the dominant economic content of the model is invariant to the ESS threshold, while the late divergence of the ESS = 30% trial indicates that the apparent KL gains from marginal additions at loose thresholds should not be interpreted as the discovery of new pricing dimensions. As documented in Table 1.4, the additional factors admitted at low ESS thresholds map almost entirely to themes already represented in higher-ESS selected sets.

Table 1.4 reports the thematic composition of the selected models aggregated across the seven pooled trials. Column (3) gives the share of trials in which at least one factor from each theme enters the selected set. Columns (4) and (5) report the median KL divergence of the theme’s factors on iteration 1 and pooled across all iterations, respectively. These two columns admit distinct interpretations. The iteration-1 median captures the theme’s unconditional pricing information, that is, the divergence obtainable when no other factor has been selected. The all-iterations median captures the theme’s average marginal contribution, which is mechanically smaller because once one representative is admitted, remaining candidates within the theme compete for residual conditional entropy. A large gap between columns (4) and (5) indicates a theme whose pricing content is concentrated in a single dimension, so that once the first representative enters, within-theme substitutes add little. A small gap indicates a theme whose factors carry pricing information along distinct economic

Table 1.4: Selected Factor Themes

Theme	# Factors (1)	Sel. Rate (2)	Models (3)	Median D_{KL} (It.1) (4)	Median D_{KL} (5)
Debt Issuance	7	14.3%	85.7%	0.061	0.010
Quality	16	7.1%	85.7%	0.009	0.005
Profit Growth	10	12.9%	71.4%	0.032	0.010
Seasonality	11	5.2%	57.1%	0.019	0.009
Value	14	7.1%	57.1%	0.013	0.009
EQP	1	42.9%	42.9%	0.035	0.035
Size	5	8.6%	42.9%	0.003	0.004
Short-Term Reversal	6	4.8%	28.6%	0.043	0.009
Profitability	10	5.7%	28.6%	0.016	0.005
Momentum	8	3.6%	28.6%	0.015	0.012
Investment	22	1.3%	28.6%	0.012	0.007
Low Leverage	10	2.9%	28.6%	0.005	0.006
Low Risk	18	0.8%	14.3%	0.001	0.003
Accruals	6	0.0%	0.0%	0.004	0.003

Table 1.4: This table shows (1) the number of factors in each theme, (2) how often an individual factor $\times$ trial results in a selection, (3) the percentage of models the theme is represented in, (4) the median KL divergence on iteration 1, and (5) the median KL divergence across all iterations.

dimensions and continue to clear the Wald threshold even after a representative is selected.

Debt issuance and quality are the most stable themes in the aggregated results, each entering in six of the seven trials. Their near-uniform selection despite very different within-theme structures is informative about the mechanism producing theme stability. Debt issuance enters with a high iteration-1 median KL of 0.061 that collapses to 0.010 when pooled across all iterations, a pattern consistent with a theme whose pricing content is concentrated in a first-order representative. A possible explanation follows from the external financing literature of Bradshaw et al. (2006) and Pontiff and Woodgate (2008): debt issuance aggregates information about managerial incentives, limits to arbitrage, and future operating performance into a small set of highly correlated signals, so that once one representative enters, the incremental KL from additional debt issuance factors falls sharply. The single trial in which debt issuance fails to enter is the ESS = 90% specification, under which quality is the only theme represented in the selected set.

Quality is the mirror image of debt issuance. It contains sixteen candidate factors and enters in six of seven trials, but its iteration-1 median KL of 0.009 is close to its all-iteration median of 0.005. The narrow gap suggests that quality is not dominated by a single informative representative. Several factors appear to capture distinct but economically related pricing dimensions, including profitability persistence, asset composition, and the earnings quality construction of Sloan (1996) as well as the composite quality measure of Asness et al. (2019). The single configuration in which quality fails to enter is $ESS = 60\%$. A possible explanation is that at this specific threshold, the admissible candidate set permits profit growth or profitability factors to attain the highest unconditional KL in an early iteration, absorbing pricing content that quality would otherwise have conveyed. Because quality shares substantial information with profitability and profit growth (Novy-Marx, 2013), the order of admission under $ESS = 60\%$ may foreclose quality through between-theme crowding rather than reflect any absolute weakness in its marginal contribution. This interpretation is consistent with the fact that quality re-enters at both lower and higher ESS thresholds, which corresponds to different admissible sets and therefore to different between-theme orderings.

Profit growth enters in five of seven trials, with an iteration-1 median KL that again substantially exceeds the all-iteration median. The pattern resembles debt issuance. The economic interpretation is that profit growth factors carry concentrated information about earnings dynamics, along the lines of Chan et al. (1996), with a single representative capturing most of the theme’s pricing content. Seasonality contains eleven factors with a median iteration-1 KL of 0.019 that falls to 0.009 in the all-iterations pool, consistent with a theme whose factors capture related but partially distinct seasonal patterns. The pattern is consistent with the evidence in Heston and Sadka (2008) and Keloharju et al. (2016) that monthly return seasonalities constitute a robust anomaly. Value enters with a small iteration-1-to-all-iterations gap, which suggests that the multiple book-to-market and earnings-yield variants in the theme carry overlapping but not identical pricing information.

Themes in the middle of the distribution include EQP and size, which each enter in three of seven trials. The contrast between them is informative. EQP contains a single factor, and its pricing content does not decrease conditional on other factors having been

selected. This suggests that the risk-signal in aggregate market returns is distinct from any other theme. Size contains five factors but enters with a very low median KL, consistent with the documented attenuation of the size premium since (Banz, 1981) and with the view in Asness et al. (2018) that size becomes informative primarily when interacted with other characteristics. Short-term reversal, profitability, momentum, investment, and low leverage all enter in two of seven trials, which indicates that under most configurations these themes are crowded out by earlier entrants with which they share pricing content.

Low risk and accruals rarely enter or not at all. Low risk enters in only one of seven trials despite containing eighteen candidate factors, and its iteration-1 median KL is 0.001. A possible explanation is that the eighteen candidates function as near-substitutes along a common beta and volatility dimension, so that whichever factor carries the theme's pricing information loses most of it conditional on the market-related representative that typically enters prior. The premia documented in Frazzini and Pedersen (2014) and Baker et al. (2011) load heavily on the first principal component of characteristic-sorted returns, which any broad market factor captures. A complementary explanation concerns the construction of the test-asset universe. Because the test assets are the first 73 principal components of 4,600+ decile-sorted portfolios, pricing content associated with narrow volatility anomalies may load on principal components beyond the truncation threshold, in which case a portion of the theme's pricing information is excluded from the test-asset representation by construction.

Accruals fail to enter in any of the seven trials despite the prominence of the accruals anomaly following Sloan (1996). A possible explanation is that accruals-based pricing information overlaps substantially with the quality theme, which enters in six of seven trials. The accruals anomaly is itself an earnings-quality signal, and the composite quality construction of Asness et al. (2019) explicitly incorporates accruals-related components, so that once a quality representative enters the conditional pricing content of the accruals factors can fall below the Wald threshold. An alternative and not mutually exclusive interpretation is that the accruals premium has attenuated in more recent samples, consistent with the post-publication decay documented in McLean and Pontiff (2015). Neither interpretation implies that accruals carry no pricing information. The observed absence reflects the specific

identifying variation available in the PCA-based test-asset representation and the presence of correlated themes in the candidate set.

Taken together, the pooled results support the claim that the set of themes represented in the selected model is broadly stable across ESS configurations. A small core (debt issuance, quality, profit growth) enters reliably, and a larger set of peripheral themes enters contingent on the admissible candidate set. Path dependence operates primarily through within-theme rotation among near-substitutes and, in the case of quality at $\text{ESS} = 60\%$, through specific between-theme ordering interactions. Neither channel overturns the conclusion that the KL selection algorithm identifies a stable economic content rather than an arbitrary subset of factors.

1.5.5 Stopping Criteria in Action

Now I examine what each stopping criterion is actually doing in this sample. The analysis relies on the distributions of incremental KL divergences and Wald statistics across all evaluated candidates, pooled over the seven ESS trials. The two criteria were introduced separately in Section 1.3 as distinct screens on the admissible factor set, and their empirical relationship in this sample is informative about which screen is doing the binding work.

Figure 1.2 plots the median incremental KL divergence by thematic category. The left panel aggregates across all evaluated candidates; the right panel restricts to the subset of candidates that pass the Wald feasibility filter. The dashed vertical line in each panel marks the entropy bound $\bar{\eta}$ calibrated to a detection error probability of $\varepsilon = 0.10$.

Two features of Figure 1.2 warrant comment. The first is the order-of-magnitude difference in scale between the two panels. In the left panel, theme-level median KL values range from roughly 0.002 to 0.012. In the right panel they range from roughly 0.017 to 0.06, approximately five times higher. The Wald filter removes from each theme a long left tail of candidates whose incremental KL is close to zero. This scale shift quantifies the volume of uninformative candidates in the factor zoo. In every theme the median unfiltered candidate has a KL contribution that is economically small, and only 61% of candidate themes survive the significance test. The pattern is consistent with the broader characterization of the fac-




Figure 1.2: This plot shows the median D_{KL} value for each theme at all iterations. The left panel aggregates across all candidates; the right panel subsets to only feasible candidates.

tor zoo in Harvey et al. (2015) as dominated by weakly identified signals, where appropriate multiple-testing adjustments eliminate a substantial share of documented factors.

The second feature is the reordering of themes between panels. In the left panel, momentum has the highest median incremental KL, followed by debt issuance, profit growth, seasonality, and short-term reversal. In the right panel debt issuance moves to the top with a median of approximately 0.06. Momentum drops from first to among the lowest medians, and low risk rises from the bottom of the left panel to an intermediate position in the right panel. The ratio of the feasible-set median to the unfiltered median captures the within-theme distribution of information content rather than its absolute level. A possible explanation for these patterns draws on the theme structures established in the discussion on path dependence. Debt issuance contains a few highly informative candidates concentrated in the right tail of its KL distribution, with a larger mass of near-zero candidates; the Wald filter removes the near-zero mass and shifts the median toward the informative representatives. This is consistent with the concentrated-signal interpretation of the external-financing literature of Bradshaw et al. (2006) and Table 1.4. Momentum, by contrast, contains candidates whose Wald statistics are relatively uniform and typically identifiable, so conditioning on feasibility produces only a modest multiplicative shift in the median. Low risk lies at

the extreme opposite of debt issuance: a high feasibility-conditional median combined with a very low unconditional median indicates that the theme contains a small subset of well-identified representatives embedded within a large mass of noise-dominated alternatives.

Across all candidate evaluations pooled over trials, no factor fails the entropy bound without also failing the Wald test. The entropy bound is therefore never the sole reason for exclusion in this sample. Two mechanisms account for the co-movement. First, the calibration at $\varepsilon = 0.10$ produces a relatively permissive bound. An investor who attributes 10% misclassification probability to the reference model attributes pricing error to omitted priced risk rather than model misspecification. Second, with $T = 666$ monthly observations, any factor with a KL contribution large enough to exceed $\bar{\eta}$ tends to produce a risk price estimate large enough to reject the null at $\alpha = 0.005$. The co-movement is a finite-sample regularity rather than a theoretical necessity. At substantially smaller ε , the bound tightens considerably, and an investor who attributes nearly all pricing error to misspecification would admit far fewer factors. In that region of the parameter space, the entropy bound would begin to bind independently of the Wald test.

Figure 1.3 presents the same thematic decomposition for Wald statistics. The left panel plots the median Wald statistic by theme for all evaluated candidates; the right panel restricts to feasible candidates only. The dashed vertical line marks the critical value $W_{\alpha=0.005} = 7.88$ for the one-degree-of-freedom chi-square test.

The most notable feature of the left panel is that every theme-level median falls well below the critical value. The highest medians are less than 60% of the threshold, and the lower tail of the distribution sits at or below 2. The interpretation is not that every individual factor in these themes fails significance. The median candidate in every theme produces a conditional risk price estimate that does not clear the multiple-testing-adjusted threshold. This further strengthens the conclusions from Figure 1.2 and Harvey et al. (2015) that the factor zoo is full of noise factors that are difficult to identify.

The right panel shows that once the Wald filter conditions on feasibility, debt issuance is distinctly separated from the remaining themes. Its median Wald statistic of approximately 29 is roughly 80% higher than the next-highest theme. The remaining themes cluster between




Figure 1.3: This plot shows the median Wald statistic for each theme at all iterations. The left panel aggregates across all candidates; the right panel subsets to only feasible candidates.

roughly 10 and 17, producing a relatively uniform distribution above the critical value. Again this points to debt issuance admitting a few highly identifiable factors whose pricing content concentrates along a single economic dimension, so that once the filter selects them, the feasible distribution retains unusually high t -statistics. For other themes the feasible set is more diffuse. Several candidates contribute moderate but broadly comparable pricing signal. This interpretation aligns with the iteration-1 versus all-iterations KL pattern documented in Table 1.4.

Comparing the two figures, the Wald statistic at $\alpha = 0.005$ separates included from excluded themes at least as cleanly as the entropy bound, and for the peripheral themes the Wald test rules out candidates well before the entropy bound is approached. This reflects a general property of short time-series data. At $T = 666$, the Wald test is powerful enough to detect any factor whose KL contribution is economically material but the factors sample KL divergence may exceed its population level, leaving the entropy bound as a redundant screen in the interior of the parameter space. The entropy bound retains its economic interpretation. However, at these sample lengths, the Wald test is the operative filter. At shorter longer lengths, or at substantially tighter ε values, the two criteria would be expected to diverge, and the entropy bound would begin to bind for candidates whose Wald statistics

clear the significance threshold but whose KL contributions are too small to survive the ambiguity-aversion adjustment.

1.5.6 Wald v.s. ESS

The ESS and Wald criteria are not independent. Both guard against overfitting, but through different channels, and their interaction determines which marginal factors survive.

The Wald test screens factors whose risk price estimates are imprecise relative to sampling variation. The ESS criterion screens factors whose inclusion would concentrate the tilted probability measure on too few observations. A tighter Wald threshold may partially substitute for a tighter ESS threshold. Factors that trigger the ESS criterion often do so because their pricing content is driven by a few extreme months, and those same factors tend to have imprecise risk prices once attention is restricted to the remainder of the sample. The converse also holds. A tighter ESS threshold may reject factors that pass the Wald test but whose risk prices are supported by concentrated observation weights. The two criteria therefore operate as partial substitutes within the interior of the parameter space, and the question of interest is the degree of that substitution.

Table 1.5 reports a line search holding $\alpha = 0.005$ and $\varepsilon = 0.10$ fixed and varying the ESS threshold from $\text{ESS}/T = 0.10$ to $\text{ESS}/T = 0.90$. The two lowest thresholds illustrate why an effective-sample-size floor is necessary in the first place. At $\text{ESS}/T = 0.10$, the selected SDF implies an annualized Sharpe ratio of 19.3, with an individual factor contributing a Sharpe ratio of 7.42 along the selection path. At $\text{ESS}/T = 0.20$, the implied Sharpe ratio is 10.25, with a within-path maximum of 5.83. Sharpe ratios of this magnitude are roughly an order of magnitude larger than the robust long-sample estimates of aggregate market Sharpe ratios in Fama and French (2002). The interpretation is that at $\text{ESS}/T < 0.30$ the algorithm is fitting a handful of extreme observations rather than identifying pervasive cross-sectional pricing. I therefore treat $\text{ESS}/T = 0.30$ as a hard lower bound throughout. The line search results for $\text{ESS}/T < 0.30$ are reported for completeness and separated from the main specifications by a horizontal rule.

Table 1.5 also documents the monotone decline of the implied Sharpe ratio as the ESS

Table 1.5: Results of the ESS Line Search

ESS/ T (1)	# Factors (2)	# Themes (3)	Baseline themes (4)	$\sum D_{KL}$ (5)	SR (6)	Max SR (7)	Act. ESS/ T (8)
10%	10	8	78%	15.519	19.30	7.42	13.2%
20%	14	10	89%	4.381	10.25	5.83	40.2%
30%	14	9	100%	2.268	7.38	2.63	46.6%
40%	10	9	78%	0.901	4.65	2.48	43.3%
50%	9	7	67%	0.625	3.87	2.08	53.3%
60%	9	7	56%	0.538	3.59	2.08	60.7%
70%	5	4	33%	0.394	3.08	1.80	70.1%
80%	6	5	44%	0.269	2.54	1.40	80.1%
90%	1	1	11%	0.051	1.11	1.11	90.3%

Table 1.5: This table shows a line search of different minimum ESS thresholds from 10% to 90% of T . Column (1) shows the minimum ESS threshold, (2) and (3) show the number of factors selected and themes represented in the final model, respectively, (4) shows the percentage of themes represented in the baseline model at ESS = 30%, (5) and (6) are the cumulative KL divergence and implied annualized Sharpe ratio, respectively, (7) is the maximum Sharpe ratio achieved by any included factor, and (8) is the ESS of the final model. The Wald significance and error detection probability are set to the baseline level of $\alpha = 0.005$ and $\epsilon = 0.1$, respectively.

floor tightens. The SDF at ESS/ $T = 0.30$ achieves an annualized Sharpe ratio of 7.38 with nine themes represented. At ESS/ $T = 0.90$, the SDF retains a single factor with an annualized Sharpe ratio of 1.11, which aligns with consensus estimates of the efficient-frontier Sharpe ratio under broadly tradable risk factors. The monotone relationship between the ESS threshold and the implied Sharpe ratio is consistent with the interpretation of ESS as a sample-breadth filter rather than a convergence criterion. Tightening the filter rejects factors whose pricing content is concentrated, and the rejected factors contribute both to the cumulative KL and to the tails of the SDF.

To identify which themes generate these concentrated violations, I examine the distribution of ESS outcomes at the candidate-evaluation level across the full grid of trials. Table 1.6 reports, for each theme, the fraction of candidate evaluations that fail the ESS criterion along with the mean realized KL divergence for the failing and passing subsets.

Three themes exhibit the highest rates of ESS violation. Seasonality leads the distribution, failing in 35.4% of candidate evaluations, followed by momentum at 33.8% and profit

Table 1.6: ESS Violations by Theme

Theme	Fail (1)	# Fail (2)	ESS Fail (3)	ESS Pass (4)	$\bar{D}_{KL}$ Fail (5)	$\bar{D}_{KL}$ Pass (6)
Seasonality	35.4%	226	39.7%	61.3%	0.498	0.018
Momentum	33.8%	162	49.3%	61.9%	0.279	0.017
Profit Growth	30.6%	162	47.2%	61.8%	0.298	0.018
Short-Term Reversal	27.3%	97	48.9%	61.8%	0.154	0.019
Value	26.8%	220	54.4%	60.5%	0.045	0.018
Low Risk	24.9%	272	56.0%	62.6%	0.046	0.011
Low Leverage	24.7%	149	53.9%	61.7%	0.037	0.012
Debt Issuance	23.9%	90	55.9%	59.8%	0.093	0.025
FF	23.6%	82	49.8%	63.5%	0.156	0.017
Investment	23.5%	312	50.0%	63.7%	0.050	0.013
Profitability	22.4%	133	58.7%	62.2%	0.044	0.014
Size	20.7%	62	51.3%	64.0%	0.039	0.013
Quality	16.2%	153	58.3%	63.3%	0.036	0.013
Accruals	14.3%	52	50.5%	63.9%	0.126	0.009

Table 1.6: This table reports, for each theme, (1) the fraction of candidate evaluations that fail the ESS criterion, (2) the number of failing evaluations, (3) the mean realized ESS/ T conditional on failure, (4) the mean realized ESS/ T conditional on passing, (5) the mean realized D_{KL} conditional on ESS failure, and (6) the mean realized D_{KL} conditional on ESS passing. Evaluations are pooled across the seven trials and all iterations.

growth at 30.6%. The remaining themes cluster between 14% and 28%. The pattern is not limited to violation frequency. Conditional on failure, seasonality candidates produce a mean $\bar{D}_{KL}$ of 0.498, roughly double the corresponding values for momentum and profit growth. The latter two are indistinguishable from each other and substantially larger than the failure-conditional means of the remaining themes. Conditional on passing, all three themes produce $\bar{D}_{KL}$ values that are indistinguishable from the cross-theme average. The failing-to-passing ratio is between 16 and 28 for these themes, an order of magnitude larger than the ratios observed in the remaining themes. This two-regime behavior is the empirical signature of overfitting on concentrated observations. When a seasonality, momentum, or profit growth factor clears the ESS filter, it behaves as an ordinary well-identified factor. When it fails, it does so by producing an apparent pricing contribution that traces to a small subset of months.

Economic intuition supports this division. Seasonality factors, as constructed in Heston and Sadka (2008) and Keloharju et al. (2016), exploit same-calendar-month return patterns whose informational content in any given month is supported only by the subset of assets in phase with the signal. When the full history of same-calendar-month observations contains a few extreme periods, the tilted measure concentrates on those periods and the effective sample size collapses. The mechanical narrowness of this conditioning set accounts for seasonality’s failing-conditional $\bar{D}_{KL}$ being roughly double that of the other two themes. Momentum factors exhibit an analogous concentration through momentum crashes, documented in Daniel and Moskowitz (2016) and Barroso and Santa-Clara (2015), during which cumulative returns drop sharply within one or two months following market rebounds. Profit growth operates through the same mechanism in a different guise. Novy-Marx (2015) argues that price momentum is largely a manifestation of earnings momentum, which rationalizes the near-identical ESS violation patterns of the two themes in Table 1.6. All three themes load on signals that refresh at discrete intervals and produce outsized spreads in a small subset of observations, precisely the configuration that causes the tilted measure to collapse onto a narrow support.

The policy implication for the selection algorithm is that seasonality, momentum, and profit growth factors should not be excluded on principle. All three themes enter the baseline model at $\text{ESS}/T = 0.30$ through representatives that pass the ESS filter, and when they pass they behave as the rest of the factor zoo. The ESS filter is required precisely because a substantial fraction of the candidates in these themes would otherwise enter on the strength of concentrated-observation fits. Table 1.6 therefore supplies a quantitative rationale for the qualitative argument that the ESS criterion is not redundant with the Wald test. The two criteria target partially overlapping but distinguishable failure modes, and the themes most affected by ESS screening are precisely those whose economic construction makes concentrated-observation fits most likely.

1.5.7 Higher-Order-Moments

The KL divergence between the MCE-tilted measure and the physical measure decomposes into a Gaussian component, which captures the mean-variance pricing content a standard SDF would price, and a residual higher-order component, which captures skewness, kurtosis, and other non-Gaussian structure in the factor's pricing information.

For an exponential-family tilt, the decomposition takes the form

$$D_{KL}(\mathcal{Q}^{\mathbb{S}}||\mathcal{P}) = D_{KL}^{\text{Gauss}} + D_{KL}^{\text{HOM}}, \quad (1.24)$$

where $D_{KL}^{\text{Gauss}} = \frac{1}{2} \frac{\hat{\theta}^2}{\sigma_f^2}$ is the KL divergence that would obtain if the factor were normally distributed with the same mean and variance, and $D_{KL}^{\text{HOM}} \geq 0$ is the residual attributable to the higher-order cumulants of the factor's return distribution. The Gaussian component is fully summarized by the factor's Sharpe ratio and coincides with the pricing content available to a mean-variance investor; the higher-order component reflects the additional information an ambiguity-averse investor extracts from asymmetry and tail events. The non-negativity of the residual follows from the convexity of the cumulant generating function.

As in the ESS analysis, the decomposition is cleanest at iteration 1. Once the algorithm has selected multiple factors, the tilt $\mathcal{Q}^{\mathbb{S}}$ is a joint function of all k included factors, and the marginal higher-order contribution of the k -th factor is confounded with the higher-order structure of the preceding tilts. At iteration 1 the mapping is one-to-one: the decomposition reflects the higher-order structure of a single factor's pricing information, unconditional on any prior selections. I report theme-level averages of the two components, averaging across all candidate factors within a theme at iteration 1 across the ESS trials.

Figure 1.4 plots the average Gaussian component against the average total KL divergence, with one point per thematic category. The 45° line is the locus of equal Gaussian and total KL divergence. Points above the line have a higher-order component that exceeds the Gaussian component, and points on the line have approximately balanced components. A point lying far above the line indicates that the factor's cross-sectional pricing power derives substantially from non-Gaussian features rather than from mean-variance compensation.

The most striking feature of the figure is that higher-order dominance is the modal

Figure 1.4: Higher-Order Moment Structure




Figure 1.4: This plot shows the average Gaussian (x-axis) and KL Divergence (y-axis) for each theme at iteration one of the algorithm. Evaluations are pooled across the seven ESS trials restricted to the subset of feasible candidates.

pattern. The majority of themes sit above the 45° line, and no theme sits materially below it once sampling noise is taken into account. The cross-section of factor pricing information is therefore not well summarized by Sharpe ratios alone, consistent with the cross-sectional evidence on non-Gaussian features in characteristic-sorted portfolios (Harvey and Siddique, 2000; Dittmar, 2002) and with the structural claim in Section 1.3 that the robust investor’s concern about model misspecification maps into sensitivity to higher cumulants of the factor return distribution.

Among the selected themes, profitability, profit growth, value, and the equity premium all sit clearly above the line, so their inclusion in the SDF is driven in part by pricing information that a mean-variance investor would not register. Quality, by contrast, sits very close to the 45° line, indicating that its pricing content is nearly Gaussian. Its selection reflects mean-variance compensation that a standard SDF would also capture, consistent with the characterization of the quality premium as compensation for persistent firm fundamentals rather than for tail risk (Asness et al., 2019; Novy-Marx, 2013).

Profit growth is the most extreme higher-order case. It carries the largest total KL divergence in the sample, paired with only a middling Gaussian component, so its pricing content is dominated by non-Gaussian structure. The economic interpretation is the earnings-news channel: expectations errors and post-announcement drift concentrate returns in episodic, skewed events tied to earnings surprises, and the cross-section of profit-growth-sorted portfolios inherits this non-Gaussian structure (Chan et al., 1996; Bouchaud et al., 2019). Debt issuance sits at the opposite end of the selected themes. Its Gaussian and higher-order components are both large, placing it high on both axes and yielding the second-largest total KL divergence. Issuance activity generates a sizeable mean-variance premium and a sizeable tail component simultaneously, reflecting both a persistent cross-sectional spread in expected returns and the episodic, conditionally skewed corporate events that drive it (Pontiff and Woodgate, 2008).

Profitability and value occupy intermediate positions above the line. Profitability’s higher-order content is consistent with the flight-to-safety behaviour of profitable firms in bad states, and value’s is consistent with the episodic negative skewness of value returns doc-

umented in the value-crash literature (Daniel and Moskowitz, 2016). The equity premium sits noticeably above the line despite a middling Sharpe ratio. The Gaussian component is modest because the Sharpe ratio is modest, but the higher-order component is non-trivial, consistent with the large literature documenting fat tails and rare-disaster exposure in aggregate market returns (Barro, 2006; Gabaix, 2012). The robust investor extracts pricing information from this tail structure that a Sharpe-ratio-based selection criterion would miss.

Accruals and size sit on or very close to the 45° line. In population the residual is non-negative, so the natural interpretation of these points is that the higher-order component is effectively zero and the pricing information these themes do carry is primarily Gaussian.⁷

The higher-order structure documented here connects directly to the ESS line search results. Factors whose KL divergence is higher-order dominated achieve their information gain through non-Gaussian concentration of the probability tilt, which is precisely the concentration the ESS criterion penalizes. Themes lying far above the 45° line in Figure 1.4 correspond to the contingent themes in Table 1.5 that drop out first as the ESS threshold is tightened.

1.5.8 Summary

The greedy KL-maximization algorithm selects a parsimonious stochastic discount factor spanning roughly nine thematic categories drawn from the candidate universe. Three themes constitute a stable core. Debt issuance and quality each enter in six of seven ESS trials, and profit growth enters in five. Seasonality, value, the equity premium, and size form a middle tier that enters in three to four trials. Short-term reversal, momentum, profitability, investment, low leverage, and low risk enter in one or two trials, and accruals is absent from every specification despite containing six candidate factors. The set of themes represented is stable across the interior of the parameter space, and the peripheral themes are excluded on statistical or economic rather than on overfitting grounds.

⁷Equation (1.24) is an exact population identity, but the empirical decomposition is a plug-in estimator with two sources of sampling error. Points whose distance below the 45° line in Figure 1.4 is small relative to the cross-sectional spread in KL divergences should therefore be read as approximately Gaussian with sampling noise in the residual.

Theme-level stability coexists with substantial factor-level variation. Of the thirty-five factors selected in at least one trial, only one—`cop_at`, the quality theme’s cash-profits-to-book representative—is selected in more than four trials. The remaining themes distribute their selections across multiple representatives. Profit growth enters through six distinct factors across seven trials, seasonality through four, and value through four. This within-theme rotation reflects the greedy algorithm’s treatment of near-substitute candidates. Between-theme ordering is more consequential and more stable. The dominant themes enter in approximately the same sequence across all trials, which indicates that the algorithm recovers a consistent ranking of the cross-sectional pricing structure even when the within-theme representative rotates.

The entropy bound does not bind singularly at the baseline calibration $\varepsilon = 0.10$, and its role in the final specification is interpretive rather than operative. The Wald test and the ESS criterion jointly govern admissibility, and they operate as partial substitutes within the interior of the parameter space. The ESS floor at $\text{ESS}/T = 0.30$ is imposed to rule out specifications whose implied Sharpe ratios are inconsistent with economically defensible risk pricing. Below this threshold the algorithm fits concentrated-observation anomalies rather than pervasive cross-sectional variation, and the tilted probability measure places disproportionate weight on a small subset of months. The themes most exposed to this failure mode are seasonality, momentum, and profit growth, whose economic construction makes concentrated-observation fits mechanically likely.

Table 1.7 provides the complete factor-level selection roster. For each factor selected in at least one ESS trial, the table reports its mnemonic, full name, thematic classification, the number of trials in which it was selected, the corresponding selection frequency, the mean iteration at which the factor entered the selection path, the mean incremental KL divergence at selection, the implied annualized Sharpe ratio of the selected factor, and the primary reference documenting its empirical construction.

Table 1.7: Final Selected Factor Roster

Factor	Full name	Trials Sel. (1)	Freq. (2)	Iter. (3)	$\overline{\Delta D}_{KL}$ (4)	SR (5)	Citation
<i>Debt Issuance</i>							
nfna_gr1a	Δ Net Financial Assets	4	57%	1.5	0.172	2.03	Richardson et al. (2005)
fnl_gr1a	Δ Financial Liabilities	2	29%	3.5	0.086	1.37	Richardson et al. (2005)
capex_abn	Abnormal Corporate Investment	1	14%	3.0	0.040	0.98	Titman et al. (2003)
<i>Quality</i>							
cop_at	Cash Operating Profits-to-Book	6	86%	3.7	0.068	1.23	Ball et al. (2016)
at_turnover	Capital Turnover	1	14%	12.0	0.200	2.19	Haugen and Baker (1996)
sale_bev	Asset Turnover	1	14%	5.0	0.033	0.88	Soliman (2008)
<i>Profit Growth</i>							
dsale_drec	Δ Sales - Δ Receivables	3	43%	1.7	0.073	1.33	Abarbanell and Bushee (1997)
dsale_dinv	Δ Sales - Δ Inventory	2	29%	1.0	0.256	2.48	Abarbanell and Bushee (1997)
ocf_at_chg1	Δ Operating Chasflows-to-Assets	1	14%	3.0	0.067	1.27	Bouchaud et al. (2019)
sale_emp_gr1	Labor Force Efficiency	1	14%	4.0	0.050	1.09	Abarbanell and Bushee (1997)
saleq_su	Standardized Revenue Surprise	1	14%	4.0	0.082	1.40	Jegadeesh and Livnat (2006)
tax_gr1a	Tax Expense Surprise	1	14%	7.0	0.070	1.30	Thomas and Zhang (2011)
<i>Seasonality</i>							
dbnetis_at	Net Debt Issuance	1	14%	1.0	0.133	1.78	Bradshaw et al. (2006)
seas_11_15an	Years 11–15 Lagged Ret.	1	14%	3.0	0.132	1.78	Heston and Sadka (2008)
seas_16_20an	Years 16–20 Lagged Ret.	1	14%	5.0	0.046	1.05	Heston and Sadka (2008)
seas_2_5an	Years 2–5 Lagged Ret.	1	14%	6.0	0.068	1.28	Heston and Sadka (2008)
<i>Value</i>							
sale_me	Sales-to-Market	3	43%	6.0	0.117	1.53	Barbee et al. (1996)
ival_me	Intrinsic Value-to-Market	2	29%	8.5	0.044	1.01	Frankel and Lee (1998)
fcf_me	Free Cashflow-to-Market	1	14%	5.0	0.024	0.76	Lakonishok et al. (1994)
ni_me	Earnings-to-Price	1	14%	11.0	0.288	2.63	Basu (1983)
<i>Size</i>							
ami_126d	Amihud Measure	1	14%	13.0	0.214	2.26	Amihud (2002)
dolvol_126d	Dollar Trading Volume	1	14%	4.0	0.032	0.88	Brennan et al. (1998)
market_equity	Market Equity	1	14%	9.0	0.031	0.86	Banz (1981)
<i>Equity Premium</i>							
EQP	Equity Premium	3	43%	6.0	0.062	1.16	Sharpe (1964)
<i>Investment</i>							
ncoa_gr1a	Δ Noncurrent Operating Assets	1	14%	6.0	0.022	0.73	Richardson et al. (2005)
nncoa_gr1a	Δ Net Noncurrent Operating Assets	1	14%	6.0	0.063	1.23	Richardson et al. (2005)
<i>Low Leverage</i>							
aliq_mat	Liquidity of Market Assets	2	29%	6.0	0.084	1.40	Ortiz-Molina and Phillips (2014)
<i>Profitability</i>							
ocf_at	Operating Cashflow-to-Assets	2	29%	7.5	0.120	1.55	Bouchaud et al. (2019)
dolvol_var_126d	Variation in Dollar Trading Volume	1	14%	7.0	0.037	0.94	Chordia et al. (2001)
o_score	Ohlson O-Score	1	14%	14.0	0.044	1.03	Dichev (1998)
<i>Momentum</i>							
prc_highprc_252d	Current Price-to-Last Year High Price	1	14%	10.0	0.034	0.90	George and Hwang (2004)
resff3_12_1	6 Month Residual Momentum of FF3	1	14%	3.0	0.031	0.87	Blitz et al. (2011)
<i>Short-Term Rev.</i>							
ret_1_0	Short-Term Reversal	1	14%	9.0	0.145	1.87	Jegadeesh (1990)
rskew_21d	Total Skewness	1	14%	8.0	0.055	1.15	Bali et al. (2017)
<i>Low Risk</i>							
beta_dimson_21d	Dimson Beta	1	14%	7.0	0.076	1.35	Dimson (1979)

Table 1.7: This table reports, for each factor selected in at least one ESS trial, (1) the number of trials in which the factor enters the model, (2) the selection frequency as a percentage of all 7 trials, (3) the average iteration of entry (averaged only over trials in which the factor is selected), (4) the average incremental KL divergence ΔD_{KL} at the iteration of entry, and (5) the average annualized implied Sharpe ratio at entry. Trials span $ESS/T \in \{0.3, \dots, 0.9\}$ with $\alpha = 0.005$ and $\varepsilon = 0.1$.

1.6 Conclusion

In this paper, I develop an information-geometric framework for selecting priced risk factors from a high-dimensional candidate set. The stochastic discount factor is cast as a point on a statistical manifold of exponential-affine pricing kernels, and the manifold’s dual flatness admits an exact Pythagorean decomposition of the Kullback-Leibler divergence between nested models. This decomposition supplies the ranking criterion for a greedy forward-selection algorithm, and the Hansen-Jagannathan distance emerges as the leading term of a curvature expansion of the KL divergence, coinciding with the geometric criterion locally and departing from it in directions governed by the higher cumulants of the factor return distribution. I provide an economic foundation for the procedure through a representative investor with multiplier preferences in the sense of Hansen and Sargent (2001, 2020). The investor’s worst-case distortion is an exponential tilt that lands on the same SDF manifold traversed by the geometric algorithm, and the investor’s ranking of candidate factors by worst-case welfare loss coincides with the purely statistical ranking by KL divergence. The entropy of the worst-case distortion provides an economic stopping threshold that, combined with a Wald test on the estimated risk price, produces a complete inclusion rule with separate statistical and economic failure modes.

On the estimation side, I implement the full nested Minimum Cross-Entropy optimization over both the Lagrange multipliers and the exponential-affine risk prices under the pricing restrictions. To my knowledge, this is the first implementation of the Kitamura and Stutzer (1997) estimator that solves the complete problem rather than a restricted subproblem. The dual architecture enforces SDF positivity by construction, penalizes concentration of probability mass on extreme observations, and delivers an estimator whose finite-sample behavior reflects the theoretical properties that motivate the exponential-affine specification.

Applying the methodology to 73 principal components of 4,686 characteristic-sorted portfolios against a candidate factor zoo of 157 factors yields several findings. First, the algorithm selects a parsimonious SDF spanning approximately nine thematic categories. Debt issuance, quality, and profit growth form a stable core that enters in six or more of seven parameter configurations; seasonality, value, the equity premium, and size form a middle

tier; and short-term reversal, momentum, profitability, investment, low leverage, and low risk enter contingent on the admissible set. Second, accruals is absent from every specification despite containing six candidate factors, a pattern attributable to overlap with the quality theme. Third, path dependence operates primarily within themes through rotation among near-substitute representatives rather than across themes, indicating that the algorithm recovers a stable economic content rather than an arbitrary subset of factors. Fourth, the selected factors derive a substantial share of their pricing content from non-Gaussian features of the return distribution, consistent with the robust investor's concern about model misspecification mapping into sensitivity to skewness and tail exposure.

Taken together, these results suggest that the factor zoo can be organized by a principled criterion that integrates statistical and economic content rather than by a significance threshold alone. The information-geometric framework supplies the ranking, the robust-control foundation supplies the stopping rule, and the two converge on the same equilibrium selection of priced risks.

Chapter 2

Selling Distress: The Economics of Failed Bank Resolutions

2.1 Introduction

Bank failures are rare, consequential, and poorly understood. When a bank fails, the Federal Deposit Insurance Corporation (FDIC) must rapidly resolve the institution, allocating losses among depositors, the insurance fund, and private acquirers. While the aggregate fiscal costs of these failures are well documented, the specific transaction mechanisms have largely remained opaque to researchers. This paper introduces a novel dataset of FDIC resolution terms constructed from over forty years of administrative press releases, enabling the first systematic analysis of how resolution mechanisms vary across failure types, economic conditions, and acquiring bank characteristics. The data reveal that resolution structure is far from mechanical: acquirer pricing, loss-sharing arrangements, and deposit franchise valuation reflect deliberate economic choices whose determinants this paper identifies and quantifies.

The stakes of efficient resolution are high. During the Savings and Loan (S&L) crisis of the 1980s, the collapse of the thrift industry generated resolution costs in excess of 150 billion dollars. More recently, the Global Financial Crisis and the regional banking failures of 2023 demonstrated that even a relatively small number of failures can impose significant

stress on the Deposit Insurance Fund (DIF). Despite the magnitude of these events, the existing academic literature faces a significant data limitation. Most research relies either on aggregate cost estimates or on privileged internal data that is unavailable to the broader scientific community. Standard public regulatory filings such as Call Reports provide a detailed view of the bank prior to failure, but the data trail typically ends the moment the charter is revoked. This leaves researchers unable to observe the specific premiums paid by acquirers, the precise volume of assets retained by the receiver, or the structure of loss-sharing agreements that define the modern resolution.

This paper bridges this gap by introducing a hand-collected dataset constructed from the universe of FDIC press releases spanning the period from 1983 to 2024. By systematically encoding over forty years of resolution announcements, I transform qualitative legal disclosures into a structured economic dataset containing 1,860 observations. This data collection effort quantifies critical resolution terms that are absent from standard regulatory reports. These variables include the premia paid for deposit franchises, the magnitude of assets retained by the FDIC in receivership, and the precise structure of loss-sharing agreements. This granular view allows for the decomposition of resolution costs and provides a new lens through which to evaluate the determinants of franchise value in failed bank auctions.

The institutional framework governing these resolutions has evolved significantly over the sample period. The passage of the Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA) marked a pivotal shift in the regulatory landscape. Designed to correct the incentive problems and regulatory forbearance that exacerbated the S&L crisis, FDICIA mandated a least-cost resolution test. This requirement forces the FDIC to select the resolution method that imposes the minimum expense on the insurance fund barring a systemic risk exception. Consequently, the resolution mix shifted dramatically from a landscape where forbearance was common to one dominated by Purchase and Assumption (P&A) transactions. In a P&A transaction, a healthy acquiring institution purchases some or all of the failed bank's assets and assumes its deposit liabilities. This method is generally preferred as it preserves the franchise value of the failed institution and maintains the continuity of banking services for the local community.

The historical arc of these failures is visualized in Figure 2.1 which plots the frequency and resolution method of every bank failure in the sample. The data reveals three distinct epochs of distress: the S&L crisis of the late 1980s, the GFC of 2008, and the concentrated regional banking shock of 2023. While the sheer volume of failures in the earlier period dwarfs modern episodes, the visual composition of the bars highlights a profound shift in regulatory technology. The 1980s exhibit a fragmented approach where the FDIC frequently relied on direct Deposit Payoffs and Insured Deposit Transfers to close small, chaotic institutions. By contrast, the modern era is characterized by the monolithic dominance of the P&A transaction.




Figure 2.1: Annual bank failures and resolution mix, 1983–2024. The top panel shows the count of failures with shaded bands for the Savings & Loan crisis (blue), Global Financial Crisis (pink), and 2023 regional bank crisis (tan). The bottom panel gives the share of each FDIC resolution type (PA, IDT, PO, PI, DINB) by year.

This uniformity in the later years is deceptive. While Figure 2.1 suggests that the resolution process has become standardized, the underlying transaction structures have actually become more complex. The disappearance of the deposit payoff is not merely a function of

bidder interest but a result of financial engineering. As the empirical results of this paper demonstrate, the ubiquity of the P&A transaction in the post-2008 period is largely sustained by the FDIC's increasing willingness to retain toxic assets. In the 1980s, a bank with a poor loan portfolio might have been liquidated. Today, that same institution is sold via a P&A transaction where the acquirer purchases the franchise but leaves the distressed assets behind with the receiver. Thus, the clean blue bars of the modern era mask a significant transfer of risk that can only be observed by analyzing the specific asset retention terms.

To interpret the variation in the pricing of these transactions, I develop a structural model of the FDIC auction process that includes both the asset and liability side of the failed bank's balance sheet. Each bidder's marginal value for the failed bank's assets declines in the share acquired, reflecting portfolio heterogeneity, regulatory capital constraints, and the diminishing returns of operational absorption. Bidders optimally select an asset take-up share by equating their marginal asset valuation to the FDIC's own expected recovery value, partitioning the portfolio into a "transferable" tranche managed more efficiently by the private sector and a "retained" tranche that the FDIC works out in receivership. On the liability side, the deposit franchise is a wasting asset where value is driven by the stability and interest-rate sensitivity of the deposit base. The FDIC's statutory least-cost mandate induces a linear scoring rule that compresses the three-dimensional bid—deposit premium, asset purchase price, and asset take-up share—into a single scalar, preserving the full tractability of first-price auction theory. The model's central prediction is that the retained asset share and the deposit premium are positively related through two reinforcing channels. A mechanical composition effect, whereby a lower take-up share shifts the bid toward the premium line, and an equilibrium private-value effect, whereby shedding negative-surplus assets raises the winning bidder's willingness to pay for the cleansed franchise.

The empirical analysis utilizes the new dataset to test these theoretical predictions through three distinct channels. The first is the deposit pricing channel, which examines the bid premium paid for the deposit franchise conditional on a successful sale. The second is the FDIC cost channel, which analyzes the fiscal expense absorbed by the Deposit Insurance Fund. The third is the P&A selection channel, which models the probability that a

bank is resolved via purchase and assumption rather than a direct deposit payoff.

Consistent with the theoretical framework, I find evidence that the stability of the deposit base matters. Banks that experience accelerated deposit runoff in the quarter prior to failure impose significantly higher costs on the FDIC. This finding supports the view that the destruction of franchise value prior to closure directly translates into fiscal losses. Furthermore, the analysis of deposit betas suggests a negative relationship between interest rate sensitivity and bid premiums, although this result is statistically noisier than the runoff effect. This implies that while acquirers theoretically value low-beta deposits, the chaotic environment of a failed bank auction may obscure this pricing signal relative to more immediate concerns about asset quality.

The headline finding across all three predictions is that the proportion of the failed bank's assets retained by the FDIC is the dominant quantitative lever. Because the retained asset share and the deposit premium are jointly determined at the auction, I instrument the realized toxic ratio with the pre-failure loan-loss provision (LLP) ratio, a balance-sheet measure determined before any bidding occurs. The instrumented toxic ratio raises bid premiums in the pricing channel, raises the probability of a P&A sale in the selection channel, and the raw provision ratio is the dominant driver of FDIC costs. A one percentage point increase in the loan-loss provision ratio raises the cost to the Deposit Insurance Fund by approximately 86 basis points of total assets, making asset quality the single most economically significant determinant of fiscal exposure. Consistent with the theoretical framework, accelerated deposit runoff and high interest-rate sensitivity reduce bid premiums and increase resolution costs in the directions the model predicts, though with more attenuated statistical significance relative to the asset-quality channel. The most robust selection predictor is log asset size: larger institutions attract P&A acquirers at rates disproportionate to their deposit-franchise quality, a pattern consistent with the Too-Big-to-Fail dynamic embedded in the resolution process.

This paper contributes to several strands of the banking literature. First, it extends the early work on resolution costs by Bovenzi and Murton (1988) and James (1991) by providing a more granular decomposition of the cost structure using data that covers the post-GFC and

2023 crisis periods. While those studies established the determinants of aggregate losses, this paper isolates the specific pricing of the deposit franchise and the cost of asset retention. Second, it complements the recent literature on fire sales and bidder behavior by Granja et al. (2017). While their work focuses on the geographic and portfolio fit of acquirers, this paper focuses on the explicit contract terms and the allocation of risk between the state and the private sector. Finally, it contributes to the literature on deposit franchise value and monetary policy transmission. By linking deposit betas to realized auction premiums, I provide a market-based validation of the theory that core deposits represent a valuable intangible asset even in the context of distress.

The construction of this dataset serves a broader purpose beyond the specific empirical tests presented here. By making these transaction-level details public, this paper opens the door for future research into the political economy of bank resolutions. The variation in loss-sharing agreements, the timing of closures, and the treatment of uninsured depositors across different regulatory regimes can now be scrutinized with greater precision. The data reveals a stark evolution in the protection of uninsured depositors where nearly all uninsured depositors were protected during the systemic events of 2008 and 2023 while non-crisis periods show a much higher incidence of losses. This cyclical in depositor protection has profound implications for market discipline and moral hazard which are themes that are central to the ongoing debate over deposit insurance reform and the “Too-Big-To-Fail” doctrine.

The remainder of the paper is organized as follows. Section 2.2 provides the institutional background on the FDIC resolution process and reviews the relevant literature. Section 2.3 details the data collection methodology and describes the construction of the novel variables. Section 2.4 presents the theoretical model of the failed bank auction. Section 2.5 reports the empirical results regarding the pricing, cost, and selection channels. Section 2.6 concludes.

2.2 Institutional Details

This section provides an overview of the Federal Deposit Insurance Corporation’s (FDIC) bank resolution process, including the typical timeline and the methods used to resolve failed banks. It begins with an introduction to how resolutions unfold under the Federal Deposit

Insurance Corporation Improvement Act of 1991 (FDICIA), outlining the key steps from regulatory intervention through closure and bidding. It then describes in detail each major resolution method employed by the FDIC: Purchase and Assumption transactions (P&As), Insured Deposit Transfers (IDTs), Insured Deposit Payoffs (POs), Purchase of Insured deposits (PIDs), and the use of a Deposit Insurance National Bank (DINB); explaining what each method entails, how it works, the cost implications for the FDIC and acquiring institutions, and the treatment of insured versus uninsured depositors. The section concludes with a brief literature review of empirical studies of United States' bank failures and resolution costs, highlighting findings from academic research and FDIC analyzes that provide context on the outcomes and costs associated with different resolution strategies.

2.2.1 S&L Crisis and the FDICIA

The Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA) 102nd Congress of the United States (1991) is best understood as a legislative response to the intertwined thrift and commercial bank crises of the 1980s, and in particular to the policy failures revealed by the savings and loan (S&L) debacle. The S&L crisis exposed the fragility of a regulatory regime that combined generous, underpriced deposit insurance with rapid deregulation, weak supervision, and extensive regulatory forbearance. FDICIA's core provisions: prompt corrective action, least-cost resolution, and risk-based insurance premiums were explicitly designed to correct those incentive problems and to limit the fiscal and political costs of future failures Benston and Kaufman (1997).

A brief review of the S&L crisis highlights why such reforms became politically unavoidable. In 1980, roughly 4,000 federally insured S&Ls held portfolios concentrated in long-term, fixed-rate mortgages that became unsustainable when short-term interest rates spiked¹, leav-

¹A key macroeconomic backdrop to the S&L crisis was the sharp rise in short-term interest rates associated with the Federal Reserve's disinflation campaign under Chair Paul Volcker. In October 1979, the Fed shifted its operating procedure to target non-borrowed reserves to bring down double-digit inflation; the resulting monetary tightening pushed the effective federal funds rate from roughly 10–11% in mid-1979 to peaks above 17% in late 1979-early 1980, and again to nearly 19–20% in 1981. From 1980 to 1982, the funds rate averaged in the low to mid-teens, while yields on short-term Treasury bills also moved into double-digit territory. This abrupt increase in funding costs, driven by anti-inflation policy rather than thrift-specific factors, devastated the traditional S&L business model of funding long-term, fixed-rate mortgages with short-term deposits.

ing hundreds of institutions economically insolvent and ultimately generating resolution costs in excess of \$150 billion, much of it borne by taxpayers (Curry and Shibut (2000); Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998)). Deregulatory measures, notably the Depository Institutions Deregulation and Monetary Control Act of 1980 removed interest rate ceilings, raised deposit insurance limits, expanded asset powers, and formalized capital forbearance, allowing weak thrifts to attract insured funds and “gamble for resurrection” while the Federal Savings and Loan Insurance Corporation (FSLIC) lacked the resources and political backing to close them promptly McCoy (2007). Congress eventually responded with the Financial Institutions Reform, Recovery, and Enforcement Act of 1989 (FIRREA), which abolished FSLIC, created the Resolution Trust Corporation to resolve failed thrifts, and transferred thrift insurance to a new fund administered by the FDIC; however, the contemporaneous wave of commercial bank failures made clear that the broader design of deposit insurance and resolution policy required more fundamental reform.

Before FDICIA, the FDIC enjoyed broad discretion over failure resolution, choosing among deposit payoffs, purchase-and-assumption transactions, and open-bank assistance without a statutory least-cost test and with considerable latitude to protect uninsured depositors. High-profile interventions, such as the rescue of Continental Illinois in 1984, signaled a willingness to shield all depositors and many creditors of large banks, reinforcing expectations of “too-big-to-fail” protection and encouraging risk-taking Davison (1997). Supervisors also sometimes delayed closure, engaging in regulatory forbearance that allowed losses to grow before resolution Benston and Kaufman (1997).

FDICIA sought to harden these soft constraints; Prompt Corrective Action linked supervisory interventions to objective capital categories, aiming to force earlier, rules-based responses to deterioration and to limit forbearance. The least-cost resolution requirement constrained the FDIC’s use of open-bank assistance and broad creditor protection, subject to a narrow systemic-risk exception. At the same time, risk-based insurance premiums and target reserve ratios shifted a greater share of the cost of failures onto insured institutions. Complementary provisions, such as more frequent examinations for weaker banks and limits on brokered deposits, reinforced these incentives. In aggregate, FDICIA represents a codified

lesson of the S&L era: generous insurance and deregulation must be paired with credible, timely intervention and loss-allocation rules if public exposure to banking crises is to be contained.

2.2.2 FDIC Resolution Process

Timeline

Under FDICIA's Prompt Corrective Action framework, regulators follow a structured timeline when dealing with a failing bank. In general, a failing institution is closed by its chartering authority once it becomes insolvent or "critically undercapitalized" (typically defined as tangible equity $\leq 2\%$ of assets). At that point, the FDIC is appointed as receiver and must resolve the bank in a manner that is "least costly" to the deposit insurance fund, barring a systemic risk exception. The resolution timeline usually unfolds over about 90 days for a capital insolvency scenario, during which the FDIC prepares for an orderly failure resolution. Key steps in this process include:

1. **Regulatory Intervention Period:** Once the chartering agency determines a bank is likely to fail, the bank is given up to 90 days to attempt corrective actions or find a private-sector solution (recapitalization or a buyer) under FDICIA's prompt corrective action provisions. During this period, the FDIC begins contingency planning for resolution in case the bank cannot recover.
2. **Failure Determination and Receivership:** If the institution cannot be saved by the end of the 90-day window, the chartering authority (state or federal) formally closes the bank and appoints the FDIC as receiver. This typically occurs once it is clear that no viable recapitalization or merger is forthcoming and the bank's capital is critically deficient.
3. **Marketing to Potential Acquirers:** The FDIC quickly moves to market the failing bank to potential bidders. Eligible bidder institutions are identified, usually well-capitalized banks in good regulatory standing capable of absorbing the failed bank.

After signing confidentiality agreements, these qualified bidders are given access to due diligence information about the failing bank's assets, liabilities, and operations. The FDIC simultaneously develops a marketing strategy and resolution structure for the sale.

4. **Auction and Bid Submission:** The FDIC conducts a sealed-bid auction for the failed bank. Bidders typically submit bids comprising three components: (a) an offer (if any) to assume the failed bank's deposits, often with a deposit premium they are willing to pay for the deposit franchise, (b) an offer for some or all of the failed bank's assets, often at a discount to book value, and (c) an indication of whether the bid is for a whole-bank transaction: all deposits, including uninsured, or only the insured deposits. Bidders may submit multiple bid options². In some cases, the FDIC may offer optional loss-sharing agreements or other structures to encourage bids on bad assets.
5. **Least Cost Bid Selection:** Once bids are received, the FDIC evaluates them using a proprietary Least Cost Test to determine which option would minimize the cost to the Deposit Insurance Fund (DIF). FDICIA mandates that the FDIC must choose the resolution alternative that is expected to be the least costly to the DIF, unless a systemic risk exception is invoked. In practice, this means that bids resulting in lower estimated losses to the FDIC, after accounting for assets recovered and any liabilities assumed by the acquirer, will be favored. The least-cost requirement also generally implies that uninsured depositors and other general creditors should absorb losses rather than be protected, unless protecting them is necessary to secure a bid or prevent systemic harm.
6. **Final Resolution:** Once the winning bid is selected, the FDIC prepares the legal closing documents, and the bank is closed. By the next business day, the FDIC carries out the resolution: either transferring operations to the acquiring institution or paying out deposits, depending on the method used. The entire resolution process from initial

²For example, a bid assuming all deposits versus a bid assuming only insured deposits.

regulatory warnings to final closing usually occurs within about 90 days for capital failures. However, in liquidity crises, this timeline may be compressed.

Resolution Types

The FDIC employs several resolution methods to address failed banks, each with different implications for the continuity of banking services, FDIC costs, and outcomes for depositors. The primary methods include Purchase and Assumption transactions (P&As), Insured Deposit Transfers (IDTs), Insured Deposit Payoffs (POs), Purchase of Insured deposits (PIs), and, in rare cases, the creation of a Deposit Insurance National Bank (DINB). In all cases, the goal is to protect depositors up to the insurance limit and dispose of the failed bank's assets and liabilities in an orderly way, while minimizing the FDIC's losses. The choice of resolution method depends on factors such as the bank's size and condition, bidders' interests, and the results of the least-cost test. Below, I describe each method in detail:

1. **Purchase and Assumption:** A P&A transaction is the most common resolution method, used in roughly 77.6% of bank failures historically. In a P&A, the "assuming institution" purchases some or all of the failed bank's assets and assumes some or all of the liabilities, especially the deposits. The FDIC acts as receiver to facilitate this transfer. P&A agreements are flexible and can be structured in various ways. In a whole-bank P&A, the acquirer assumes all deposits, including uninsured deposits, and takes on essentially all assets, even troubled loans, leaving very little with the FDIC. In other P&A variants, the acquirer may assume only the insured deposits, leaving uninsured depositors as creditors of the receivership, and may also purchase a portion of the assets. The FDIC typically sweetens the deal to ensure it meets the least-cost test: it will reimburse the acquirer for the difference between liabilities assumed and assets purchased by injecting cash or retaining bad assets, so that the acquiring bank does not take on a net loss at closing. Essentially, the acquirer buys the failed bank's good assets and assumes the deposits. At the same time, the FDIC absorbs the loss on bad assets by paying out receivership funds to balance the transaction. Often, the acquirer pays the FDIC a deposit premium, a cash payment in exchange for acquiring the

deposit franchise, if the franchise has value, which offsets the FDIC's cost. P&A deals allow for continuity of banking services: the failed bank's branches typically reopen as branches of the acquiring institution immediately, and all insured depositors automatically become customers of the acquirer with uninterrupted access to their accounts. If the P&A is a whole-bank transaction, even uninsured depositors are fully protected, and the acquirer assumes their accounts. However, the FDICIA tightened the rules governing the protection of uninsured depositors. Since 1994, the FDIC has been prohibited from covering uninsured deposits in a resolution except when a systemic risk exemption is invoked. Thus, in modern P&A transactions under the least-cost mandate, uninsured deposits are assumed only by the acquirer if doing so does not increase the FDIC's cost or if regulators determine that not protecting them would threaten financial stability. Overall, P&A transactions tend to be less costly to the FDIC than straight liquidations because they preserve franchise value, the premium paid by an acquirer, and they help avoid asset fire-sales. The acquiring bank benefits by gaining new customers and assets, often at a discount or with loss protection. Indeed, during the Great Financial Crisis of 2008, many P&A deals included FDIC loss-sharing agreements (where the FDIC agreed to absorb a large percentage of future losses on certain loan portfolios), which further encouraged bidders to participate. These arrangements greatly limit the acquirer's risk, effectively providing a subsidy. Empirical evidence shows that acquirers in FDIC P&A auctions during the crisis realized significant gains. For example, one study found that winning bidders' stock prices jumped on the announcement of the acquisition, reflecting the advantageous terms and FDIC support they received (Cowan and Salotti (2015)).

2. **Insured Deposit Transfers:** An IDT is a resolution method typically used when no open bank is willing to acquire the entire failing institution, but the FDIC can at least arrange to transfer the insured deposit accounts to another bank. In an IDT, the FDIC transfers all insured deposits of the failed bank to a healthy bank, which agrees to act as the paying agent for those deposits. The receiving bank does not purchase the failed bank as a going concern; rather, it simply opens accounts on its books

for the failed bank's insured depositors so that those customers can quickly access their insured funds, usually at full-service bank branches in their community Barth and Prabha (2013). The FDIC, in its receivership capacity, provides cash (or quickly sellable assets) equal to the total insured deposit amounts to the acquiring institution to cover the transferred deposits. Essentially, this is a facilitated payout: the insured depositors receive their money by seamlessly becoming depositors of the agent bank, rather than receiving a check from the FDIC. Uninsured depositors in this scenario are not transferred; they remain creditors of the receivership for the amount of their uninsured balances, to be paid out from asset liquidation recoveries (if any) over time. The agent bank typically receives a fee or other considerations for assisting, or sometimes pays a modest premium if the deposit franchise is attractive. However, this method is usually employed when franchise value is low.

The IDT method was more common in earlier decades and has rarely been used since the mid-1990s; in fact, the FDIC has not employed an IDT since 1995. It is considered similar to a direct payoff in that the failed bank is liquidated, with no ongoing operations sold. However, it has the advantage of giving insured depositors uninterrupted access to their funds through a new bank, thereby avoiding service disruptions that a payoff by mailed checks would entail. For the FDIC, the cost of an IDT is comparable to an insured deposit payoff. The FDIC still bears the loss for the difference between the failed bank's assets and insured liabilities. However, it may reduce resolution costs if an acquiring institution pays a small premium for the deposits or if having a ready agent bank streamlines the claims process. From the depositor's perspective, insured customers are made whole and can quickly access their funds at the acquiring bank. In contrast, uninsured depositors incur losses, receiving only receivership certificates and, if asset recoveries permit, eventual dividends.

3. **Insured Deposit Payoffs:** A deposit payoff (PO), also known as a deposit payout or liquidation, is the resolution method of last resort when no assuming institution can be found. In a payoff, the FDIC, as receiver, pays insured depositors directly the amount of their insured funds and winds down the failed bank's operations. The bank is closed

and liquidated: all assets of the failed bank are retained by the FDIC receivership for sale over time, and the proceeds are used to reimburse the FDIC for insured payouts and to pay any remaining claims by uninsured depositors or other creditors, in order of priority. Uninsured depositors in a payoff do not get an assuming bank; instead, they receive a receivership certificate and may eventually recover a portion of their uninsured balances through liquidation dividends if the asset sales yield sufficient funds. In most cases, uninsured depositors suffer losses (a partial recovery) because failed bank assets often liquidate for less than the total liabilities. From the FDIC's perspective, a straight payoff is usually the most costly resolution method, because it forgoes any franchise value. The FDIC must immediately draw on the insurance fund to pay all insured deposits, and then bear the risk and expense of liquidating the assets over time. Payoffs tend to result in higher FDIC losses as a percentage of the failed bank's assets compared to P&A transactions. POs also disrupt banking services: the failed bank does not reopen, and customers must wait to receive their insured funds (often by check or electronic transfer) and then find a new bank on their own. Because of these drawbacks, deposit payoffs have been relatively infrequently used in roughly 7% of failures in past decades and generally only occur when no acceptable bid is offered for a P&A or transfer. Notably, even in a payoff, the FDIC may use an IDT or DINB as a vehicle to facilitate the payout, as described next, to reduce disruption.

4. **Purchase of Insured Deposits:** A PI resolution is essentially a hybrid of a P&A and an IDT/PO. In a PI transaction, an acquiring bank agrees to purchase only the insured deposit liabilities of the failed bank, without taking on the uninsured deposits. The acquirer usually also takes a portion of the assets equal to those deposits, often mostly cash and investment securities, or the FDIC may provide cash to fund the assumed deposits. In effect, the acquirer bank is stepping into the FDIC's shoes to pay insured depositors, but through a P&A agreement for the insured deposit portion. The acquirer may pay the FDIC a small premium for the insured deposit franchise since it gains a base of new customer accounts, which can reduce the FDIC's net cost compared to a payoff. However, the acquirer leaves behind all uninsured deposits and

any unwanted assets to be handled by the FDIC's receivership. From the depositor's standpoint, insured depositors are protected and become customers of the acquiring bank, just as in a P&A or IDT, they retain access to their accounts, now at the new bank. Uninsured depositors, by contrast, are not assumed by the acquirer; they suffer an interruption and must seek recovery from the receivership for the uninsured portion of their funds, likely recovering only a fraction after asset liquidation. PI transactions can be viewed as a partial P&A: they preserve some insured deposit relationships but not the entire franchise. They have been used in cases where an acquirer is willing to take on the deposit base but perhaps not the riskier assets or the uninsured liabilities. This method was employed in a minority of failures, around 6% historically. The cost to the FDIC in a PI is typically lower than a full payoff because the FDIC might receive a premium and avoid the logistical costs of directly paying insured deposits. Still, because the FDIC must cover losses on assets not taken by the acquirer and handle the uninsured claims, the DIF will incur losses on a PI, albeit hopefully less than in a straight liquidation. Acquiring banks in PI deals gain insured deposit accounts without exposing themselves to the failed bank's credit risks since they can cherry-pick assets or take none at all beyond cash, making it an attractive proposition for the acquirer with minimal risk.

5. **Deposit Insurance National Bank:** A DINB is a temporary national bank chartered by the FDIC to maintain banking services for depositors of a failed bank when no immediate acquirer is found and a straight payoff would be disruptive. In this resolution, the FDIC itself obtains a new national bank charter from the Office of the Comptroller of the Currency and operates the DINB solely to provide depositors with access to their insured funds while they transition to other banks. The DINB takes over the insured deposits of the failed bank; depositors can access their insured balances by doing business with the DINB, which might operate out of the failed bank's existing offices for a limited period. By law, a DINB can operate for up to 2 years (with possible extensions), but in practice, the FDIC usually keeps a DINB open for only a short time (a few weeks to a few months). During that time, depositors are expected to move

their funds to other institutions. The DINB does not make loans or conduct a full banking business; it exists solely to facilitate the smooth payout of insured deposits. It has no capitalization requirements because the FDIC backs it. The DINB does not assume all uninsured deposits of the failed bank; initially, those uninsured depositors would receive receivership certificates and stand to recover funds later as assets are sold. In some cases, if a buyer for the whole bank emerges shortly after failure, the DINB can be bridged into a purchase transaction. For example, in the March 2023 failure of Silicon Valley Bank, the FDIC first created a DINB to pay insured depositors and was preparing to handle uninsured claims through receivership. However, within days, it transferred all deposits, including uninsured ones, into a newly created bridge bank after invoking a systemic risk exception. DINBs are exceedingly rare; they have accounted for less than 1% of failures because the FDIC usually prefers to either find an acquirer or, if possible, pay off deposits directly. However, the DINB tool is available for situations where the sudden loss of banking services would significantly harm the community or where an orderly payoff requires maintaining an interim bank. The FDIC's cost for a DINB is similar to a payoff; it still ultimately pays out insured deposits and bears the costs of asset liquidation and the administrative expenses of operating the temporary bank. Nevertheless, the benefit is that it gives depositors time to make alternative arrangements and can preserve a semblance of normalcy in the immediate aftermath of a bank failure, potentially preventing local panic.

In summary, the FDIC's resolution toolkit under FDICIA is designed to protect insured depositors while minimizing the insurance fund's losses. Purchase and Assumption transactions are the preferred method in most cases, as they maintain continuity and often yield the lowest cost by leveraging an acquiring institution. Insured deposit transfers and DINBs are fallbacks to ensure insured depositors are made whole when no outright buyer is available. A deposit payoff is the least favored due to higher costs and disruption, and is used only when necessary. The treatment of uninsured depositors depends on the method and the least-cost calculus: post-FDICIA, uninsured depositors are generally at risk of loss unless the FDIC can arrange a whole-bank acquisition or invoke a systemic risk exception to protect them.

2.2.3 Related Literature

The empirical literature on bank failures and FDIC resolutions is large and spans several distinct strands: early work that measures resolution losses and their determinants at the bank level; studies that embed these experiences in broader evaluations of deposit insurance and insolvency regimes; and more recent research that exploits detailed transaction-level data from FDIC auctions, often focusing on bidder behavior, acquirer outcomes, and the incidence of losses across claimants. This subsection briefly reviews the main contributions in each strand, with an emphasis on U.S. evidence and on studies that explicitly analyze the FDIC’s choice of resolution method and its implications for the deposit insurance fund and depositors.

Early empirical work by FDIC staff and affiliated researchers documents basic facts about resolution costs across failure episodes and resolution methods. Bovenzi and Murton (1988) and Bovenzi and Muldoon (1990) use FDIC Banking Review data to estimate the average costs of alternative resolution methods, discussed in Section 2.2.2, and show that P&A transactions generally reduce losses relative to straight liquidations when suitable acquirers exist. James (1991) measures realized losses as the difference between book liabilities and discounted recovery values of failed bank assets, documenting that losses are strongly related to pre-failure asset quality, capital ratios, and local economic conditions. Brown and Epstein (1992) updates the FDIC’s historical loss model, while subsequent work by McDill (2004) links resolution costs to the business cycle and macroeconomic stress. Parallel work on the thrift sector, such as the analysis of the savings and loan crisis in Curry and Shibut (2000), quantifies the aggregate fiscal burden of large failure waves and emphasizes that taxpayers ultimately bore a substantial share of losses via the deposit insurance funds. These studies establish the now-standard definition of “resolution cost” as the shortfall between the failed bank’s liabilities and the discounted market value of its assets net of resolution expenses, and they highlight systematic differences in this cost across resolution technologies and institutional environments.

A second strand embeds these experiences within a broader evaluation of deposit insurance design and the legal framework for bank resolution. FDIC historical volumes, notably

Managing the Crisis: The FDIC and RTC Experience, 1980–1994 Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998), provide a narrative and quantitative account of the 1980s banking and thrift crises, including the extensive use of P&A transactions, bridge banks, and open-bank assistance, as well as the evolution of FDIC practice in response to large-scale failures. Benston and Kaufman (1997) interpret the FDICIA as a response to moral hazard problems created by mispriced deposit insurance and forbearance during the 1980s, and they emphasize two innovations with direct implications for resolution policy: prompt corrective action and the statutory least-cost resolution requirement. Bliss and Kaufman (2006) compare U.S. corporate bankruptcy and bank receivership regimes, stressing that bank resolution is an administrative process run by the FDIC under the Federal Deposit Insurance Act rather than a judicial Chapter 11 proceeding, with different objectives (maintaining confidence and limiting systemic risk), different priority rules (including depositor preference), and much faster timelines for closing, valuing, and reprivatizing failed institutions. Related work on depositor preference, secondary deposit insurance, and changes in creditor priority (e.g., Hirschhorn and Zervos (1990); Marino and Bennett (1999); Osterberg (1996)) clarifies how statutory and contractual priority rules govern the distribution of resolution losses across insured and uninsured depositors, general creditors, and the FDIC.

More recent research exploits detailed FDIC data to dissect the components of resolution cost and to study the auction process used to sell failed banks. Bennett and Unal (2015) construct a reconciled time series of cash flows for more than 1,200 FDIC receiverships between 1986 and 2007, decomposing total resolution costs into losses on assets, direct receivership expenses, and indirect FDIC overhead, and documenting how these components vary with resolution method, bank size, legislative period (pre-FIRREA, FIRREA³, FDICIA), and fraud. They show, for example, that resolution costs as a share of assets decline with bank size and are lower in the post-FDICIA period, consistent with earlier closures and the operation of the least-cost test. Complementary FDIC publications, such as Failed Bank Cost Analysis: 1986–2002 Corporation (2003) and subsequent staff studies, extend

³The Financial Institutions Reform, Recovery, and Enforcement Act of 1989

this line of work by formalizing FDIC cost estimation methodologies and relating realized costs to ex-ante projections used in the least-cost decision. A parallel literature, starting with Granja et al. (2017), uses FDIC auction data to study which potential acquirers win failed bank auctions and how bidder characteristics—geographic proximity, portfolio similarity, and capitalization—shape both allocation and FDIC losses, interpreting these findings through a fire-sale framework.

Finally, there is a growing body of event-study and contract-level work on FDIC-assisted transactions during the 2008–2013 crisis and its aftermath. Cowan and Salotti (2015) analyzes stock returns for winning bidders in FDIC-assisted acquisitions and shows that acquirers earn positive abnormal returns that are increasing in the generosity of FDIC guarantees, indicating that resolution design redistributes surplus between the FDIC and acquiring banks. Heitz (2022) focuses on loss-sharing agreements in failed-bank sales and finds that acquiring banks entering into loss-sharing arrangements experience significantly lower abnormal returns than otherwise similar auction participants, implying that loss-sharing shifts more downside risk back to acquirers and away from the FDIC. Recent work also examines externalities on rival and losing bidders (e.g., Molyneux and Zhou (2022)) and compares FDIC-assisted acquisitions to open-market mergers in terms of long-run profitability and risk (Cowan et al. (2022); Spokeviciute et al. (2025)). FDIC staff studies contrast the 1980–1994 SNL and 2008 GFC crises, emphasizing that, although far fewer banks failed in the latter episode, the average size of failures and the complexity of resolution strategies, including greater use of loss sharing and structured transactions, were much larger.

2.3 Press Release Data

The empirical analysis relies on a novel dataset constructed by merging administrative failure records with hand-collected details from the FDIC’s failure-specific press releases. The foundational population is drawn from the FDIC’s BankFind Suite: Bank Failures and Assistance Data, which provides the universe of federally insured depository institution failures. This base file is augmented with granular transaction-level data extracted directly from the FDIC’s press release archive and matched to quarterly Call Report data from Drechsler et al.

(2017, 2021). The final sample covers the period from January 21, 1983, to October 18, 2024, and contains 1,860 observations. The panel consists primarily of commercial banks, national and state-chartered, and 68 savings banks.

This section has two purposes. First, it documents the construction and coverage of the hand-collected dataset. Second, it presents the descriptive cross-sectional and time-series variation in the six variables that the structural model in Section 2.4 treats as primitives or outcomes: deposit runoff, the deposit beta of Drechsler et al. (2021), the loan-loss-provision ratio, the uninsured share, the realized retained-asset share, and the auction outcomes. Taken together, Tables 2.1 and 2.2 and the accompanying figures show that there is substantial cross-sectional dispersion in each primitive, and that simple bivariate patterns in the data already motivate the structural links that the model later formalizes.

2.3.1 Construction and Coverage

For each failed bank, the FDIC issues a standardized public announcement that describes the closing authority, the assuming institution, the timing of branch reopening, and the key balance-sheet and resolution terms. An example of the common template appears in Appendix B.4. From each release I code the volume of assets and deposits at the moment of resolution, the total number of deposit accounts, the subset of accounts and deposit balances above the \$100,000 (later \$250,000) insurance limit, the premium paid by the acquirer on deposits and on assets, the volume of assets retained by the FDIC in receivership, any cash advance from the FDIC to the acquirer, the notional amount of assets covered by a loss-sharing agreement, the FDIC's share of those losses, and the FDIC's projected resolution cost. A full list of the variables contained in the data set can be found in Table B.3. These variables are not available in standard FDIC statistical publications or call report data but are central to allocating resolution outcomes across acquirers, the DIF, and uninsured depositors. The Call Report match, following Bennett and Unal (2015), provides the pre-failure (QBF) balance sheet, from which I construct the loan-loss-provision ratio and the bank's size.

Table 2.1 documents the coverage of all variables used in the paper. Four remarks are in

Table 2.1: Coverage of FDIC Press Release Data		
Data Field	Count	% Missing
Panel A: Balance Sheet Inputs		
QBF Assets	1,860	0.00
QBF Deposits	1,860	0.00
Final Assets	1,860	0.00
Final Deposits	1,860	0.00
Deposit Runoff	1,858	0.11
Log Asset Size	1,860	0.00
Panel B: Theoretical Primitives		
LLP Ratio	1,785	4.03
Deposit Beta	373	79.95
Expense Beta	373	79.95
Uninsured Share	414	77.74
Core Deposit Share	414	77.74
Panel C: Auction Outcomes		
Recovery Category	1,860	0.00
Recovery %	1,860	0.00
Premium on Deposits	1,860	0.00
Premium on Assets	1,856	0.22
Retained-Asset Share	1,860	0.00
FDIC Advance	1,860	0.00
Loan-Loss Provision (press release)	1,860	0.00
Loan-Loss FDIC Share	7	99.62
Cost	1,860	0.00

Table 2.1: Coverage of variables used in the paper ($N = 1,860$ failures). Panel A lists balance-sheet inputs from call reports and press releases. Panel B lists the theoretical primitives used as regressors. Panel C lists auction outcomes from the press-release dataset.

order. First, QBF and final balance-sheet items are available for essentially the full sample, as are the core auction outcomes: premiums on deposits and assets, retained assets, cost, and uninsured-depositor recovery. Second, the loan-loss provision from the Call Report is available for 1,785 of 1,860 failures. Third, the matched Drechsler et al. (2021) deposit betas cover 373 failures, concentrated in post-2000 closures. Fourth, the number and balance of accounts above the insurance limit are available for 414 failures, primarily during the 1980s, when the FDIC reported this granularity in its releases. I therefore treat uninsured exposure as a restricted-sample covariate and flag the sample restriction explicitly in the regression tables in Section 2.5.

2.3.2 Failure Landscape

Figure 2.1 in Section 2.1 traces the annual count and resolution-method mix of failures over the sample. Three features of that figure are worth carrying into the rest of this section. First, failures cluster in three waves: the S&L crisis of 1986–1995, the Global Financial Crisis (GFC) of 2008–2013, and the small but large-asset 2023 regional banking episode. Second, the share of resolutions structured as P&As rises from roughly three-quarters in the 1980s to nearly one in the Dodd-Frank period, with pure payoffs and insured deposit transfers becoming rare. Third, as noted above, average failed-bank asset size rises by nearly two orders of magnitude across the sample. However, medians remain in the community-bank range, a pattern already documented in Curry and Shibut (2000) and Bennett and Unal (2015).

Figure 2.2 adds a spatial dimension by mapping the location of each failed bank by regulatory regime. Failures cluster persistently in Texas and the broader Southwest, the agricultural Midwest, California, and parts of the Northeast, with the concentration shifting across regions as local real estate and commodity cycles evolve. Many S&L-era failures involve small community banks in rural or semi-rural counties with limited branch competition. Morgan et al. (2016) shows that post-GFC branch closures are associated with the emergence of “banking deserts” in which households and small firms face reduced access to basic deposit and credit services. The spatial pattern in Figure 2.2 is consistent with that concern. It underscores that resolution policy is not only about minimizing DIF losses but also about managing localized disruptions to relationship banking.

2.3.3 Balance Sheet Primitives

The model in Section 2.4 treats four bank-level objects as exogenous primitives that determine franchise value and the equilibrium bid: the deposit attrition hazard λ , the deposit rate sensitivity β , the pre-resolution loan-loss provision ratio that proxies for asset toxicity γ , and the uninsured share of deposits. This subsection presents the cross-sectional variation in each.

The model’s attrition hazard λ is not directly observable, but failing banks approach

Figure 2.2: Location of Domestic Bank Failures




Figure 2.2: Locations of failed banks by regulatory regime. Each point is a failed institution colored by regime: Pre-FIRREA, FIRREA, FDICIA/PCA, or Dodd-Frank.

resolution over a short window of days to weeks. I use the log change in reported deposits between the last Call Report and the resolution date as a one-shot realization of cumulative attrition up to $t = 0$,

$$\text{Deposit Runoff}_i := \ln \left(\frac{\text{Final Deposits}_i}{\text{QBF Deposits}_i} \right), \quad (2.1)$$

where $t = -1$ is the last QBF Call Report and $t = 0$ is the resolution date. The timing convention matches the model.

Figure 2.3 plots the distribution of Deposit Runoff_i for the full sample alongside the subsamples of failures that occurred during the three crisis windows. The full-sample distribution is centered near zero, with a pronounced left tail: a nontrivial share of banks lose more than 20 log points of deposits in the quarters leading up to closure. The 2023 subsample shifts the entire distribution to the left, consistent with the rapid deposit outflows documented in the March 2023 regional episode, while the S&L and GFC subsamples display

heavier left tails than the full sample but retain substantial mass near zero. This variation is what identifies the coefficient on Deposit Runoff in the deposit-premium regressions of Section 2.5.

Figure 2.3: Deposit Runoff Distribution




Figure 2.3: Histogram of $\ln(\text{Final Deposits}/\text{QBF Deposits})$ for the full sample, overlaid with density step plots for the S&L (1986–1995), GFC (2008–2013), and 2023 subsamples. Runoff values are clipped at $[-1.0, 0.5]$ for display.

Drechsler et al. (2021) estimate bank-level sensitivities of deposit rates to short-term policy rates and argue that low-beta deposits are the core source of franchise value because they insulate net interest income from rate shocks. I match each failed bank to the deposit beta and interest-expense beta posted on Philipp Schnabl’s website; the match rate is 373 banks, concentrated in the post-2000 period. Figure 2.4 plots the log deposit premium from the FDIC auction against the matched deposit beta. The raw scatter is noisy, but the decile means on the right panel are monotonically decreasing: acquirers pay systematically higher premiums for banks whose deposits reprice slowly. This bivariate pattern is the empirical counterpart of the mapping from β to franchise value developed in Section 2.4 and is what identifies the coefficient on deposit beta in the premium regressions.




Figure 2.4: Left: scatter of log deposit premiums against bank-level deposit betas matched from Drechsler et al. (2021). Right: mean deposit premium by beta decile.

The model's asset-valuation wedge γ is proxied by the ratio of loan-loss provisions to QBF assets,

$$\text{LLP}_i := \frac{\text{Loan-Loss Provision}_i}{\text{QBF Assets}_i}, \quad (2.2)$$

computed from the last Call Report before closure and winsorized at the 1% and 99% tails. The identification logic is straightforward. Loan-loss provisions are an accounting recognition of expected credit losses that is booked before the FDIC opens the auction. They are therefore predetermined with respect to the auction's equilibrium outcome and can be treated as exogenous in regressions of cost and the retained-asset share, in the spirit of James (1991), who shows that realized resolution losses are strongly related to pre-failure asset quality.

Figure 2.5 plots the cross-sectional distribution of the LLP ratio. The distribution has substantial right skewness: the sample mean is 1.8%, but the 90th percentile exceeds 5%.

Figure 2.6 shows the first-stage relationship that motivates the use of LLP_i as an instrument for the realized retained-asset share τ_i in the cost regressions. Restricting to P&A transactions, I bin the sample into 20 quantile bins of the LLP ratio and plot the mean

Figure 2.5: Loan Loss Provision Ratio




Figure 2.5: Panel (a): histogram of LLP scaled by QBF assets, winsorized at 1% and 99%, with the sample mean and median. Panel (b): annual mean and interquartile band; shaded regions mark the S&L, GFC, and 2023 crisis windows.

retained share in each bin, together with the OLS fit on the underlying observations. The slope is 1.18, so a one-percentage-point increase in the pre-failure LLP ratio is associated with a 1.2-percentage-point increase in the realized retained share. This pattern is consistent with the least-cost logic: banks with more troubled assets enter the auction with a larger share of assets that acquirers are unwilling to bid for at book value, and the FDIC retains the residual in receivership. It is the identifying variation behind the asset-side instrumentation in the analysis of Section 2.5.

For the subsample of failures in which the FDIC release reports over-the-limit balances, I construct the uninsured share as

$$\text{Unins}_i = \frac{\text{Amount Over Limit}_i}{\text{Final Deposits}_i} \quad (2.3)$$

and the complementary core-deposit share as $1 - \text{Unins}_i$. Coverage is heavily tilted toward the 1980s and early 1990s, and I flag this restriction whenever these variables enter the regression analysis. In line with the classical bank-run logic of Diamond and Dybvig (1983), uninsured balances are the claim most exposed to loss in a payoff, and variation




Figure 2.6: Realized retained-asset share vs. pre-failure LLP ratio, P&A sample. Points are means within 20 quantile bins of LLP/QBF Assets; the line is the OLS fit on the underlying observations. The slope and HC1 standard error are reported in the title.

in the uninsured share across failures is a natural shifter of the option value of running in resolution.

Table 2.2 collects the cross-sectional moments of the six primitives. The interquartile range of deposit runoff indicates that half of failed banks experience at most modest attrition in the quarter before closure, but a meaningful minority experience 5–15% deposit attrition. Deposit and expense betas range from roughly 0.2 to 0.7 across the matched subsample, so a wide cross-section identifies the model’s rate sensitivity parameter. The retained-asset share has a mean of 0.31 but a median of 0.23, with 30% of failures recording $\tau_i = 0$, a whole-bank P&A or straight payoff.

2.3.4 Resolution Outcomes

I now turn to the outcome variables that the model’s equilibrium bid determines: the premiums paid on deposits and assets, the realized retained-asset share τ , the projected cost to the

Table 2.2: In-Sample Bank Failure Summary Statistics

	N	Mean	Std. Dev.	p10	p25	Median	p75	p90
Deposit Runoff	1,858	-0.040	0.288	-0.123	-0.057	-0.002	0.014	0.058
Deposit Beta	373	0.416	0.100	0.297	0.354	0.408	0.481	0.540
Expense Beta	373	0.373	0.093	0.261	0.316	0.362	0.427	0.486
LLP Ratio	1,785	0.018	0.031	0.000	0.000	0.004	0.023	0.052
Retained-Asset Share	1,860	0.307	0.321	0.000	0.000	0.228	0.544	0.794
Uninsured Share	414	0.713	13.843	0.001	0.004	0.015	0.043	0.075

Table 2.2: Summary statistics for the theoretical primitives used in the empirical analysis. Deposit Runoff is the log change in deposits between the last call report and the resolution date. Deposit and Expense Betas are matched from bank-level estimates following Drechsler et al. (2021). LLP Ratio is the pre-failure loan-loss provision scaled by QBF assets, winsorized at the 1% and 99% tails. Retained-Asset Share is the ratio of assets retained by the FDIC to QBF assets. Uninsured Share is deposits above the insurance limit as a fraction of final deposits.

DIF, and the recovery paid to uninsured depositors. Before documenting the unconditional shape of these outcomes, I first describe how failures and resolution methods vary across the legislative and macro-financial environment in which they occur. The institutional regime fixes the menu of resolution technologies available to the FDIC and the statutory tests the agency must satisfy, and the macro-financial regime shapes both the size and the asset composition of the banks that resolve. Both margins matter for interpreting the pricing moments reported below.

Table 2.3 summarizes the distribution of failures across legislative regime (Panel A) and macro-financial environment (Panel B). Two shifts are first-order for interpreting the auction outcomes. First, mean QBF assets rise by nearly two orders of magnitude across regimes, from \$43 million under pre-FIRREA to \$1.9 billion under Dodd-Frank, with a further order-of-magnitude jump to \$106 billion for the five banks in the 2023 crisis subsample. The contrast between the exploding mean and the community-bank-sized median documented in Bennett and Unal (2015) is therefore entirely driven by a few large failures concentrated in the FDICIA and Dodd-Frank windows. Second, the composition of resolution methods changes discontinuously at each legislative boundary. In the pre-FIRREA period, 15% of resolutions are structured as insured deposit transfers and roughly 10% as straight pay-offs, so that, by construction, more than one in four failures impose losses on uninsured

Table 2.3: Resolution Outcomes by Period

Panel A: Regulatory Periods								
	# Failures	Assets	Deposits	P&A	IDT	PO	PI	DINB
Pre-FIRREA (1983–1988)	714	42.53	38.62	0.749	0.153	0.098	0.000	0.000
FIRREA (1989–1991)	420	151.81	132.47	0.840	0.110	0.050	0.000	0.000
FDICIA (1992–2009)	363	1,443.62	1,015.84	0.573	0.047	0.080	0.292	0.008
Dodd-Frank (2010–2024)	363	1,906.31	1,593.72	0.956	0.000	0.022	0.008	0.014
Panel B: Crisis Periods								
S&L Crisis (1986–1995)	1,052	99.439	89.346	0.718	0.144	0.072	0.067	0.000
Global Financial Crisis (2008–2013)	471	1,340.25	977.03	0.930	0.000	0.034	0.019	0.017
2023 Regional Crisis (2023)	5	106,445.68	88,122.01	1.000	0.000	0.000	0.000	0.000
No Crisis (1983–1985, 1996–2007, 2013–2022, 2024)	332	126.68	107.53	0.738	0.063	0.108	0.090	0.000

Table 2.3: Summary of failed banks by regulatory regime (Panel A) and crisis environment (Panel B). Columns report counts of failures, mean quarter-before assets and deposits (in millions of dollars), and the share of resolutions by type: purchase and assumption (P&A), insured deposit transfer (IDT), payout (PO), purchase of insured deposits (PI), and deposit payoff with DINB bridge bank (DINB).

depositors. Under FIRREA, the IDT and PO shares fall. Under FDICIA, the purchase of insured deposits (PI) appears and accounts for 29% of resolutions, reflecting the least-cost logic of Benston and Kaufman (1997): the PI preserves the insured-deposit franchise for an acquirer while leaving uninsured depositors and distressed assets in the receivership. The Dodd-Frank regime then returns to a near-exclusive reliance on whole-bank P&A transactions, with 96% of failures resolved by a single acquiring institution, consistent with the evidence in Spokeviciute et al. (2025) that the FDIC has increasingly favored franchise-preserving resolutions over the last three decades.

Panel B of Table 2.3 cuts the sample by macro-financial environment. Three features are worth emphasizing. First, the S&L crisis accounts for 1,052 of the 1,860 failures in the

sample but for only a small share of aggregate failed-bank assets, confirming the “many small failures” character of the episode narrated in Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998) and quantified in Curry and Shibut (2000). Second, the GFC and the 2023 crisis invert this pattern: 471 and 5 failures, respectively, but mean assets an order of magnitude (GFC) and three orders of magnitude (2023) larger than in the S&L episode. The 2023 subsample is essentially the mass of the large regional banks resolved under systemic-risk exceptions. Third, the P&A share is monotone in crisis severity: 72% in the S&L crisis, 93% in the GFC, and 100% in 2023. This pattern is consistent with the view in Granja et al. (2017) that, conditional on a sufficiently deep pool of well-capitalized bidders, whole-bank P&As dominate the least-cost calculus even under significant asset-quality stress, and with the Cowan and Salotti (2015) evidence that FDIC loss-sharing and cash advances were used aggressively during the GFC to maintain this preference for franchise-preserving resolutions. The “no crisis” subsample, by contrast, looks more like a low-intensity extension of the S&L pattern: a modestly larger average bank, a slightly lower P&A share, and a small but persistent use of IDTs and POs for the smallest and least attractive institutions.

The heterogeneity across the rows of Table 2.3 is the institutional backdrop against which the pricing moments below are formed. In particular, (i) the asset-premium, deposit-premium and cost distributions pool banks resolved under four different statutory regimes, and (ii) the unconditional means are heavily influenced by the post-FDICIA era, in which the P&A has become the dominant technology and in which resolution losses are allocated between the FDIC, acquirers, and uninsured depositors through a least-cost-scored auction rather than a discretionary negotiation.

Three moments in Table 2.4 summarize the pricing of failed banks. First, the Cost/Assets ratio has a right tail but is tightly bunched near 22%, close to the figure cited in Bennett and Unal (2015) for the 1986–2007 period. Second, the average premium paid on assets is near zero and strictly negative when concentrated discounts are included. In contrast, the premium on deposits is strictly positive on average, reinforcing the observation of James (1991) that value in a P&A is a deposit-side phenomenon. Third, the retained-asset share τ_i has mass at zero (whole-bank P&As) and a continuous right tail (partial retention). This

Table 2.4: Resolution Outcome Summary Statistics

	N	Mean	Std. Dev.	Min	q25%	Median	q75%	Max
Cost / QBF Assets	1,860	0.230	0.138	0.000	0.132	0.218	0.312	0.937
Premium on Deposits (%)	1,860	1.053	2.358	-1.500	0.000	0.075	1.160	48.146
Premium on Assets (%)	1,856	-2.010	6.406	-78.226	0.000	0.000	0.000	5.791
Retained-Asset Share τ	1,860	0.307	0.321	0.000	0.000	0.228	0.544	3.234
Uninsured Deposit Recovery (%)	1,860	90.935	37.843	0.000	100	100	100	100
FDIC Advance (\$m)	1,860	9.459	33.623	0.000	0.000	0.000	6.425	410.500

Table 2.4: Summary statistics for FDIC resolution outcomes. Cost is the projected DIF loss scaled by QBF assets. Premiums are reported as a percentage of deposits (assets). The retained-asset share is the ratio of retained assets to QBF assets. Recovery is the realized payoff to uninsured depositors.

bimodality is the quantity-side analog of the mass at zero in the asset-premium distribution: the FDIC either hands the entire asset book to the acquirer at a market price, or retains a nontrivial share in receivership.

Figure 2.7 plots the distribution of the realized retained-asset share τ , the single quantity in the press-release data most directly tied to the structural model. Thirty percent of failures have $\tau = 0$: either the acquirer takes the whole balance sheet or the resolution is a pure payoff. Among the remaining two-thirds, τ is continuously distributed with a mean of 0.43 and a right tail extending beyond 1.0 for cases in which the FDIC retains assets valued above QBF book value after write-downs.

Figure 2.8 tracks the uninsured recovery rate over time. The unconditional mean is 91%, with 78% of failures delivering full recovery, 5% partial recovery (mean 58%), and 17% zero recovery. The time series reveals the institutional shift traced in Section 2.2: outside major crises, recoveries are lower and more volatile, consistent with FDICIA's limits on protecting uninsured deposits absent a systemic-risk determination. In the GFC and 2023 episodes, the combination of whole-bank P&As and explicit systemic-risk exceptions pushes recoveries close to 100%. The ex-ante distribution of recoveries is a mixed incentive for uninsured depositors to run, as emphasized by Diamond and Dybvig (1983), and is empirically consistent with the stock-return evidence in Cowan and Salotti (2015) on the incidence of FDIC guarantees.

Before turning to the structural model, Figure 2.9 previews the three regressions es-




Figure 2.7: Distribution of the realized retained-asset share $\tau = \text{Retained Assets} / \text{QBF Assets}$. Histogram restricted to positive retention; the fraction of failures with $\tau = 0$ is reported in the title.

timated in Section 2.5 with binned scatter plots of the reduced-form relationships they formalize. Panel (A) plots the log deposit premium against deposit runoff, restricted to P&A transactions: premiums decline monotonically with larger runoff, consistent with the franchise-value interpretation of the deposit premium. Panel (B) plots Cost/Assets against the LLP ratio over the full sample: the slope is strongly positive, consistent with James (1991) and Bennett and Unal (2015). Panel (C) plots the share of resolutions structured as P&As across deciles of log QBF assets: the probability of a P&A is near 0.9 in the top deciles and falls toward 0.5 for the smallest banks, reflecting the interaction of least-cost pricing with the geography of potential acquirers documented in Granja et al. (2017). These three patterns, in that order, preview Tables 2.5, 2.6, and 2.7.




Figure 2.8: Uninsured deposit recovery over time. The series shows average recovery rates with shaded crisis periods: Savings & Loan crisis (blue), Global Financial Crisis (pink), and 2023 regional bank crisis (tan).

2.4 Theory

This section develops a structural model of FDIC failed-bank auctions that jointly determines the deposit premium, the asset premium (or discount), and the fraction of assets retained by the FDIC. I begin with a continuous time cash flow framework for the deposit franchise, then introduce the balance-sheet objects that govern asset valuation, and embed both sides in a sealed-bid first-price scoring auction. The central analytical result is that the FDIC's least-cost mandate collapses a multidimensional bid, comprising a deposit premium, an asset purchase price, and an asset take-up share, into a scalar score, thereby preserving the tractability of standard first-price auction theory while endogenizing the composition of the resolution transaction. The model yields directly testable comparative statics linking deposit-base quality, asset toxicity, and bidder heterogeneity to observed premiums, retained asset shares, and sale probabilities.

2.4.1 Environment

Time t is continuous, with $t = 0$ denoting the auction date. Failed banks are indexed by $i \in \{1, \dots, I\}$ and potential acquiring banks by $j \in \{1, \dots, J\}$. Each failed bank i holds a stock of core deposits D_{i0} and a portfolio of assets with aggregate book value A_{i0} at the date

Figure 2.9: Motivating Bivariates




Figure 2.9: Panel (A): binned scatter of log deposit premium on deposit runoff, P&A sample. Panel (B): binned scatter of Cost/Assets on the LLP ratio, full sample. Panel (C): mean probability of a P&A across deciles of log QBF assets. In each panel, the line is the OLS fit (or the decile mean in Panel C) on the underlying observations.

of auction. Deposits exhibit exponential decay, at any time $t \geq 0$, the remaining deposit balance is

$$D_{it} := D_{i0}e^{-\lambda_i t}, \quad (2.4)$$

where $\lambda_i \in \mathbb{R}_{++}$ is the deposit-attribution hazard rate for bank i . A smaller λ_i indicates a “stickier” deposit base that runs off more slowly, whereas a larger λ_i implies rapid outflows following failure. This exponential specification is a parsimonious form for the rich heterogeneity in depositor behavior documented in the literature: some depositors leave immediately upon resolution, while others remain with the acquiring institution indefinitely. The hazard rate λ_i summarizes the net effect across the depositor population.

Let r_t^m denote the market short-term interest rate and let r_{jt} denote acquirer j ’s marginal return on earning assets. Each acquirer also incurs an operating cost $c_j \in \mathbb{R}_+$ per dollar of deposits serviced. The deposit rate paid to depositors of failed bank i is

$$r_{it}^d := \alpha_i + \beta_i r_t^m, \quad (2.5)$$

where α_i is a bank-specific intercept capturing the base level of deposit compensation and $\beta_i \in [0, 1]$ measures the pass-through of market rates to deposit costs. A low β_i characterizes

a deposit base that is insensitive to the prevailing rate environment, the hallmark of valuable deposit franchise. A high β_i indicates deposits that reprice rapidly with the market, eroding the funding-cost advantage that makes the franchise worth acquiring in the first place. Let $r_f \in \mathbb{R}_+$ denote the acquirer's discount rate, and define $\mu_x := \mathbf{E}[x_t]$ for any stochastic process x , so that μ_{r^m} and μ_{r_j} are the expected market rate and the expected asset return, respectively.

2.4.2 Valuing the Deposit Franchise

I first derive the economic value of the failed bank's deposit franchise to a potential acquirer. Because the deposit base declines over time at hazard rate λ_i ⁴, the franchise is a wasting asset that generates a stream of cost-saving benefits until deposits run off. At each instant, the acquirer can invest the remaining deposits in earning assets at rate r_{jt} , while paying depositors the rate r_{it}^d and bearing operating cost c_j . A single dollar of deposits therefore produces instantaneous net profit

$$\Pi_t := (r_{jt} - r_{it}^d - c_j) D_t. \quad (2.6)$$

I impose $\mathbf{E}[\Pi_t] > 0$, consistent with the FDIC's requirement that acquiring institutions be in sound financial condition and capable of profitably deploying the acquired deposit base.

The present value of the deposit franchise is the expected discounted flow of net spreads, accounting for exponential deposit attrition

$$V_{ij}^D(\beta_i, \lambda_i) := \mathbf{E} \left[\int_0^\infty e^{-r_f t} \Pi_t dt \right] = \frac{(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j) D_{i0}}{r_f + \lambda_i}. \quad (2.7)$$

This closed-form expression distills the franchise value into a transparent function of three economic forces. The numerator captures the expected net interest margin per dollar of deposits: the spread between asset returns and the all-in cost of deposit funding. The denominator is an effective discount rate that combines the acquirer's time preference r_f with the deposit-attrition rate λ_i . Intuitively, deposits that decay faster are worth less because

⁴The franchise is a wasting asset because the acquirer inherits a fixed, run-off deposit portfolio, no new deposits are originated to replace attrition.

the acquirer has a shorter window over which to harvest the spread.

Two parameters are central to the empirical analysis that follows. First, a lower β_i preserves the acquirer's funding advantage even as market rates rise, directly increasing the numerator of (2.7). Second, a lower λ_i reduces the effective discount rate in the denominator, amplifying the present value of every dollar of expected spread. Both channels raise V_{ij}^D and, as I show below, feed through monotonically into equilibrium bid premiums.

2.4.3 Asset Valuation

While the deposit franchise constitutes the primary source of going-concern value in most FDIC resolutions, the asset portfolio of the failed bank also enters the acquirer's valuation. In practice, FDIC P&A transactions allow bidders to acquire some or all of the failed bank's assets alongside its deposits. Assets that are not transferred to the winning bidder are retained by the FDIC as receiver and liquidated or resolved over time. To capture this institutional feature, I introduce a minimal asset-valuation structure that generates endogenous variation in the fraction of assets transferred.

Let $r_i^{\text{FDIC}} \in (0, 1)$ denote the FDIC's expected recovery value per dollar of book assets if bank i 's assets are retained and worked out in receivership. I treat r_i^{FDIC} as a reduced-form primitive that summarizes the receiver's liquidation technology, the composition of the asset portfolio (e.g., performing loans versus distressed commercial real estate), and the time and cost of the workout process. Higher r_i^{FDIC} indicates that the FDIC is relatively efficient at recovering value from the asset pool.

Let $s_{ij} \in [0, 1]$ denote the fraction of the failed bank i 's total asset book that a bidder j proposes to acquire, so that $s_{ij}A_{i0}$ dollars of assets are transferred and $(1 - s_{ij})A_{i0}$ are retained by the FDIC. I model each bidder j 's marginal valuation for the x -th fractional unit of bank i 's assets as

$$v_{ij}(x) := \bar{v}_{ij} - \gamma_i x, \quad \gamma_i > 0. \quad (2.8)$$

Here $\bar{v}_{ij} \in (0, 1]$ is bidder j 's marginal valuation for the first tranche of assets, and $\gamma_i > 0$ governs the rate at which marginal value declines as the bidder absorbs a larger share of the

portfolio.

The declining marginal value specification in (2.8) is motivated by three economic forces. First, failed-bank asset portfolios are heterogeneous: performing residential mortgages, consumer auto loans, commercial real estate, and distressed or non-performing assets coexist on the same balance sheet. A rational bidder will cherry-pick the highest-quality tranches first, so marginal value falls as the bidder moves deeper into the portfolio. Second, acquiring banks face regulatory capital constraints; each additional dollar of risk-weighted assets demands incremental capital, raising the effective cost of the marginal asset. Third, absorbing a failed bank's assets requires operational capacity, loan servicing infrastructure, workout expertise, local market knowledge, that is subject to diminishing returns at scale. The slope parameter γ_i captures the combined force of these frictions. When γ_i is large, the portfolio is highly heterogeneous or the bidder's capacity constraints bind quickly, so marginal value deteriorates rapidly. When γ_i is small, the asset pool is relatively homogeneous and the bidder can absorb a large share without significant value erosion.

Integrating (2.8) over the take-up share s_{ij} yields the bidder's total value from acquiring fraction s_{ij} of the asset book

$$V_{ij}^A(s_{ij}) := A_{i0} \int_0^{s_{ij}} v_{ij}(x) dx = A_{i0} \left(\bar{v}_{ij} s_{ij} - \frac{1}{2} \gamma_i s_{ij}^2 \right). \quad (2.9)$$

The quadratic form in (2.9) has a natural interpretation; the first term $\bar{v}_{ij} s_{ij}$ is the value the bidder would obtain if all acquired assets were worth $\bar{v}_{ij}$ per dollar of book, while the second term $\frac{1}{2} \gamma_i s_{ij}^2$ is the cumulative “toxicity penalty” that reflects the declining quality of successive asset tranches. For small s_{ij} , the penalty is negligible and the bidder captures close to $\bar{v}_{ij}$ per dollar; for s_{ij} near one, the penalty can be substantial, potentially rendering the marginal asset a net liability. This curvature is the minimal device needed to generate interior solutions for asset take-up; without it, bidders would either acquire all assets or none, counterfactual to the rich variation in retained asset shares observed in the data.

2.4.4 Auction Mechanism

The FDIC resolves the majority of failed banks through P&A transactions structured as sealed-bid first-price auctions. I extend the standard single-dimensional format to accommodate joint bidding on deposits and assets, then show that the FDIC's statutory least-cost mandate reduces the multidimensional bid to a scalar score, preserving the tractability of classical auction theory.

I restrict attention to whole-bank P&A transactions in which the acquirer assumes all deposits, including uninsured deposits, and no loss-share arrangement is in place. Each bidder j submits a sealed bid for failed bank i consisting of three components:

- (4) a *deposit premium rate* $p_{ij} \geq 0$, expressed per dollar of assumed deposits, so that total premium dollars are $P_{ij} = p_{ij}D_{i0}$,
- (4) an *asset take-up share* $s_{ij} \in [0, 1]$, specifying the fraction of the failed bank's asset book the bidder proposes to acquire, and
- (4) an *asset purchase price* $q_{ij} \in [0, 1]$ per dollar of book value for the acquired assets, so that total asset payment is $Q_{ij} = q_{ij}s_{ij}A_{i0}$.

Under this bid, the acquirer pays a total of $p_{ij}D_{i0} + q_{ij}s_{ij}A_{i0}$ to the FDIC, assumes all D_{i0} of deposits as liabilities, takes ownership of $s_{ij}A_{i0}$ of assets, and the FDIC retains $(1 - s_{ij})A_{i0}$ for disposal in receivership. The three-dimensional bid (p_{ij}, q_{ij}, s_{ij}) nests the standard deposit-only auction as the special case $s_{ij} = 0$, in which no assets are transferred, and the entire premium is paid on the deposit line.

The joint bid structure reflects institutional practice: in actual FDIC resolutions, acquirers routinely negotiate over both the deposit premium and the terms on which assets are transferred. Modeling these as separate but simultaneous choices allows us to endogenize the composition of the transaction and, in particular, the fraction of assets retained by the FDIC.

FDICIA requires the FDIC to resolve each bank failure using the method that is least costly to the DIF. I formalize this mandate as a scoring rule over bids. Consider a bid (p, q, s) for failed bank i . The FDIC's cash inflows consist of the deposit premium pD_{i0} and

the asset sale proceeds qsA_{i0} . Against these, the FDIC must honor the full D_{i0} of assumed deposit liabilities and expects to recover $r_i^{\text{FDIC}}(1-s)A_{i0}$ from the retained assets. The net cost to the DIF is therefore

$$C(p, q, s) := D_{i0} - \underbrace{(pD_{i0} + qsA_{i0})}_{\text{cash received from acquirer}} - \underbrace{r_i^{\text{FDIC}}(1-s)A_{i0}}_{\text{expected recovery on retained assets}}. \quad (2.10)$$

The first term is the deposit liability the FDIC must cover; the second is the total cash the acquirer pays; and the third is the present value the FDIC expects to realize by working out the assets it keeps. A higher bid, whether through a larger premium, a higher asset price, or a greater asset take-up share, reduces the DIF's out-of-pocket cost.

Rearranging (2.10) yields

$$C(p, q, s) = \underbrace{(D_{i0} - r_i^{\text{FDIC}}A_{i0})}_{\text{baseline shortfall}} - [pD_{i0} + (q - r_i^{\text{FDIC}})sA_{i0}]$$

Assumption 2. *The failed bank is insolvent in expectation: $D_{i0} > r_i^{\text{FDIC}}A_{i0}$.*

Under Assumption 2, the baseline shortfall $D_{i0} - r_i^{\text{FDIC}}A_{i0} > 0$, is common across all bids, and reflects the fundamental insolvency of the failed bank: its deposit liabilities exceed the expected liquidation value of its assets. The FDIC's bid-selection problem therefore reduces to maximizing the second term, which I denote the *score*

$$\mathcal{S}_i(p, q, s) := pD_{i0} + (q - r_i^{\text{FDIC}})sA_{i0}. \quad (2.11)$$

This is the key tractability result of the model. Despite the three-dimensional bid space, the FDIC's least-cost mandate induces a linear scoring rule that compresses all bid components into a single scalar. I refer to $\mathcal{S}_i$ as the score—the object the FDIC evaluates when comparing bids. The score has a clean economic interpretation. The first term, pD_{i0} , is the direct cash benefit from the deposit premium: every dollar of premium reduces DIF cost dollar-for-dollar. The second term, $(q - r_i^{\text{FDIC}})sA_{i0}$, measures the incremental benefit of asset transfer relative to the FDIC's outside option of retaining and resolving those assets itself. Transferring

an additional dollar of book assets to the acquirer benefits the FDIC if and only if the acquirer's purchase price q exceeds the FDIC's own expected recovery r_i^{FDIC} . When $q < r_i^{\text{FDIC}}$, the FDIC would recover more by keeping the assets, so the asset-transfer component of the score is negative, and the bidder must compensate with a higher deposit premium. This scoring structure nests the pure deposit-premium auction as the limiting case: when $s = 0$ or $q = r_i^{\text{FDIC}}$, the score collapses to pD_{i0} and the FDIC selects the highest deposit premium.

Each bidder j values the acquisition of failed bank i along two margins. On the deposit side, the bidder derives franchise value V_{ij}^D as defined in (2.7), scaled by a private synergy parameter $\theta_{ij} \in (0, 1]$ that captures heterogeneity in strategic fit, geographic overlap, or operational capacity. Note that V_{ij}^D carries a j subscript because it depends on acquirer-specific parameters c_j and μ_{r_j} ; the franchise value of the same deposit base differs across bidders. On the asset side, the bidder obtains value $V_{ij}^A(s)$ from acquiring share s of the asset book, as defined in (2.9). The bidder also bears an idiosyncratic integration cost $\bar{K}_{ij} \in \mathbb{R}_+$ encompassing systems' conversion, regulatory compliance, personnel retention, and other merger-related expenses.

If bidder j wins bank i 's auction with bid (p, q, s) , the bidder's payoff is

$$U_{ij}(p, q, s) := \theta_{ij}V_{ij}^D(\beta_i, \lambda_i) + V_{ij}^A(s) - \bar{K}_{ij} - pD_{i0} - qsA_{i0}. \quad (2.12)$$

The first term is the deposit franchise value scaled by synergy; the second is the gross asset value from (2.9); the third is the integration cost; and the final two terms are the cash payments to the FDIC. This payoff structure makes transparent the fundamental trade-off the bidder faces: paying more (higher p or q) increases the probability of winning but reduces the surplus conditional on winning.

The private-information structure is as follows. The parameters θ_{ij} and $\bar{K}_{ij}$ are private to bidder j ; rival bidders and the FDIC observe only their joint distribution. The deposit-franchise parameters (β_i, λ_i) and the asset-portfolio parameters $(\bar{v}_{ij}, \gamma_i)$ are common knowledge in the sense that they characterize the failed bank's observable balance sheet, though bidders may differ in $\bar{v}_{ij}$ due to heterogeneous asset-management capabilities.

I now show that the combination of the FDIC's linear scoring rule and the bidder's payoff

structure reduces the multidimensional auction to a standard first-price auction on the scalar score $\mathcal{S}_i$. This reduction is the central analytical device of the model. The first step is to express the bidder's total cash payment in terms of the score. Using (2.11),

$$pD_{i0} + qsA_{i0} = \mathcal{S}_i + r_i^{\text{FDIC}} sA_{i0}. \quad (2.13)$$

Substituting (2.13) into the bidder payoff (2.12) and collecting terms yields

$$U_{ij} = \underbrace{\left[\theta_{ij} V_{ij}^D(\beta_i, \lambda_i) - \bar{K}_{ij} + A_{i0} \left((\bar{v}_{ij} - r_i^{\text{FDIC}}) s_{ij} - \frac{1}{2} \gamma_i s_{ij}^2 \right) \right]}_{:= W_{ij}(s_{ij})} - \mathcal{S}_i. \quad (2.14)$$

The bracketed expression, which I denote $W_{ij}(s)$, is the bidder's *private value*: the total surplus the bidder would earn if it won the auction at a score of zero. The payoff is then the private value minus the score, which is the canonical structure of a first-price auction with private value W_{ij} and bid $\mathcal{S}_i$.

The economic content of equation (2.14) is as follows. The term $\theta_{ij} V_{ij}^D - \bar{K}_{ij}$ is the net deposit-side surplus. The asset-side contribution $A_{i0}[(\bar{v}_{ij} - r_i^{\text{FDIC}})s - 1/2\gamma_i s^2]$ measures the surplus the bidder creates by acquiring assets relative to the FDIC's own recovery. The relevant wedge is not the bidder's absolute valuation $\bar{v}_{ij}$ but the comparative advantage $\bar{v}_{ij} - r_i^{\text{FDIC}}$: the bidder creates surplus on the asset transfer only to the extent that it can extract more value from the assets than the FDIC receiver could through its own workout process. This comparative-advantage structure ensures that asset transfers occur precisely when they are socially efficient, a direct consequence of the least-cost mandate.

The reduction to a scalar auction preserves the entire apparatus of first-price auction theory. In particular, the equilibrium analysis of Milgrom and Weber (1982) applies directly: each bidder submits a score bid $\mathcal{S}_i$ that is a monotonically increasing function of their private value W_{ij} , with bid shading that depends on the distribution of rival private values. The highest-private-value bidder wins and pays a score strictly below their value, earning positive expected surplus. All comparative statics derived for standard first-price auctions carry over immediately.

2.4.5 Retained Share of Assets

The asset take-up share s is not merely a component of the bid vector; it is a strategic choice that each bidder optimizes as part of its bid preparation. Because the private value $W_{ij}(s)$ is the bidder's "budget" for bidding aggressively in the scoring auction, each bidder selects s to maximize $W_{ij}(s)$; a higher private value translates directly into a higher equilibrium score bid, which raises the probability of winning.

Taking the first-order condition of $W_{ij}(s)$ with respect to s and solving yields the optimal asset take-up share

$$s_{ij}^* := \min \left\{ \frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i}, 1 \right\}, \quad (2.15)$$

where I impose $s_{ij}^* \geq 0$, a bidder cannot be forced to return assets, and $s_{ij}^* \leq 1$, a bidder cannot acquire more than the full portfolio. The second-order condition $\partial_s^2 W_{ij} = -A_{i0}\gamma_i < 0$ confirms that this is a global maximum.

Equation (2.15) encodes a sharp economic logic. The bidder acquires assets up to the point where the marginal asset's private value $v_{ij}(s_{ij}^*) = \bar{v}_{ij} - \gamma_i s_{ij}^*$ exactly equals the FDIC's recovery value r_i^{FDIC} . Assets with private value above r_i^{FDIC} generate positive surplus when transferred, so acquiring them raises the bidder's valuation. Assets with private value below r_i^{FDIC} are worth more to the FDIC in receivership than to the bidder, so acquiring them would reduce the bidder's valuation. The optimal take-up share thus partitions the asset portfolio into a "transferable" tranche, for which private-sector management dominates FDIC technology, and a "retained" tranche, for which the reverse is true.

I define the model-implied *retained asset share* as the fraction of assets kept by the FDIC under the winning bidder's optimal take-up

$$\tau_i := 1 - s_{ij^*}^*, \quad (2.16)$$

where j^* denotes the auction winner. The retained share is an equilibrium object, jointly determined by the bidder's asset-management comparative advantage $(\bar{v}_{ij^*} - r_i^{\text{FDIC}})$, the heterogeneity of the asset portfolio (γ_i) , and the FDIC's resolution technology (r_i^{FDIC}) .

The model predicts a positive association between the retained share τ_i and the deposit premium p , operating through two reinforcing channels. The first is mechanical: using (2.11), the deposit premium can be expressed as $p = [\mathcal{S}_i - (q - r_i^{\text{FDIC}})sA_{i0}]/D_{i0}$, so for a given score $\mathcal{S}_i$, a lower take-up share (higher τ) shifts the bid's composition toward the premium line. The second is an equilibrium channel: when marginal assets satisfy $v_{ij}(x) < r_i^{\text{FDIC}}$, forcing the bidder to acquire them would lower W_{ij} and depress the equilibrium score bid. Allowing the FDIC to retain these negative-surplus assets removes a drag on the bidder's private value, enabling more aggressive bidding that raises both $\mathcal{S}_i$ and the deposit premium.

Crucially, both channels generate the same positive sign. This means that the model's prediction for the relationship between the retained share and the premium is robust to the joint determination of τ_i and p : regardless of whether one conditions on the score or allows the score to adjust endogenously, a higher τ_i is associated with a higher p . In the empirical analysis that follows, I construct an exogenous proxy for the retained share using pre-failure balance-sheet characteristics of the failed bank. The model provides the structural basis for this approach: cross-sectional variation in asset-portfolio quality shifts the primitives $\bar{v}_{ij}$ and γ_i , which determine s_{ij}^* and hence τ_i through (2.15). Because these balance-sheet characteristics are determined before the resolution auction, they are not a function of the bidding equilibrium.

2.4.6 Comparative Statics

I now derive comparative statics along three dimensions: deposit-franchise fundamentals, asset-side parameters, and the interaction between the two. These results generate the model's testable predictions.

The sensitivity of the deposit franchise value to the rate-sensitivity parameter is

$$\frac{\partial V_{ij}^D}{\partial \beta_i}(\beta_i, \lambda_i) = -\frac{\mu_{r^m} D_{i0}}{r^f + \lambda_i} < 0. \quad (2.17)$$

A higher β_i reduces franchise value because deposits that reprice rapidly with the market compress the net interest margin in rising-rate environments. When deposit costs track

market rates one-for-one ($\beta_i \rightarrow 1$), the acquirer's funding-cost advantage vanishes and the franchise reduces to a costly liability net of operating expenses. The magnitude of the effect is scaled by the expected market rate μ_{r^m} ; in high-rate environments, the penalty for rate-sensitive deposits is proportionally larger.

The sensitivity to the attrition rate is

$$\frac{\partial V_{ij}^D}{\partial \lambda_i}(\beta_i, \lambda_i) = -\frac{\mathbf{E}[\Pi_t]D_{i0}}{(r^f + \lambda_i)^2} < 0. \quad (2.18)$$

Faster deposit attrition reduces franchise value by shortening the horizon over which the acquirer can harvest the net spread. λ_i acts as an additive premium to the discount rate, and increasing it shrinks the present value of the entire cash-flow stream. Note that the numerator is the expected net margin, which we have assumed positive; the sign of the derivative is therefore unambiguously negative.

The cross-partial derivative reveals an important interaction:

$$\frac{\partial^2 V_{ij}^D}{\partial \beta_i \partial \lambda_i}(\beta_i, \lambda_i) = \frac{\mu_{r^m} D_{i0}}{(r^f + \lambda_i)^2} > 0. \quad (2.19)$$

This positive cross-derivative indicates that the marginal cost of rate sensitivity is smaller when deposits are short-lived. Intuitively, if deposits are going to run off quickly regardless, their rate sensitivity matters less because the acquirer's exposure to interest-rate risk is brief. Conversely, when deposits are very sticky, rate sensitivity becomes acutely important because the acquirer bears the compressed spread for a long horizon. Stickiness and rate insensitivity are substitutes in generating franchise value: strength in one dimension partially compensates for weakness in the other.

The optimal take-up share responds to the bidder's asset-management advantage and the portfolio's heterogeneity

$$\frac{\partial s_{ij}^*}{\partial \bar{v}_{ij}} = \frac{1}{\gamma_i} > 0, \quad \frac{\partial s_{ij}^*}{\partial \gamma_i} = -\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i^2} < 0. \quad (2.20)$$

A bidder with a higher baseline valuation $\bar{v}_{ij}$ acquires a larger share of the portfolio because more tranches clear the r_i^{FDIC} threshold. Conversely, greater portfolio heterogeneity (higher

γ_i) causes marginal value to decline more steeply, so the bidder reaches the indifference point sooner and takes up less.

The response to the FDIC's recovery technology is

$$\frac{\partial s_{ij}^*}{\partial r_i^{\text{FDIC}}} = -\frac{1}{\gamma_i} < 0. \quad (2.21)$$

When the FDIC is relatively efficient at working out assets (r_i^{FDIC} high), the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$ narrows, and the bidder optimally takes up a smaller share. In the limit $r_i^{\text{FDIC}} \rightarrow \bar{v}_{ij}$, the bidder has no comparative advantage and the entire asset book is retained. This result has a clear policy implication: investments in the FDIC's receivership infrastructure (e.g., loan-workout expertise, technology for managing distressed portfolios) endogenously shift the equilibrium toward greater retention, which may or may not reduce total resolution cost depending on the deposit-side response.

Substituting the interior optimal take-up s_{ij}^* from (2.15) into the private value W_{ij} yields the maximized private value

$$W_{ij}(s_{ij}^*) = \theta_{ij} V_{ij}^D(\beta_i, \lambda_i) - \bar{K}_{ij} + \frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i}. \quad (2.22)$$

The three terms correspond, respectively, to the net deposit-side surplus, the integration cost, and the asset-side option value, the surplus the bidder extracts from acquiring those asset tranches for which it has a comparative advantage over the FDIC. The asset-side option value is always non-negative and is quadratic in the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$, reflecting the fact that a larger wedge both increases the surplus on each infra-marginal tranche and expands the set of tranches worth acquiring.

Equation (2.22) reveals that the deposit and asset channels are additively separable in the private value. This separability has two important implications. First, improvements in deposit-franchise quality increase the private value without altering the optimal asset take-up, because s_{ij}^* depends only on asset-side parameters. Second, improvements in asset-side comparative advantage raise the private value without altering the deposit-side valuation. The two channels reinforce each other in driving up equilibrium score bids and premiums,

but they operate through independent mechanisms. This additive structure is a consequence of the linear scoring rule and the separability of the bidder's payoff across deposit and asset margins.

2.4.7 Probability of P&A

Not every failed bank is resolved through a P&A transaction. When no bidder submits a bid that clears the FDIC's least-cost threshold, the FDIC instead executes a deposit payoff (PO): insured depositors are paid directly from the DIF, uninsured depositors receive a receivership claim, and all assets are liquidated by the FDIC as receiver. The model generates predictions about which failures result in a P&A sale and which do not.

A P&A transaction is least-cost relative to a payoff if and only if the winning bid's DIF cost from (2.10) is below the cost of the payoff alternative. Let C_i^{payoff} denote the DIF cost under a payoff resolution. A sale occurs if and only if

$$\max_j \mathcal{S}_i(p_{ij}, q_{ij}, s_{ij}) \geq \underbrace{(D_{i0} - r_i^{\text{FDIC}} A_{i0}) - C_i^{\text{payoff}}}_{:= \underline{\mathcal{S}}_i}, \quad (2.23)$$

where $\underline{\mathcal{S}}_i$ is the minimum score required for the P&A to be cost-effective. When the best available score falls below this threshold, no sale occurs.

In equilibrium, each bidder's score bid is a monotonically increasing function of their maximized private value $W_{ij}(s_{ij}^*)$ from (2.22). The probability that at least one bidder clears the threshold $\underline{\mathcal{S}}_i$ therefore depends on the distribution of private values across potential acquirers.

Several predictions follow directly. First, failures with higher deposit-franchise value are more likely to attract bids that clear the threshold, because V_{ij}^D enters positively in every bidder's private value. Attractive deposit franchises are, in a precise sense, easier to sell. Second, failures with more toxic asset portfolios present a subtler prediction. The direct effect of asset toxicity is to reduce the asset-side option value in (2.22), which lowers private values and reduces sale probability. However, the endogenous retention mechanism partially offsets this: by allowing the FDIC to retain the most toxic tranches, the model removes the

negative-surplus assets from the bidder's calculus, sustaining a positive private value even when the overall portfolio is impaired. The net effect on sale probability depends on the magnitude of the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$ relative to γ_i . Third, the number of bidders J affects sale probability through the standard order-statistic channel: the maximum score bid among J draws from the private-value distribution is stochastically increasing in J . Greater competition raises the probability that at least one bidder clears the threshold, consistent with the well-established empirical finding that FDIC failures attracting more interested parties are more likely to be resolved through P&A transactions.

Among failures that result in a P&A sale, the model generates a positive conditional association between the retained share τ_i and the deposit premium p . Failures that clear the sale threshold despite high asset toxicity must have particularly valuable deposit franchises or particularly well-matched acquirers. Conditional on sale, the high- τ observations are therefore selected to have strong deposit-side fundamentals, which manifest as high premiums. This selection effect reinforces the mechanical and equilibrium channels identified in Section 2.4.5.

2.4.8 Testable Predictions

The model generates six testable predictions that organize the empirical analysis in the following section.

The first two predictions concern the deposit franchise. The winning deposit premium p is decreasing in the deposit-rate sensitivity β_i , because rate-insensitive deposits preserve the acquirer's funding-cost advantage in rising-rate environments, raising franchise value and equilibrium bids. The premium is also decreasing in the deposit-attrition rate λ_i , because stickier deposits extend the horizon over which the acquirer earns the net spread, amplifying the present value of the franchise through the effective discount rate $r_f + \lambda_i$.

The next three predictions concern the asset side and its interaction with the premium. The retained asset share τ_i is increasing in asset-portfolio heterogeneity γ_i and decreasing in the bidder's asset-management advantage $\bar{v}_{ij}$: the FDIC retains a larger fraction of assets when the portfolio is more heterogeneous or when bidders have less comparative advantage in

managing distressed assets. Conditional on a P&A sale, the deposit premium p is positively associated with the retained share τ_i , a prediction that holds through both the mechanical and equilibrium channels identified in Section 2.4.5 and that the model's structure signs unambiguously. The retained share is also increasing in the FDIC's expected recovery value r_i^{FDIC} , because a more efficient FDIC receiver raises the threshold for asset transfer, leading to greater equilibrium retention.

The final prediction concerns the extensive margin of resolution. The probability of a P&A transaction, relative to a deposit payoff, is increasing in the deposit-franchise value V_{ij}^D , the asset-side comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$, and the number of potential bidders J . Each operates through the private value $W_{ij}(s_{ij}^*)$ and the order-statistic structure of the auction: more valuable franchises, larger asset-side surpluses, and greater competition all raise the probability that the maximum score bid exceeds the least-cost threshold $\underline{S}_i$.

The model is intentionally parsimonious. It abstracts from several features of actual FDIC resolutions: loss-share agreements, partial deposit assumptions, dynamic renegotiation between the FDIC and potential acquirers, and the possibility that the FDIC's recovery value r_i^{FDIC} is itself endogenous to the resolution method. These omissions are deliberate: the goal is to isolate the first-order interaction between deposit-franchise quality and asset-side composition in a framework that is transparent enough to generate sharp, sign-identified predictions. The additive separability of the deposit and asset channels in the maximized private value (2.22) is the key structural result, it implies that deposit-side and asset-side fundamentals can be tested independently in the cross-section, and that their contributions to the premium can be decomposed without requiring a fully specified joint distribution of bidder types. The empirical specifications in the next section map each prediction to an observable regressor, exploiting pre-failure balance-sheet variation to identify the model's comparative statics.

2.5 Empirics

The theoretical framework of Section 2.4 maps six comparative statics onto three estimating tests: a deposit pricing channel, an FDIC cost channel, and a selection channel. The analysis

centers on three constructed variables. Deposit runoff proxies for the attrition parameter λ_i and is defined in (2.1). Deposit beta measures the interest-rate sensitivity of deposit costs and is matched from Drechsler et al. (2021). The toxic ratio proxies for the retained asset share τ_i and equals the fraction of total assets held by the FDIC receivership rather than transferred to the acquirer.

Two identification threats require explicit attention before turning to the results. The first is simultaneity in the toxic ratio: because the premium and the retained asset share are jointly determined at the auction, OLS estimates of the toxic ratio coefficient in the pricing regression are biased toward zero. I address this by instrumenting the realized toxic ratio with the pre-failure loan loss provision ratio, defined in (2.2).⁵ The first-stage coefficient on the provision ratio is 0.897 with a standard error of 0.231, yielding an F -statistic of 15.0, which satisfies the Staiger and Stock (1997) threshold for instrument relevance. The first stage reflects a mechanical relationship: banks with larger pre-failure provisioning generate a larger pool of impaired assets that acquirers decline to purchase, so more of the asset book remains with the FDIC. The full first-stage regression is reported in Table A1 of Appendix B.2.2.

The second identification threat is sample selection on P&A transactions: a deposit premium is observed only when a sale occurs, so the pricing regression is estimated on a subsample that is not a random draw from the failure universe. I address this with the two-step correction of Heckman et al. (1997), using the log number of concurrent failures in the same calendar quarter as the exclusion restriction. The argument for exclusion follows James (1991) and Granja et al. (2017): resolution bandwidth and acquirer attention are finite, so aggregate quarterly congestion reduces the probability of a timely P&A bid without affecting the conditional price that a successful acquirer pays for the specific target. The Heckman correction leaves the pricing coefficients essentially unchanged; the selection-corrected estimates and full diagnostics are reported in Appendix B.2.2.

The headline result across all three tests is that the toxic ratio is the dominant quantita-

⁵The loan loss provision ratio is the ratio of total loan loss provisions to assets in the quarter before failure. It is determined on the balance sheet before any bidding occurs and is therefore unaffected by the auction outcome.

tive lever. Deposit runoff and deposit beta enter with the signs predicted by the model but with attenuated statistical significance.

2.5.1 Data and Sample

The sample spans 1,860 failure observations from 1983 through 2024 across four distinct regulatory epochs. Sample coverage varies systematically across the three tests. The full failure universe supports the cost and selection regressions, with between 1,784 and 1,860 observations depending on variable availability. The P&A subsample, restricted to transactions for which an explicit deposit premium is observed, provides approximately 1,488 observations for the pricing regressions. The deposit-beta-matched subsample further restricts to banks for which Drechsler et al. (2021) supply a matched rate-sensitivity estimate, yielding 373 observations. An uninsured-share subsample of 133 observations, concentrated in the pre-1990 period, isolates the compositional effects of liability structure.

Year fixed effects appear in every specification and robust standard errors are used throughout. State fixed effects are omitted on power grounds given the sample size.⁶

2.5.2 Deposit Pricing Channel

The pricing regression implements the comparative statics of the deposit franchise value in equations (2.17) and (2.18). The estimating equation is

$$\ln \left(1 + \frac{\text{Premium}_{it}}{\text{Deposits}_{it}} \right) = \alpha + \gamma_1 \text{Deposit Runoff}_i + \gamma_2 \text{Deposit Beta}_i + \Gamma X_i + \delta_t + \varepsilon_{it}, \quad (2.24)$$

where X_i includes the loan loss provision ratio and log asset size and δ_t are year fixed effects. The model predicts $\gamma_1 < 0$ and $\gamma_2 < 0$. Table 2.5 reports the results.

The primary finding of the pricing regressions is the narrow dispersion of the dependent variable. The sample standard deviation of $\ln(1+p)$ on the P&A subsample is approximately

⁶State fixed effects would ideally absorb idiosyncratic regional shocks such as the collapse of oil prices in the 1980s, which disproportionately affected community banks in Texas and Oklahoma. The sample does not have sufficient within-state variation to support them.

Table 2.5: Deposit Premium Regressions

	Deposit Premium: $\ln(1 + p)$				
	(1)	(2)	(3)	(4)	(5)
Deposit Runoff	0.00763 (0.00518)	0.01628 (0.01177)	0.01589 (0.01176)	0.06707*** (0.02066)	0.01623 (0.01182)
Uninsured Share				0.06238*** (0.01696)	
Deposit Beta		-0.01468 (0.00900)			-0.02052 (0.01408)
Expense Beta			-0.01221 (0.00931)		0.00723 (0.01413)
LLP Ratio	-0.00508 (0.01331)	0.04591 (0.04634)	0.04559 (0.04712)	0.07518 (0.09179)	0.04592 (0.04608)
Log Asset Size	0.00033 (0.00035)	-0.00088 (0.00056)	-0.00088 (0.00056)	-0.00029 (0.00143)	-0.00091 (0.00056)
Year FE	X	X	X	X	X
N	1488	359	359	133	359
R^2	0.30	0.45	0.45	0.46	0.45

Table 2.5: Deposit Side: Pricing Channel. OLS regressions of the log deposit premium on deposit-franchise characteristics and the exogenous LLP ratio proxy. Sample restricted to P&A transactions. All specifications include year fixed effects. Robust standard errors in parentheses. Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

0.015 log-points, and year fixed effects absorb most of the cross-sectional variation, yielding $R^2 = 0.30$ in column (1). This compression reflects the competitive structure of the auction: the winning bid is tightly concentrated around a year-specific median, leaving little residual variation for deposit fundamentals to explain.

In column (1), deposit runoff enters with a point estimate of +0.0076 that is not statistically distinguishable from zero. Columns (2) and (3) introduce deposit beta and expense beta on the matched subsample of 359 observations. Deposit beta enters at -0.015 and expense beta at -0.012 , both carrying the sign predicted by equation (2.17): higher rate sensitivity reduces the premium an acquirer is willing to pay. Neither achieves conventional significance. The most likely explanation is classical measurement error in rate-sensitivity

estimates for small failing banks, which typically have limited exposure to the branch-level deposit pricing data underlying the Drechsler et al. (2021) betas.

Column (4), which introduces the uninsured-share subsample, generates the sharpest deposit-side pricing result in the table. Deposit runoff enters at +0.067 and uninsured share enters at +0.062, both significant at the one-percent level. The sign on deposit runoff in this column requires interpretation. When uninsured share is controlled, residual runoff conditional on a given uninsured fraction reflects sophisticated-depositor sorting rather than generalized panic. Acquirers value the insured core franchise; runoff concentrated among uninsured depositors who bear losses in any resolution is a qualitatively different event from runoff among insured depositors whose claims the acquirer assumes. The positive coefficient is therefore consistent with the insured franchise being more valuable precisely in banks where the uninsured depositor base has already self-selected out.

The loan loss provision ratio enters with a negative point estimate in column (1). This sign reflects the selection and bargaining channel: banks that require a larger loss-sharing pool to attract any bidder are precisely the banks whose residual franchise commands a lower premium. The structural positive effect of asset cleanliness on the premium emerges once simultaneity is purged by the instrumental-variable strategy described above. The second-stage IV coefficient on the instrumented toxic ratio is larger in magnitude than the corresponding OLS coefficient, consistent with simultaneity bias attenuating OLS estimates toward zero.

2.5.3 FDIC Cost Channel

The FDIC resolution cost is the mirror image of the deposit premium. Every dollar the acquirer does not pay, the insurance fund absorbs. The estimating equation is

$$\frac{\text{Cost to FDIC}_{it}}{\text{Assets}_{it}} = \alpha + \psi_1 \text{Deposit Runoff}_i + \psi_2 \text{Deposit Beta}_i + \Psi X_i + \delta_t + \varepsilon_{it}, \quad (2.25)$$

where X_i includes the loan loss provision ratio and log asset size. The model predicts

$\psi_1 > 0$ and $\psi_3 > 0$, with all signs opposite to those in the pricing channel. Table 2.6 reports the results.

Table 2.6: FDIC Cost Regressions				
	Cost/Assets			
	(1)	(2)	(3)	(4)
Deposit Runoff	0.16784*** (0.04336)	0.22858* (0.12515)	0.22833* (0.12540)	0.10202 (0.08042)
Uninsured Share				-0.00034*** (0.00008)
Deposit Beta		0.05810 (0.08943)		
Expense Beta			0.06870 (0.09640)	
LLP Ratio	0.85942*** (0.10276)	0.62922** (0.29927)	0.62754** (0.29891)	0.81910*** (0.23310)
Log Asset Size	-0.01557*** (0.00267)	-0.02609*** (0.00541)	-0.02636*** (0.00551)	-0.00487 (0.00604)
Year FE	X	X	X	X
N	1784	372	372	402
R^2	0.16	0.21	0.21	0.26

Table 2.6: OLS regressions of the cost-to-assets ratio on deposit-franchise characteristics and the exogenous LLP ratio. All specifications include year fixed effects. Robust standard errors in parentheses. Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

The cost regressions deliver the sharpest results in the paper. In column (1) all three primary regressors enter with the predicted signs at one-percent significance. Deposit runoff carries a coefficient of +0.168: a one-standard-deviation increase in runoff raises the cost-to-assets ratio by roughly 4.9 percentage points relative to a sample median of approximately 22 percent.⁷ The loan loss provision ratio carries a coefficient of +0.859: a one-standard-deviation increase in the provision pool raises the cost-to-assets ratio by approximately 20 percentage points, making it the dominant quantitative driver across all specifications. Log asset size enters at -0.016 : larger banks are resolved at lower proportional cost, consistent

⁷The one-standard-deviation increase corresponds to approximately 0.29 log-points of deposit runoff.

with scale economies in resolution and with the size effects documented in the selection regression below.

These results confirm the theoretical trade-off at the core of the paper. The same provision ratio that enters negatively and insignificantly in the deposit pricing regressions, enters positively and significantly in the cost regressions at a magnitude of roughly +0.86. The same variable, measured from the same balance sheet, takes opposite signs on the two sides of the auction ledger in exactly the pattern that the model's scoring mechanism in equation (2.11) predicts. The FDIC's asset-retention decision simultaneously raises the acquirer's willingness to pay and raises the insurance fund's ultimate loss.

Column (4) introduces the uninsured-share subsample. Uninsured share enters at -0.00034 with high significance. This is a mechanical accounting channel distinct from the franchise-value channel in the pricing regressions: in a liquidation or partial assumption, uninsured depositors absorb losses before the insurance fund, directly reducing the DIF's out-of-pocket cost. Deposit beta enters in columns (2) and (3) with a positive point estimate on the cost side, directionally consistent with the model, but neither coefficient achieves conventional significance in the smaller matched subsample.

2.5.4 P&A Selection Channel

The extensive margin of resolution, whether the FDIC finds a P&A acquirer at all, is governed by equation (2.23) of the model. The probability that the highest private value clears the least-cost threshold $\underline{S}_i$ is increasing in the deposit franchise value V_{ij}^D , the asset-side comparative-advantage wedge, and the number of bidders. The estimating equation is

$$\begin{aligned} \mathbf{P}\{\text{Resolution Type}_i = \text{P\&A}\} = & \alpha + \phi_1 \text{Deposit Runoff}_i + \phi_2 \text{Core Deposit Share}_i \\ & + \phi_3 \text{LLP Ratio}_i + \Phi X_i + \delta_t + \varepsilon_{it}, \end{aligned} \quad (2.26)$$

where the dependent variable equals one if the failure was resolved via P&A or Insured Deposit Transfer and zero for deposit payoff or DINB resolutions. Table 2.7 reports linear

probability model and logit estimates.

Table 2.7: Probability of P&A Regressions

	Linear Probability Model			Logit
	(1)	(2)	(3)	(4) ME
Deposit Runoff	-0.28891*** (0.10179)	-0.11663 (0.10990)	0.07859 (0.22185)	-0.25095*** (0.08850)
Core Deposit Share			0.00048*** (0.00011)	
LLP Ratio	0.24152 (0.32063)	0.30955 (0.33020)	-1.58490** (0.62798)	0.21558 (0.31572)
Log Asset Size	0.03587*** (0.00544)	0.00958 (0.00745)	0.12576*** (0.01493)	0.04112*** (0.00686)
Year FE		X		
N	1784	1784	402	1784
R^2 / Pseudo R^2	0.03	0.07	0.15	0.04

Table 2.7: Columns (1)–(3) report Linear Probability Model estimates. Column (4) reports logit marginal effects evaluated at the mean. Robust standard errors in parentheses. Significance: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Column (1) of Table 2.7 presents the baseline linear probability model without year fixed effects. Deposit runoff enters at -0.289 and log asset size enters at $+0.036$, both significant at the one-percent level. Faster pre-failure deposit attrition materially reduces the probability that any bidder submits a P&A-viable offer, consistent with the franchise-value channel in the model.

Column (2) adds year fixed effects. The runoff coefficient loses significance. The runoff-to-P&A relationship is primarily a cross-episode phenomenon driven by aggregate deposit flight during crisis years. Year effects absorb the macroeconomic component of deposit flight common to all failing banks in a given year, leaving only idiosyncratic runoff variation whose marginal explanatory power for resolution type is limited.

Column (3) introduces core deposit share, defined as insured deposits as a fraction of total deposits, on the uninsured-share subsample. Core deposit share enters at $+0.00048$ with high significance, confirming that the stickiness and insurance-coverage properties of the liability

base emphasized by Drechsler et al. (2021) translate directly into resolution marketability. The loan loss provision ratio reverses sign to -1.585 , significant at five percent. Without core deposit share as a control, the provision ratio proxies jointly for asset-quality deterioration and for the FDIC's willingness to engineer a deal through loss-sharing arrangements. Once the franchise-quality component is absorbed by core deposit share, the provision ratio isolates the pure bad-asset component, which reduces P&A marketability. Column (4) reports logit marginal effects evaluated at the sample mean and confirms the linear probability model sign on runoff at -0.251 , significant at the one-percent level.

Log asset size is the single most robust selection predictor across all columns. Larger banks are disproportionately more likely to attract P&A acquirers, consistent with scale economies in due diligence and with the historical record that the FDIC has expended substantial effort to avoid liquidating large institutions during systemic episodes.

2.5.5 Heckman Corrected Deposit Pricing Channel

As described in the opening section, the pricing regression is estimated on the P&A subsample and the Heckman correction assesses whether this conditioning generates material bias in the pricing coefficients. Table 2.8 presents the key results. Column (1) reproduces the baseline OLS pricing regression for comparison; Columns (2) through (5) report the second-stage estimates of the Heckman et al. (1997) correction under alternative specifications, with case-bootstrap standard errors.

The columns differ in the exclusion restriction used to identify the first-stage probit and, in Column (5), the estimation sample. Column (2) (Congestion IV) uses the log number of concurrent failures in the same calendar quarter as the sole exclusion variable. The identification argument follows James (1991) and Granja et al. (2017): when many banks are failing simultaneously, the FDIC's finite resolution bandwidth and potential acquirers' finite attention reduce the probability that a specific bank attracts a timely P&A bid, but conditional on a sale occurring, the transaction price for that specific franchise is unaffected by how many other auctions are running in the same period. Column (3) (Cycle IV) replaces the congestion measure with the log number of days elapsed since the nearest macroeconomic

Table 2.8: Heckman Second Stage

	ln(1 + p) on the P&A Subsample				
	(1) OLS	(2) Congestion IV	(3) Cycle IV	(4) Over-Id	(5) β
Deposit Runoff	0.00763 (0.00518)	0.00648 (0.00757)	0.00819 (0.00871)	0.00685 (0.00714)	0.01676 (0.01463)
LLP Ratio	-0.00508 (0.01331)	-0.00249 (0.01801)	-0.00622 (0.02463)	-0.00335 (0.01500)	0.05407 (0.05832)
Log Asset Size	0.00033 (0.00035)	0.00040 (0.00046)	0.00029 (0.00055)	0.00038 (0.00041)	-0.00087 (0.00069)
Deposit Beta					-0.01387 (0.00941)
Inverse Mills Ratio		0.00628 (0.02359)	-0.00303 (0.04005)	0.00422 (0.01615)	0.03048 (0.05428)
Implied $\hat{\rho}$		0.419	-0.207	0.286	1.794
Implied $\hat{\sigma}_\epsilon$		0.015	0.015	0.015	0.017
Year FE	X	X	X	X	X
N	1488	1488	1488	1488	359
R^2	0.297	0.297	0.297	0.297	0.453

Table 2.8: OLS and Heckman two-step second-stage estimates of the log deposit premium on deposit-franchise controls, asset-quality proxy, and the inverse Mills ratio. Column (1) reports OLS on the P&A subsample without selection correction. Columns (2)–(5) report the second stage of a Heckman et al. (1997) two-step estimator, including the inverse Mills ratio computed from a first-stage probit of P&A selection estimated on the universe of failures. Standard errors in parentheses are heteroskedasticity-consistent in column (1) and case-bootstrapped (200 replications) in columns (2)–(5) to account for first-stage generated-regressor variability. Significance * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

crisis trough, a continuous measure of where any given failure falls in the credit cycle. The cycle variable captures a different dimension of resolution-market conditions: during troughs, acquirer balance sheets are most impaired and bidder participation is thinnest, reducing the probability of a P&A outcome for reasons unrelated to the target's own pricing fundamentals. Column (4) includes both exclusion variables simultaneously in an over-identified specification; the pricing coefficients serve as a robustness check on whether the two instruments are capturing the same or independent variation in selection. Column (5) restricts to the deposit-beta matched subsample of 359 observations and adds deposit beta as both a pricing

covariate and a selection covariate in the first stage.

The pricing coefficients are essentially unchanged by the correction. The deposit runoff coefficient moves from 0.0076 in OLS to 0.0065 in the primary corrected specification (column 2); the provision ratio coefficient moves from -0.0051 to -0.0025 ; log asset size is stable at approximately 0.0004. The inverse Mills ratio coefficient is 0.0063 with a bootstrap standard error of 0.0236, failing to reject the null of no selection bias at any conventional level. The key robustness check is the stability of the pricing coefficients between columns (2) and (3): despite the two exclusion variables operating through different economic channels—within-year congestion versus cross-year cycle position—the pricing estimates barely move, which strengthens the case that the conditional pricing equation is genuinely orthogonal to resolution-market conditions. The implied correlation between the pricing and selection residuals is approximately 0.42 in the primary specification, but the corresponding bias in the pricing coefficients is bounded by a quantity of order 0.008, which is comparable to the standard errors on the pricing coefficients themselves.⁸

The economic reason the correction is small is the narrow dispersion of the dependent variable. The sample standard deviation of the log deposit premium is approximately 0.015 log-points. Even a moderate selection correlation translates into a negligible absolute bias when the outcome variable has little total variance to begin with. This narrow dispersion is direct evidence that the FDIC auction delivers a competitive outcome in which the winning bid extracts most of the franchise surplus. The Heckman correction and the instrumental-variable strategy are complementary: the former addresses selection into the P&A sample, the latter addresses the simultaneity of the premium and the retained asset share at the auction. Full diagnostics, including the first-stage probit estimates and the decomposition of the selection correlation, are reported in Appendix B.2.2.

2.5.6 Summary

The three tests converge on a coherent description of how asset retention mediates the FDIC's resolution problem. The toxic ratio is the mechanical lever. Retaining distressed assets in

⁸The bound is computed as $|\hat{\rho}|\hat{\sigma}_\varepsilon\bar{\lambda} \approx 0.42 \times 0.015 \times 1.3 \approx 0.008$, where $\bar{\lambda}$ is the average inverse Mills ratio on the P&A subsample.

the receivership raises the deposit premium that acquirers are willing to pay, confirmed by the 2SLS estimate in the deposit pricing channel. It raises the probability that a buyer is found at all, through the loan loss provision and core deposit share channels in the P&A selection channel. And it raises the fiscal cost borne by the Deposit Insurance Fund, with a loan loss provision coefficient of +0.859 in the FDIC cost channel. All three channels carry the predicted sign at economically large magnitudes. The Heckman correction confirms that the pricing results are not an artifact of the P&A-only sample restriction.

2.6 Conclusion

This paper opens the black box of bank resolution by introducing a novel dataset of transaction terms from over forty years of FDIC press releases, transforming qualitative disclosures into a structured economic analysis of 1,860 failures. The central finding is that the apparent standardization of the modern resolution process masks a fundamental shift in regulatory technology. The transition from the chaotic liquidations of the 1980s to the monolithic dominance of the Purchase and Assumption transaction is not merely a function of industry consolidation or improved bidder appetite. It is the product of financial engineering. By retaining the most impaired assets in receivership regulators systematically cleanse distressed institutions of their troubled assets and convert them into marketable franchises. A structural model that encompasses both the liability and asset sides of the auction formalizes this mechanism, showing that the efficient partition of the asset portfolio between private and public management is a direct implication of the FDIC's least-cost mandate. The efficiency gain, however, is not free, the insurance fund absorbs the credit risk that private acquirers decline to carry.

The empirical record confirms this trade-off with consistency across all three channels of the analysis. Asset quality, as measured by the pre-failure loan-loss provision ratio, is the dominant quantitative determinant of resolution outcomes. It raises the cost borne by the Deposit Insurance Fund, depresses bid premiums unless offset by asset stripping, and reduces the probability of attracting a P&A acquirer at all. An instrumental-variable strategy that exploits the predetermined nature of the provision ratio relative to the auction confirms that

the positive relationship between asset retention and the deposit premium is causal rather than spurious. Deposit franchise quality matters in the direction the model predicts, but its quantitative influence is secondary to the asset-quality channel. The selection margin is dominated by bank size, with larger institutions resolving via purchase and assumption at rates disproportionate to their franchise fundamentals.

These findings carry direct implications for the design of the regulatory safety net. The reliance on asset retention reveals that the least-cost requirement operates with significant elasticity, effectively trading immediate marketability for retained credit risk. The fiscal cost of stripping distressed assets from a failing bank is not borne at closure; it accumulates silently on the FDIC's receivership balance sheet, socialized over the workout horizon. To the extent that market participants anticipate that the FDIC will engineer a sale for any sufficiently impaired institution by stripping its bad assets, this practice creates a distress-subsidy expectation that weakens market discipline and encourages risk-taking. The dominance of size in the selection regression adds another dimension to this moral hazard: large institutions enjoy implicit P&A protection at rates their deposit-franchise quality alone would not justify, a quantitative manifestation of the Too-Big-to-Fail doctrine.

Mitigating these distortions requires intervention earlier in the failure process. The evidence here shows that accelerated deposit attrition prior to closure destroys franchise value in a manner that is predictable *ex ante* and costly *ex post*. An eroded deposit base forces the FDIC to offer larger asset concessions to attract any bid at all. Intervening before the liability base is hollowed out would maximize auction premiums and reduce the necessity of offering expensive asset guarantees. The cyclical pattern of uninsured depositor protection compounds this problem. Systemic episodes yield near-total recovery through whole-bank P&A transactions, creating an expectation of *de facto* insurance coverage that persists into normal times and minimizes the market discipline that uninsured depositors are designed to provide. As the banking system enters a new era of digital deposit flight, where information spreads instantly and withdrawals can be executed in seconds, recalibrating these intervention thresholds is not a peripheral concern. It is the central challenge of resolution policy.

Chapter 3

Depositing Trust: Banking Demand and Too-Big-to-Fail

3.1 Introduction

Commercial banks are the engines of credit creation, but they run on a specific fuel... deposits. Deposits comprise over 90% of liabilities and three-quarters of total funding for the average bank. Consequently, the stability and cost of this funding source are the primary determinants of bank profitability and solvency. The United States banking sector is uniquely segmented in how it gathers this fuel, split between thousands of small “community banks” or “local banks” serving single towns, larger regional institutions, and a handful of Globally Systemically Important Banks. This paper investigates a central question of modern financial architecture: Did the government interventions during the Global Financial Crisis (GFC) distort the competitive landscape of this market, effectively subsidizing the largest banks at the expense of their smaller competitors?

The Financial Stability Board annually designates a set of Global Systemically Important Banks (GSIBs) on the basis of five criteria reflecting systemic risk: size, cross-jurisdictional activity, interconnectedness, substitutability, and complexity Financial Stability Board (2025). The rationale is that the disorderly failure of such an institution would disrupt critical financial services and impose cascading losses on the broader economy. In the

United States, four domestic commercial banks currently hold this designation: JPMorgan Chase, Bank of America, Citigroup, and Wells Fargo. Collectively, these four institutions hold approximately \$10 trillion in total assets Federal Reserve System (2026). Wells Fargo is the smallest GSIB at \$1.8 trillion in assets which is three times bigger than the next largest bank, US Bank.

The broader U.S. commercial banking sector is organized in a pronounced size hierarchy. Beneath the four GSIBs sits a tier of large regional banks, institutions with assets typically ranging from roughly \$100 billion to \$500 billion, followed by thousands of community and local banks whose operations are largely confined to single metropolitan areas or rural counties. As of recent years, over 4,000 federally insured commercial banks operate across the United States (Egan et al., 2017). This extreme size dispersion defines the competitive landscape studied in this paper: institutions of vastly different scale draw from the same national deposit base yet face markedly different regulatory obligations and perceived safety profiles.

Figure 3.1: GSIB Deposit Rate Spread




Figure 3.1: The top panel shows the deposit weighted average of the savings deposit rate for the four domestic GSIBs and all other commercial banks in the sample. The bottom panel shows the spread between the savings deposit rates of all other commercial banks minus the GSIBs scaled by the Federal Funds rate. Additionally, the right axis of the bottom panel shows the path of the Federal Funds rate. The data is collected quarterly between 1987 and 2022. The savings deposit rate is calculated from Call Reports by dividing the total interest expense on savings deposits by the total amount of savings deposits.

Figure 3.1 gives motivating evidence for this project. The top panel plots the aggregate savings deposit rates for GSIBs versus all other commercial banks. While a spread existed pre-GFC, the relationship fundamentally shifted following the crisis. To control for the dramatic decline in the interest rate environment, the bottom panel plots the “Savings Rate Ratio”, the spread between non-GSIB and GSIB rates scaled by the Federal Funds Rate.

The data reveals a stark regime change; prior to 2008, the scaled spread was relatively stable. Post-GFC, however, the ratio spikes dramatically, indicating that non-GSIB banks were forced to pay a significantly higher premium for deposits relative to GSIBs. Crucially, this advantage appears counter-cyclical: the premium paid by small banks spikes during periods of intense systemic distress, most notably during the GFC and the onset of the COVID-19 pandemic—and mean-reverts during periods of calm. While one might argue this dispersion is a mechanical artifact of the Zero Lower Bound (ZLB), the ratio exhibits significant variation between 2012 and 2016 even as the policy rate remained pinned at zero, suggesting the driver is economic, not mechanical.

This paper argues that this widening spread represents the price of safety. The architecture of the modern banking system is predicated on a fundamental tension, the transformation of liquid, short-term liabilities into illiquid, long-term assets. As formalized by Diamond and Dybvig (1983), this maturity mismatch exposes institutions to self-fulfilling runs where the mere expectation of insolvency creates actual insolvency. To mitigate this, governments have erected safety nets, implicitly guaranteeing that GSIBs are “Too-Big-To-Fail” (TBTF). The stylized facts in Figure 3.1 suggest that this guarantee distorts market discipline. During times of crisis, uninsured depositors flee to the safety of GSIBs, allowing these giants to lower their funding costs exactly when the rest of the market faces a liquidity squeeze.

The distinction between insured and uninsured depositors is central to this dynamic. The Federal Deposit Insurance Corporation (FDIC) guarantees deposits up to \$250,000 per depositor per institution, effectively eliminating credit risk for the majority of retail account holders. Below this threshold, a depositor has no incentive to withdraw preemptively; their principal is protected regardless of the bank’s condition. Uninsured depositors face a materially different problem. Their claims above \$250,000 are subject to potential losses in

the FDIC receivership process, exposing them to genuine default risk. It is this cohort that responds to deteriorating bank fundamentals by reallocating funds toward perceived safe havens Egan et al. (2017), and their capacity for coordinated withdrawal constitutes the primary source of banking system fragility.

The structural fragility of the U.S. banking system is perhaps best illustrated by the counterfactual analysis of Wachovia Bank during the 2008 financial crisis. Egan et al. (2017) estimate a structural model of the U.S. banking sector that confirms the existence of multiple equilibria. In March 2008, the market-implied probability of default for Wachovia, based on its fundamental asset quality and leverage, was a manageable 3.3%. However, the authors identified a second, “fragile” equilibrium consistent with those exact same fundamentals where the default probability skyrocketed to 52%. This stark divergence highlights the central thesis of this paper: large banks, inoculated by implicit government guarantees, persistently live in the “good” equilibrium. Even if their fundamentals deteriorate, the expectation of a bailout prevents the coordinated withdrawal of uninsured funds. Conversely, smaller banks, lacking this credible guarantee, remain perpetually exposed to the “bad” equilibrium.

This disparity has significant implications for the political economy of bank stability. Following the GFC, regulatory frameworks such as the Dodd-Frank Act explicitly aimed to end the TBTF regime. However, the effectiveness of these reforms remains debated. Berndt et al. (2025) estimate that while creditors now price in some risk of failure, the market still assigns an overwhelming 87% probability that the government will save a GSIB in a crisis. The implicit guarantee has not been eliminated; it has merely been repriced.

This persistent expectation of support drives the “Flight to Safety” dynamic investigated here. Testing the magnitude of this distortion, however, presents a significant identification challenge. A naive comparison of deposit rates is confounded by non-regulatory factors, such as the “sophistication paradox” identified by d’Avernas et al. (2023), where wealthy consumers accept lower rates at large banks in exchange for superior technology (e.g., mobile apps, expansive ATM network). This paper resolves this ambiguity by exploiting the temporal dynamics of the safety premium. Technological advantages are static over short

horizons, whereas the TBTF guarantee is highly counter-cyclical. By using a difference-in-differences methodology that isolates the change in deposit spreads during stress events, I disentangle the safety premium from the technology premium.

To formalize the channel through which TBTF guarantees distort deposit market competition, I develop a structural model that integrates the spatial differentiation framework of Hotelling (1929) with the endogenous default mechanics of Merton (1974) and Leland (1994). In this economy, banks compete for funding not only on the basis of offered rates and spatial convenience but also on their perceived safety. The model posits a fundamental asymmetry: while the large “National” bank effectively enjoys a regulatory backstop that insulates its depositors from loss, the smaller “Local” bank is subject to default risk, forcing it to pay a credit risk premium to retain uninsured depositors. This framework generates a sharp, testable implication; the deposit spread between large and small banks should not be constant, but rather a convex function of the smaller bank’s financial health. Specifically, as a small bank’s probability of default rises, the model predicts a “flight-to-safety” dynamic where the small bank must aggressively raise rates to compensate for risk, while the TBTF incumbent strategically widens its net interest margin, effectively monetizing the systemic distress. This theoretical prediction provides a structural foundation for my empirical tests, which aim to identify whether the sensitivity of deposit flows to bank solvency is indeed magnified for unprotected institutions during periods of aggregate stress.

Operationalizing these model implications requires a reliable, continuous measure of default probability for the universe of small banks. Unlike GSIBs, whose risk is priced in liquid Credit Default Swap (CDS) markets, most regional banks lack market-based risk indicators. To overcome this, I construct a novel proxy for bank distress using a supervised machine learning framework. Building on the bankruptcy prediction literature, I employ a weighted ensemble model that aggregates predictions from multiple machine learning architectures. This approach addresses the limitations of traditional linear models by capturing the complex, non-linear interactions inherent in banking data.

The model is trained on a rich dataset of quarterly regulatory filings spanning 2002 to 2023. The target variable is the market-implied probability of default derived from the CDS

spreads of the largest U.S. banks. By mapping the accounting fundamentals of these large banks to their market-priced risk, the model learns a transfer function that can be applied to smaller banks to estimate their default probability. This methodological contribution allows for the first comprehensive analysis of deposit elasticity across the entire size spectrum of the U.S. banking system.

The empirical analysis validates the flight-to-safety mechanism predicted by the theoretical model, particularly through the quantity channel. As presented in Table 3.3, I find a stark hierarchy in how uninsured depositors respond to bank distress in the post-GFC regime. With local banks as the reference category, the post-GFC triple interaction terms for both regional banks and GSIBs are positive and statistically significant, indicating that larger institutions attract rather than repel uninsured capital as their default risk rises. The effect is largest for GSIBs: a one-percentage-point increase in default probability is associated with an approximately 48 percent increase in a GSIB's uninsured deposit share relative to the outside option, while local banks experience an essentially flat or slightly negative relationship. Crucially, this sensitivity is effectively absent for insured deposits, isolating the mechanism to those liabilities most exposed to credit risk, and the result survives a within-bank triple-difference that controls for time-varying bank quality.

The results for the pricing channel, presented in Table 3.4, offer corroborating but statistically weaker evidence. While the model predicts that GSIBs should be able to lower their deposit rates during distress, the empirical identification of this effect is constrained by the structural characteristics of the deposit market. As shown by the high R^2 values exceeding 0.99 in the spread regressions, the vast majority of variation in deposit rates is absorbed by time and bank fixed effects, leaving little residual variation to be explained by idiosyncratic risk. This lack of statistical power is consistent with the well-documented “stickiness” of deposit rates Drechsler et al. (2017); because banks are hesitant to adjust rates frequently due to administrative costs and customer inertia, the primary margin of adjustment during stress events is deposit flows rather than interest rates. Consequently, while the signs of the coefficients align with the TBTF distortion, suggesting GSIBs enjoy a pricing advantage, the flight-to-safety is principally observed through the massive reallocation of volume shown in

Table 3.3.

The remainder of this paper is organized as follows. Section 3.2 reviews the related literature on banking stability, market discipline, and TBTF distortions. Section 3.3 develops the continuous-time structural model of deposit competition and derives the equilibrium pricing conditions. Section 3.4 details the empirical strategy, including the construction of the machine learning default model, the difference-in-differences specification, and the main results. Section 3.5 concludes with a discussion of policy implications.

3.2 Related Literature

The theoretical foundation of banking fragility rests on the seminal model of Diamond and Dybvig (1983), who characterize bank runs as equilibrium outcomes driven by self-fulfilling panic rather than solely by fundamental insolvency. While this framework supports multiple equilibria, Goldstein and Pauzner (2005) refine the theory using global games to demonstrate that runs are triggered when fundamentals deteriorate below a specific threshold, re-establishing the link between asset quality and depositor flight. This transmission of distress is not purely financial; Kelly and Gráda (2000) and Iyer and Puri (2012) document that social networks amplify contagion, a dynamic accelerated by modern technology as evidenced by the “digital runs” documented by Cookson et al. (2023).

A structural assessment of this fragility is provided by Egan et al. (2017), who confirm that the U.S. banking sector is characterized by multiple equilibria. Their work reveals a stark bifurcation in market discipline: while insured depositors are insensitive to risk, the demand for uninsured deposits is highly elastic. Their estimates suggest that a one percentage point increase in a bank’s risk-neutral probability of default leads to a 12% decline in its market share of uninsured deposits. This elasticity creates a “bad” equilibrium where fear of default becomes self-fulfilling. Their counterfactual analysis of Wachovia Bank in 2008 illustrates this mechanism, while Wachovia’s fundamentals implied a default probability of only 3.3%, the authors identified an alternative equilibrium consistent with the same fundamentals where the default probability was 52%. This finding is critical to this paper, as it suggests that smaller banks without implicit guarantees remain structurally exposed to

this “fragile” equilibrium.

While government guarantees stabilize the system against such runs, they introduce profound moral hazard. Stern and Feldman (2004) argue that TBTF subsidies distort risk pricing, encouraging banks to leverage up and capture the upside of risky bets while taxpayers bear the tail risk. This subsidy is empirically robust across asset classes. In the bond market, Santos (2014) provides evidence that large banks enjoy a significant funding cost advantage that is counter-cyclical; the spread between TBTF and non-TBTF costs widens dramatically during crises when the guarantee is most valuable. Acharya et al. (2016) corroborate this, finding that prior to the Dodd-Frank Act, bond spreads for large financial institutions were remarkably insensitive to risk, implying a near-total implicit guarantee.

This distortion extends to equity markets and corporate strategy. Gandhi and Lustig (2015) identify a size anomaly where large bank stocks yield significantly lower risk-adjusted returns than small bank stocks. They argue this reflects an “insurance premium” investors pay for the government’s tail-risk protection, which accrues to the bank as a lower cost of equity capital. Consequently, banks are incentivized to pursue scale not for efficiency, but to secure this subsidy; Penas and Unal (2004) document that bond returns following mergers reflect the attainment of “TBTF status” rather than purely economic synergies.

Recent literature debates whether post-2008 regulatory reforms have successfully eliminated these distortions. Berndt et al. (2025) provide a current assessment, utilizing a reduced-form asset pricing model to decompose credit spreads into insolvency risk and bailout probability. They find that while the risk-neutral probability of not getting bailed out rose from approximately 5% pre-crisis to 13% post-crisis, the market still assigns an overwhelming 87% probability that the government will save a GSIB. They estimate the present value of this remaining subsidy at 1.6% of GDP. This persistent expectation confirms that despite statutory changes, the flight to safety dynamic remains a potent competitive force, driving uninsured deposits toward GSIBs during stress events.

Identifying this safety premium is complicated by non-regulatory factors. d’Avernas et al. (2023) argue that the funding advantage of large banks may stem from a sophistication paradox, where wealthy clients accept lower rates in exchange for superior mobile technology and

liquidity services. This paper addresses this identification challenge by exploiting the temporal variation in distress; while technological advantages are static, the safety premium is counter-cyclical. Finally, to operationalize this analysis for smaller banks lacking market data, we build upon the bankruptcy prediction literature, evolving from the static discriminant analysis of Altman (1968) and Ohlson (1980) to modern Machine Learning techniques that better capture non-linear distress signals.

3.3 Model

In this section, we develop a structural model of deposit competition that integrates the spatial differentiation framework of Hotelling (1929) with the endogenous default risk mechanics of Merton (1974) and Leland (1994). The primary objective is to formalize the mechanism through which TBTF guarantees distort deposit pricing. By moving to a continuous-time setting, we can analyze how the equilibrium deposit spread evolves dynamically as a function of the smaller bank’s distance-to-default.

3.3.1 Model Setup

Let $(\Omega, \mathbb{F})$ be the given probability space and canonical filtration. We consider a continuous-time, infinite-horizon economy populated by a unit mass of risk-neutral agents. The banking sector consists of two competing depository institutions: a National Bank (N) and a Local Bank (L). Following the standard spatial competition literature, these banks are located at the extremes of a unit interval $[0, 1]$, with bank N at $x = 0$ and bank L at $x = 1$.

The fundamental asymmetry in the model arises from the regulatory treatment of these institutions. The National Bank is designated as a Systemically Important Financial Institution (SIFI). Consistent with the “implicit guarantee” literature (e.g., O’Hara and Shaw (1990); Gandhi and Lustig (2015)), we assume this designation renders the National Bank “Too-Big-To-Fail.” Consequently, its deposits are perceived by agents as risk-free, regardless of the bank’s underlying asset performance. Conversely, the Local Bank does not enjoy this protection. Its deposits are subject to credit risk, and in the event of default, uninsured

depositors suffer a loss of principal.

The Local Bank holds a portfolio of risky assets, V_t , which evolves according to a Geometric Brownian Motion (GBM) under the risk-neutral measure $\mathcal{Q}$. Following Merton (1974), the dynamics of the bank's asset value are given by

$$\frac{dV_t}{V_t} = (r_f - \delta)dt + \sigma dZ_t, \quad (3.1)$$

where r_f is the constant net risk-free rate, δ represents the total payout rate (dividends and coupons) to claimholders, σ is the constant volatility of the asset portfolio, and Z_t is a standard Brownian motion. δ prevents the bank from growing indefinitely and ensures a stationary leverage ratio in the absence of shocks, consistent with the stationary debt structure in Leland (1994). The initial asset value is V_0 , and we assume V_0 is sufficiently high such that the bank is initially solvent.

A critical departure from the simple static formulation is the treatment of default. Rather than a fixed exogenous probability, default is an endogenous stopping time driven by the solvency decisions of the Local Bank's shareholders. Following Black and Cox (1976), default occurs the first time the bank's asset value V_t touches a lower barrier V_B . We define the default time τ as

$$\tau := \inf \{t \geq 0 : V_t \leq V_B\} \quad (3.2)$$

The barrier V_B is not imposed by debt covenants as in Merton (1974) but is chosen endogenously by equity holders to maximize the value of their claim, a feature central to the trade-off models of Leland (1994) and Leland and Toft (1996). We derive the optimal V_B in Section C.1.

3.3.2 Valuing Default Risk

Before analyzing the deposit market equilibrium, we must establish how risk-neutral depositors value the contingent loss associated with the Local Bank. Depositors face a potential loss given default (LGD) of $\alpha \in (0, 1]$. If the Local Bank defaults at time τ , depositors

recover $(1 - \alpha)$ of their deposit face value.

To price this risk, we utilize the standard pricing kernel for a “first-passage-time” security. Let $P(V_t)$ denote the present value of a claim that pays \$1 at the moment default occurs. Specifically, $P(V_t) = \mathbf{E}_{\mathcal{Q}}[e^{-r_f \tau} \mid V_t]$. For the GBM process specified in Equation (3.1), the closed-form solution for this security is

$$P(V_t) = \left(\frac{V_t}{V_B} \right)^{-\gamma}, \quad (3.3)$$

where $-\gamma$ is the positive root of the fundamental quadratic equation derived from the generator of the diffusion process

$$\frac{1}{2}\sigma^2\gamma(\gamma + 1) + (\delta - r_f)\gamma - r_f = 0.$$

The term $P(V_t)$ can be interpreted as the risk-neutral probability of default discounted at the risk-free rate. It captures the “distance-to-default” in value space. As $V_t \rightarrow \infty$, $P(V_t) \rightarrow 0$, and as $V_t \rightarrow V_B$, $P(V_t) \rightarrow 1$. This valuation serves as the basis for the “credit premium” depositors demand from the Local Bank.

3.3.3 Deposit Market

We model the deposit market using a continuous-time extension of the Hotelling framework. There is a unit mass of uninsured depositors uniformly distributed along the line $x \in [0, 1]$. Depositors seek to maximize their instantaneous flow utility, which consists of the interest income net of transport costs and expected credit losses.

Let r_N and r_L denote the instantaneous deposit rates offered by the National and Local Banks, respectively. Depositors at the local bank have their transport cost normalized to 1 and the national bank has a convenience yield $\phi \in (0, 1]$. This parameter ϕ proxies for the degree of differentiation between banks, capturing factors such as branch density, digital app quality, or brand loyalty for the national bank d’Avernas et al. (2023).

For a depositor located at x , the flow utility from depositing at the risk-free National Bank is simply the net yield

$$U_N(x) = r_N - \phi x. \quad (3.4)$$

For the Local Bank, the depositor faces credit risk. The flow utility must account for the instantaneous intensity of the expected loss. We define the instantaneous credit spread required to compensate for the default claim $P(V_t)$ as $\Psi(V_t, V_B)$. Thus, the utility from the Local Bank is

$$U_L(x) = r_L - \alpha \Psi(V_t, V_B) - (1 - x). \quad (3.5)$$

Here, $\alpha \Psi(V_t, V_B)$ represents the flow cost of credit risk. Intuitively, this is the insurance premium a depositor would need to pay in a perfectly competitive market to insure their deposit against the loss α .

The marginal depositor, located at x^* , is indifferent between the two banks when $U_N(x^*) = U_L(x^*)$. Solving this equality for x^* yields the market share of deposits at the national bank,

$$M_N(V_t, V_B) = \frac{r_N - r_L + \alpha \Psi(V_t, V_B) + 1}{1 + \phi}, \quad (3.6)$$

with the local bank taking the other $M_L = 1 - M_N$ mass of deposits.

Equation (3.6) reveals the central tension of the model. The Local Bank's market share is determined by three forces: (1) the baseline spatial split ($1/(1 + \phi)$), (2) the interest rate differential ($r_N - r_L$), and (3) the TBTF distortion term $\alpha \Psi(V_t, V_B)$. As the Local Bank's asset value V_t declines, the risk term $\Psi(V_t, V_B)$ rises, forcing the bank to either accept a lower market share or compress its net interest margin by raising r_L .

3.3.4 Bank Competition

We assume the Local Bank maximizes shareholder value. This involves two distinct optimization problems: a static decision regarding the deposit rates r_N and r_L to offer at each instant, and a dynamic decision regarding when to exercise the option to default, for the local bank.

Shareholders hold a residual claim on the bank's assets. Let $E(V_t)$ denote the value of

the Local Bank's equity. The bank pays a total flow of coupons to depositors, denoted C . For tractability, we approximate the total interest expense as a constant proportional to the steady-state deposit base, $C = r_L \bar{D}$.

The default problem for the shareholders is an optimal stopping problem. In the continuation region $V_t > V_B$, the value of equity must satisfy the standard asset pricing differential equation

$$\frac{1}{2}\sigma^2 V^2 E''(V) + (r_f - \delta)V E'(V) - rE(V) + \delta V - C = 0. \quad (3.7)$$

The term $(\delta V - C)$ represents the net cash flow to equity holders. To solve this second-order ordinary differential equation (ODE), we impose three boundary conditions standard in the structural credit risk literature Leland (1994)

- Limited Liability: At the default barrier, equity value is zero

$$E(V_B) = 0.$$

- Solvency: As asset values grow arbitrarily large, the probability of default vanishes, and equity value approaches the risk-free value of assets net of debt

$$\lim_{V \rightarrow \infty} E(V) = V - \frac{C}{r_f}.$$

- Smooth Pasting: Shareholders choose the default barrier V_B optimally. This requires the derivative of the equity value to be continuous at the boundary

$$E'(V_B) = 0.$$

Solving the Euler-Cauchy ODE subject to the first two boundary conditions yields the general solution

$$E(V) = V - \frac{C}{r_f} + \left(\frac{C}{r_f} - V_B \right) \left(\frac{V}{V_B} \right)^{-\gamma}.$$

Applying the smooth pasting condition, we derive the closed-form solution for the endogenous default barrier

$$V_B = \frac{C}{r_f} \left(\frac{\gamma}{\gamma + 1} \right). \quad (3.8)$$

Equation (3.8) provides a micro-founded threshold for default. The term $\frac{\gamma}{\gamma+1} < 1$ indicates that shareholders will keep the bank operating even when the value of assets falls below the face value of liabilities. This “gambling for resurrection” behavior arises from the option value of limited liability.

We assume banks compete in rates continuously to maximize instantaneous expected profits, taking the asset value V_t and the default barrier V_B as given. The profit maximization problem for bank $k \in \{N, L\}$ is

$$\max_{r_k \in \mathbb{R}_{++}} \Pi_k \equiv (R - r_k)M_k \quad (3.9)$$

where R is the expected gross return on assets.

The FOCs for the bank optimization problem relates the marginal benefit of attracting more depositors and the marginal cost of the additional cost of capital that each depositor earns.

3.3.5 Equilibrium

Combining the demand side and supply side, we arrive at the main theoretical result of the paper, the Nash equilibrium deposit spread. In a symmetric Hotelling model without risk, the equilibrium spread only depends on the amenity value that the national bank offers. However, the presence of the TBTF distortion increases the deposit spread. The equilibrium condition requires the Local Bank to fully compensate the marginal depositor for the expected credit loss to maintain its market share.

Proposition 4 (Equilibrium Deposit Rates).

In the Nash equilibrium, the optimal deposit rates set by the National Bank (r_N^) and the Local Bank (r_L^*) are given by:*

$$r_N^*(V_t, V_B) = R - \frac{2 + \phi}{3} - \frac{\alpha}{3}\Psi(V_t, V_B) \quad (3.10)$$

$$r_L^*(V_t, V_B) = R - \frac{1 + 2\phi}{3} + \frac{\alpha}{3}\Psi(V_t, V_B) \quad (3.11)$$

where $\Psi(V_t, V_B)$ is the instantaneous credit yield derived from the structural model

$$\Psi(V_t, V_B) = r \left(\frac{(V_t/V_B)^{-\gamma}}{1 - (V_t/V_B)^{-\gamma}} \right). \quad (3.12)$$

As V_t approaches V_B , the term $(V_t/V_B)^{-\gamma}$ approaches 1, causing the denominator in Ψ to vanish and the spread $\Psi(V_t, V_B)$ to diverge. This convexity is a key prediction of the model; the TBTF subsidy is relatively small during good economic times but expands non-linearly during periods of financial distress.

The equilibrium deposit rates exhibit a distinct “scissors” dynamic that reflects the strategic exploitation of the TBTF subsidy. In the absence of credit risk, both banks set rates based on the marginal product of capital, R , markdown by the spatial differentiation parameter, ϕ , which serves as a measure of the National Bank’s convenience yield. However, as the Local Bank’s distance-to-default shortens, the pricing strategies diverge. The Local Bank is forced to raise its rate, r_L , to compensate the marginal depositor for the escalating intensity of expected credit losses, $\alpha\Psi(V_t, V_B)$. Simultaneously, the National Bank optimally lowers its rate, r_N . This markdown occurs because the Local Bank’s increasing risk profile reduces its substitutability, effectively softening the competitive constraint on the National Bank; the TBTF incumbent exploits this reduced demand elasticity to widen its net interest margin, monetizing the flight to safety by offering lower yields while still retaining market share.

Corollary 2 (Equilibrium Market Structure and Profitability). *Given the Nash equilibrium deposit rates r_N^* and r_L^* , the resulting equilibrium deposit spread S^* , market shares M_k^* , and*

profit functions Π_k^* are characterized as follows:

(ii) The equilibrium deposit rate spread between the local and national banks is

$$S^*(V_t, V_B) := r_L^* - r_N^* = \frac{1 - \phi}{3} + \frac{2}{3}\alpha\Psi(V_t, V_B). \quad (3.13)$$

(ii) The mass of depositors retained by the National Bank (M_N^*) and the Local Bank (M_L^*) are:

$$M_N^*(V_t, V_B) = \frac{2 + \phi + \alpha\Psi(V_t, V_B)}{3(1 + \phi)} \quad M_L^*(V_t, V_B) = \frac{1 + 2\phi - \alpha\Psi(V_t, V_B)}{3(1 + \phi)}.$$

(ii) The equilibrium profits for each bank are given by

$$\begin{aligned} \Pi_N^*(V_t, V_B) &= (1 + \phi) \left(\frac{2 + \phi + \alpha\Psi(V_t, V_B)}{3(1 + \phi)} \right)^2 \\ \Pi_L^*(V_t, V_B) &= (1 + \phi) \left(\frac{1 + 2\phi - \alpha\Psi(V_t, V_B)}{3(1 + \phi)} \right)^2 \end{aligned}$$

It is important to note that the model dynamics permit corner solutions during periods of extreme systemic stress. As the Local Bank's probability of default rises sufficiently high, the flight-to-safety becomes absolute, driving the National Bank's market share toward unity, $M_N^* \rightarrow 1$. In this limit, the Local Bank is effectively crowded out of the funding market; its deposit base evaporates, and its expected profits collapse to zero, precipitating an immediate default and market exit. Simultaneously, having captured the entire market, the National Bank transitions to an unconstrained monopoly pricing strategy. Crucially, the model assumes that depositors face a prohibitive cost to accessing outside options, such as money market funds or physical storage (the “mattress”). Without a viable alternative store of value to enforce a participation constraint, the National Bank can optimally set deposit rates to zero or negative levels to extract maximal consumer surplus from the captive depositor base, monetizing the systemic instability.

The decomposition of the equilibrium deposit spread in (i) reveals that the pricing differential between GSIBs and smaller institutions is driven by two distinct economic forces: a structural convenience yield and a state-dependent risk premium. The first component, $(1 - \phi)/3$, represents the baseline “convenience premium” that the National Bank enjoys solely due to its superior network density and technological sophistication; even in the absence of default risk, the Local Bank must pay a premium to overcome the National Bank’s non-pecuniary amenity value. The second component, $\frac{2}{3}\alpha\Psi$, captures the TBTF distortion. Notably, the pass-through of credit risk to the spread is incomplete (two-thirds rather than unity) because the Nash equilibrium involves mutual adjustment; the Local Bank absorbs one-third of the shock through margin compression to retain depositors, while the National Bank strategically accommodates the remaining two-thirds by widening the spread from the bottom, lowering its own rate to exploit its rival’s distress.

This decomposition provides a structural interpretation for the behavior of insured versus uninsured deposit spreads. For fully insured deposits, the depositor faces no credit loss in the event of default, implying an effective loss-given-default of zero ($\alpha = 0$). In this limiting case, the risk premium term vanishes $\frac{2}{3}\alpha\Psi \equiv 0$, and the deposit spread collapses to the constant convenience yield, $(1 - \phi)/3$. This result aligns with the findings of d’Avernas et al. (2023), who argue that the persistent spread between large and small banks is driven by differences in “banking services” and technology rather than credit risk. Our model nests this service quality view within the convenience parameter ϕ but demonstrates that for uninsured depositors, who cannot ignore Ψ , the TBTF subsidy drives a wedge beyond what quality differences alone can explain, particularly during periods of financial distress.

Item (ii) quantifies the flight-to-safety dynamic, demonstrating that the National Bank’s market share expands linearly with the Local Bank’s credit intensity, Ψ , despite the National Bank offering lower deposit rates. This result highlights a breakdown in the standard price elasticity of demand during periods of stress; as the probability of default rises, the safety attribute of the National Bank becomes the dominant factor in depositor utility, overwhelming the yield differential. Consequently, depositors effectively become price-insensitive to the National Bank’s rate cuts, allowing the TBTF incumbent to hoard liquidity precisely when

its competitor is experiencing a runoff. This mechanism formalizes the empirical observation that deposit flows during financial crises are driven not by the pursuit of yield, but by the avoidance of tail risk, granting TBTF institutions a distinct funding advantage that accelerates as systemic stress deepens.

Finally, the profit functions in (iii) illustrate the convex nature of the TBTF subsidy. Because equilibrium profits are proportional to the square of market share, the National Bank acts as the beneficiary of its competitor's fragility, experiencing a compounding increase in profitability as the Local Bank approaches the default barrier. This double-dipping, gaining both higher deposit volume and wider net interest margins, implies that TBTF protections do not merely shield large banks from their own failures, but actively transfer franchise value from distressed smaller banks to the protected sector. Conversely, the Local Bank faces a convex penalty; as it attempts to defend its deposit base with higher rates, its net interest margin collapses faster than its market share, potentially accelerating the depletion of equity capital and creating a self-reinforcing death spiral that justifies the endogenous decision to default earlier than book-value insolvency would suggest.

The following comparative statics elucidate the economic mechanisms driving the flight-to-safety dynamics and the competitive distortion introduced by TBTF guarantees.

The most direct determinant of the deposit spread is the financial health of the Local Bank, proxied by its current asset value V_t . Differentiating the equilibrium spread equation with respect to V_t yields

$$\frac{\partial S^*}{\partial V_t}(V_t, V_B) = \frac{2}{3}\alpha \frac{\partial \Psi}{\partial V_t}(V_t, V_B) < 0.$$

Since the credit intensity Ψ is proportional to the first passage time probability $(V_t/V_B)^{-\gamma}$, and the exponent $-\gamma$ is negative, this derivative is strictly negative. Mathematically, this confirms that as the Local Bank's asset value declines towards the endogenous default barrier V_B , the deposit spread widens non-linearly. This result generates the flight-to-safety dynamic observed in the empirical data. In “good times” (when $V_t \gg V_B$), the default probability is negligible, and the spread is driven almost entirely by the structural convenience yield, $(1-\phi)/3$. However, as the bank enters distress, the TBTF distortion activates. The convexity

of the function implies that a small deterioration in asset value for a healthy bank has little impact on funding costs, whereas the same deterioration for a distressed bank causes a sharp spike in the spread. Importantly, in our Nash equilibrium, this widening is driven by two simultaneous reactions: the distressed Local Bank raising rates to retain wary depositors, and the TBTF National Bank opportunistically lowering rates to exploit its rival's weakness.

While the analysis of deposit spreads illustrates how prices adjust to distress, the sensitivity of market shares, M_k^* , reveals the flow of funds and the resulting concentration of systemic liquidity. Differentiating the equilibrium market share of the National Bank, M_N^* , with respect to the Local Bank's asset value yields

$$\frac{\partial M_N^*}{\partial V_t}(V_t, V_B) = \frac{\alpha}{3(1+\phi)} \frac{\partial \Psi}{\partial V_t}(V_t, V_B) < 0.$$

Since $\partial_{V_t} \Psi < 0$ the inverse relationship implies that as the Local Bank's asset value falls (distress increases), the National Bank's market share expands. This result demonstrates that the price mechanism is insufficient to clear the market during a crisis. In a standard competitive market, the Local Bank would raise rates until demand stabilizes. However, in our Nash equilibrium, despite the Local Bank raising rates and the National Bank lowering them, the Local Bank still suffers a net outflow of deposits. This indicates that the demand elasticity with respect to safety overwhelms the elasticity with respect to yield. The National Bank acts as a liquidity sink, absorbing deposits purely due to its TBTF status. This dynamic suggests that TBTF distortions lead to procyclical deposit concentration: precisely when the banking system is most fragile, funding flees the periphery and concentrates in the systemic core, potentially exacerbating TBTF risks by making the largest institutions even larger during downturns.

The parameter ϕ captures the convenience yield or non-pecuniary amenity value of the National Bank, such as a superior branch network or digital infrastructure. We examine the sensitivity of both the equilibrium spread and market share to this structural friction

$$\frac{\partial S^*}{\partial \phi}(V_t, V_B) = -\frac{1}{3} < 0 \quad \frac{\partial M_N^*}{\partial \phi}(V_t, V_B) = -\frac{1 + \alpha \Psi(V_t, V_B)}{3(1 + \phi)^2} < 0.$$

The negative derivative of the spread isolates the structural component of the National Bank's funding advantage; a lower ϕ (indicating higher convenience) increases the baseline spread $(1 - \phi)/3$, allowing the National Bank to pay consistently lower rates independent of credit risk. Simultaneously, the sensitivity of market share reveals a crucial interaction effect: the presence of the risk term $\alpha\Psi$ in the numerator implies that the volume gains from superior technology are magnified during periods of distress. When depositors are sensitized to default risk, they become hyper-sensitive to frictions. Consequently, large banks that invest in friction-reducing technologies secure a double dividend, capturing a disproportionately larger share of flight-to-safety flows while maintaining a wider structural wedge in their cost of funds.

The volatility of the bank's underlying assets, σ , affects the spread through two opposing channels: the kinetic effect, which increases the probability of hitting the barrier, and the strategic effect, which lowers the optimal default barrier V_B as shareholders delay default to preserve option value. The net effect on the spread is positive

$$\frac{\partial S^*}{\partial \sigma}(V_t, V_B) > 0.$$

This static highlights how aggregate uncertainty exacerbates TBTF distortions. During periods of high market volatility, such as the Global Financial Crisis or the COVID-19 pandemic, the TBTF National Bank gains a competitive advantage even if the Local Bank's expected asset value remains constant. Depositors, fearing that higher variance increases the likelihood of a left-tail event, demand a higher risk premium from the unprotected Local Bank. Meanwhile, the TBTF bank, shielded by the implicit guarantee, sees its funding costs remain insensitive to the volatility shock. This suggests that TBTF policies distort competition most severely during turbulent market regimes.

The parameter α represents the Loss Given Default (LGD) faced by uninsured depositors at the Local Bank. The sensitivity is straightforward

$$\frac{\partial S^*}{\partial \alpha}(V_t, V_B) = \frac{2}{3}\Psi(V_t, V_B) > 0.$$

This static formalizes the direct link between regulatory credibility and market pricing. The magnitude of the TBTF distortion scales linearly with the expected loss severity. If regulators signal a credible commitment to impose losses on uninsured depositors at small banks, increasing α , the spread widens. Conversely, if a blanket guarantee is effectively extended to all institutions, as seen in the immediate aftermath of the Silicon Valley Bank failure FDIC (2023) $\alpha \rightarrow 0$, the risk-driven component of the spread collapses. This implies that ambiguity regarding the scope of deposit insurance is a primary driver of the cost of capital differential between large and small banks.

This discussion elucidates the economic mechanisms driving the flight-to-safety. The model is calibrated using an insured deposit spread of 17 basis points and a historical uninsured recovery rate of $1 - \alpha = 91\%$ found in Perry (2026). In this analysis, the probability of default proxies for aggregate sector distress. Crucially, the model features asymmetric risk exposure: while the Local Bank’s solvency is tethered to this systemic risk, the National Bank’s TBTF designation effectively inoculates it from asset volatility.

The bifurcation in equilibrium pricing strategies is vividly illustrated in Figure 3.2. The data reveals a distinct “scissors” dynamic predicted by Proposition 4. As the probability of default rises, the Local Bank is compelled to aggressively raise its offered rate to compensate uninsured depositors for the escalating expected loss. Conversely, the National Bank responds to this same distress by lowering its deposit rate. This strategic markdown occurs because the Local Bank’s increasing risk profile reduces its substitutability, softening the competitive constraint on the incumbent. Consequently, in the region of extreme distress where the default probability approaches unity, the National Bank can optimally effectively zero-out its deposit rate. In this regime, the safety attribute dominates the depositor’s utility function, allowing the TBTF bank to retain funds while paying negligible pecuniary returns.

Figure 3.3 translates these pricing dynamics into equilibrium market shares and profitability, highlighting the quantity channel of the distortion. The left panel demonstrates that despite the Local Bank’s efforts to defend its funding base through higher rates, it suffers a monotonic loss of market share as risk rises. This implies that demand elasticity with respect to safety overwhelms elasticity with respect to yield; during crises, price adjustments

Figure 3.2: Equilibrium Deposit Rates




Figure 3.2: The top panel plots the equilibrium offered deposit rates in basis points for the Local Bank (blue line) and National Bank (yellow line) as a function of the Local Bank's probability of default. The bottom panel displays the resulting deposit rates for uninsured deposits (blue line) and insured deposits (yellow dashed line) across the range of default probabilities.

fail to clear the market for unprotected banks. The right panel underscores the destabilizing implications for profitability. While the Local Bank's profits collapse toward zero as it faces both volume outflows and margin compression, the National Bank's profits exhibit a convex, counter-cyclical profile. By capturing cheap flight-to-safety flows, the TBTF institution enjoys a profit boost precisely when the sector is failing, effectively transferring franchise value from the periphery to the core.

To examine these mechanics in a continuous setting, Figure 3.4 presents a single simulation path of the sector's asset value.

Figure 3.4 confirms that the safety premium is not static but behaves as a highly volatile, state-dependent variable. During periods of relative calm (e.g., years 0–1), the spread remains compressed, primarily reflecting the baseline convenience yield ϕ . However, as as-

Figure 3.3: Equilibrium Deposit Market Shares




Figure 3.3: The left panel plots the equilibrium deposit market shares for the Local Bank and National Bank as a function of the probability of default, with the dotted line indicating an equal 50% share. The right panel plots the scaled equilibrium profit functions for both institutions against the same domain of default probabilities.

set values deteriorate and the probability of default spikes, the spread widens non-linearly. This convexity highlights the option-like nature of the implicit guarantee; the cost of being small is negligible in benign economic environments but expands dramatically during distress episodes, forcing the spread to widen exactly when the Local Bank can least afford it.

The raw time-series dynamics presented in Figure 3.5 provide a granular view of the scissors mechanism in operation. As economic conditions deteriorate, evidenced by the spikes in default probability, we observe a simultaneous divergence in bank behavior: the National Bank lowers rates and gains deposit inflows, while the Local Bank raises rates and suffers outflows. This inverse correlation confirms that the theoretical predictions hold dynamically; the National Bank effectively monetizes the systemic instability by widening its net interest margin during the very periods where its competitor is fighting for liquidity.

Finally, Figure 3.6 scales profits to reveal a positive feedback loop. Because the Na-

Figure 3.4: Simulation of Bank Fundamentals




Figure 3.4: A single stochastic simulation of the economy over a five-year horizon. The top panel displays the path of the banking sector’s asset value (V_t) relative to the constant default barrier (V_B , dashed line). The middle panel shows the corresponding instantaneous probability of default derived from the structural model. The bottom panel plots the resulting equilibrium spread for uninsured deposits in basis points.

tional Bank generates windfall profits during “bad times” the TBTF incumbent accumulates capital at a faster rate over the business cycle. This dynamic suggests that the current regulatory architecture creates an evolutionary pressure toward industry consolidation. Over time, the repeated profit boosts enjoyed by TBTF banks during downturns could drive the market toward an equilibrium where all surviving banks are systemically important, not due to superior efficiency, but because the implicit guarantee allows them to cannibalize the profitability of smaller rivals during every contraction. This mechanism is exemplified by historical consolidations, such as Wells Fargo’s acquisition of Wachovia during the 2008 Great Financial Crisis and J.P. Morgan Chase’s purchase of First Republic Bank in 2023, where TBTF incumbents absorbed distressed competitors and expanded their market dominance amidst systemic fragility.

Figure 3.5: Simulation of Bank Deposit Spreads and Shares




Figure 3.5: Time-series realizations of equilibrium pricing and quantity metrics corresponding to the simulation path in Figure 3.4. The top-left panel plots the uninsured deposit rates (bps) for the Local and National Banks. The bottom-left panel plots the deposit market shares for both banks. The right column displays the scaled versions of the deposit rates (top-right) and market shares (bottom-right) over the same five-year period. Scaled by the corresponding metric when the probability of default is zero.

3.4 Empirics

To estimate the elasticity of deposit demand with respect to bank failure risk, we require a continuous measure of the probability of default for every commercial bank in the United States. While the default probabilities of Globally Systemically Important Banks (GSIBs) can be inferred directly from the market prices of Credit Default Swaps (CDS), no such market exists for the vast majority of regional and local banks. Consequently, we must construct a proxy for these unobservable default probabilities.

We approach this as a supervised learning problem. We operate under the structural assumption that the mapping between observable bank accounting variables, derived from regulatory filings, and default risk is consistent across bank sizes, conditional on a sufficiently

Figure 3.6: Simulation of Bank Profits




Figure 3.6: The plot displays the time-series evolution of scaled profits for the Local Bank (blue line) and National Bank (yellow line) over the five-year simulation period shown in Figure 3.4. Scaled by the profit when the probability of default is zero.

rich feature set. The estimation procedure follows a multi-stage framework designed to balance the bias-variance trade-off inherent in high-dimensional financial data. First we filter the feature space using Elastic Net regularization, then we utilize a Weighted Ensemble model that aggregates predictions from a diverse set of base learners, including Random Forests and Gradient Boosted Decision Trees (GBM).

3.4.1 Data

The primary dataset is constructed at a quarterly frequency spanning from 2002 to 2023. This period encompasses several distinct phases of the banking cycle, including the pre-crisis expansion, the Global Financial Crisis (GFC), the zero-lower-bound (ZLB) era, and the post-COVID inflationary period. We integrate data from three primary sources to construct the training and prediction sets.

First, to proxy for default risk, we utilize 5-year CDS spreads sourced from Markit. The raw data consists of daily spread observations for the 17 largest U.S. commercial banks. To align this high-frequency market data with quarterly regulatory filings, we take the average CDS spread of the final month of each quarter. This “end-of-quarter” alignment ensures that the market signal reflects the information set available at the time of the regulatory report.

Second, granular balance sheet and income statement data are sourced from the Federal Financial Institutions Examination Council’s (FFIEC) Reports of Condition and Income (“Call Reports”) and the Uniform Bank Performance Report (UBPR). The UBPR serves as a cleaned and extended version of the Call Reports, providing pre-calculated regulatory ratios and peer group averages that are critical for benchmarking bank performance.

Third, to ensure our analysis covers the full spectrum of the U.S. banking system, we incorporate data from the Statistics on Depository Institutions (SDI), which covers 6,980 savings institutions. To ensure consistency in the unit of observation, all data—including the 2,464 commercial banks in the training set—are aggregated to the highest Bank Holding Company (BHC) level. This aggregation aligns with the regulatory treatment of TBTF designations, which are applied at the holding company level.

Following Begenau and Stafford (2023), we define the transaction deposit rate as the interest expense on non-time domestic deposits divided by the quarterly average of non-time domestic deposits. This precise definition isolates the cost of funding for the most liquid, flight-prone liabilities, which are the primary focus of our elasticity analysis.

The target variable for our training set is the implied probability of default derived from the CDS spreads of the 17 largest banks. Following the methodology of Hull (2015), we extract the risk-neutral default probability from the credit spread.

A CDS is priced so that, at inception, the expected present value (EPV) of the premium leg equals the EPV of the protection leg under the risk-neutral measure. Assuming continuous discounting, exponential default density, and payment dates of $0 < t_1 < \dots < t_n = T$ the pricing relationship is

$$S \left[\sum_{i=1}^n (t_i - t_{i-1}) e^{-(r_f + \lambda)t_i} + \sum_{i=1}^n \int_{t_{i-1}}^{t_i} (t - t_{i-1}) \lambda e^{-(r_f + \lambda)t} dt \right] = (1 - R) \int_0^T \lambda e^{-(r_f + \lambda)t} dt$$

where S is the CDS spread and R is the recovery rate for the insurance contract. The bracketed term is the PV of contractual premium payments plus the PV of accrued premium paid at default; the right-hand side is the PV of the default-contingent protection payment $(1 - R)$.

We solve the equation above numerically for λ . The implied risk-neutral default probability over horizon T is then

$$\rho(T) = 1 - e^{-\lambda T}.$$

This delivers a market-implied, risk-neutral measure of default risk that we use as the supervised learning target. Over the sample the average 5-year market implied default probability is 1.4%.

Table 3.1: In-Sample Banking Data Summary Statistics

Observations: 1124					
Variable	Mean	Std. Dev.	Min	Median	Max
Total Deposits (\$Bn)	334.792	439.185	2.684	153.289	2243.484
Insured Deposits (\$Bn)	188.581	235.899	2.148	87.376	1250.422
Uninsured Deposits (\$Bn)	146.21	215.058	0.536	63.502	1240.76
CDS Spread (bps)	89.19	90.008	4.63	66.885	1297.946
Probability of Default	1.481%	1.5%	0.075%	1.112%	21.707%
Deposit Rate (bps)	39.323	39.678	1.066	23.614	175.232

Table 3.1: Summary statistics for the in-sample training data, covering the 17 largest U.S. commercial banks from 2002 to 2023. Variables reported are total deposits, insured and uninsured deposits (in billions of dollars), CDS spreads (basis points), implied probability of default, and the quarterly deposit rate (basis points).

Table 3.1 reports summary statistics for the 1,124 bank-quarter observations that comprise the in-sample training data. The deposit distribution is highly right-skewed: the median institution holds approximately \$153 billion in total deposits, while the mean exceeds \$335

billion, reflecting balance sheet concentration among the GSIBs. Insured deposits account for roughly 56% of total deposits on average, though this share varies across institutions and over time. The CDS spread and implied default probability distributions exhibit comparable skewness; the median default probability of 1.1% lies well below the sample mean of 1.5%, and the maximum of 21.7% reflects acute GFC-era stress. The average quarterly deposit rate of 39 basis points captures the full span of monetary policy regimes in the sample, from the zero lower bound through the post-COVID tightening cycle.

A critical contribution of this paper is the construction of a high-dimensional feature space capable of capturing the non-linear dynamics of bank failure. We begin with the universe of 2,464 commercial banks that have complete coverage of all 523 features available in the UBPR. While these level variables provide a static snapshot of bank health, they fail to capture the trajectory of financial distress. To address this, we augment the feature space by calculating first order lags and percent changes. First-order lags are included to capture the persistence and mean-reversion in bank performance; while percent changes capture rapid deteriorations in asset quality or liquidity that may not yet be reflected in static capital ratios. Multiple lag and percent change horizons were tested with one order being the most performant.

This augmentation expands the initial dimensionality to 1,569 features. Given the scale variance among banking metrics (e.g., Total Assets vs. Return on Equity), we standardize all features prior to training using Z-score normalization. For every feature j , the observation x_{ij} is transformed as

$$z_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j},$$

ensuring that the optimization algorithms in subsequent steps are not biased by the magnitude of the underlying variables.

3.4.2 Default Model

The initial feature space 1,569 suffers from severe multicollinearity, a pervasive issue in accounting data where stock and flow variables are mechanically linked. A standard OLS

approach in this environment would yield high-variance estimators and poor generalization.

To address this, we employ Elastic Net regularization as a filtering mechanism. Elastic Net is chosen over Lasso or Ridge regression individually because it combines the sparsity-inducing properties of the ℓ_1 penalty (Lasso) with the grouping effect of the ℓ_2 penalty (Ridge). In banking data, where groups of liquidity and capital ratios are highly correlated, Lasso tends to select a single variable from the group arbitrarily. Elastic Net encourages a grouping effect, retaining correlated predictors that collectively inform default risk.

We model the relationship as a linear equation $\mathbf{y} = X\boldsymbol{\beta} + \boldsymbol{\epsilon}$. The estimator $\hat{\boldsymbol{\beta}}$ is obtained by minimizing the residual sum of squares subject to the weighted sum of the ℓ_1 and ℓ_2 norms

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \left\{ \frac{1}{2N} \|\mathbf{y} - X\boldsymbol{\beta}\|_2^2 + \xi \left(\zeta \|\boldsymbol{\beta}\|_1 + \frac{1 - \zeta}{2} \|\boldsymbol{\beta}\|_2^2 \right) \right\} \quad (3.14)$$

where the first term is the standard OLS goodness-of-fit, $\|\boldsymbol{\beta}\|_1$ is the ℓ_1 (“Taxi Cab”) norm, and $\|\boldsymbol{\beta}\|_2^2$ is the ℓ_2 (“Euclidean”) norm. The taxi cab norm enforces sparsity by driving coefficients of non-predictive features to exactly zero, while the Euclidean norm stabilizes the solution path and handles collinearity.

The hyperparameters ξ and ζ control the overall strength and mixing of the regularization. We tune these hyperparameters via 5-fold cross-validation with an 85/15 train test-split. Post-estimation, we perform feature selection by retaining only the subset of variables for which $\hat{\beta}_j \neq 0$. This selection process distills the feature space down to 170 distinct variables, composed of: 75 level variables, 73 lagged variables, and 22 percent change variables. The most predictive features identified by the model include Non-Current Loans to Tier 1 Capital, the Lag of Total Servicing with Recourse to Total Assets, and the Percent Change in Allowance for Credit Losses.

While the Elastic Net provides interpretability and efficient feature selection, it assumes a linear relationship between bank fundamentals and default risk. However, financial distress is inherently non-linear; the marginal impact of a liquidity shock increases exponentially as capital buffers deplete. To capture these non-linearities and high-order interactions, we deploy a Weighted Ensemble model using the AutoGluon framework.

The ensemble aggregates predictions from diverse base learners: Random Forests (RF) and Gradient Boosted Machines (GBM), specifically XGBoost and LightGBM. The core of our ensemble consists of decision tree-based models, which are uniquely suited to capturing threshold effects in financial data (e.g., a bank is safe until its capital ratio drops below a specific regulatory minimum).

The Random Forest component addresses the variance inherent in individual decision trees. We utilize a bagging approach. For a forest of B trees, we generate B bootstrap samples from the training data. A decision tree T_b is grown on each sample. Unlike standard CART trees, the Random Forest introduces additional randomness by considering only a random subset of features at each split candidate. The final prediction is the average of the individual tree predictions

$$\hat{f}_{RF}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B T_b(\mathbf{x}; \Theta_b).$$

To reduce bias, we employ Gradient Boosted Decision Trees (GBDT). Unlike Random Forests, which train trees in parallel, GBDTs train trees sequentially. Each new tree attempts to correct the residuals of the previous ensemble.

We define the ensemble model at step K as the sum of K additive functions

$$\hat{y}_i = \eta(\mathbf{x}_i) = \sum_{k=1}^K g_k(\mathbf{x}_i)$$

for $g_k \in \mathbb{G}$ is the space of regression trees. The objective is to minimize a regularized loss function

$$L(\eta) := \sum_i \ell(\hat{y}_i, y_i) + \sum_k \Omega(g_k).$$

Here, ℓ is the mean-square-error loss function measuring the difference between the predicted default probability $\hat{y}_i$ and the target y_i . The term $\Omega(g_k)$ penalizes model complexity to prevent overfitting.

Both XGBoost and LightGBM optimize this objective using a second-order Taylor expansion of the loss function. At iteration t , we seek the function h_t that minimizes

$$\Xi^{(t)} \approx \sum_{i=1}^N \left(\ell(y_i, \hat{y}_i^{(t-1)}) + \partial_{\hat{y}_i^{(t-1)}} \ell(y_i, \hat{y}_i^{(t-1)}) h_t(\mathbf{x}_i) + \frac{1}{2} \partial_{\hat{y}_i^{(t-1)}}^2 \ell(y_i, \hat{y}_i^{(t-1)}) h_t^2(\mathbf{x}_i) \right) + \Omega(h_t). \quad (3.15)$$

This formulation allows the model to utilize the curvature of the loss function, resulting in faster convergence and more precise estimates of the probability of default compared to standard gradient descent.

The final stage of the estimation is the aggregation of the base learners into a single predictive model. We employ a stacking strategy via the Weighted Ensemble algorithm. Let $\hat{g}_{ij}$ denote the prediction of base learner $j \in \{1, \dots, M\}$ for observation i . The final predicted default probability $\hat{\rho}_i$ is a convex combination of the base learners

$$\hat{\rho}_i = \sum_{j=1}^M w_j \hat{g}_{ij}.$$

The optimal weights vector $\mathbf{w} \in \mathbb{R}^M$ is determined by solving a constrained least squares optimization problem on a hold-out validation set with N_v observations. We minimize the MSE

$$\min_{\mathbf{w} \in \mathbb{R}^M} \frac{1}{N_v} \sum_{i=1}^{N_v} \left(y_i - \sum_{j=1}^M w_j \hat{g}_{ij} \right)^2 \quad \text{subject to} \quad \sum_{j=1}^M w_j = 1, \quad w_j \geq 0 \quad \forall j.$$

The non-negativity constraint $w_j \geq 0$ and the summation constraint $\sum w_j = 1$ ensure models are weighted by their validation performance. This reduces the generalization error by averaging out the idiosyncratic errors of diverse models. For instance, if the Random Forest overestimates risk in high-volatility environments while the LightGBM underestimates it, the ensemble average will converge closer to the true probability.

To ensure the stability of the estimated default probabilities, we implement a rigorous validation protocol. The model training utilizes an 85/15 train/test split. Within the training phase, we employ 5-fold cross-validation to tune the hyperparameters of the Elastic Net and the base learners. Furthermore, we bootstrap the training data for the ensemble. By training on repeated bootstrap samples, we reduce the variance of the final estimator and mitigate

the risk of overfitting to specific idiosyncratic bank failures in the training set.

Table 3.2: Estimated Default Probabilities Summary Statistics

	N	Mean	Std. Dev	Min	p50	p99	Max
$\rho_{k,t}$	1,076	0.0141	0.0139	0.0008	0.0108	0.0659	0.2003
$\hat{\rho}_{k,t}$	39,527	0.0143	0.0122	0.0014	0.0116	0.0620	0.1282

Table 3.2: Summary statistics for observed default probabilities $\rho_{k,t}$ for the in-sample banks and estimated default probabilities $\hat{\rho}_{k,t}$ for the out-of-sample local banks.

Table 3.2 compares the distribution of market-implied default probabilities for the in-sample banks with the model’s out-of-sample predictions for the local bank universe. The central moments align closely, the means differ by less than two basis points and the medians and standard deviations are similarly matched. The primary discrepancy arises in the right tail, where the observed distribution extends to 20% while the estimated distribution reaches approximately 13%.

Figure 3.7: Distribution of Estimated Default Probabilities




Figure 3.7: Distribution of observed default probabilities for the in-sample large regional banks and GSIBs and estimated default probabilities for the out-of-sample local banks.

Figure 3.7 presents the distributional comparison as overlapping histograms. The two distributions share a common modal range and broadly similar shape, consistent with the summary statistics in Table 3.2.

Figure 3.8: Wells Fargo Time Series of Estimated Default Probabilities




Figure 3.8: Estimated and implied default probabilities for Wells Fargo under a hold-one-out specification. Out-of-sample predictions are compared against market-implied default probabilities from 2004 to 2024.

Figure 3.8 presents a hold-one-out validation for Wells Fargo. The model is re-estimated on all remaining banks and the resulting out-of-sample predictions are compared against the market-implied default probabilities over the full 2004–2024 period. The two series track each other closely, with a correlation of 0.95, and the model captures both the sharp increase in default risk during the GFC and the more moderate elevation observed around 2020. This performance across distinct stress episodes, estimated entirely without information from the held-out institution, supports the view that the model generalizes beyond the training sample.

The weighted ensemble achieves an R^2 of 0.8636 on the held-out test set, compared with 0.4164 for the baseline Elastic Net, and a root mean squared error of 0.0058, approximately

half that of the linear benchmark. The ratio of validation RMSE to training RMSE is 23.5%, indicating some in-sample overfit that likely reflects the limited number of institutions available for training. The resulting $\hat{\rho}_{k,t}$ provides a continuous, risk-sensitive measure of default probability for the universe of small and regional banks, enabling the subsequent difference-in-differences analysis of deposit demand elasticity.

3.4.3 Causal Tests

To identify the causal mechanism through which TBTF guarantees distort deposit market competition, we must estimate the elasticity of deposit demand with respect to bank failure risk. Our approach follows the structural identification strategy of Egan et al. (2017), identifying demand parameters by exploiting cross-sectional variation in bank default probabilities. We extend their framework by introducing a difference-in-differences (DiD) specification that isolates the safety premium generated by government interventions following the Global Financial Crisis (GFC).

We model the deposit market as a discrete choice problem where depositors choose between depositing funds at a specific commercial bank k or retaining them in a risk-free outside option (e.g., physical cash or government treasuries). We define the market share of bank k at time t for a given deposit type l (insured or uninsured) as

$$s_{k,t}^l := \frac{D_{k,t}^l}{\sum_{j=1}^K D_{j,t}^l} \quad (3.16)$$

where $D_{k,t}^l$ represents the total dollar amount of deposits of type l held by bank k . Crucially, the inclusion of the outside option allows for aggregate contractions in the banking sector's deposit base. Following Egan et al. (2017), we define the deposit share of the outside option, $s_{O,t}^l$, as the difference between the total deposits in the market and the sum of all in-sample bank deposits.

To disentangle the TBTF subsidy from fundamental quality differences, we employ a difference-in-differences methodology. We estimate the sensitivity of deposit market shares to the estimated probability of default $\hat{\rho}_{k,t}$, derived from our machine learning model described

in Section 3.4.2, conditional on the bank’s regulatory status, TBTF or not-TBTF, and the macroeconomic regime, pre/post GFC and between GSIBs and Local Banks.

The baseline specification for the quantity channel is given by the following linear regression:

$$\begin{aligned} \ln \left(\frac{s_{k,t}}{s_{O,t}} \right) = & \beta_0 + \beta_1 \hat{\rho}_{k,t} + \beta_2 \times \text{Regional} \times \hat{\rho}_{k,t} + \beta_3 \times \text{GSIB} \times \hat{\rho}_{k,t} \\ & + \beta_4 \times \text{Post GFC} \times \text{GSIB} \times \hat{\rho}_{k,t} + \beta_5 \times \text{Post GFC} \times \text{Regional} \times \hat{\rho}_{k,t} \\ & + \beta_6 \times \text{Post GFC} \times \hat{\rho}_{k,t} + \beta_7 \times \text{Post GFC} \times \text{GSIB} \\ & + \beta_8 \times \text{Post GFC} \times \text{Regional} \\ & + \delta_k + \delta_t + \Gamma X_{k,t} + \epsilon_{k,t} \end{aligned} \quad (3.17)$$

where $\hat{\rho}_{k,t}$ is the continuous measure of default probability for bank k at time t estimated via the Weighted Ensemble model. “Regional” and “GSIB” are indicator variables for large regional commercial banks and Globally Systemically Important Banks, respectively, with local banks serving as the omitted reference group. “Post GFC” is a time indicator set to one for periods following the implementation of the Dodd-Frank Act, post-2012. δ_k and δ_t represent bank and time fixed effects to control for unobserved heterogeneity and aggregate shocks. $X_{k,t}$ includes time-varying bank-level controls as in Egan et al. (2017), in particular the number of domestic branches is included.

The coefficients of primary interest are the triple interaction terms β_4 and β_5 . These coefficients capture the differential change in depositor sensitivity to default risk for GSIBs and regional banks relative to local banks in the post-GFC period. A positive β_4 indicates that GSIBs gain deposit share as their measured risk rises in the post-reform era, consistent with the TBTF distortion. A positive β_5 indicates that regional banks similarly attract deposits as their risk rises relative to local banks, reflecting a partial flight-to-safety effect that extends across the size hierarchy.

We estimate a parallel specification for the pricing channel, where the dependent variable is the deposit spread, defined as the 3-month U.S. T-Bill rate minus the bank’s offered deposit

rate. This allows us to test the scissors dynamic predicted by our theoretical model: whether distressed local banks are forced to compress spreads to retain funds while protected GSIBs widen them.

3.4.4 Results

Table 3.3 presents estimates from the deposit share regression on an unbalanced panel of 2,464 commercial banks observed from 2003 to 2023. Throughout, local banks serve as the omitted reference category. All four specifications include bank and time fixed effects and control for the number of domestic branches; standard errors are clustered at the bank level.

I estimate four specifications that progressively isolate the credit-risk channel from alternative explanations. Column (1) reports the aggregate effect on total deposits. Columns (2) and (3) decompose this into insured and uninsured deposits, respectively. Column (4) estimates the within-bank difference between the uninsured and insured responses, a triple-differences specification following Egan et al. (2017) that removes the contribution of unobserved, time-varying bank quality from the identified effect.

Column (2) provides the critical identification test for the credit-risk mechanism. Because the FDIC guarantees deposits up to \$250,000 per depositor per institution, insured depositors bear no credit exposure to the issuing bank. Their decisions should therefore be insensitive to the probability of default, and the post-GFC triple interaction terms for insured deposits should be statistically indistinguishable from zero. If, instead, the documented post-GFC shifts in deposit flows were driven by changing bank quality or service provision, insured depositors would exhibit behavior symmetric to uninsured depositors, since quality affects both cohorts equally.

The data support the identification restriction. The coefficient on $\text{Post GFC} \times \text{Regional} \times \hat{\rho}_{k,t}$ and the coefficient on $\text{Post GFC} \times \text{GSIB} \times \hat{\rho}_{k,t}$ in Column (2) are both statistically insignificant. Insured depositors do not discriminate between bank types in response to rising default risk, consistent with the FDIC guarantee eliminating the incentive to monitor solvency. The large post-GFC effects documented in the uninsured specification cannot be attributed to a general shift in bank attractiveness or quality.

Table 3.3: Bank Deposit Share Regressions

	Total (1)	Insured (2)	Uninsured (3)	Uninsured - Insured (4)
$\hat{\rho}_{k,t}$	4.048*** (0.605)	4.708*** (0.674)	2.972*** (0.993)	-1.736* (0.971)
Regional $\times \hat{\rho}_{k,t}$	1.012* (3.465)	5.494 (3.765)	5.494 (3.765)	5.943*** (1.737)
GSIB $\times \hat{\rho}_{k,t}$	-15.430* (8.523)	-17.483** (7.685)	-10.327 (9.794)	7.156* (3.802)
Post GFC $\times \hat{\rho}_{k,t}$	-3.400*** (1.022)	-4.091*** (1.052)	-3.762** (1.612)	0.330 (1.496)
Post GFC $\times$ Regional	-0.138 (0.143)	-0.281 (0.180)	-0.005 (0.120)	0.276* (0.145)
Post GFC $\times$ GSIB	-0.379 (0.348)	-0.443 (0.407)	-0.473 (0.368)	-0.029 (0.286)
Post GFC $\times$ Regional $\times \hat{\rho}_{k,t}$	10.279** (4.764)	8.229 (5.610)	14.549** (5.754)	6.320 (6.345)
Post GFC $\times$ GSIB $\times \hat{\rho}_{k,t}$	29.144* (16.485)	17.422 (20.698)	50.590*** (16.911)	33.168* (18.931)
Bank FE	X	X	X	X
Time FE	X	X	X	X
N	40,603	40,603	40,603	40,603
R^2	0.985	0.980	0.960	0.727

Table 3.3: Deposit Share Regressions by Bank Type and Deposit Class. This table presents estimates from the regression of a bank's log deposit share relative to the outside option on its probability of default, estimated on an unbalanced panel of 2,464 commercial banks from 2003 to 2023. Local Banks serve as the omitted reference category. All specifications include bank and time fixed effects and control for the number of domestic branches. Standard errors are clustered at the bank level.

Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

The pre-GFC coefficients establish the deposit-risk relationship before crisis-era government interventions created sharp asymmetries across bank types. Local and regional banks exhibit a directionally similar but statistically indistinguishable pre-GFC relationship. For GSIBs, the interaction $\text{GSIB} \times \hat{\rho}_{k,t}$ is -10.327 and marginally significant at the 10% level, suggesting that uninsured depositors at the largest banks were somewhat less likely to increase allocations in response to rising risk even in the pre-crisis period. The pre-GFC period therefore reveals no dramatic hierarchy in the deposit-risk relationship across bank types.

The TBTF guarantee had not yet been made credible through the experience of a systemic bailout, and the structural break emerges sharply only in the post-GFC period.

The triple interaction terms with the Post GFC indicator reveal a fundamental reorganization of the relationship between bank default risk and deposit flows. For regional banks the coefficient on $\text{Post GFC} \times \text{Regional} \times \hat{\rho}_{k,t}$ in Column (3) is positive and significant at the 5% level. As regional bank default risk rises in the post-reform era, uninsured depositors reallocate toward these institutions rather than away from them. This pattern is consistent with a partial flight-to-safety effect that extends beyond the four GSIBs: depositors withdrawing from local banks may perceive large regional institutions as relatively safe havens, even absent a formal systemic designation.

The effect is largest and most precisely estimated for GSIBs. The coefficient on $\text{Post GFC} \times \text{GSIB} \times \hat{\rho}_{k,t}$ is 50.590, significant at the 1% level. The net post-GFC elasticity for GSIBs is 39.4 per unit of $\hat{\rho}_{k,t}$. Since $\hat{\rho}_{k,t}$ is measured as a decimal probability, a one-percentage-point increase in a GSIB's probability of default is associated with an approximately 48% increase in its uninsured deposit share in the post-GFC era. GSIBs have become liquidity sinks in the post-reform era, attracting rather than repelling uninsured capital as their measured default risk increases. The implicit government guarantee, reinforced by the systematic rescues during the GFC and the resolution framework established under Dodd-Frank, has decoupled credit risk from depositor response at the top of the banking hierarchy.

The flight-to-safety gradient across the size hierarchy is economically striking. Relative to local banks, GSIBs capture 50.590 additional units of log deposit share per unit increase in default probability in the post-GFC period, and 36.041 more than regional banks. The deposit reallocation is not simply a flow from small to any larger institution; it concentrates disproportionately in the four institutions perceived to carry the most credible government backstop.

Column (4) directly tests whether the uninsured deposit flows in Column (3) reflect credit-risk sensitivity or changes in unobserved bank quality. The triple-differences estimator computes the within-institution differential response of uninsured versus insured depositors to default probability. Because quality improvements generate symmetric inflows across both

deposit types, any within-bank asymmetry between insured and uninsured responses must be attributed to credit risk rather than quality.

For local banks, the baseline coefficient on $\hat{\rho}_{k,t}$ in Column (4) is negative and significant at the 10% level. Higher default probability at local banks is associated with a greater decline in the uninsured deposit share than in the insured share. This within-institution differential is consistent with market discipline: uninsured depositors monitor and respond to solvency risk at local banks, while insured depositors do not.

The coefficient of primary interest is $\text{Post GFC} \times \text{GSIB} \times \hat{\rho}_{k,t}$, significant at the 10% level. This estimate survives the within-bank quality control. Relative to insured depositors at the same GSIB, uninsured depositors disproportionately increase their allocations to GSIBs as default probability rises in the post-GFC era.

The analogous triple-difference coefficient for regional banks, $\text{Post GFC} \times \text{Regional} \times \hat{\rho}_{k,t}$, is not statistically significant. This contrast with the significant GSIB result clarifies the scope of the TBTF mechanism. Regional banks attract deposit flows on the basis of general safety perceptions relative to local banks, but the clean credit-risk channel identified by the within-bank comparison is concentrated in the four institutions with a credible systemic backstop.

Table 3.4 presents the results for the pricing channel, regressing the deposit spread on default probabilities. The results corroborate the quantity channel but with weaker statistical significance. The results are presented with regional banks as the reference group.

A critical factor driving these results is the underlying interest rate environment. In the Pre-GFC period, when benchmark rates were well above zero, the pricing mechanism functioned consistently with the theory presented in Section 3.3. The baseline coefficient for regional banks is negative, suggesting that riskier banks paid higher rates or that deposit rates were insensitive to default probabilities, while the GSIB interaction is significantly positive, suggesting an early presence of the safety subsidy. However, the lack of statistical significance in the interaction terms during the Post GFC and ZLB periods is likely a mechanical artifact of the extremely low interest rate environment that characterized the post-crisis era. When the Federal Funds Rate is pinned near zero, the variance in deposit rates collapses; banks

cannot lower nominal deposit rates below zero to widen spreads, nor can riskier banks raise rates significantly without destroying net interest margins. This compression removes the variation required to identify the pricing coefficient using interaction terms alone.

Table 3.4: Bank Deposit Spread Regressions

Time Period	Pre GFC 2003Q1-2007Q4	GFC 2008Q1-2009Q4	ZLB 2010Q1-2019Q4	Covid 2020Q1-2023Q1
$\hat{\rho}_{k,t}$	-0.064 (0.072)	0.001 (0.007)	0.026** (0.011)	-0.036 (0.023)
$\hat{\rho}_{k,t} \times \text{GSIB}$	0.266*** (0.048)	0.008 (0.015)	-0.020 (0.018)	0.219*** (0.041)
$\hat{\rho}_{k,t} \times \text{Local}$	0.036 (0.073)	-0.010 (0.008)	-0.027** (0.011)	0.032 (0.024)
<i>Bank FE</i>	Yes	Yes	Yes	Yes
<i>Quarter FE</i>	Yes	Yes	Yes	Yes
R^2	0.987	0.974	0.992	0.998
N	12,737	5,982	16,620	4,912
$\hat{\rho}_{k,t} + \hat{\rho}_{k,t} \times \text{GSIB}$	0.202*** (0.051)	0.010 (0.014)	0.006 (0.016)	0.183*** (0.035)
$\hat{\rho}_{k,t} + \hat{\rho}_{k,t} \times \text{Local}$	-0.028*** (0.009)	-0.008* (0.004)	-0.001 (0.002)	-0.004 (0.005)

Table 3.4: This table shows the regression of a bank's deposit spread on its probability of default. A bank's deposit spread is defined as the 3-month U.S. T-Bill minus the bank's offered deposit rate. Each regression is estimated using an unbalanced panel of 2,464 commercial banks in the United States over the period 2003-2023. All specifications control for the number of bank branches. Standard errors are clustered at the bank level.

Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

To rigorously evaluate the net impact of these distortions across regimes, we calculate the total marginal effect of default risk for each bank type, presented in the bottom two rows of Table 3.4. These total effect rows sum the baseline coefficient with the respective interaction terms $\hat{\rho}_{k,t} + \hat{\rho}_{k,t} \times \text{Type}$, providing a clearer picture of the absolute elasticity. The results confirm a persistent divergence. For GSIBs, the total effect remains positive and statistically significant across the Pre-GFC and Covid periods. This indicates that even during the pandemic, GSIBs retained the ability to widen spreads as their risk profile deteriorated. In sharp contrast, the total effect for local banks is consistently negative or near zero in the Pre-GFC period and during the GFC. This negative total elasticity confirms

that local banks are subject to market discipline, forcing them to compress their spreads and pay a premium to retain funding when their solvency is threatened, while their systemic competitors effectively monetize the distress.

3.5 Conclusion

This study provides a structural and empirical characterization of the safety premium in the U.S. deposit market, demonstrating that implicit government guarantees fundamentally distort the transmission of market discipline. The theoretical framework, combining Hotelling’s spatial competition with Merton’s credit risk mechanics, reveals that the elasticity of deposit demand with respect to safety overwhelms the elasticity with respect to yield during periods of stress. This generates a convex “scissors” dynamic in equilibrium pricing: unprotected local banks face a self-reinforcing death spiral where rising default risk necessitates higher deposit rates that further deplete capital, while TBTF institutions exploit the inelasticity of flight-to-safety flows to widen net interest margins. Crucially, the model demonstrates that this distortion is not static but highly counter-cyclical, expanding non-linearly when the banking sector is most fragile and the marginal utility of safety is highest.

Empirically, this paper overcomes the challenge of unobservable default risk for non-publicly traded institutions by deploying a Weighted Ensemble machine learning model, which maps the accounting fundamentals of the universe of commercial banks to market-implied probabilities of default. The subsequent difference-in-differences analysis validates the theoretical predictions via the quantity channel, confirming a regime change following the Global Financial Crisis. While the Zero Lower Bound environment constrained the identification of the pricing channel, the flow of funds data provides distinct evidence of bifurcation: uninsured depositors exert severe discipline on local banks while treating GSIBs as effectively risk-free substitutes for government obligations. This asymmetry confirms that TBTF status grants large banks a put option on the taxpayer that is valued most highly by the market during tail events.

The policy implications of these findings suggest that the current regulatory architecture inadvertently creates evolutionary pressure toward industry consolidation, not through

efficiency gains, but through regulatory arbitrage. By shielding GSIBs from the market discipline that constrains their smaller rivals, the TBTF regime allows systemic institutions to accumulate capital and market share at an accelerated rate during downturns, effectively cannibalizing the periphery of the banking system. Furthermore, the sensitivity of the deposit spread to the loss given default parameter implies that ambiguity regarding the scope of deposit insurance remains a primary driver of cost-of-capital differentials. Consequently, unless regulatory frameworks can credibly commit to imposing losses on the unsecured creditors of systemic institutions, the flight-to-safety subsidy will continue to incentivize excessive scale and concentration, undermining the competitive neutrality of the financial system.

Bibliography

- 102nd Congress of the United States (1991). Federal deposit insurance corporation improvement act of 1991.
- Abarbanell, J. S. and Bushee, B. J. (1997). Abnormal returns to a fundamental analysis strategy. *The Accounting Review*, 72(1):1–31.
- Acharya, V. V., Anginer, D., and Warburton, A. J. (2016). The end of market discipline? investor expectations of implicit government guarantees. *SSRN Electronic Journal*.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.
- Almeida, C. and Garcia, R. (2012). Assessing misspecified asset pricing models with empirical likelihood estimators. *Journal of Econometrics*, 170(2):519–537.
- Altman, E. I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance*, 23(4):589–609.
- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.
- Anderson, E. W., Hansen, L. P., and Sargent, T. J. (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123.
- Armijo, L. (1966). Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific Journal of Mathematics*, 16(1):1–3.

- Asness, C., Frazzini, A., Israel, R., Moskowitz, T. J., and Pedersen, L. H. (2018). Size matters, if you control your junk. *Journal of Financial Economics*, 129(3):479–509.
- Asness, C. S., Frazzini, A., and Pedersen, L. H. (2019). Quality minus junk. *Review of Accounting Studies*, 24(1):34–112.
- Backus, D., Chernov, M., and Zin, S. (2014). Sources of entropy in representative agent models. *Journal of Finance*, 69(1):51–99.
- Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221.
- Baker, M., Bradley, B., and Wurgler, J. (2011). Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal*, 67(1):40–54.
- Bali, T., Engle, R., and Murray, S. (2017). *Empirical Asset Pricing: The Cross Section of Stock Returns: An Overview*, pages 1–8. John Wiley & Sons, Ltd.
- Ball, R., Gerakos, J., Linnainmaa, J. T., and Nikolaev, V. (2016). Accruals, cash flows, and operating profitability in the cross section of stock returns. *Journal of Financial Economics*, 121(1):28–45.
- Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1):3–18.
- Barbee, W. C., Mukherji, S., and Raines, G. A. (1996). Do sales—price and debt—equity explain stock returns better than book—market and firm size? *Financial Analysts Journal*, 52(2):56–60.
- Barillas, F. and Shanken, J. (2018). Comparing asset pricing models. *The Journal of Finance*, 73(2):715–754.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866.

- Barroso, P. and Santa-Clara, P. (2015). Momentum has its moments. *Journal of Financial Economics*, 116(1):111–120.
- Barth, J. R. and Prabha, A. (2013). Resolving too-big-to-fail banks in the united states. *Mercatus Center Working Paper Series*, 13(5).
- Basu, S. (1983). The relationship between earnings' yield, market value, and return for nyse common stocks: Further evidence. *Journal of Financial Economics*, 12(1):129–156.
- Begenau, J. and Stafford, E. (2023). Uniform rate setting, branch expansion, and the deposit channel. *Working Paper*.
- Bennett, R. L. and Unal, H. (2015). Understanding the components of bank failure resolution costs. *Financial Markets, Institutions & Instruments*, 24(5):349–389.
- Benston, G. J. and Kaufman, G. G. (1997). Fdicia after five years. *Journal of Economic Perspectives*, 11(3):139–158.
- Berndt, A., Duffie, D., and Zhu, Y. (2025). The decline of too big to fail. *American Economic Review*, 115(3):945–974.
- Black, F. (1972). Capital market equilibrium with restricted borrowing. *The Journal of Business*, 45(3):444–455.
- Black, F. and Cox, J. C. (1976). Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance*, 31(2):351–367.
- Black, F., Jensen, M. C., and Scholes, M. (1972). The capital asset pricing model: Some empirical tests. In *Studies in the theory of capital markets*, pages 79–121. Praeger, New York.
- Bliss, R. R. and Kaufman, G. G. (2006). A comparison of u.s. corporate and bank insolvency resolution. *Federal Reserve Bank of Chicago: Economic Perspectives*, 30(2).
- Blitz, D., Huij, J., and Martens, M. (2011). Residual momentum. *Journal of Empirical Finance*, 18(3):506–521.

- Bouchaud, J.-P., Krüger, P., Landier, A., and Thesmar, D. (2019). Sticky expectations and the profitability anomaly. *Journal of Finance*, 74(2):639–674.
- Bovenzi, J. F. and Muldoon, M. E. (1990). Failure-resolution methods and policy considerations. *FDIC Banking Review*, 3(1):1–11.
- Bovenzi, J. F. and Murton, A. J. (1988). Resolution costs of bank failures. *FDIC Banking Review*, 1(1):1–13.
- Boyd, S. and Vandenberghe, L. (2009). *Convex Optimization*. Cambridge University Press.
- Bradshaw, M. T., Richardson, S. A., and Sloan, R. G. (2006). The relation between corporate financing activities, analysts’ forecasts and stock returns. *Journal of Accounting and Economics*, 42(1):53–85. Conference Issue on Implications of Changing Financial Reporting Standards.
- Bregman, L. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217.
- Brennan, M. J., Chordia, T., and Subrahmanyam, A. (1998). Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. *Journal of Financial Economics*, 49(3):345–373.
- Brown, R. A. and Epstein, S. (1992). Resolution costs of bank failures: An update of the fdic historic al loss model. *FDIC Banking Review*, 5(1):1–16.
- Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, 52(1):57–82.
- Chamberlain, G. and Rothschild, M. (1983). Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*, 51(5):1281–1304.
- Chan, L. K. C., Jegadeesh, N., and Lakonishok, J. (1996). Momentum strategies. *The Journal of Finance*, 51(5):1681–1713.

- Chen, A. Y. and Zimmermann, T. (2022). Open source cross-sectional asset pricing. *Critical Finance Review*, 11(2):207–264.
- Chordia, T., Subrahmanyam, A., and Anshuman, R. (2001). Trading activity and expected stock returns. *Journal of Financial Economics*, 59(1):3–32.
- Cochrane, J. C. (2011). Presidential address: Discount rates. *The Journal of Finance*, 66(4):1047–1108.
- Connor, G. and Korajczyk, R. A. (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of Financial Economics*, 15(3):373–394.
- Connor, G. and Korajczyk, R. A. (1988). Risk and return in an equilibrium apt: Application of a new test methodology. *Journal of Financial Economics*, 21(2):255–289.
- Cookson, J. A., Fox, C., Gil-Bazo, J., Imbet, J. F., and Schiller, C. (2023). Social media as a bank run catalyst. *Université Paris-Dauphine Research Papers*, (4422754). Available at SSRN: <https://ssrn.com/abstract=4422754> or <http://dx.doi.org/10.2139/ssrn.4422754>.
- Corporation, F. D. I. (2003). Failed bank cost analysis: 1986-2002. *FDIC Publication*.
- Cowan, A. R. and Salotti, V. (2015). The resolution of failed banks during the crisis: Acquirer performance and fdic guarantees, 2008-2013. *Journal of Banking & Finance*, 54(C):222–238.
- Cowan, A. R., Salotti, V., and Schenck, N. A. (2022). The long-term impact of bank mergers on stock performance and default risk: The aftermath of the 2008 financial crisis. *Finance Research Letters*, 48:102925.
- Curry, T. and Shibut, L. (2000). The cost of the savings and loan crisis: Truth and consequences. *FDIC Banking Review*.
- Daniel, K., Hirshleifer, D., and Sun, L. (2020). Short- and long-horizon behavioral factors. *The Review of Financial Studies*, 33(4):1673–1736.

- Daniel, K. and Moskowitz, T. J. (2016). Momentum crashes. *Journal of Financial Economics*, 122(2):221–247.
- d’Avernas, A., Eisfeldt, A. L., Huang, C., Stanton, R., and Wallace, N. (2023). The deposit business at large vs. small banks. *NBER Working Paper Series*, (31865).
- Davis, G., Mallat, S., and Avellaneda, M. (1997). Adaptive greedy approximations. *Constructive Approximation*, 13(1):57–98.
- Davison, L. K. (1997). *Continental Illinois and Continental Illinois and "Too Big to Fail"*, chapter 7, pages 236–258. Federal Deposit Insurance Corporation.
- Dembo, R. S., Eisenstat, S. C., and Steihaug, T. (1982). Inexact Newton methods. *SIAM Journal on Numerical Analysis*, 19(2):400–408.
- Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419.
- Dichev, I. D. (1998). Is the risk of bankruptcy a systematic risk? *The Journal of Finance*, 53(3):1131–1147.
- Dimson, E. (1979). Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics*, 7(2):197–226.
- Dittmar, R. F. (2002). Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *The Journal of Finance*, 57(1):369–403.
- Donsker, M. D. and Varadhan, S. R. S. (1975). Asymptotic evaluation of certain markov process expectations for large time, i. *Communications on Pure and Applied Mathematics*, 28(1):1–47.
- Drechsler, I., Savov, A., and Schnabl, P. (2017). The deposits channel of monetary policy. *The Quarterly Journal of Economics*, 132(4):1819–1876.
- Drechsler, I., Savov, A., and Schnabl, P. (2021). Banking on deposits: Maturity transformation without interest rate risk. *The Journal of Finance*, 76(3):1091–1143.

- Efron, B. (1975). Defining the curvature of a statistical problem (with applications to second order efficiency). *The Annals of Statistics*, 3(6):1189–1242.
- Egan, M., Hortaçsu, A., and Matvos, G. (2017). Deposit competition and financial fragility: Evidence from the us banking sector. *American Economic Review*, 107(1):169–216.
- Engle, R. F. (1984). Wald, likelihood ratio, and lagrange multiplier tests in econometrics. In Griliches, Z. and Intriligator, M. D., editors, *Handbook of Econometrics*, volume 2, pages 775–826. North-Holland, New York.
- Fama, E. F. and French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465.
- Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.
- Fama, E. F. and French, K. R. (2002). The equity premium. *The Journal of Finance*, 57(2):637–659.
- Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.
- Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *The Journal of Political Economy*, 81(3):607–636.
- FDIC (2023). Fdic acts to protect all depositors of the former silicon valley bank, santa clara, california. Online.
- Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998). *Managing the Crisis: The FDIC and RTC Experience, 1980–1994*. Federal Deposit Insurance Corporation, Washington, DC. 2 vols.
- Federal Reserve System (2026). U.s. domestically chartered commercial banks. <https://www.federalreserve.gov/releases/lbr/current/>. Accessed: April 20, 2026.

- Feng, G., Giglio, S., and Xiu, D. (2020). Taming the factor zoo: A test of new factors. *The Journal of Finance*, 75(3):1327–1370.
- Financial Stability Board (2025). 2025 list of global systemically important banks (g-sibs). <https://www.fsb.org/2025/11/2025-list-of-global-systemically-important-banks-g-sibs/>. Accessed: April 20, 2026.
- Frankel, R. and Lee, C. M. (1998). Accounting valuation, market expectation, and cross-sectional stock returns. *Journal of Accounting and Economics*, 25(3):283–319.
- Frazzini, A. and Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1):1–25.
- Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance *. *The Quarterly Journal of Economics*, 127(2):645–700.
- Gandhi, P. and Lustig, H. (2015). Size anomalies in u.s. bank stock returns. *The Journal of Finance*, 70(2):733–768.
- George, T. J. and Hwang, C.-Y. (2004). The 52-week high and momentum investing. *The Journal of Finance*, 59(5):2145–2176.
- Ghosh, A., Julliard, C., and Taylor, A. P. (2016). What is the consumption-capm missing? an information-theoretic framework for the analysis of asset pricing models. *The Review of Financial Studies*, 30(2):442–504.
- Gibbons, M. R., Ross, S. A., and Shanken, J. (1989). A test of the efficiency of a given portfolio. *Econometrica*, 57(5):1121–1152.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.
- Goldstein, I. and Pauzner, A. (2005). Demand-deposit contracts and the probability of bank runs. *The Journal of Finance*, 60(3):1293–1327.

- Gram, J. P. (1883). Ueber die entwicklung reeller funktionen in reihen mittels der methode der kleinsten quadrate. *Journal für die reine und angewandte Mathematik*, 94:41–73.
- Granja, J., Matvos, G., and Seru, A. (2017). Selling failed banks. *The Journal of Finance*, 72(4):1723–1784.
- Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273.
- Gu, S., Kelly, B., and Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics*, 222(1, Part B):429–450. Annals Issue: Financial Econometrics in the Age of the Digital Economy.
- Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.
- Hansen, L. P. and Jagannathan, R. (1997). Assessing specification errors in stochastic discount factor models. *The Journal of Finance*, 52(2):557–590.
- Hansen, L. P. and Sargent, T. J. (2001). Robust control and model uncertainty. *The American Economic Review*, 91(2):60–66.
- Hansen, L. P. and Sargent, T. J. (2010). Fragile beliefs and the price of uncertainty. *New York University*.
- Hansen, L. P. and Sargent, T. J. (2020). Structured ambiguity and model misspecification. Technical report, New York University.
- Harvey, C. R., Liu, Y., and Zhu, H. (2015). ... and the cross-section of expected returns. *The Review of Financial Studies*, 29(1):5–68.
- Harvey, C. R. and Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, 55(3):1263–1295.
- Haugen, R. A. and Baker, N. L. (1996). Commonality in the determinants of expected stock returns. *Journal of Financial Economics*, 41(3):401–439.

- Heckman, J. J., Ichimura, H., and Todd, P. E. (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The Review of Economic Studies*, 64(4):605–654.
- Heitz, A. (2022). The long-run impact of government asset guarantees. *FDIC Working Paper Series*.
- Heston, S. L. and Sadka, R. (2008). Seasonality in the cross-section of stock returns. *Journal of Financial Economics*, 87(2):418–445.
- Hirschhorn, E. and Zervos, D. (1990). Policies to change the priority of claimants: The case of depositor preference laws. *Journal of Financial Services Research*, 4:111–125.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30.
- Hotelling, H. (1929). Stability in competition. *The Economic Journal*, 39(153):41–57.
- Hou, K., Xue, C., and Zhang, L. (2014). Digesting anomalies: An investment approach. *The Review of Financial Studies*, 28(3):650–705.
- Hull, J. C. (2015). *Options, Futures, and Other Derivatives*. Pearson Education, 9th edition.
- Amari, S. (2020). *Information Geometry and Its Applications*, volume 194 of *Applied Mathematical Sciences*. Springer.
- Iyer, R. and Puri, M. (2012). Understanding bank runs: The importance of depositor-bank relationships and networks. *The American Economic Review*, 102(4):1414–1445.
- James, C. (1991). The losses realized in bank failures. *The Journal of Finance*, 46(4):1223–1242.
- Jegadeesh, N. (1990). Evidence of predictable behavior of security returns. *The Journal of Finance*, 45(3):881–898.
- Jegadeesh, N. and Livnat, J. (2006). Revenue surprises and stock returns. *Journal of Accounting and Economics*, 41(1-2):147–171.

- Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91.
- Jensen, T. I., Kelly, B., and Pedersen, L. H. (2023). Is there a replication crisis in finance? *The Journal of Finance*, 78(5):2465–2518.
- Jin, Q., Jiang, R., and Mokhtari, A. (2025). Non-asymptotic global convergence analysis of bfgs with the armijo-wolfe line search.
- Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3):501–524.
- Kelly, M. and Gráda, C. Ó. (2000). Market contagion: Evidence from the panics of 1854 and 1857. *The American Economic Review*, 90(5):1110–1124.
- Keloharju, M., Linnainmaa, J. T., and Nyberg, P. (2016). Return seasonalities. *The Journal of Finance*, 71(4):1557–1590.
- Kessy, A., Lewin, A., and Strimmer, K. (2018). Optimal whitening and decorrelation. *The American Statistician*, 72(4):309–314.
- Kitamura, Y. and Stutzer, M. (1997). An information-theoretic alternative to generalized method of moments estimation. *Econometrica*, 65(4):861–874.
- Kitamura, Y. and Stutzer, M. (2002). Connections between entropic and linear projections in asset pricing estimation. *Journal of Econometrics*, 107(2):159–174.
- Kiwiel, K. C. (1997). Proximal minimization methods with generalized bregman functions. *SIAM Journal on Control and Optimization*, 35(4):1142–1168.
- Kozak, S., Nagel, S., and Santosh, S. (2020). Shrinking the cross-section. *Journal of Financial Economics*, 135(2):271–292.
- Lakonishok, J., Shleifer, A., and Vishny, R. W. (1994). Contrarian investment, extrapolation, and risk. *The Journal of Finance*, 49(5):1541–1578.

- Laloux, L., Cizeau, P., Bouchaud, J.-P., and Potters, M. (1999). Noise dressing of financial correlation matrices. *Physical Review Letters*, 83(7):1467–1470.
- Leland, H. E. (1994). Corporate debt value, bond covenants, and optimal capital structure. *The Journal of Finance*, 49(4):1213–1252.
- Leland, H. E. and Toft, K. B. (1996). Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads. *The Journal of Finance*, 51(3):987–1019.
- Lettau, M. (2023). High-dimensional factor models and the factor zoo. *NBER Working Paper Series*, (31719).
- Lettau, M. and Pelger, M. (2020). Estimating latent asset-pricing factors. *Journal of Econometrics*, 218(1):1–31.
- Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168.
- Litner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37.
- Liu, D. C. and Nocedal, J. (1989). On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45:503–528.
- Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498.
- Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *The Review of Financial Studies*, 17(4):951–983.
- Mallat, S. G. and Zhang, Z. (1993). Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415.
- Marchenko, V. A. and Pastur, L. A. (1967). Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457–483.

- Marino, J. A. and Bennett, R. L. (1999). The consequences of national depositor preference. *FDIC Banking Review*, 12(2):19–38.
- Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441.
- Martino, L., Elvira, V., and Louzada, F. (2017). Effective sample size for importance sampling based on discrepancy measures. *Signal Processing*, 131:386–401.
- McCoy, P. A. (2007). *The Moral Hazard Implications of Deposit Insurance: Theory and Evidence*, chapter 16, pages 417–442. International Monetary Fund.
- McCullagh, P. (1987). *Tensor Methods in Statistics*, volume 29 of *Monographs on Statistics and Applied Probability*. Chapman and Hall/CRC, London.
- McDill, K. (2004). Resolution costs and the business cycle. *FDIC Working Paper Series*.
- McLean, R. D. and Pontiff, J. (2015). Does academic research destroy stock return predictability? *The Journal of Finance*, 71(1):5–32.
- Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2):449–470.
- Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122.
- Molyneux, P. and Zhou, T. M. (2022). Banking market reaction to auctions of failed banks. *International Journal of Finance & Economics*, 27(1):518–534.
- Morgan, D. P., Pinkovskiy, M. L., and Yang, B. (2016). Banking deserts, branch closings, and soft information. Liberty Street Economics 20160307, Federal Reserve Bank of New York.
- Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, 34(4):768–783.
- Natarajan, B. K. (1995). Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234.

- Newey, W. K. (2009). Two-step series estimation of sample selection models. *The Econometrics Journal*, 12(s1):S217–S229.
- Nielsen, F. (2024). Divergences induced by the cumulant and partition functions of exponential families and their deformations induced by comparative convexity. *Entropy*, 26(3).
- Nishiyama, T. (2020). Convex optimization on functionals of probability densities.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer Series in Operations Research. Springer, 2 edition.
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28.
- Novy-Marx, R. (2015). Fundamentally, momentum is fundamental momentum. NBER Working Papers 20984, National Bureau of Economic Research, Inc.
- O’Hara, M. and Shaw, W. (1990). Deposit insurance and wealth effects: The value of being “too big to fail”. *The Journal of Finance*, 45(5):1587–1600.
- Ohlson, J. A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research*, 18(1):109–131.
- Ortiz-Molina, H. and Phillips, G. M. (2014). Real asset illiquidity and the cost of capital. *The Journal of Financial and Quantitative Analysis*, 49(1):1–32.
- Osterberg, W. P. (1996). The impact of depositor preference laws. *Economic Review*, None(Q III):2–11.
- Pati, Y. C., Rezaiifar, R., and Krishnaprasad, P. S. (1993). Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of the 27th Asilomar Conference on Signals, Systems and Computers*, pages 40–44.
- Penas, M. F. and Unal, H. (2004). Gains in bank mergers: Evidence from the bond markets. *Journal of Financial Economics*, 74(1):149–179.

- Perry, A. (2026). Selling distress: The economics of failed bank resolutions. *Working Paper*.
- Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.
- Pontiff, J. and Woodgate, A. (2008). Share issuance and cross-sectional returns. *The Journal of Finance*, 63(2):921–945.
- Puhani, P. (2000). The heckman correction for sample selection and its critique. *Journal of Economic Surveys*, 14(1):53–68.
- Richardson, S. A., Sloan, R. G., Soliman, M. T., and Tuna, I. (2005). Accrual reliability, earnings persistence and stock prices. *Journal of Accounting and Economics*, 39(3):437–485.
- Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360.
- Santos, J. A. C. (2014). Evidence from the bond market on banks’ “too-big-to-fail” subsidy. *Economic Policy Review*, pages 29–39.
- Schmidt, E. (1907). Zur theorie der linearen und nichtlinearen integralgleichungen. i. entwicklung willkürlicher funktionen nach systemen vorgeschriebener. *Mathematische Annalen*, 63:433–476.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442.
- Sloan, R. G. (1996). Do stock prices fully reflect information in accruals and cash flows about future earnings? *The Accounting Review*, 71(3):289–315.
- Soliman, M. T. (2008). The use of dupont analysis by market participants. *The Accounting Review*, 83(3):823–853.

- Spokeviciute, L., Jahanshahloo, H., Keasey, K., and Vallascas, F. (2025). Three decades of failed bank acquisitions. *Journal of Banking and Finance*, 170.
- Staiger, D. and Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586.
- Stambaugh, R. F. and Yuan, Y. (2017). Mispricing factors. *The Review of Financial Studies*, 30(4):1270–1315.
- Stern, G. H. and Feldman, R. J. (2004). *Too Big to Fail: The Hazards of Bank Bailouts*. Brookings Institution Press.
- Stutzer, M. (1995). A bayesian approach to diagnosis of asset pricing models. *Journal of Econometrics*, 68(2):367–397.
- Thomas, J. and Zhang, F. X. (2011). Tax expense momentum. *Journal of Accounting Research*, 49(3):791–821.
- Titman, S., Wei, K. J., and Xie, F. (2003). Capital investments and stock returns. *NBER Working Paper Series*, 9951.
- Tropp, J. A. (2004). Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242.
- Tropp, J. A. (2011). User-friendly tail bounds for sums of random matrices. *The Journal of the Society for the Foundations of Computational Mathematics*, 12:389–434.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57(2):307–333.
- Wald, A. (1943). Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations Is Large. *Transactions of the American Mathematical Society*, 54(3):426–482.
- Wolfe, P. (1969). Convergence conditions for ascent methods. *SIAM Review*, 11(2):226–235.

- Zehna, P. W. (1966). Invariance of maximum likelihood estimators. *The Annals of Mathematical Statistics*, 37(3):744–744.

Appendix A

Appendix to Chapter 1

A.1 Proofs

This appendix provides formal proofs for all propositions, corollaries, and formal results. We adopt the notation from the main text. Let $(\Omega, \mathfrak{F}, \mathcal{P})$ be the probability space of the physical economy, and let $\mathbb{M}$ be the manifold of admissible SDFs parameterized by $\boldsymbol{\theta} \in \mathbb{R}^K$ using the normalized exponential-affine SDF.

A.1.1 Proposition 1

Under the exponential-affine SDF with factor $\mathbf{g}^\perp$ as the orthogonal component of $\mathbf{g}$ with respect to $\mathbb{S}$, the Wald statistic $W_T := \hat{\boldsymbol{\theta}}_{\mathbf{g}}^2 / \hat{\sigma}_{\boldsymbol{\theta}_{\mathbf{g}}}^2$ satisfies:

1. Under $H_0 : \boldsymbol{\theta}_{\mathbf{g}} = 0$, $W_T \rightarrow \chi_1^2$.
2. Under $H_1 : \boldsymbol{\theta}_{\mathbf{g}} \neq 0$, $W_T \rightarrow \chi_1^2(\delta_{\mathbf{g}})$, with noncentrality

$$\delta_{\mathbf{g}} = \frac{T(\boldsymbol{\theta}_{\mathbf{g}}^*)^2}{\sigma_{\boldsymbol{\theta}_{\mathbf{g}}, \infty}^2} \tag{A.1}$$

where $\sigma_{\boldsymbol{\theta}_{\mathbf{g}}, \infty}^2$ denotes the long-run asymptotic variance of $\sqrt{T}\hat{\boldsymbol{\theta}}_{\mathbf{g}}$.

3. To second order in $\boldsymbol{\theta}_{\mathbf{g}}^*$, the noncentrality parameter coincides with twice the KL diver-

gence between the augmented and restricted models scaled by the long run variance,

$$\delta_{\mathbf{g}} = \frac{T}{\Omega_{\mathbf{gg}}} \frac{2D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})}{(\nabla_{\theta_{\mathbf{g}}} \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}))^2}. \quad (\text{A.2})$$

Proof. Throughout the main text, $\hat{\sigma}_{\theta_{\mathbf{g}}}^2$ denotes the estimated variance of the estimator $\hat{\theta}_{\mathbf{g}}$ itself, which is of order T^{-1} . Writing $\sigma_{\theta_{\mathbf{g}}, \infty}^2$ for the long-run asymptotic variance of $\sqrt{T}\hat{\theta}_{\mathbf{g}}$, the two are related by $\hat{\sigma}_{\theta_{\mathbf{g}}}^2 = \hat{\sigma}_{\theta_{\mathbf{g}}, \infty}^2/T$. The factor of T in the noncentrality parameter is therefore absorbed into the standard-error normalization appearing in the denominator of W_T , and no explicit T scaling is required when computing the statistic in practice.

Under H_0 , the MCE estimator satisfies the asymptotic normality established by Kitamura and Stutzer (1997)

$$\sqrt{T}\hat{\theta}_{\mathbf{g}} \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_{\theta_{\mathbf{g}}, \infty}^2), \quad \sigma_{\theta_{\mathbf{g}}, \infty}^2 = (D_{\mathbf{gg}}^\top \Omega_{\mathbf{gg}}^{-1} D_{\mathbf{gg}})^{-1}, \quad (\text{A.3})$$

where $D_{\mathbf{gg}} = \nabla_{\theta_{\mathbf{g}}} \mathbf{E}[m_t(\boldsymbol{\theta}) \mathbf{r}_t^e]$ and $\Omega_{\mathbf{gg}}$ is the long-run covariance of the orthogonalized moment condition. The orthogonalization of $\mathbf{g}$ against $\mathbb{S}$ renders the Fisher information matrix block diagonal, so $\hat{\theta}_{\mathbf{g}}$ is asymptotically independent of $\hat{\theta}_{\mathbb{S}}$.

Given any consistent estimator $\hat{\sigma}_{\theta_{\mathbf{g}}, \infty}^2$ of $\sigma_{\theta_{\mathbf{g}}, \infty}^2$, the Wald statistic admits the scale-invariant representation

$$W_T = \frac{\hat{\theta}_{\mathbf{g}}^2}{\hat{\sigma}_{\theta_{\mathbf{g}}, \infty}^2} = \frac{(\sqrt{T}\hat{\theta}_{\mathbf{g}})^2}{\hat{\sigma}_{\theta_{\mathbf{g}}, \infty}^2} \xrightarrow{\mathcal{D}} \frac{Z^2}{1} = \chi_1^2, \quad Z \sim \mathcal{N}(0, 1), \quad (\text{A.4})$$

by the continuous mapping theorem. This is the standard Wald result of Wald (1943); see Engle (1984) for a unified treatment.

Under H_1 , local asymptotic normality Kitamura and Stutzer (1997) gives

$$\sqrt{T}(\hat{\theta}_{\mathbf{g}} - \theta_{\mathbf{g}}^*) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_{\theta_{\mathbf{g}}, \infty}^2), \quad (\text{A.5})$$

so that $\sqrt{T}\hat{\theta}_{\mathbf{g}} \xrightarrow{\mathcal{D}} \mathcal{N}(\sqrt{T}\theta_{\mathbf{g}}^*, \sigma_{\theta_{\mathbf{g}}, \infty}^2)$. It follows that

$$W_T = \left(\frac{\sqrt{T} \hat{\theta}_{\mathbf{g}}}{\hat{\sigma}_{\theta_{\mathbf{g}}, \infty}} \right)^2 + o_{\mathcal{P}}(1) \xrightarrow{\mathcal{D}} \chi_1^2(\delta^2), \quad \delta^2 := \frac{T(\theta_{\mathbf{g}}^*)^2}{\sigma_{\theta_{\mathbf{g}}, \infty}^2}. \quad (\text{A.6})$$

We now relate the noncentrality parameter to the KL divergence between the augmented and restricted models. By the Bregman-divergence representation (1.7),

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}}) = \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) - \psi(\boldsymbol{\theta}^{\mathbb{S}}) - \nabla \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})^\top (\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}). \quad (\text{A.7})$$

Orthogonalization ensures $\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}} = (0, \dots, 0, \theta_{\mathbf{g}}^*)^\top$, so only the g -block of the Hessian of ψ contributes. A Taylor expansion of ψ around $\boldsymbol{\theta}^{\mathbb{S}}$ yields

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}}) = \frac{1}{2}(\theta_{\mathbf{g}}^*)^2 [\nabla^2 \psi(\boldsymbol{\theta}^{\mathbb{S}})]_{\mathbf{gg}} + \mathcal{O}_{\mathcal{P}}((\theta_{\mathbf{g}}^*)^3). \quad (\text{A.8})$$

By the information-matrix identity for exponential families, the Hessian of the CGF equals the Fisher information, which in turn coincides (up to the long-run adjustment $\Omega_{\mathbf{gg}}$ that vanishes in the i.i.d. case) with the inverse asymptotic variance

$$[\nabla^2 \psi(\boldsymbol{\theta}^{\mathbb{S}})]_{\mathbf{gg}} = \mathcal{I}_{\mathbf{gg}} = \sigma_{\theta_{\mathbf{g}}, \infty}^{-2}. \quad (\text{A.9})$$

Combining (A.8) and (A.9),

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}}) = \frac{(\theta_{\mathbf{g}}^*)^2}{2\sigma_{\theta_{\mathbf{g}}, \infty}^2} + \mathcal{O}_{\mathcal{P}}((\theta_{\mathbf{g}}^*)^3), \quad (\text{A.10})$$

so that

$$\delta^2 = \frac{T(\theta_{\mathbf{g}}^*)^2}{\sigma_{\theta_{\mathbf{g}}, \infty}^2} = \frac{T}{\Omega_{\mathbf{gg}}} \frac{2D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \| \mathcal{Q}^{\mathbb{S}})}{(\nabla_{\theta_{\mathbf{g}}} \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}))^2} + \mathcal{O}_{\mathcal{P}}((\theta_{\mathbf{g}}^*)^3). \quad (\text{A.11})$$

The identity (A.11) is a *local* second-order result; it equates the Wald noncentrality with twice the KL divergence up to cubic terms in the displacement $\theta_{\mathbf{g}}^*$. In the Gaussian i.i.d. baseline, the CGF ψ is quadratic in $\boldsymbol{\theta}$, the higher-order remainder vanishes identically, and one recovers the exact equality $\delta^2 = T(\theta_{\mathbf{g}}^*)^2 \sigma_{\mathbf{g}}^2 = 2T \cdot D_{KL}$. Away from the Gaussian case, (A.11) should be read as a local approximation whose accuracy is governed by the

curvature of ψ along the direction of the orthogonalized factor. $\square$

A.1.2 Proposition 2

For any fixed portfolio $\mathbf{w}$, the worst case density distortion that solves the inner problem in (1.11) is

$$h^*(\mathbf{w}) = \frac{\exp(-\vartheta^{-1}u(W_1))}{\mathbf{E}_{\mathcal{P}}[\exp(-\vartheta^{-1}u(W_1))]} \quad (\text{A.12})$$

The worst-case measure $\mathcal{Q}_{\mathbf{w}}^* := h^*(\mathbf{w})d\mathcal{P}$ is an exponential tilt of the reference model in the direction of low utility.

Proof. We solve the constrained variational problem

$$V(\mathbf{w}) := \min_{h \geq 0, \mathbf{E}_{\mathcal{P}}[h] = 1} \underbrace{\{\mathbf{E}_{\mathcal{P}}[hu(W_1)] + \vartheta \mathbf{E}_{\mathcal{P}}[h \ln(h)]\}}_{:= J(h)}. \quad (\text{A.13})$$

Let $\varsigma \in \mathbb{R}$ be the Lagrange multiplier for the normalization constraint $\mathbf{E}_{\mathcal{P}}[h] = 1$

$$\mathcal{L}(h; \varsigma) := \mathbf{E}_{\mathcal{P}}[hu(W_1)] + \vartheta \mathbf{E}_{\mathcal{P}}[h \ln(h)] - \varsigma (\mathbf{E}_{\mathcal{P}}[h] - 1). \quad (\text{A.14})$$

Note: The constraint $h \geq 0$ will be satisfied automatically by the exponential form of the solution. Taking the Gâteaux derivative w.r.t. the functional h in the direction of an arbitrary perturbation δh

$$\frac{\delta \mathcal{L}}{\delta h}(h; \varsigma) = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (\mathcal{L}(h + \epsilon \delta h; \varsigma) - \mathcal{L}(h; \varsigma)) = \mathbf{E}_{\mathcal{P}}[\delta h u(W_1) + \vartheta \delta h (1 + \ln(h)) - \varsigma \delta h] \quad (\text{A.15})$$

where I used the fact

$$\left. \frac{d}{d\epsilon} ((h + \epsilon \delta h) \ln(h + \epsilon \delta h)) \right|_{\epsilon=0} = \delta h (1 + \ln(h)). \quad (\text{A.16})$$

The F.O.C. for all admissible perturbations δh yields the pointwise condition

$$\delta h u(W_1) + \vartheta \delta h (1 + \ln(h)) - \varsigma \delta h = 0 \Rightarrow h^*(\mathbf{w}) = \exp\left(\frac{\varsigma - \vartheta}{\vartheta}\right) \exp(-\vartheta^{-1}u(W_1)) \quad (\text{A.17})$$

As promised this expression satisfies the non-negativity constraint for all realizations of W_1 . To impose the normalization constraint $\mathbf{E}_{\mathcal{P}}[h] = 1$ we require that

$$\exp\left(\frac{\varsigma - \vartheta}{\vartheta}\right) = \frac{1}{\mathbf{E}_{\mathcal{P}}[\exp(-\vartheta^{-1}u(W_1))]} \quad (\text{A.18})$$

Substituting this back into (A.17) gives the canonical Gibbs tilt

$$h^*(\mathbf{w}) = \frac{\exp(-\vartheta^{-1}u(W_1))}{\mathbf{E}_{\mathcal{P}}[\exp(-\vartheta^{-1}u(W_1))]} \quad (\text{A.19})$$

To show this functional is unique note that the second Gâteaux derivative of the entropy penalty is

$$\frac{\delta^2}{\delta h^2} \mathbf{E}_{\mathcal{P}}[h \ln(h)] = \mathbf{E}_{\mathcal{P}} \left[\frac{(\delta h)^2}{h} \right] > 0 \quad (\text{A.20})$$

for all perturbations $\delta h \neq 0$ with $h > 0$. Since the entropy penalty is strictly convex and the distorted expected utility is linear in the distortion h , the objective function $J(h)$ is strictly convex in h when $\vartheta > 0$. By Kiwiel (1997); Nishiyama (2020), because the constraint set $\{h \geq 0 : \mathbf{E}_{\mathcal{P}}[h] = 1\}$ is convex; a strictly convex functional on a convex constraint set has at most one minimizer, so $h^*(\mathbf{w})$ is unique.

Finally, I show that $h^*(\mathbf{w})$ is truly a minimizer and not a maximizer or saddle point. For any feasible perturbation with $\mathbf{E}_{\mathcal{P}}[\delta h] = 0$, a second order expansion gives

$$J(h^* + \epsilon \delta h) - J(h^*) = \frac{\epsilon^2 \vartheta}{2} \mathbf{E}_{\mathcal{P}} \left[\frac{(\delta h)^2}{h^*} \right] + \mathcal{O}_{\mathcal{P}}(\epsilon^3) > 0, \quad (\text{A.21})$$

confirming that $h^*(\mathbf{w})$ is a strict minimizer. □

A.1.3 Lemma 1

Under the SDF (1.14) the expected excess return on asset i decomposes

$$\mathbf{E}_{\mathcal{P}}[r_i^e] = - \left(\frac{\text{Cov}_{\mathcal{P}}(m^{\mathbb{S}}, r_i^e)}{\mathbf{E}_{\mathcal{P}}[m^{\mathbb{S}}]} + \frac{\text{Cov}_{\mathcal{P}}(h^*, m^{\mathbb{S}} r_i^e)}{\mathbf{E}_{\mathcal{P}}[h^* m^{\mathbb{S}}]} \right) \quad (\text{A.22})$$

Proof. Under the misspecification adjusted SDF

$$m := h^*(\mathbf{w}^*) \frac{u'(W_1)}{R_f \mathbf{E}_{\mathcal{P}}[u'(W_1)]} = h^*(\mathbf{w}^*) m^{\mathbb{S}} \quad (\text{A.23})$$

the pricing restriction for asset i is $\mathbf{E}_{\mathcal{P}}[h^* m^{\mathbb{S}} r_i^e] = 0$. Using the covariance decomposition

$$\mathbf{E}[ABC] = \mathbf{E}[A]\mathbf{E}[BC] + \text{Cov}(A, BC) \quad (\text{A.24})$$

and the normalization $\mathbf{E}_{\mathcal{P}}[h^*] = 1$ we write the pricing anomaly as

$$\mathbf{E}_{\mathcal{P}}[m^{\mathbb{S}} r_i^e] = -\text{Cov}_{\mathcal{P}}(h^*(\mathbf{w}^*), m^{\mathbb{S}} r_i^e). \quad (\text{A.25})$$

Applying the same covariance decomposition to $\mathbf{E}_{\mathcal{P}}[m^{\mathbb{S}} r_i^e]$ gives

$$\mathbf{E}_{\mathcal{P}}[r_i^e] = -\frac{\text{Cov}_{\mathcal{P}}(m^{\mathbb{S}}, r_i^e)}{\mathbf{E}_{\mathcal{P}}[m^{\mathbb{S}}]} - \frac{\text{Cov}_{\mathcal{P}}(h^*(\mathbf{w}^*), m^{\mathbb{S}} r_i^e)}{\mathbf{E}_{\mathcal{P}}[m^{\mathbb{S}}]}. \quad (\text{A.26})$$

□

A.1.4 Proposition 3

Under the multiplier preferences (1.11) and the exponential-affine SDF specification, the worst-case welfare loss from misspecification in the direction of factor $\mathbf{g}$, to leading order, satisfies

$$|\Delta V_{\mathbf{g}}| \leq \sqrt{2\vartheta \mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(u(W_1^*))} \sqrt{D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})}. \quad (\text{A.27})$$

Moreover, the ranking of candidates by worst-case welfare loss is identical to the ranking by KL divergence

$$|\Delta V_{\mathbf{g}}| > |\Delta V_{\mathbf{g}'}| \iff D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) > D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g}')} || \mathcal{Q}^{\mathbb{S}}). \quad (\text{A.28})$$

Proof. Let $\mathcal{Q}^{\mathbb{S}}$ denote the baseline risk-neutral measure with Radon-Nikodym derivative

$$\frac{d\mathcal{Q}^{\mathbb{S}}}{d\mathcal{P}} = m_t(\boldsymbol{\theta}^{\mathbb{S}}) = \exp\left(-(\boldsymbol{\theta}^{\mathbb{S}})^\top \mathbf{f}_t - \psi(\boldsymbol{\theta}^{\mathbb{S}})\right). \quad (\text{A.29})$$

Similarly, let $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ denote the augmented measure with the augmented risk price vector $\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}$. The parameter vectors $\boldsymbol{\theta}^{\mathbb{S}}, \boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} \in \mathbb{R}^K$; the components corresponding to factors not in $\mathbb{S}$ are zero in $\boldsymbol{\theta}^{\mathbb{S}}$, and the component corresponding to the candidate $\mathbf{g}$ takes the value $\theta_{\mathbf{g}}^*$ in $\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}$.

Let

$$\mathbf{g}_t^\perp := \mathbf{g}_t - \text{Proj}_{\text{span}(\mathbb{S})}^{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}_t) \quad (\text{A.30})$$

be the orthogonalized factor defined in Lemma 3. By Assumption 1 the moment generating function $\mathbf{E}_{\mathcal{P}}[\exp(-\boldsymbol{\theta}^\top \mathbf{f}_t)]$ exists and is finite on an open set containing $\Theta \subseteq \mathbb{R}^K$, the factors are non-degenerate under $\mathcal{P}$, and $u(W_1^*)$ has finite variance under $\mathcal{Q}^{\mathbb{S}}$.

The indirect value function under model $\mathcal{Q}$ is

$$V(\mathbf{w}; \mathcal{Q}) := -\vartheta \ln(\mathbf{E}_{\mathcal{Q}}[\exp(-\vartheta^{-1}u(W_1))]). \quad (\text{A.31})$$

The welfare damage from misspecification in the direction of $\mathbf{g}$ is

$$\Delta V_{\mathbf{g}} := V(\mathbf{w}^*; \mathcal{Q}^{(\mathbb{S}, \mathbf{g})}) - V(\mathbf{w}^*; \mathcal{Q}^{\mathbb{S}}). \quad (\text{A.32})$$

The proof proceeds by deriving the upper bound, the equality condition, and the ranking equivalence.

The Donsker-Varadhan sets the upper bound on the welfare damage. The Donsker-Varadhan variational representation Donsker and Varadhan (1975) states that for any two probability measures $\mathcal{Q}_1, \mathcal{Q}_2$ with $\mathcal{Q}_1 \ll \mathcal{Q}_2$ and any measurable function φ with finite moment generating function under $\mathcal{Q}_2$,

$$\ln(\mathbf{E}_{\mathcal{Q}_2}[e^\varphi]) = \sup_{\mathcal{Q} \ll \mathcal{Q}_2} \{\mathbf{E}_{\mathcal{Q}}[\varphi] - D_{KL}(\mathcal{Q}||\mathcal{Q}_2)\}. \quad (\text{A.33})$$

Applying (A.33) to $t\varphi$ with the sub-Gaussian bound on the cumulant generating function

$$\ln \left(\mathbf{E}_{\mathcal{Q}_2} \left[e^{t(\varphi - \mathbf{E}_{\mathcal{Q}_2}[\varphi])} \right] \right) \leq \frac{1}{2} t^2 \mathbf{V}_{\mathcal{Q}_2}(\varphi) \quad (\text{A.34})$$

see Hoeffding (1963), and optimizing the resulting inequality over $t \in \mathbb{R}$, one obtains

$$|\mathbf{E}_{\mathcal{Q}}[\varphi] - \mathbf{E}_{\mathcal{Q}_2}[\varphi]| \leq \sqrt{2\mathbf{V}_{\mathcal{Q}_2}(\varphi) D_{KL}(\mathcal{Q} \parallel \mathcal{Q}_2)}. \quad (\text{A.35})$$

Applying (A.35) with $\mathcal{Q} = \mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$, $\mathcal{Q}_2 = \mathcal{Q}^{\mathbb{S}}$, and $\varphi = u(W_1^*)$:

$$|\mathbf{E}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}[u(W_1^*)] - \mathbf{E}_{\mathcal{Q}^{\mathbb{S}}}[u(W_1^*)]| \leq \sqrt{2\mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(u(W_1^*)) D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} \parallel \mathcal{Q}^{\mathbb{S}})}. \quad (\text{A.36})$$

Relating (A.36) to the welfare damage (A.32). For the augmented model $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ we have

$$\mathbf{E}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}} \left[\exp(-\vartheta^{-1} u(W_1^*)) \right] = \mathbf{E}_{\mathcal{Q}^{\mathbb{S}}} \left[\frac{d\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}{d\mathcal{Q}^{\mathbb{S}}} \exp(-\vartheta^{-1} u(W_1^*)) \right]. \quad (\text{A.37})$$

On the exponential-affine manifold, the likelihood ratio between the augmented and selected measures is

$$\frac{d\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}{d\mathcal{Q}^{\mathbb{S}}} = \frac{m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})}{m_t(\boldsymbol{\theta}^{\mathbb{S}})} = \exp(-\boldsymbol{\theta}_{\mathbf{g}}^* \mathbf{g}_t^\perp - (\psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) - \psi(\boldsymbol{\theta}^{\mathbb{S}}))) \quad (\text{A.38})$$

The orthogonality of $\mathbf{g}_t^\top$ w.r.t. $\text{span}(\mathbb{S})$ under $\mathcal{Q}^{\mathbb{S}}$ ensures that the likelihood ratio depends on $\mathbf{g}^\perp$ only. Substituting (A.38) into (A.37)

$$\mathbf{E}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}} \left[\exp(-\vartheta^{-1} u(W_1^*)) \right] = \mathbf{E}_{\mathcal{Q}^{\mathbb{S}}} \left[\exp(-\boldsymbol{\theta}_{\mathbf{g}}^* \mathbf{g}_t^\perp - \Delta\psi) \exp(-\vartheta^{-1} u(W_1^*)) \right] \quad (\text{A.39})$$

$$= e^{-\Delta\psi} \mathbf{E}_{\mathcal{Q}^{\mathbb{S}}} \left[\exp(-\boldsymbol{\theta}_{\mathbf{g}}^* \mathbf{g}_t^\perp - \vartheta^{-1} u(W_1^*)) \right] \quad (\text{A.40})$$

where $\Delta\psi := \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) - \psi(\boldsymbol{\theta}^{\mathbb{S}})$. Define the joint cumulant generating

$$\Upsilon(a, b) := \ln \left(\mathbf{E}_{\mathcal{Q}^{\mathbb{S}}} \left[\exp(a(-u(W_1^*)) + b(-\mathbf{g}_t^\perp)) \right] \right). \quad (\text{A.41})$$

Υ is well-defined and finite in the neighborhood of the origin by 1. Taking logarithms of (A.40) and multiplying by $-\vartheta$ gives

$$V(\mathbf{w}^*; \mathcal{Q}^{(\mathbb{S}, \mathbf{g})}) = -\vartheta \ln \left(e^{-\Delta\psi} \exp \left(\Upsilon(\vartheta^{-1}, \theta_{\mathbf{g}}^*) \right) \right) = \vartheta \Delta\psi - \vartheta \Upsilon(\vartheta^{-1}, \theta_{\mathbf{g}}^*). \quad (\text{A.42})$$

For the selected model $b = 0$,

$$V(\mathbf{w}^*; \mathcal{Q}^{\mathbb{S}}) = -\vartheta \Upsilon(\vartheta^{-1}, 0). \quad (\text{A.43})$$

The welfare loss is

$$\Delta V_{\mathbf{g}} = \vartheta \left(\Delta\psi + \Upsilon(\vartheta^{-1}, \theta_{\mathbf{g}}^*) - \Upsilon(\vartheta^{-1}, 0) \right). \quad (\text{A.44})$$

Expanding (A.44) to leading order. The CGF of $-u(W_1^*)$ under measure $\mathcal{Q}$ admits the cumulant expansion (McCullagh, 1987, Equation 2.4)

$$-\vartheta \ln \left(\mathbf{E}_{\mathcal{Q}} \left[e^{-\vartheta^{-1} u(W_1^*)} \right] \right) = -\vartheta \sum_{n=1}^{\infty} \frac{(-\vartheta^{-1})^n}{n!} \kappa_{\mathcal{Q}}^{(n)}(u(W_1^*)) \quad (\text{A.45})$$

$$= \sum_{n=1}^{\infty} \frac{(-1)^{n+1} \vartheta^{1-n}}{n!} \kappa_{\mathcal{Q}}^{(n)}(u(W_1^*)) \quad (\text{A.46})$$

$$= \mathbf{E}_{\mathcal{Q}}[u(W_1^*)] - \frac{\vartheta^{-1}}{2} \mathbf{V}_{\mathcal{Q}}(u(W_1^*)) + \dots \quad (\text{A.47})$$

where $\kappa_{\mathcal{Q}}^{(n)}(u(W_1^*))$ is the n -th cumulant of under $\mathcal{Q}$. Differencing between $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ and $\mathcal{Q}^{\mathbb{S}}$ gives

$$\begin{aligned} \Delta V_{\mathbf{g}} &= \mathbf{E}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}[u(W_1^*)] - \mathbf{E}_{\mathcal{Q}^{\mathbb{S}}}[u(W_1^*)] - \frac{\vartheta^{-1}}{2} \left(\mathbf{V}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}(u(W_1^*)) - \mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(u(W_1^*)) \right) \\ &\quad + \mathcal{O}_{\mathcal{P}} \left(\vartheta^{-2} \Delta \kappa^{(n)} \right) \end{aligned} \quad (\text{A.48})$$

where $\Delta \kappa^{(n)} := \kappa_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}^{(n)}(u(W_1^*)) - \kappa_{\mathcal{Q}^{\mathbb{S}}}^{(n)}(u(W_1^*))$. The change of measure from $\mathcal{Q}^{\mathbb{S}}$ to $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ is an exponential tilt with parameter $\theta_{\mathbf{g}}^*$, so each cumulant is a smooth function of $\theta_{\mathbf{g}}^*$ and

$\Delta\kappa^{(n)} = \mathcal{O}_{\mathcal{P}}(|\theta_{\mathbf{g}}^*|)$, $\forall n \in \mathbb{N}_+$. The leading term is bounded by (A.36).

The ranking equivalence result (A.28) asserts that $D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})$ provides a sufficient statistic for ordering candidates by welfare damage. We establish this in three parts: expressing the KL divergence as a Bregman divergence, showing that both $|\Delta V_{\mathbf{g}}|$ and D_{KL} are strictly increasing functions of the same pair of sufficient statistics.

From the closed-form derivation in (1.7), the KL divergence on the exponential-affine manifold equals the Bregman divergence generated by ψ ichi Amari (2020). Since $\boldsymbol{\theta}^{\mathbb{S}}$ and $\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}$ differ only in the $\mathbf{g}$ -component, and the $\mathbf{g}$ -component of $\boldsymbol{\theta}^{\mathbb{S}}$ is zero

$$\nabla\psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})^{\top}(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}) = \theta_{\mathbf{g}}^* \frac{\partial\psi}{\partial\theta_{\mathbf{g}}}(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}). \quad (\text{A.49})$$

Substituting

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) = \psi(\boldsymbol{\theta}^{\mathbb{S}}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) + \theta_{\mathbf{g}}^* \frac{\partial\psi}{\partial\theta_{\mathbf{g}}}(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}). \quad (\text{A.50})$$

Differentiating w.r.t. $\theta_{\mathbf{g}}^*$

$$\frac{\partial D_{KL}}{\partial\theta_{\mathbf{g}}^*}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) = \theta_{\mathbf{g}}^* \frac{\partial^2\psi}{\partial\theta_{\mathbf{g}}^2}(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) = \theta_{\mathbf{g}}^* \mathbf{V}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}(\mathbf{g}^{\perp}) \quad (\text{A.51})$$

where $\partial_{\theta_{\mathbf{g}}}^2\psi = I_{\mathbf{g}\mathbf{g}}(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) = \mathbf{V}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}(\mathbf{g}^{\perp}) > 0$ by the positive-definiteness of the Fisher information 1. $\partial_{\theta_{\mathbf{g}}} D_{KL} > 0$ when $\theta_{\mathbf{g}} > 0$ and negative when $\theta_{\mathbf{g}} < 0$; so D_{KL} is strictly increasing in $|\theta_{\mathbf{g}}^*|$ with a minimum at $\theta_{\mathbf{g}}^* = 0$. Also note that the divergence is strictly increasing in the residual variance $\mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}^{\perp})$.

Now, I show that the absolute welfare loss is also monotonic in $|\theta_{\mathbf{g}}^*|$ and $\mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}^{\perp})$. Differentiating (A.44) with respect to $\theta_{\mathbf{g}}^*$

$$\frac{\partial\Delta V_{\mathbf{g}}}{\partial\theta_{\mathbf{g}}^*} = \vartheta \left(\frac{\partial\Delta\psi}{\partial\theta_{\mathbf{g}}^*} + \frac{\partial\Upsilon}{\partial\theta_{\mathbf{g}}^*}(\vartheta^{-1}, \theta_{\mathbf{g}}^*) \right). \quad (\text{A.52})$$

The first term is

$$\frac{\partial\Delta\psi}{\partial\theta_{\mathbf{g}}^*} = \frac{\partial\psi}{\partial\theta_{\mathbf{g}}}(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) = -\mathbf{E}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}[\mathbf{g}_t^{\perp}]. \quad (\text{A.53})$$

The second term is given by

$$\frac{\partial \Upsilon}{\partial \theta_{\mathbf{g}}^*}(\vartheta^{-1}, \theta_{\mathbf{g}}^*) = -\mathbf{E}_{\Xi(\vartheta^{-1}, \theta_{\mathbf{g}}^*)}[\mathbf{g}_t^\perp] \quad (\text{A.54})$$

where $\Xi(\vartheta^{-1}, \theta_{\mathbf{g}}^*)$ is the exponential tilt of $\mathcal{Q}^{\mathbb{S}}$ with Radon-Nikodym derivative

$$\frac{d\Xi}{d\mathcal{Q}^{\mathbb{S}}}(\vartheta^{-1}, \theta_{\mathbf{g}}^*) = \frac{\exp(-\vartheta^{-1}u(W_1^*) - \theta_{\mathbf{g}}^* \mathbf{g}_t^\perp)}{\mathbf{E}_{\mathcal{Q}^{\mathbb{S}}}[\exp(-\vartheta^{-1}u(W_1^*) - \theta_{\mathbf{g}}^* \mathbf{g}_t^\perp)]}. \quad (\text{A.55})$$

Evaluate at $\theta_{\mathbf{g}}^* = 0$. By the orthogonalization construction, $\mathbf{E}_{\mathcal{Q}^{\mathbb{S}}}[\mathbf{g}_t^\perp] = 0$. The measure $\Xi(\vartheta^{-1}, 0)$ is the tilt in the $-\vartheta^{-1}u(W_1^*)$ direction alone; when $\theta_{\mathbf{g}}^* = 0$ the optimal portfolio does not load on $\mathbf{g}^\perp$, so $u(W_1^*)$ and $\mathbf{g}^\perp$ are uncorrelated under $\mathcal{Q}^{\mathbb{S}}$ and

$$\left. \frac{\partial \Delta V_{\mathbf{g}}}{\partial \theta_{\mathbf{g}}^*} \right|_{\theta_{\mathbf{g}}^*=0} = 0. \quad (\text{A.56})$$

For the second derivative, the map $b \mapsto \Upsilon(\vartheta^{-1}, b)$ is the CGF of $-\mathbf{g}_t^\perp$ under $\Xi(\vartheta^{-1}, 0)$, and CGFs of non-degenerate random variables are strictly convex

$$\frac{\partial^2 \Upsilon}{\partial b^2}(\vartheta^{-1}, b) = \mathbf{V}_{\Xi(\vartheta^{-1}, b)}(\mathbf{g}_t^\perp) > 0. \quad (\text{A.57})$$

The variance is strictly positive because $\mathbf{g}_t^\perp$ is non-degenerate under $\mathcal{Q}^{\mathbb{S}}$ and exponential tilting preserves non-degeneracy.

The strict convexity of Υ in b implies that $\Upsilon(\vartheta^{-1}, \theta_{\mathbf{g}}^*) - \Upsilon(\vartheta^{-1}, 0)$ is strictly convex in $\theta_{\mathbf{g}}^*$ with minimum zero at the origin. Similarly, $\Delta\psi$ is strictly convex in $\theta_{\mathbf{g}}^*$ with minimum zero because $I_{\mathbf{g}\mathbf{g}} > 0$. The sum $\Delta\psi + \Upsilon(\vartheta^{-1}, \theta_{\mathbf{g}}^*) - \Upsilon(\vartheta^{-1}, 0)$ is therefore strictly convex with minimum zero at $\theta_{\mathbf{g}}^* = 0$. A strictly convex function with minimum at the origin is strictly increasing in the absolute value of its argument, so $|\Delta V_{\mathbf{g}}|$ is strictly increasing in $|\theta_{\mathbf{g}}^*|$. The dependence on $\mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}^\perp)$ follows by the same logic: increasing $\mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}^\perp)$ increases the curvature $\partial_b^2 \Upsilon$ and increases $\Delta\psi$ through $I_{\mathbf{g}\mathbf{g}}$.

Define the maps

$$F : (|\theta_{\mathbf{g}}^*|, \sigma_{\mathbf{g}}^2) \mapsto D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}), \quad (\text{A.58})$$

$$H : (|\theta_{\mathbf{g}}^*|, \sigma_{\mathbf{g}}^2) \mapsto |\Delta V_{\mathbf{g}}|, \quad (\text{A.59})$$

where $\sigma_{\mathbf{g}}^2 := \mathbf{V}_{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}^{\perp})$. Both F and H are strictly increasing in each argument, and both equal zero if and only if $\theta_{\mathbf{g}}^* = 0$.

When the baseline $\mathbb{S}$ is held fixed, the pair $(|\theta_{\mathbf{g}}^*|, \sigma_{\mathbf{g}}^2)$ constitutes the only quantities that vary across candidates $\mathbf{g} \notin \mathbb{S}$. The risk price $\theta_{\mathbf{g}}^*$ is determined by the m -projection of the physical measure $\mathcal{P}$ onto $\mathbb{M}^{(\mathbb{S}, \mathbf{g})}$, and $\sigma_{\mathbf{g}}^2$ is determined by the covariance structure of the factors under $\mathcal{Q}^{\mathbb{S}}$. Both are population quantities independent of the investor's preferences.

For each candidate $\mathbf{g}$, the value $\sigma_{\mathbf{g}}^2$ is fixed. Along the one-dimensional submanifold $\mathbb{M}^{(\mathbb{S}, \mathbf{g})}$ parameterized by $\theta_{\mathbf{g}}$ alone, both $\Delta V_{\mathbf{g}}$ and $D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}})$ are functions of the single scalar $\theta_{\mathbf{g}}^*$. The maps F and H , restricted to this path, are strictly increasing functions of the same scalar, so the composition $H \circ F^{-1}$ is strictly increasing on $[0, \infty)$. It follows that

$$|\Delta V_{\mathbf{g}}| > |\Delta V_{\mathbf{g}'}| \iff H(|\theta_{\mathbf{g}}^*|, \sigma_{\mathbf{g}}^2) > H(|\theta_{\mathbf{g}'}^*|, \sigma_{\mathbf{g}'}^2) \quad (\text{A.60})$$

$$\iff F(|\theta_{\mathbf{g}}^*|, \sigma_{\mathbf{g}}^2) > F(|\theta_{\mathbf{g}'}^*|, \sigma_{\mathbf{g}'}^2) \quad (\text{A.61})$$

$$\iff D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) > D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g}')} || \mathcal{Q}^{\mathbb{S}}). \quad (\text{A.62})$$

□

A.1.5 Lemma 2

The forward selection algorithm is invariant to affine reparameterization of the exponential family.

Proof. Let $\mathcal{Q}(\boldsymbol{\theta})$ be an exponential family of distributions induced by the SDF

$$m_t(\boldsymbol{\theta}) := \exp(-\boldsymbol{\theta}^{\top} \mathbf{f}_t - \psi(\boldsymbol{\theta})) \quad (\text{A.63})$$

where $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^K$ are the canonical coordinates.

Consider an affine change of coordinates $\eta : \Theta \rightarrow \Xi$ defined by $\eta = A\theta + c$, where A is an invertible matrix and c is a constant vector. This transformation defines a new parameterization of the same family of distributions. The selection criterion is the KL divergence between the augmented and restricted models,

$$\Delta D_{KL} = D_{KL}(\mathcal{Q}^{(\mathbb{S}, g)} || \mathcal{Q}^{\mathbb{S}}). \quad (\text{A.64})$$

which depends only on the Radon-Nikodym derivative $d\mathcal{Q}^{(\mathbb{S}, g)}/d\mathcal{Q}^{\mathbb{S}}$ and hence only on the distributions themselves, not on their parameterization. Because the MCE estimator is equivariant under invertible parameter transformations Zehna (1966), the fitted distributions $\mathcal{Q}^{(\mathbb{S}, g)}$ and $\mathcal{Q}^{\mathbb{S}}$ are identical across coordinate systems. It follows that ΔD_{KL} , and therefore the ranking of candidate factors $\mathbf{g}$, is invariant to the choice of coordinates. $\square$

A.1.6 Lemma 3

Replacing a candidate factor $\mathbf{g}$ with its orthogonal residual $\mathbf{g}^\perp := \mathbf{g} - \text{Proj}_{\text{span}(\mathbb{S})}^{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g})$ under the current risk-neutral measure $\mathcal{Q}^{\mathbb{S}}$ leaves the augmented SDF manifold and the KL divergence unchanged.

Proof. Let $\mathbb{S} = \{\mathbf{f}_1, \dots, \mathbf{f}_k\}$ be the set of currently selected factors. The manifold of SDFs spanned by $\mathbb{S}$ is

$$\mathbb{M}^{\mathbb{S}} := \left\{ \mathcal{Q}(\boldsymbol{\theta}) : \frac{d\mathcal{Q}}{d\mathcal{P}} = \exp \left(- \sum_{i=1}^k \theta_i f_{i,t} - \psi(\boldsymbol{\theta}) \right) \right\} \quad (\text{A.65})$$

Consider a new candidate factor $\mathbf{g}_t$. The augmented manifold $\mathbb{M}^{(\mathbb{S}, \mathbf{g})}$ is spanned by the linear combination of $\mathbb{S} \cup \{\mathbf{g}\}$. We define the orthogonalized factor $\mathbf{g}_t^\perp$ as the residual of $\mathbf{g}_t$ after projecting onto the span of $\mathbb{S}$

$$\mathbf{g}_t^\perp = \mathbf{g}_t - \text{Proj}_{\text{span}(\mathbb{S})}^{\mathcal{Q}^{\mathbb{S}}}(\mathbf{g}_t) = \mathbf{g}_t - \sum_{i=1}^k \alpha_i f_{i,t} \quad (\text{A.66})$$

where α_i are the projection coefficients under $\mathcal{Q}^{\mathbb{S}}$. For any vector $\mathbf{v}$ in the augmented span

$$\mathbf{v} = \sum_{i=1}^k \beta_i \mathbf{f}_i + \lambda \mathbf{g} = \sum_{i=1}^k \beta_i \mathbf{f}_i + \lambda \left(\mathbf{g}^{\perp} + \sum_{j=1}^k \alpha_j \mathbf{f}_j \right) = \sum_{i=1}^k (\beta_i + \lambda \alpha_i) \mathbf{f}_i + \lambda \mathbf{g}^{\perp}. \quad (\text{A.67})$$

Since the set of reachable vectors is identical, the set of admissible probability distributions is identical

$$\mathbb{M}^{(\mathbb{S}, \mathbf{g})} \equiv \mathbb{M}^{(\mathbb{S}, \mathbf{g}^{\perp})}. \quad (\text{A.68})$$

The information projection is unique. The estimation of the SDF parameters constitutes an m -projection of the true data-generating process $\mathcal{Q}^{\dagger}$ onto the model manifold. Since the two manifolds coincide, the projections coincide $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} = \mathcal{Q}^{(\mathbb{S}, \mathbf{g}^{\perp})}$. Therefore, the KL divergence contribution is identical

$$D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}}) = D_{KL}(\mathcal{Q}^{(\mathbb{S}, \mathbf{g}^{\perp})} || \mathcal{Q}^{\mathbb{S}}). \quad (\text{A.69})$$

□

A.2 Derivations

A.2.1 Closed Form Kullback-Leibler Divergence

This appendix details the derivation of equation (1.7), which provides the closed-form expression for the Kullback-Leibler (KL) divergence between two exponential-affine pricing models.

Let the exponential-affine SDF be defined as

$$m_t(\boldsymbol{\theta}) = \exp(-\boldsymbol{\theta}^{\top} \mathbf{f}_t - \psi(\boldsymbol{\theta})) \quad (\text{A.70})$$

where $\psi(\boldsymbol{\theta})$ is the cumulant generating function defined as

$$\psi(\boldsymbol{\theta}) = \ln \left(\mathbf{E}_{\mathcal{P}} \left[\exp \left(-\boldsymbol{\theta}^\top \mathbf{f}_t \right) \right] \right), \quad (\text{A.71})$$

ensuring that $\mathbf{E}_{\mathcal{P}}[m_t(\boldsymbol{\theta})] = 1$.

We consider the divergence between a restricted model $\mathcal{Q}^{\mathbb{S}}$ (parameterized by $\boldsymbol{\theta}^{\mathbb{S}}$) and an augmented model $\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}$ (parameterized by $\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}$). Both parameter vectors $\boldsymbol{\theta} \in \mathbb{R}^K$ are defined over the full factor space, where coefficients for non-selected factors are zero.

The KL divergence is defined as

$$D_{KL} \left(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}} \right) = \mathbf{E}_{\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}} \left[\ln \left(\frac{m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})}{m_t(\boldsymbol{\theta}^{\mathbb{S}})} \right) \right]. \quad (\text{A.72})$$

Switching the expectation to the physical measure $\mathcal{P}$, we have

$$D_{KL} \left(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}} \right) = \mathbf{E}_{\mathcal{P}} \left[m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) \ln \left(\frac{m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})}{m_t(\boldsymbol{\theta}^{\mathbb{S}})} \right) \right]. \quad (\text{A.73})$$

Substituting the exponential functional form into the log ratio

$$\ln \left(\frac{m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})}{m_t(\boldsymbol{\theta}^{\mathbb{S}})} \right) = \left(-(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})^\top \mathbf{f}_t - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) \right) - \left(-(\boldsymbol{\theta}^{\mathbb{S}})^\top \mathbf{f}_t - \psi(\boldsymbol{\theta}^{\mathbb{S}}) \right) \quad (\text{A.74})$$

$$= \psi(\boldsymbol{\theta}^{\mathbb{S}}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) - \mathbf{f}_t^\top (\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}) \quad (\text{A.75})$$

Substituting this back into the expectation

$$D_{KL} \left(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}} \right) = \mathbf{E}_{\mathcal{P}} \left[m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) \left(\psi(\boldsymbol{\theta}^{\mathbb{S}}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) - \mathbf{f}_t^\top (\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}) \right) \right]. \quad (\text{A.76})$$

Using the linearity of expectations and the normalization condition $\mathbf{E}_{\mathcal{P}}[m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})] = 1$

$$D_{KL} \left(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}} \right) = \psi(\boldsymbol{\theta}^{\mathbb{S}}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) - \mathbf{E}_{\mathcal{P}} \left[m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) \mathbf{f}_t^\top \right] (\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}). \quad (\text{A.77})$$

Next, we relate the expected factor payoff to the gradient of the cumulant generating

function. Taking the gradient of $\psi(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$

$$\nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \ln (\mathbf{E}_{\mathcal{P}} [\exp (-\mathbf{f}_t^{\top} \boldsymbol{\theta})]) \quad (\text{A.78})$$

$$= \mathbf{E}_{\mathcal{P}} [\exp (-\mathbf{f}_t^{\top} \boldsymbol{\theta})]^{-1} \mathbf{E}_{\mathcal{P}} [-\mathbf{f}_t^{\top} \exp (-\mathbf{f}_t^{\top} \boldsymbol{\theta})] \quad (\text{A.79})$$

$$= -\exp (-\psi(\boldsymbol{\theta})) \mathbf{E}_{\mathcal{P}} [m_t(\boldsymbol{\theta}) \mathbf{f}_t^{\top}] \exp (\psi(\boldsymbol{\theta})) \quad (\text{A.80})$$

$$= -\mathbf{E}_{\mathcal{P}} [m_t(\boldsymbol{\theta}) \mathbf{f}_t^{\top}] \quad (\text{A.81})$$

Plugging this back into the KL divergence expression gives

$$D_{KL} \left(\mathcal{Q}^{(\mathbb{S}, \mathbf{g})} || \mathcal{Q}^{\mathbb{S}} \right) = \psi(\boldsymbol{\theta}^{\mathbb{S}}) - \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) + \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})}) (\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})} - \boldsymbol{\theta}^{\mathbb{S}}). \quad (\text{A.82})$$

A.2.2 Minimum Cross Entropy Estimator

This appendix details the derivation of the Minimum Cross-Entropy (MCE) estimator used in Section 1.4. We follow the information-theoretic approach of Kitamura and Stutzer (1997), transforming the constrained primal problem into a computationally tractable unconstrained dual problem.

Let $p(\boldsymbol{\gamma}, \boldsymbol{\theta}) := \frac{d\mathcal{P}_{\boldsymbol{\gamma}, \boldsymbol{\theta}}}{d\mathcal{P}^{\dagger}}$; the constrained minimization uses a Lagrangian with multipliers ξ for $\int_{\Omega} p(\boldsymbol{\gamma}, \boldsymbol{\theta}) d\mathcal{P}^{\dagger} = 1$ and $\boldsymbol{\gamma} \in \mathbb{R}^N$ for the N pricing conditions

$$L(p(\boldsymbol{\gamma}, \boldsymbol{\theta}); \xi, \boldsymbol{\gamma}) := \int_{\Omega} (p(\boldsymbol{\gamma}, \boldsymbol{\theta}) \ln(p(\boldsymbol{\gamma}, \boldsymbol{\theta})) - \xi(p(\boldsymbol{\gamma}, \boldsymbol{\theta}) - 1) - \boldsymbol{\gamma}^{\top} m_t(\boldsymbol{\theta}) \mathbf{r}_t^e p(\boldsymbol{\gamma}, \boldsymbol{\theta})) d\mathcal{P}^{\dagger} \quad (\text{A.83})$$

Taking the Gateaux derivative w.r.t. $p(\boldsymbol{\gamma}, \boldsymbol{\theta})$ gives the F.O.C.

$$\frac{\partial L}{\partial p(\boldsymbol{\gamma}, \boldsymbol{\theta})} (p(\boldsymbol{\gamma}, \boldsymbol{\theta}); \xi, \boldsymbol{\gamma}) = \ln(p(\boldsymbol{\gamma}, \boldsymbol{\theta})) + 1 - \xi - \boldsymbol{\gamma}^{\top} m_t(\boldsymbol{\theta}) \mathbf{r}_t^e = 0. \quad (\text{A.84})$$

Which will have a solution of the form

$$p(\boldsymbol{\gamma}, \boldsymbol{\theta}) \propto \exp(\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e) \quad (\text{A.85})$$

Normalizing A.85, gives the exponential/Gibbs tilt

$$\frac{d\mathcal{P}_{\boldsymbol{\gamma}, \boldsymbol{\theta}}}{d\mathcal{P}^\dagger} = \frac{\exp(\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e)}{\mathbf{E}_{\mathcal{P}^\dagger}[\exp(\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e)]} \quad (\text{A.86})$$

where we choose $\boldsymbol{\gamma}$ so that the pricing conditions hold under $\mathcal{P}_{\boldsymbol{\gamma}, \boldsymbol{\theta}}$, $\mathbf{E}_{\mathcal{P}_{\boldsymbol{\gamma}, \boldsymbol{\theta}}}[m_t(\boldsymbol{\theta}) \mathbf{r}_t^e] = \mathbf{0}$. This is a challenging problem to solve in function space, but under the empirical measure, it becomes a finite-dimensional convex optimization problem.

The dual objective is the cumulant-generating-function of $\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e$

$$\phi_T(\boldsymbol{\gamma}; \boldsymbol{\theta}) := \ln \left(\frac{1}{T} \sum_{t=1}^T \exp(\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e) \right). \quad (\text{A.87})$$

This can be seen by plugging in the Gibbs tilt probability density $\mathcal{P}_{\boldsymbol{\gamma}, \boldsymbol{\theta}}$ back into the Lagrangian A.83 and minimizing over densities.

Define softmax weights

$$w_t(\boldsymbol{\gamma}, \boldsymbol{\theta}) := \frac{\exp(\boldsymbol{\gamma}^\top m_t(\boldsymbol{\theta}) \mathbf{r}_t^e)}{\sum_{\tau=1}^T \exp(\boldsymbol{\gamma}^\top m_\tau(\boldsymbol{\theta}) \mathbf{r}_\tau^e)}. \quad (\text{A.88})$$

Computing the gradient and Hessian of ϕ_T w.r.t. $\boldsymbol{\gamma}$ gives

$$\nabla_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta}) = \sum_{t=1}^T w_t(\boldsymbol{\gamma}, \boldsymbol{\theta}) m_t(\boldsymbol{\theta}) \mathbf{r}_t^e \quad (\text{A.89})$$

and

$$\nabla_{\boldsymbol{\gamma}}^2 \phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta}) = \sum_{t=1}^T w_t(\boldsymbol{\gamma}, \boldsymbol{\theta}) (m_t(\boldsymbol{\theta}) \mathbf{r}_t^e - \nabla_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta})) (m_t(\boldsymbol{\theta}) \mathbf{r}_t^e - \nabla_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta}))^\top. \quad (\text{A.90})$$

The optimization problem in $\boldsymbol{\gamma}$ is globally convex; the Hessian $\nabla_{\boldsymbol{\gamma}}^2 \phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta})$ is a weighted covariance, hence semidefinite and when $\{m_t(\boldsymbol{\theta}) \mathbf{r}_t^e\}$ spans $\mathbb{R}^N$ it is positive definite. I solve for $\boldsymbol{\gamma}$ by optimizing the moment conditions $\nabla_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\gamma}^*, \boldsymbol{\theta}) = \mathbf{0}$, i.e., the tilted moments are

zero. When $\nabla_{\boldsymbol{\gamma}}^2 \phi_T(\boldsymbol{\gamma}, \boldsymbol{\theta})$ is positive semidefinite, ϕ_T is convex in $\boldsymbol{\gamma}$, the minimizer exists and is unique Boyd and Vandenberghe (2009). Under standard smoothness assumptions and a Wolfe line search the Limited Memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) optimizer converges rapidly Liu and Nocedal (1989); Jin et al. (2025); Politis and Romano (1994).

After optimizing for $\boldsymbol{\gamma}^*$ we optimize over $\boldsymbol{\theta}$; define the profiled objective

$$J_T(\boldsymbol{\theta}) := \phi_T(\boldsymbol{\gamma}^*(\boldsymbol{\theta}), \boldsymbol{\theta}). \quad (\text{A.91})$$

By the Envelope Theorem,

$$\nabla_{\boldsymbol{\theta}} J_T(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \phi_T(\boldsymbol{\gamma}^*(\boldsymbol{\theta}), \boldsymbol{\theta}) = \sum_{t=1}^T w_t(\boldsymbol{\gamma}^*(\boldsymbol{\theta}), \boldsymbol{\theta}) \boldsymbol{\gamma}^{\top} \nabla_{\boldsymbol{\theta}} m_t(\boldsymbol{\theta}) \mathbf{r}_t^e. \quad (\text{A.92})$$

Taking the gradient of the pricing restrictions

$$\nabla_{\boldsymbol{\theta}} m_t(\boldsymbol{\theta}) \mathbf{r}_t^e = -m_t(\boldsymbol{\theta}) \mathbf{r}_t^e \left(\mathbf{f}_t - \frac{1}{T} \sum_{\tau=1}^T m_{\tau}(\boldsymbol{\theta}) \mathbf{f}_{\tau} \right)^{\top}, \quad (\text{A.93})$$

So the objective for the outer- $\boldsymbol{\theta}$ loop is

$$\nabla_{\boldsymbol{\theta}} J_T(\boldsymbol{\theta}) = - \sum_{t=1}^T w_t(\boldsymbol{\gamma}^*(\boldsymbol{\theta}), \boldsymbol{\theta}) m_t(\boldsymbol{\theta}) \boldsymbol{\gamma}^*(\boldsymbol{\theta})^{\top} \mathbf{r}_t^e \left(\mathbf{f}_t - \frac{1}{T} \sum_{\tau=1}^T m_{\tau}(\boldsymbol{\theta}) \mathbf{f}_{\tau} \right) \quad (\text{A.94})$$

Each loop completes different parts of the optimization; the inner $\boldsymbol{\gamma}$ loop chooses tilted probabilities $w_t(\boldsymbol{\gamma}, \boldsymbol{\theta})$ that make the pricing restrictions hold, while the outer $\boldsymbol{\theta}$ loop changes the SDF weights to reduce the Kullback-Leibler divergence between the candidate pricing kernel and the true pricing kernel.

A.2.3 Asymptotic Standard Errors

This appendix details the derivation of the standard errors for the parameter estimates and the Kullback-Leibler (KL) divergence, discusses the testing procedure for nested models, and establishes the conditions for parameter stability. We assume throughout that the stochastic process generating the data is stationary and ergodic and that the relevant moment

conditions satisfy the standard regularity conditions for identification and differentiability.

We employ the estimator of Kitamura and Stutzer (1997), which minimizes the Kullback-Leibler (KL) divergence between the estimated probability measure and the empirical measure. This is done under the exponential tilting.

Let $\mathbf{x}_t$ be a stationary, ergodic sequence of random vectors for $t = 1, \dots, T$. Let $\mathbf{h}(\mathbf{x}_t, \boldsymbol{\theta}) := m_t(\boldsymbol{\theta})\mathbf{r}_t^e$ denote the N -dimensional vector of moment conditions dependent on parameters $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^k$. The estimator is defined by the saddle-point problem

$$\hat{\boldsymbol{\theta}} := \arg \min_{\boldsymbol{\theta} \in \Theta} \max_{\boldsymbol{\gamma} \in \mathbb{R}^N} \phi_T(\boldsymbol{\theta}, \boldsymbol{\gamma}) \equiv \ln \left(\frac{1}{T} \sum_{t=1}^T \exp(\boldsymbol{\gamma}^\top \mathbf{h}(\mathbf{x}_t, \boldsymbol{\theta})) \right) \quad (\text{A.95})$$

where $\boldsymbol{\gamma}$ is the vector of Lagrange multipliers. Kitamura and Stutzer (1997) establish that the asymptotic variance of $\hat{\boldsymbol{\theta}}$ is identical to that of the optimal two-step efficient GMM estimator

$$\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \xrightarrow{\mathcal{D}} \mathcal{N}(0, (D^\top \Omega^{-1} D)^{-1}) \quad (\text{A.96})$$

where $D := \mathbf{E}_{\mathcal{P}}[\nabla_{\boldsymbol{\theta}} \mathbf{h}(\mathbf{x}_t, \boldsymbol{\theta}_0)]$ is the Jacobian of the moment conditions, and Ω is the long-run covariance matrix of the moments. To account for serial correlation and heteroskedasticity in the data, Ω is estimated using a HAC estimator with bandwidth L

$$\hat{\Omega} := \hat{\Gamma}_0 + \sum_{j=1}^L w(j, L)(\hat{\Gamma}_j + \hat{\Gamma}_j^\top) \quad (\text{A.97})$$

where

$$\hat{\Gamma}_j := \frac{1}{T} \sum_{t=j+1}^T \mathbf{h}(\mathbf{x}_t, \hat{\boldsymbol{\theta}}) \mathbf{h}(\mathbf{x}_t, \hat{\boldsymbol{\theta}})^\top \quad (\text{A.98})$$

and $w(j, L)$ are kernel weights.

We derive the standard errors for the minimized KL divergence statistic itself. The sample objective function be the one derived in Equation A.95. Let $(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}})$ denote the sample estimates, and $(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)$ denote the pseudo-true population parameters that minimize the population KL divergence. We seek the asymptotic distribution of $\hat{K} = \phi_T(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}})$. A first

order expansion of $\hat{K}$ around the pseudo-true values $(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)$

$$\hat{K} - K_0 \approx \phi_T(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*) - K_0 + \nabla_{\boldsymbol{\theta}} \phi_T(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) + \nabla_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)(\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}^*). \quad (\text{A.99})$$

By the Envelope Theorem, the variation in the objective function with respect to the optimal parameters is zero. The FOCs for the optimization problem imply that at the optimum

$$\mathbf{E}_{\mathcal{P}} [\nabla_{\boldsymbol{\theta}} \phi_T(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)] = 0 \quad \text{and} \quad \mathbf{E}_{\mathcal{P}} [\nabla_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)] = 0. \quad (\text{A.100})$$

Under the assumption that the model is pseudo-true, the gradient terms with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\gamma}$ vanish asymptotically ($o_{\mathcal{P}}(1)$). Thus, the variation in the KL divergence is driven solely by the randomness of the moment conditions at the fixed pseudo-true parameters.

Unlike standard Maximum Likelihood or GMM, where the objective function is a sum of logs the exponential tilt objective function involves the log of a sum. Therefore, we must use the delta method to derive the asymptotic standard errors.

Define $u_t(\boldsymbol{\theta}, \boldsymbol{\gamma}) := \exp(\boldsymbol{\gamma}^\top \mathbf{h}(\mathbf{x}_t, \boldsymbol{\theta}))$; the population expectation is $\mu := \mathbf{E}_{\mathcal{P}} [u_t(\boldsymbol{\theta}, \boldsymbol{\gamma})]$ with sample average $\bar{u}_T(\boldsymbol{\theta}, \boldsymbol{\gamma}) := T^{-1} \sum_{t=1}^T u_t(\boldsymbol{\theta}, \boldsymbol{\gamma})$. The estimated KL divergence is $\hat{K} = \ln(\bar{u}_T(\boldsymbol{\theta}, \boldsymbol{\gamma}))$ and the true divergence is $K_0 = \ln(\mu)$. Applying the delta method to $\ln(\bar{u}_T)$ and expanding around μ gives

$$\ln(\bar{u}_T(\boldsymbol{\theta}, \boldsymbol{\gamma})) \approx \ln(\mu) + \frac{1}{\mu}(\bar{u}_T(\boldsymbol{\theta}, \boldsymbol{\gamma}) - \mu) \quad (\text{A.101})$$

this implies that

$$\sqrt{T}(\hat{K} - K_0) = \sqrt{T}(\ln(\bar{u}_T(\boldsymbol{\theta}, \boldsymbol{\gamma})) - \ln(\mu)) \approx \frac{1}{\sqrt{T}} \sum_{t=1}^T \left(\frac{u_t(\boldsymbol{\theta}, \boldsymbol{\gamma})}{\bar{u}_T(\boldsymbol{\theta}, \boldsymbol{\gamma})} - 1 \right). \quad (\text{A.102})$$

Under the null hypothesis that the model is correctly specified, where $\boldsymbol{\gamma}^* = \mathbf{0}$, this simplifies to the sample mean of the moments. Under misspecification, it follows a central

limit theorem for the scalar random variable

$$d_t(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*) := \frac{u_t(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)}{\mathbf{E}_{\mathcal{P}}[u_t(\boldsymbol{\theta}^*, \boldsymbol{\gamma}^*)]} - 1 \quad (\text{A.103})$$

giving

$$\sqrt{T}(\hat{K} - K_0) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_K^2) \quad (\text{A.104})$$

where σ_K^2 is estimated using the HAC estimator on $d_t(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\gamma}})$

$$\hat{\sigma}_K^2 := \frac{1}{T} \sum_{j=1}^T \tilde{d}_t(\boldsymbol{\theta}, \boldsymbol{\gamma}) \tilde{d}_{t-j}(\boldsymbol{\theta}, \boldsymbol{\gamma}) + \frac{2}{T} \sum_{j=1}^L \sum_{t=j+1}^T w(j, L) \tilde{d}_t(\boldsymbol{\theta}, \boldsymbol{\gamma}) \tilde{d}_{t-j}(\boldsymbol{\theta}, \boldsymbol{\gamma}) \quad (\text{A.105})$$

where $\tilde{d}_t(\boldsymbol{\theta}, \boldsymbol{\gamma}) := d_t(\boldsymbol{\theta}, \boldsymbol{\gamma}) - T^{-1} \sum_{t=1}^T d_t(\boldsymbol{\theta}, \boldsymbol{\gamma})$ and $w(j, L)$ are kernel weights.

A.2.4 Evaluating the Significance of a Candidate Factor

We evaluate the significance of including an additional factor $\mathbf{g}$ into the benchmark model. Let the benchmark model $\mathbb{M}^{\mathbb{S}}$ be strictly nested within the augmented model $\mathbb{M}^{(\mathbb{S}, \mathbf{g})}$. Vuong (1989) establishes the limiting distribution of the likelihood ratio statistic $\Lambda_T = T(\hat{K}_1 - \hat{K}_2)$ based on the relationship between the competing models. They define the variance of the difference in the log-likelihoods as

$$\omega^2 := \mathbf{V} \left(\ln \left(\frac{d\mathcal{Q}^{(\mathbb{S}, \mathbf{g})}}{d\mathcal{Q}^{\mathbb{S}}} \right) \right) = \mathbf{V} \left(\ln \left(\frac{m_t(\boldsymbol{\theta}^{(\mathbb{S}, \mathbf{g})})}{m_t(\boldsymbol{\theta}^{\mathbb{S}})} \right) \right) \quad (\text{A.106})$$

If models are non-nested, Vuong shows

$$T^{-1/2} \Lambda_T \xrightarrow{\mathcal{D}} \mathcal{N}(0, \omega^2). \quad (\text{A.107})$$

If the models are nested $\mathbb{M}^{\mathbb{S}} \subset \mathbb{M}^{(\mathbb{S}, \mathbf{g})}$, under the null hypothesis that the restricted model is the pseudo-true DGP $\mathbb{M}^{\mathbb{S}} \equiv \mathbb{M}^{(\mathbb{S}, \mathbf{g})}$, the likelihoods are identical almost surely. Consequently, the variance degenerates $\omega^2 = 0$.

Vuong demonstrates that the normalized t-statistic is undefined in this regime. Instead, the asymptotic distribution of the likelihood ratio is a weighted sum of Chi-square distributions. Calculating the weights for the general misspecified case is computationally burdensome. However, classical asymptotic theory applies to the parameter estimates. We therefore employ the asymptotically equivalent Wald Test on the parameter estimate associated with the additional factor.

In the greedy forward selection algorithm, we orthogonalize the candidate factor $\mathbf{g}$ with respect to the span of the benchmark model factors before estimation. This orthogonality ensures that the Fisher Information matrix and the HAC covariance matrix are asymptotically block-diagonal

$$V^{(\mathbb{S}, \mathbf{g})} = \begin{pmatrix} V^{(\mathbb{S})} & 0 \\ 0 & \sigma_{\mathbf{g}}^2 \end{pmatrix} \quad (\text{A.108})$$

Consequently, the covariance between the estimates of the existing parameters and the new parameter vanishes $\text{Cov}(\hat{\boldsymbol{\theta}}^{\mathbb{S}}, \hat{\boldsymbol{\theta}}^{(\mathbb{S}, \mathbf{g})}) = 0$. This simplifies the Wald statistic from a matrix quadratic form to the square of the t-statistic for the new parameter. We test $H_0 : \theta_{\mathbf{g}} = 0$ using

$$W_T := \frac{\hat{\theta}_{\mathbf{g}}^2}{\hat{\sigma}_{\mathbf{g}}^2} \xrightarrow{\mathcal{D}} \chi_1^2. \quad (\text{A.109})$$

A.3 Numerical Implementation

This section describes the numerical techniques used to keep the MCE estimator and the greedy forward selection well-behaved in the low-signal to noise ratio regime of financial returns.

A.3.1 Bounding the Concentration Ratio

Valid estimation of Ω and especially Ω^{-1} requires the covariance matrix of the moment conditions to be well-conditioned. We rely on Random Matrix Theory (RMT) Tropp (2011)

to establish a bound on the ratio of moments N to observations T . Consider the sample covariance matrix of N i.i.d. variables with variance σ^2 . As $N, T \rightarrow \infty$ with $N/T \rightarrow c \in \mathbb{U}^o$, the Marchenko-Pastur law Marchenko and Pastur (1967) states that the eigenvalues λ are distributed with support

$$\lambda \in \left[\sigma^2 \left(1 - \sqrt{\frac{N}{T}} \right)^2, \sigma^2 \left(1 + \sqrt{\frac{N}{T}} \right)^2 \right]. \quad (\text{A.110})$$

While the assumptions for the eigenvalues to follow this law do not seem applicable, Laloux et al. (1999) find that the bulk of the eigenvalues of a financial correlation matrix fits the Marchenko-Pastur distribution (noise), while only the eigenvalues outside the upper bound represent true signal. To ensure the signal eigenvalues are distinguishable from the noise bulk and to prevent instability in matrix inversion, we impose the constraint

$$\frac{N}{T} < 0.1 \quad (\text{A.111})$$

Under this requirement the condition number of the noise distribution is

$$\kappa \approx \frac{(1 + \sqrt{0.1})^2}{(1 - \sqrt{0.1})^2} = 3.71. \quad (\text{A.112})$$

Restricting $N/T < 0.1$ keeps the condition number of the noise component low ($\kappa < 4$), ensuring that the parameter estimates $\hat{\boldsymbol{\theta}}$ and standard errors remain numerically stable. With monthly data $T = 674$ the maximum number of test assets that can be safely used is $N = 67$, for weekly data $N = 268$, and for daily data $N = 2,022$.

A.3.2 Numerical Implementation of the MCE Estimator

This subsection describes the numerical implementation of the Minimum Cross-Entropy (MCE) estimator in detail. While the information-theoretic approach to SDF estimation is theoretically well-grounded, its practical execution is non-trivial. The low signal-to-noise ratio inherent in financial return data produces an objective surface that is nearly flat over broad regions of the parameter space and riddled with shallow local minima. Standard

gradient-based solvers fail in this environment: gradients provide little directional signal away from the solution, quasi-Newton Hessian approximations deteriorate rapidly under flat curvature, and the exponential structure of the SDF can concentrate probability mass on isolated observations in a manner indistinguishable from optimization progress. The implementation addresses these challenges through a layered architecture combining data preconditioning, a structured nested solver, rigorous numerical health monitors, a probe-and-exact evaluation protocol, and a multi-start initialization strategy with local and global restart mechanisms.

Before estimation, I pre-condition the factor panel and the test-asset return panel using zero-phase component analysis (ZCA) whitening Kessy et al. (2018). For a data matrix $X \in \mathbb{R}^{T \times N}$ with sample covariance $\hat{\Sigma}_X$, the ZCA whitening operator is the symmetric inverse square root

$$\hat{\Sigma}_X^{-1/2} = U_X \Lambda_X^{-1/2} U_X^\top, \quad (\text{A.113})$$

where $\hat{\Sigma}_X = U_X \Lambda_X U_X^\top$ is the eigendecomposition. To guard against near-singular covariance matrices, eigenvalues below $\max(10^{-10}, 10^{-6}\bar{\lambda})$ are clipped before inversion, where $\bar{\lambda}$ is the mean eigenvalue. The whitened series are

$$\tilde{\mathbf{f}}_t := \hat{\Sigma}_F^{-1/2}(\mathbf{f}_t - \bar{\mathbf{f}}), \quad \tilde{\mathbf{r}}_t^e := \hat{\Sigma}_R^{-1/2} \mathbf{r}_t^e. \quad (\text{A.114})$$

The asymmetric treatment of the two panels is deliberate. The factor panel is demeaned before whitening. The cumulant generating function $\psi(\boldsymbol{\theta}) := \ln T^{-1} \sum_t \exp(-\boldsymbol{\theta}^\top \mathbf{f}_t)$ already absorbs the unconditional level of the SDF, so the factor mean contributes a redundant location component that competes with $\psi(\boldsymbol{\theta})$ and inflates the condition number of the Newton system without contributing identifying information. Removing it focuses identification on factor innovations. Test-asset returns are whitened without demeaning. The Euler condition $\mathbf{E}_{\mathcal{P}}[m_t(\boldsymbol{\theta}) r_{i,t}^e] = 0$ targets the level of the risk premium, not a mean-adjusted return; subtracting the return mean would shift the pricing target and distort inference on $\boldsymbol{\theta}$.

The symmetry of the ZCA transformation has a further advantage. Among all lin-

ear whitening transformations that map the data to unit covariance, ZCA minimizes the mean-squared distortion of the original factor directions Kessy et al. (2018). This property preserves the economic interpretability of the whitened coordinates: the whitened factors remain as close as possible to the original factors in the L^2 sense. Crucially, it ensures that the optimization landscape is spherical rather than elliptical. Because the condition number of the Hessian of the inner Newton problem scales directly with the condition number of the factor covariance matrix $\kappa(\hat{\Sigma}_F)$, pre-whitening reduces the effective condition number of the Newton system from $\kappa(\hat{\Sigma}_F)$ to approximately unity, dramatically accelerating convergence of the inner solver and reducing its susceptibility to numerical breakdown.

After estimation, the whitened risk price vector $\tilde{\boldsymbol{\theta}}$ is mapped back to the original factor space via

$$\hat{\boldsymbol{\theta}} = \hat{\Sigma}_F^{-1/2} \tilde{\boldsymbol{\theta}}, \quad (\text{A.115})$$

so that all reported estimates and standard errors are expressed in the units of the original, unwhitened factors.

The MCE estimator minimizes the profiled dual objective $J_T(\boldsymbol{\theta})$ through a two-loop nested structure. The outer loop iterates over the risk price vector $\boldsymbol{\theta} \in \mathbb{R}^K$ using the Limited-Memory BFGS (L-BFGS) algorithm Liu and Nocedal (1989), which approximates the inverse Hessian of J_T using the most recent $m = 12$ gradient-difference pairs. This limited-memory scheme avoids the $O(K^2)$ storage and $O(K^3)$ factorization costs of full BFGS while retaining its superlinear local convergence rate. For each $\boldsymbol{\theta}$, the inner loop solves for the Lagrange multipliers $\boldsymbol{\gamma} \in \mathbb{R}^N$ that satisfy the pricing constraints by minimizing the convex subproblem

$$\phi_T(\boldsymbol{\gamma}; \boldsymbol{\theta}) := \ln \left(\frac{1}{T} \sum_{t=1}^T \exp(m_t(\boldsymbol{\theta}) \mathbf{r}_t^e \boldsymbol{\gamma}) \right), \quad (\text{A.116})$$

where $m_t(\boldsymbol{\theta}) = \exp(-\boldsymbol{\theta}^\top \mathbf{f}_t - \psi(\boldsymbol{\theta}))$ is the normalized SDF. The profiled objective is $J_T(\boldsymbol{\theta}) = \phi_T(\boldsymbol{\gamma}^*(\boldsymbol{\theta}); \boldsymbol{\theta})$, where $\boldsymbol{\gamma}^*(\boldsymbol{\theta}) = \arg \min_{\boldsymbol{\gamma}} \phi_T(\boldsymbol{\gamma}; \boldsymbol{\theta})$.

Both the cumulant $\psi(\boldsymbol{\theta})$ and the inner objective $\phi_T(\boldsymbol{\gamma}; \boldsymbol{\theta})$ require computing the logarithm of a sample average of exponentials. Direct evaluation is numerically unsafe: arguments

that differ by more than the floating-point dynamic range cause overflow or underflow in the exponentials, producing $\pm\infty$ or NaN. I apply the standard log-sum-exp identity

$$\ln \left(\frac{1}{T} \sum_t \exp(z_t) \right) = \max_t \{z_t\} + \ln \left(\frac{1}{T} \sum_t \exp \left(z_t - \max_j \{z_j\} \right) \right), \quad (\text{A.117})$$

shifting all arguments by the running maximum before exponentiation. The shift reduces the largest argument to zero and keeps all exponential inputs non-positive, eliminating overflow. The shifted quantities remain representable in double precision even when the raw arguments span hundreds of natural-log units, as is common when $\boldsymbol{\theta}$ is far from the optimum during early iterations or when factor realizations are extreme.

For fixed $\boldsymbol{\theta}$, the inner subproblem is strictly convex Kitamura and Stutzer (1997), and I solve it using damped Newton's method with adaptive Levenberg-Marquardt regularization Levenberg (1944); Marquardt (1963). The gradient and Hessian of ϕ_T with respect to $\boldsymbol{\gamma}$ are

$$\nabla_{\boldsymbol{\gamma}} \phi_T = \frac{1}{T} \sum_{t=1}^T w_t m_t \mathbf{r}_t^e \equiv R^\top (\mathbf{w} \odot \mathbf{m}), \quad (\text{A.118})$$

$$\nabla_{\boldsymbol{\gamma}}^2 \phi_T = R^\top \text{diag}(\mathbf{w} \odot \mathbf{m}^2) R - (R^\top (\mathbf{w} \odot \mathbf{m})) (R^\top (\mathbf{w} \odot \mathbf{m}))^\top, \quad (\text{A.119})$$

where $w_t \propto \exp(m_t \mathbf{r}_t^{e\top} \boldsymbol{\gamma})$ are the normalized exponential tilting weights, $R \in \mathbb{R}^{T \times N}$ is the matrix of excess returns, and $\odot$ denotes elementwise multiplication. The Hessian is a rank-one downward correction to a weighted outer product sum, and its efficient evaluation and inversion depend on the relative magnitudes of T and N .

When the number of test assets is moderate ($N \leq 512$), I assemble the $N \times N$ Hessian explicitly and factorize it via Cholesky decomposition using a standard dense linear algebra routine. A Levenberg-Marquardt damping term λI_N is added to the Hessian before factorization to ensure positive definiteness. The damping parameter λ is initialized at 10^{-4} and governed by an adaptive schedule Nocedal and Wright (2006): λ increases by a factor of ten when the ratio of actual to predicted objective reduction falls below 0.1 or when three

consecutive backtracking steps occur; it decreases toward the floor 10^{-8} when two successive Newton steps achieve a gradient reduction factor of 0.7 or better. This schedule engages strong regularization precisely when the curvature is ill-determined and releases it as the solver approaches a well-conditioned region.

For Hessians that remain ill-conditioned even after damping (condition number exceeding 10^7), I compute the full eigendecomposition of the pre-damping Hessian and restrict the Newton step to the subspace spanned by eigenmodes with eigenvalues above a floor $\varepsilon_{\text{floor}} = \max(10^{-9}, 10^{-7}\lambda_{\text{max}})$, where λ_{max} is the largest eigenvalue. This truncated Newton step Dembo et al. (1982) discards directions of near-zero curvature, which would otherwise produce Newton steps of unbounded norm and cause the inner solver to overshoot the solution by large amounts. If the retained rank is too small to form a reliable direction, the solver falls back to a gradient step scaled by the reciprocal of the largest eigenvalue.

When N is large, assembling and factorizing an $N \times N$ matrix at each Newton step becomes computationally prohibitive. I instead exploit the Woodbury matrix identity to solve the damped Newton system implicitly using only $T \times T$ computations. Writing the damped Hessian as $H + \lambda I_N$, where $H = R^\top D R - \boldsymbol{\mu} \boldsymbol{\mu}^\top$ with $D = \text{diag}(\mathbf{w} \odot \mathbf{m}^2)$ and $\boldsymbol{\mu} = R^\top(\mathbf{w} \odot \mathbf{m})$, define the scaled return matrix $\tilde{R} = D^{1/2} R \in \mathbb{R}^{T \times N}$ and the $T \times T$ kernel $G = I_T + \lambda^{-1} \tilde{R} \tilde{R}^\top$. The action of the

inverse on a vector $\mathbf{v}$ is then

$$(R^\top D R + \lambda I_N)^{-1} \mathbf{v} = \frac{\mathbf{v}}{\lambda} - \frac{1}{\lambda^2} \tilde{R}^\top G^{-1} \tilde{R} \mathbf{v}, \quad (\text{A.120})$$

where G is solved by a $T \times T$ Cholesky factorization. Because $T \ll N^2$ in all empirically relevant configurations, this reduces the dominant computational cost from $O(N^3)$ to $O(T^2 N + T^3)$ per Newton step. The rank-one correction from the $\boldsymbol{\mu} \boldsymbol{\mu}^\top$ term is handled by a second application of the Sherman-Morrison formula Nocedal and Wright (2006), with a positivity guard ensuring the effective denominator does not change sign.

The inner solver terminates when the gradient infinity norm falls below the tolerance $\|\nabla_{\boldsymbol{\gamma}} \phi_T\|_\infty \leq \varepsilon_\gamma$ (default $\varepsilon_\gamma = 10^{-8}$), or when the Newton decrement

$$\delta := \sqrt{\nabla_{\gamma}^{\top}(H + \lambda I_N)^{-1}\nabla_{\gamma}} \quad (\text{A.121})$$

falls below 10^{-6} . The Newton decrement provides a scale-free stopping criterion that accounts for the local curvature of the objective and correctly measures the remaining reducible gap regardless of the problem scale Boyd and Vandenberghe (2009). A hard budget of 800 inner iterations is imposed as a fallback.

The gradient of the profiled objective with respect to $\boldsymbol{\theta}$ is computed via the envelope theorem. Because $\boldsymbol{\gamma}^*(\boldsymbol{\theta})$ satisfies the first-order condition $\nabla_{\boldsymbol{\gamma}}\phi_T = \mathbf{0}$ at optimality, the total derivative of $J_T(\boldsymbol{\theta}) = \phi_T(\boldsymbol{\gamma}^*(\boldsymbol{\theta}); \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ equals the partial derivative of ϕ_T evaluated at $\boldsymbol{\gamma}^*(\boldsymbol{\theta})$, with no contribution from the implicit dependence of $\boldsymbol{\gamma}^*$ on $\boldsymbol{\theta}$:

$$\nabla_{\boldsymbol{\theta}} J_T(\boldsymbol{\theta}) = - \sum_{t=1}^T w_t^* m_t(\boldsymbol{\theta})(\boldsymbol{\gamma}^*)^{\top} \mathbf{r}_t^e \left(\mathbf{f}_t - \frac{1}{T} \sum_{\tau=1}^T m_{\tau}(\boldsymbol{\theta}) \mathbf{f}_{\tau} \right). \quad (\text{A.122})$$

This expression is efficient to evaluate given the inner solver output $(\boldsymbol{\gamma}^*, \{w_t^*\})$ and requires no additional linear system solves beyond those already performed in the inner loop.

Each L-BFGS outer iteration computes a descent direction $d^{(\ell)}$ and selects a step size α via a strong Wolfe line search Wolfe (1969); Jin et al. (2025). The search enforces two conditions simultaneously. The Armijo sufficient-decrease condition Armijo (1966) requires

$$J_T(\boldsymbol{\theta}^{(\ell)} + \alpha d^{(\ell)}) \leq J_T(\boldsymbol{\theta}^{(\ell)}) + c_1 \alpha \nabla J_T^{(\ell)\top} d^{(\ell)}, \quad c_1 = 10^{-4}, \quad (\text{A.123})$$

and the strong curvature condition requires

$$\left| \nabla J_T(\boldsymbol{\theta}^{(\ell)} + \alpha d^{(\ell)})^{\top} d^{(\ell)} \right| \leq c_2 \left| \nabla J_T^{(\ell)\top} d^{(\ell)} \right|, \quad c_2 = 0.9. \quad (\text{A.124})$$

The search proceeds in two phases. A bracketing phase doubles the candidate step size, starting from the unit step, until the Armijo condition is violated or α reaches a ceiling of 64.0. A zoom phase then applies cubic interpolation within the bracket to locate a step satisfying both conditions, falling back to bisection if interpolation produces an iterate outside the bracket. Up to 30 zoom iterations per outer step are permitted.

The outer loop terminates when the gradient satisfies $\|\nabla_{\theta} J_T\|_{\infty} \leq \varepsilon_{\theta}$ (default $\varepsilon_{\theta} = 5 \times 10^{-4}$). A hard budget of 400 outer iterations is imposed. Because the outer problem is generally non-convex, gradient convergence does not guarantee a global minimum; the multi-start and restart procedures described below mitigate this risk.

The most distinctive feature of this implementation is the embedding of numerical health monitors—hereafter *guards*—directly into the line search protocol. A candidate step α is not committed unless it simultaneously satisfies the following conditions. If any guard triggers, the line search backtracks and attempts a shorter step, forcing the optimizer to find a direction that respects the structure of the estimation problem.

In the exponential-family setting, an observation with an extreme factor realization can attract a tilting weight near unity while all other weights collapse toward zero. This degenerate configuration reduces the effective sample size (ESS) Martino et al. (2017) to one and produces an SDF that prices assets by fitting a single period of data rather than the full distribution. I monitor the ESS

$$\text{ESS} := \frac{1}{\sum_{t=1}^T w_t^2}, \quad (\text{A.125})$$

which ranges from one (all weight on a single observation) to T (uniform weights). A step is rejected if $\text{ESS} < \max(0.15T, 0.5 \cdot \text{ESS}_{\text{prev}})$, where ESS_{prev} is the ESS at the last accepted iterate. The relative floor $0.5 \cdot \text{ESS}_{\text{prev}}$ prevents rapid weight concentration even when the absolute floor has not yet been breached.

As a complement to the ESS, I monitor the ratio of the largest tilting weight to the median weight. A step is rejected if $\max_t(w_t)/\text{median}(w_t) > 50$. This guard detects configurations in which a few observations capture a disproportionate fraction of probability mass without necessarily reducing the ESS below the absolute floor, catching concentration patterns that the ESS test alone may miss.

Low signal-to-noise ratios and multicollinear factors cause the inner Hessian $\nabla_{\gamma}^2 \phi_T$ to become numerically singular even after Levenberg-Marquardt damping. I evaluate the condition number of the inner Hessian at each candidate step and reject any step for which $\kappa(H + \lambda I) > 10^7$. I additionally require that the minimum eigenvalue satisfies $\lambda_{\min} \geq 10^{-8}$.

These conditions ensure that the Newton step returned by the inner solver is numerically well-defined and that the envelope gradient passed to the outer L-BFGS loop is not contaminated by errors that originate in an ill-conditioned Newton system.

The inner solver seeks $\boldsymbol{\gamma}^*$ such that the weighted Euler equation residuals vanish: $R^\top(\mathbf{w} \odot \mathbf{m}) = \mathbf{0}$. If the inner solver exhausts its iteration budget before convergence, the residuals may be non-negligible, producing an incorrect envelope gradient. Any step for which $\|R^\top(\mathbf{w} \odot \mathbf{m})\|_\infty > 5 \times 10^{-3}$ is rejected. The threshold is intentionally loosened from the inner solver convergence tolerance to accommodate the probe evaluations described in the next subsection.

Evaluating $J_T(\boldsymbol{\theta})$ at a candidate step requires running the inner Newton solver to tight convergence, which is expensive when the problem is near-singular. During the line search, I use a two-stage evaluation protocol to reduce this cost without sacrificing solution quality.

In the *probe* stage, each candidate step is evaluated using the inner solver with a loosened tolerance $\varepsilon_\gamma^{\text{probe}} = 10^{-3}$ and a hard budget of eight inner Newton iterations. The probe evaluation is sufficiently accurate to determine whether a step satisfies the Wolfe conditions and passes all guards. Only the step selected by the line search, the single step that survives both the Wolfe filter and the guard system, is then re-evaluated in the *exact* stage using the full inner solver at the target tolerance $\varepsilon_\gamma = 10^{-8}$. This splitting reduces the number of full inner solves per outer iteration from the number of line-search function evaluations (typically three to ten) to one, with a corresponding reduction in computation time.

Before committing any accepted step to the L-BFGS history, I apply a pre-commit filter that rejects microscopic steps offering negligible objective improvement. A step $(\Delta\boldsymbol{\theta}, \Delta J)$ is rejected if $\|\Delta\boldsymbol{\theta}\| \leq 10^{-6}(1 + \|\boldsymbol{\theta}\|)$ and $|\Delta J| \leq 10^{-12}(1 + |J|)$ simultaneously, or if the quasi-Newton curvature condition $\Delta\boldsymbol{\gamma}^\top \Delta\boldsymbol{\theta} \leq 10^{-12}\|\Delta\boldsymbol{\gamma}\|\|\Delta\boldsymbol{\theta}\|$ is violated, which would cause the L-BFGS update to introduce negative curvature information into the Hessian approximation. Rejected steps trigger an immediate trust-radius reduction rather than being appended to the L-BFGS history. This prevents the optimizer from accumulating degenerate curvature pairs that would corrupt subsequent search directions.

The strong Wolfe line search can fail to find a valid step in three ways: the Wolfe condi-

tions may be unsatisfiable along the current direction, the guards may reject all candidate steps, or the pre-commit filter may discard the selected step. A Nocedal and Wright (2006) trust-region strategy is layered on top of the line search to recover from each of these failures.

When the line search fails, the trust radius is contracted by a factor of 0.25, bounding the maximum permitted step norm and forcing the optimizer to take a shorter move. If failures persist across two consecutive outer iterations, a *local restart* is triggered: the L-BFGS curvature history is flushed, and the current iterate is displaced by a random perturbation $\delta\theta = \sigma\mathbf{v}/\|\mathbf{v}\|$, where $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, I_K)$ and $\sigma \approx 0.02\|\tilde{\theta}\|$ (roughly two percent of the current iterate norm in whitened coordinates). Flushing the L-BFGS history forces the next iteration to use a steepest-descent direction, which is always a valid descent direction regardless of how accurately the stored curvature pairs approximate the true Hessian. Up to three local restarts are permitted per optimization run. If local restarts are exhausted and the gradient norm remains above ε_θ , a single *global restart* is attempted from a fresh starting point, chosen as either $\mathbf{0}$ or a uniform random draw from $[-0.1, 0.1]^K$.

Conversely, if the trust radius collapses to its minimum floor while all guards report a healthy region, indicating that the optimizer is making negligible progress due to an artificially tight step bound rather than a genuine curvature deficit, the radius is restored to a larger value to escape the low-curvature plateau.

The non-convexity of the outer objective motivates running independent optimization trials from a diverse set of starting points. All trials are executed in parallel using a thread pool, and the global optimum is taken as the trial that achieves the lowest final dual objective value among all runs that satisfy the guard conditions at termination.

The starting points are constructed as follows. The first deterministic start is $\theta^{(0)} = \mathbf{0}$, corresponding to the risk-neutral SDF and representing the null hypothesis of the model. The second deterministic start is an informed risk premia estimate. The trimmed mean of factor realizations (trimming five percent from each tail) serves as a robust proxy for the risk premium vector, $\hat{\lambda} = \text{TrimMean}(F, 0.05)$. This is mapped to the whitened coordinate system via the safe inverse square root of the factor covariance matrix: $\theta_{\text{rp}}^{(0)} = \hat{\Sigma}_F^{-1/2}\hat{\lambda}$, where the inverse square root is computed from the clipped eigendecomposition described above.

For problems in which $\boldsymbol{\theta}_{\text{rp}}^{(0)}$ has near-zero norm (below 10^{-12}), perturbations along the top- K principal directions of the factor panel are substituted.

Supplementing these two anchors, I add coordinate-axis perturbations $\pm se_k$ for $k = 1, \dots, K$ and scales $s \in \{0.01, 0.05, 0.1\}$, applied around both the zero start and the risk premia start. Finally, I add up to two orthogonal jitter directions chosen to be maximally orthogonal to $\boldsymbol{\theta}_{\text{rp}}^{(0)}$ in the span of the principal factor directions, increasing coverage of parameter regions that the systematic starts may under-explore. If no run converges successfully, the optimizer returns a null result and a non-convergence flag is recorded.

Appendix B

Appendix to Chapter 2

B.1 Model Derivations

This appendix provides the formal derivations for the structural model presented in Section 2.4. We derive the deposit franchise valuation, the asset valuation, the FDIC's scoring rule and the reduction of the multidimensional auction to a scalar first-price format, the optimal asset take-up share, the maximized private value, the comparative statics, and the sale-threshold condition for the probability of a P&A resolution.

B.1.1 Deposit Franchise Value

We derive the economic value $V_{ij}^D(\beta_i, \lambda_i)$ of failed bank i 's deposit franchise to acquiring bank j .

Assumption 3. *The core deposit base of failed bank i follows a deterministic exponential decay process from an initial stock $D_{i0} > 0$ at the auction date $t = 0$, governed by a constant hazard rate $\lambda_i > 0$:*

$$D_{it} = D_{i0}e^{-\lambda_i t}, \quad t \geq 0. \quad (\text{B.1})$$

Assumption 4. *The market short-term interest rate r_t^m follows a stochastic process with constant unconditional mean $\mu_{r^m} := \mathbf{E}[r_t^m]$. Acquirer j earns a gross return on earning assets*

$r_{jt} > 0$ with mean $\mu_{r_j} := \mathbf{E}[r_{jt}]$ and incurs a constant marginal operating cost $c_j > 0$ per dollar of deposits serviced.

Assumption 5. *The deposit rate paid to customers of failed bank i adjusts linearly to the market rate:*

$$r_{it}^d = \alpha_i + \beta_i r_t^m, \quad (\text{B.2})$$

where $\alpha_i \geq 0$ is a bank-specific intercept and $\beta_i \in [0, 1]$ is the deposit rate pass-through.

Assumption 6. *The expected net interest margin is strictly positive:*

$$\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j > 0. \quad (\text{B.3})$$

The instantaneous profit flow generated by the deposit base is

$$\Pi_t := (r_{jt} - r_{it}^d - c_j) D_{it}. \quad (\text{B.4})$$

Substituting Assumptions 3–5, the expected instantaneous profit is

$$\mathbf{E}[\Pi_t] = (\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j) D_{i0} e^{-\lambda_i t}. \quad (\text{B.5})$$

The franchise value is the risk-neutral present value of this expected profit stream, discounted at the acquirer's cost of capital $r_f > 0$:

$$V_{ij}^D(\beta_i, \lambda_i) := \int_0^\infty \mathbf{E}[\Pi_t] e^{-r_f t} dt \quad (\text{B.6})$$

$$= (\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j) D_{i0} \int_0^\infty e^{-(r_f + \lambda_i)t} dt \quad (\text{B.7})$$

$$= \frac{(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j) D_{i0}}{r_f + \lambda_i}, \quad (\text{B.8})$$

where the integral converges because $r_f + \lambda_i > 0$.

B.1.2 Asset Valuation

Assumption 7. *The FDIC's expected recovery value per dollar of book assets for failed bank i is $r_i^{\text{FDIC}} \in (0, 1)$.*

Assumption 8. *Bidder j 's marginal valuation for the x -th fractional unit of bank i 's asset book is linear and declining:*

$$v_{ij}(x) := \bar{v}_{ij} - \gamma_i x, \quad \gamma_i > 0, \quad (\text{B.9})$$

where $\bar{v}_{ij} \in (0, 1]$ is the marginal valuation for the first tranche and γ_i governs the rate of decline.

Let $s_{ij} \in [0, 1]$ denote the fraction of the asset book that bidder j proposes to acquire. The total asset value from acquiring share s_{ij} is obtained by integrating the marginal valuation:

$$V_{ij}^A(s_{ij}) := A_{i0} \int_0^{s_{ij}} v_{ij}(x) dx \quad (\text{B.10})$$

$$= A_{i0} \int_0^{s_{ij}} (\bar{v}_{ij} - \gamma_i x) dx \quad (\text{B.11})$$

$$= A_{i0} \left(\bar{v}_{ij} s_{ij} - \frac{1}{2} \gamma_i s_{ij}^2 \right). \quad (\text{B.12})$$

B.1.3 Scoring Rule and Reduction to a Scalar Auction

Each bidder j submits a sealed bid (p_{ij}, q_{ij}, s_{ij}) for failed bank i , consisting of a deposit premium rate $p_{ij} \geq 0$, an asset purchase price $q_{ij} \in [0, 1]$ per dollar of book, and an asset take-up share $s_{ij} \in [0, 1]$.

Assumption 9. *The failed bank is insolvent in expectation: $D_{i0} > r_i^{\text{FDIC}} A_{i0}$.*

The net cost to the DIF under bid (p, q, s) is

$$C(p, q, s) = D_{i0} - (pD_{i0} + qsA_{i0}) - r_i^{\text{FDIC}}(1 - s)A_{i0}. \quad (\text{B.13})$$

Rearranging:

$$C(p, q, s) = D_{i0} - r_i^{\text{FDIC}} A_{i0} - pD_{i0} - qsA_{i0} + r_i^{\text{FDIC}} sA_{i0} \quad (\text{B.14})$$

$$= \underbrace{(D_{i0} - r_i^{\text{FDIC}} A_{i0})}_{\text{baseline shortfall } (>0)} - \underbrace{[pD_{i0} + (q - r_i^{\text{FDIC}})sA_{i0}]}_{:= \mathcal{S}_i(p, q, s)}. \quad (\text{B.15})$$

Under Assumption 9 the baseline shortfall is positive and common across bids. The FDIC's least-cost mandate therefore reduces to maximizing the score

$$\mathcal{S}_i(p, q, s) := pD_{i0} + (q - r_i^{\text{FDIC}}) sA_{i0}. \quad (\text{B.16})$$

Assumption 10. *There are $J \geq 2$ risk-neutral bidders. Each bidder j draws a private synergy parameter $\theta_{ij} \in (0, 1]$ and an integration cost $\bar{K}_{ij} \in \mathbb{R}_+$, both unobserved by rivals and the FDIC. The deposit-franchise parameters (β_i, λ_i) and asset parameters $(\bar{v}_{ij}, \gamma_i)$ are common knowledge, though $\bar{v}_{ij}$ may vary across bidders.*

The payoff to bidder j upon winning with bid (p, q, s) is

$$U_{ij}(p, q, s) = \theta_{ij} V_{ij}^D(\beta_i, \lambda_i) + V_{ij}^A(s) - \bar{K}_{ij} - pD_{i0} - qsA_{i0}. \quad (\text{B.17})$$

We now express the total payment in terms of the score. From (B.16):

$$pD_{i0} + qsA_{i0} = \mathcal{S}_i + r_i^{\text{FDIC}} sA_{i0}. \quad (\text{B.18})$$

Substituting (B.18) and (B.10) into (B.17):

$$U_{ij} = \theta_{ij} V_{ij}^D + A_{i0} \left(\bar{v}_{ij} s - \frac{1}{2} \gamma_i s^2 \right) - \bar{K}_{ij} - \mathcal{S}_i - r_i^{\text{FDIC}} sA_{i0} \quad (\text{B.19})$$

$$= \underbrace{\theta_{ij} V_{ij}^D - \bar{K}_{ij} + A_{i0} \left[(\bar{v}_{ij} - r_i^{\text{FDIC}}) s - \frac{1}{2} \gamma_i s^2 \right]}_{:= W_{ij}(s)} - \mathcal{S}_i. \quad (\text{B.20})$$

The function $W_{ij}(s)$ is the bidder's *private value*: the surplus earned from winning at

a score of zero. Equation (B.19) has the canonical structure of a first-price auction with private value W_{ij} and bid $\mathcal{S}_i$. The standard equilibrium analysis of Milgrom and Weber (1982) applies directly.

B.1.4 Optimal Asset Take-Up

Each bidder selects s to maximize $W_{ij}(s)$, which determines the “budget” for score bidding. Taking the first-order condition of $W_{ij}(s)$ with respect to s :

$$\frac{\partial W_{ij}}{\partial s} = A_{i0} (\bar{v}_{ij} - r_i^{\text{FDIC}} - \gamma_i s) = 0 \quad \implies \quad s_{ij}^* = \frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i}. \quad (\text{B.21})$$

The second-order condition confirms a global maximum:

$$\frac{\partial^2 W_{ij}}{\partial s^2} = -A_{i0} \gamma_i < 0. \quad (\text{B.22})$$

Imposing the box constraint $s \in [0, 1]$:

$$s_{ij}^* = \min \left\{ \frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i}, 1 \right\}, \quad (\text{B.23})$$

with $s_{ij}^* \geq 0$ ensured whenever $\bar{v}_{ij} \geq r_i^{\text{FDIC}}$. The model-implied retained asset share under the winning bidder j^* is

$$\tau_i := 1 - s_{ij^*}^*. \quad (\text{B.24})$$

B.1.5 Maximized Private Value

Substituting the interior solution $s_{ij}^* = (\bar{v}_{ij} - r_i^{\text{FDIC}})/\gamma_i$ into $W_{ij}(s)$:

$$W_{ij}(s_{ij}^*) = \theta_{ij}V_{ij}^D - \bar{K}_{ij} + A_{i0} \left[(\bar{v}_{ij} - r_i^{\text{FDIC}}) \frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i} - \frac{1}{2} \gamma_i \left(\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i} \right)^2 \right] \quad (\text{B.25})$$

$$= \theta_{ij}V_{ij}^D - \bar{K}_{ij} + A_{i0} \left[\frac{(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{\gamma_i} - \frac{(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i} \right] \quad (\text{B.26})$$

$$= \theta_{ij}V_{ij}^D(\beta_i, \lambda_i) - \bar{K}_{ij} + \frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i}. \quad (\text{B.27})$$

The three terms are, respectively, the net deposit-side surplus, the integration cost, and the asset-side option value. The deposit and asset channels are additively separable: s_{ij}^* depends only on asset-side parameters $(\bar{v}_{ij}, \gamma_i, r_i^{\text{FDIC}})$, and V_{ij}^D depends only on deposit-side parameters $(\beta_i, \lambda_i, \mu_{r_j}, \alpha_i, c_j)$.

B.1.6 Symmetric Equilibrium Bidding

Assumption 11. *The private values $W_{ij}(s_{ij}^*)$ are i.i.d. draws from a distribution G with support $[\underline{W}, \bar{W}]$, density $g > 0$, and CDF G . Each bidder submits a score bid $\mathcal{S}_i$ to maximize expected surplus.*

Define the normalized per-dollar private value $w_j := W_j/D_{i0}$, with CDF G and density g inherited from the distribution of W_j . The maximization problem for bidder j with value w is

$$\max_{\mathcal{S}} (w - \mathcal{S}) \mathbf{P}\{\text{Win} \mid \mathcal{S}\}. \quad (\text{B.28})$$

In a symmetric equilibrium with a strictly increasing bidding function $b(\cdot)$, the probability of winning is the probability that all $J - 1$ rivals have values below $b^{-1}(\mathcal{S})$:

$$\mathbf{P}\{\text{Win} \mid \mathcal{S}\} = G(b^{-1}(\mathcal{S}))^{J-1}. \quad (\text{B.29})$$

Substituting $\mathcal{S} = b(w)$ and taking the first-order condition:

$$b'(w) = (J-1) \frac{g(w)}{G(w)} (w - b(w)). \quad (\text{B.30})$$

This is a first-order linear ODE. With the boundary condition $b(\underline{w}) = \underline{w}$ (the lowest-type bidder earns zero surplus), the unique solution is

$$b(w) = w - \frac{\int_{\underline{w}}^w G(y)^{J-1} dy}{G(w)^{J-1}}. \quad (\text{B.31})$$

Proof. Define $h(w) := b(w)G(w)^{J-1}$. Differentiating:

$$h'(w) = b'(w)G(w)^{J-1} + b(w)(J-1)G(w)^{J-2}g(w). \quad (\text{B.32})$$

Substituting the first-order condition $b'(w)G(w)^{J-1} = (J-1)g(w)G(w)^{J-2}(w - b(w))G(w)$:

$$h'(w) = (J-1)g(w)G(w)^{J-2}(w - b(w)) + (J-1)g(w)G(w)^{J-2}b(w) \quad (\text{B.33})$$

$$= (J-1)g(w)G(w)^{J-2}w \quad (\text{B.34})$$

$$= w \frac{d}{dw} [G(w)^{J-1}]. \quad (\text{B.35})$$

Integrating from $\underline{w}$ to w with $h(\underline{w}) = \underline{w}G(\underline{w})^{J-1}$:

$$h(w) - h(\underline{w}) = [yG(y)^{J-1}]_{\underline{w}}^w - \int_{\underline{w}}^w G(y)^{J-1} dy, \quad (\text{B.36})$$

where we integrate by parts. Simplifying:

$$b(w)G(w)^{J-1} = wG(w)^{J-1} - \int_{\underline{w}}^w G(y)^{J-1} dy, \quad (\text{B.37})$$

which yields (B.31). $\square$

The winning score observed in the data corresponds to the bid of the value leader $w^{(1)} :=$

$\max_j \{w_j\}$:

$$\mathcal{S}^* = b(w^{(1)}). \quad (\text{B.38})$$

The bid-shading term $\int_{\underline{w}}^w G(y)^{J-1} dy / G(w)^{J-1}$ is strictly positive for $w > \underline{w}$, ensuring that the winning bidder earns a positive surplus.

B.1.7 Deposit-Side Comparative Statics

We derive the partial derivatives of the franchise value V_{ij}^D in (B.6) with respect to the deposit-side primitives β_i and λ_i .

Differentiating (B.6) with respect to β_i :

$$\frac{\partial V_{ij}^D}{\partial \beta_i}(\beta_i, \lambda_i) = \frac{-\mu_{r^m} D_{i0}}{r_f + \lambda_i} < 0. \quad (\text{B.39})$$

The inequality follows from $\mu_{r^m} > 0$, $D_{i0} > 0$, and $r_f + \lambda_i > 0$.

Differentiating (B.6) with respect to λ_i :

$$\frac{\partial V_{ij}^D}{\partial \lambda_i}(\beta_i, \lambda_i) = -\frac{(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j) D_{i0}}{(r_f + \lambda_i)^2} < 0. \quad (\text{B.40})$$

The inequality follows from Assumption 6.

Differentiating (B.39) with respect to λ_i , or equivalently (B.40) with respect to β_i :

$$\frac{\partial^2 V_{ij}^D}{\partial \beta_i \partial \lambda_i}(\beta_i, \lambda_i) = \frac{\mu_{r^m} D_{i0}}{(r_f + \lambda_i)^2} > 0. \quad (\text{B.41})$$

Rate sensitivity and attrition are substitutes: the marginal cost of a higher β_i is decreasing in λ_i .

B.1.8 Asset-Side Comparative Statics

We derive the responses of the optimal take-up share s_{ij}^* in (B.23) and the maximized private value $W_{ij}(s_{ij}^*)$ in (B.27) to the asset-side primitives, restricting to interior solutions $s_{ij}^* < 1$.

Differentiating (B.23):

$$\frac{\partial s_{ij}^*}{\partial \bar{v}_{ij}} = \frac{1}{\gamma_i} > 0. \quad (\text{B.42})$$

$$\frac{\partial s_{ij}^*}{\partial \gamma_i} = -\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i^2} < 0. \quad (\text{B.43})$$

$$\frac{\partial s_{ij}^*}{\partial r_i^{\text{FDIC}}} = -\frac{1}{\gamma_i} < 0. \quad (\text{B.44})$$

Differentiating (B.27) with respect to $\bar{v}_{ij}$:

$$\frac{\partial W_{ij}}{\partial \bar{v}_{ij}}(s_{ij}^*) = \frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})}{\gamma_i} = A_{i0}s_{ij}^* > 0, \quad (\text{B.45})$$

where the second equality uses (B.23). An envelope-theorem argument confirms the same result: since s_{ij}^* maximizes $W_{ij}(s)$, the indirect effect through $\partial s^*/\partial \bar{v}_{ij}$ vanishes at the optimum, and only the direct effect $A_{i0}s_{ij}^*$ remains.

Differentiating (B.27) with respect to γ_i :

$$\frac{\partial W_{ij}}{\partial \gamma_i}(s_{ij}^*) = -\frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i^2} < 0. \quad (\text{B.46})$$

Differentiating (B.27) with respect to r_i^{FDIC} :

$$\frac{\partial W_{ij}}{\partial r_i^{\text{FDIC}}}(s_{ij}^*) = -\frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})}{\gamma_i} = -A_{i0}s_{ij}^* < 0. \quad (\text{B.47})$$

B.1.9 Premium–Retention Relationship

We establish the positive association between the retained share τ_i and the deposit premium p through two channels.

From (B.16), the deposit premium consistent with a given score $\mathcal{S}_i$ and asset terms (q, s) is

$$p = \frac{\mathcal{S}_i - (q - r_i^{\text{FDIC}})sA_{i0}}{D_{i0}}. \quad (\text{B.48})$$

Fix $\mathcal{S}_i$ and q . Then $\partial p / \partial s = -(q - r_i^{\text{FDIC}})A_{i0}/D_{i0}$. At the optimal take-up, $q = v_{ij}(s_{ij}^*) = r_i^{\text{FDIC}}$ at the margin, and for infra-marginal units $q > r_i^{\text{FDIC}}$, so $(q - r_i^{\text{FDIC}}) \geq 0$. A lower s (higher τ_i) therefore weakly increases p for a given score.

From (B.27), the maximized private value is increasing in the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$. When marginal assets satisfy $v_{ij}(x) < r_i^{\text{FDIC}}$, retaining them (increasing τ_i) removes a drag on W_{ij} , raising the equilibrium score bid through the monotonicity of $b(\cdot)$ in (B.31). The higher score translates into a higher premium via (B.48). Both channels sign the τ_i - p relationship as positive.

B.1.10 Probability of P&A

Let C_i^{payoff} denote the DIF cost under a deposit payoff. A P&A sale occurs if and only if

$$\max_j \mathcal{S}_i(p_{ij}, q_{ij}, s_{ij}) \geq \underbrace{(D_{i0} - r_i^{\text{FDIC}} A_{i0}) - C_i^{\text{payoff}}}_{:= \underline{\mathcal{S}}_i}. \quad (\text{B.49})$$

In the symmetric equilibrium, the score bid $b(w)$ is strictly increasing. The event $\{\text{sale}\}$ is therefore equivalent to $\{w^{(1)} \geq b^{-1}(\underline{\mathcal{S}}_i)\}$. Let $\underline{w}_i := b^{-1}(\underline{\mathcal{S}}_i)$ denote the threshold private value. The probability of a P&A is

$$\mathbf{P}\{\text{P\&A}\} = 1 - G(\underline{w}_i)^J. \quad (\text{B.50})$$

We verify three comparative statics:

- (x) *Franchise value.* An increase in V_{ij}^D shifts the distribution of W_{ij} to the right (by (B.27)), increasing $G(\underline{w}_i)^J$'s complement.
- (x) *Asset-side comparative advantage.* An increase in $\bar{v}_{ij} - r_i^{\text{FDIC}}$ raises the asset-side option value in (B.27), shifting W_{ij} to the right.
- (x) *Number of bidders.* Differentiating (B.50) with respect to J :

$$\frac{\partial}{\partial J} [1 - G(\underline{w}_i)^J] = -G(\underline{w}_i)^J \ln G(\underline{w}_i) > 0, \quad (\text{B.51})$$

since $G(\underline{w}_i) \in (0, 1)$ implies $\ln G(\underline{w}_i) < 0$.

All three channels increase the probability that the maximum score bid exceeds the least-cost threshold $\underline{\mathcal{S}}_i$, completing the derivation of the model's testable predictions.

B.2 Robustness

B.2.1 IV First-Stage Regression

Table A1: LLP IV First-Stage	
	Toxic Ratio
LLP Ratio	0.89711*** (0.23143)
Log Asset Size	0.00691 (0.00599)
Year FE	X
N	1785
R^2	0.25
First-stage F	15.0

Table A1: OLS regression of the realized toxic ratio on the pre-failure loan loss provision ratio and controls. The instrument is the loan loss provision ratio (LLP / assets) in the quarter before failure. Year fixed effects included. F -statistic for instrument relevance reported. Robust standard errors in parentheses. Significance levels: $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

B.2.2 Heckman Selection-Corrected Pricing

This appendix provides the complete econometric derivation, first-stage probit estimates, bias decomposition, and robustness analysis for the Heckman two-step correction summarized in Section 2.5.

Table A2 reports the first-stage probit estimates. The dependent variable is the P&A indicator. Column (1) presents the primary specification, which includes deposit runoff, the LLP Ratio, log asset size, and the quarterly concurrent-failure exclusion variable with year fixed effects. Column (2) replaces the congestion measure with log days elapsed since

Table A2: Heckman First Stage				
	$\mathbf{1}\{\text{Resolution}_i = \text{P\&A}\}$			
	(1)	(2)	(3)	(4)
Deposit Runoff	-0.46428 (0.40928)	-0.44134 (0.41929)	-0.43073 (0.42101)	0.26531 (1.75178)
LLP Ratio	1.02656 (1.17418)	0.98665 (1.17059)	1.01374 (1.17450)	1.90658 (7.14635)
Log Asset Size	0.03607 (0.03437)	0.03733 (0.03427)	0.03573 (0.03440)	-0.00374 (0.10845)
concurrent_failures	0.29254* (0.15959)		0.29006* (0.15980)	0.33925 (0.41185)
cr_gap_log		-0.02524 (0.05922)	-0.02030 (0.05927)	
Year FE	X	X	X	X
N	1784	1784	1784	372
Pseudo R^2	0.098	0.096	0.098	0.110
Log-likelihood	-723.0	-724.6	-722.9	-50.1

Table A2: Probit regressions of an indicator for P&A resolution on deposit and asset controls plus exclusion restrictions. Column (1) uses log concurrent failures in the resolution quarter; column (2) uses log days from the nearest crisis trough; column (3) includes both exclusion variables; column (4) restricts to the deposit-beta matched subsample and adds the deposit expense beta. Probit marginal index coefficients. Robust standard errors in parentheses. Significance * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

the nearest macroeconomic crisis trough. Column (3) includes both exclusion variables simultaneously to form an over-identified specification. Column (4) restricts to the deposit-beta matched subsample and adds deposit beta as an additional selection covariate.

The probit estimates are consistent with the linear probability model of Section 2.5. Asset size enters with a large positive coefficient in all specifications: larger institutions are substantially more likely to attract a P&A acquirer. Deposit runoff carries a negative point estimate that is imprecisely estimated once year effects absorb the macroeconomic flight component. The LLP Ratio enters positively across all columns, consistent with the asset-cleansing narrative: banks for which a larger loss-sharing pool has been negotiated are precisely those for which sufficient acquirer interest exists to complete a full P&A.

The coefficient on log concurrent failures in Column (1) is positive and marginally significant. This sign is the opposite of the ex ante capacity-constraint prediction, but it is consistent with the FDIC's pro-cyclical resolution infrastructure: when failures cluster in a single quarter, fixed costs of assembling bidder lists and coordinating with chartering authorities are amortized across many institutions, and the agency has stronger incentives to clear the resolution queue quickly through P&A rather than payoff. This pattern is consistent with the historical narrative of the S&L crisis resolution documented in Curry and Shibut (2000). For the identification argument, the sign of the exclusion coefficient is immaterial; what matters is that the variable enters the selection equation with a non-zero coefficient and is plausibly absent from the conditional pricing equation.

In Column (4), deposit beta enters with a negative coefficient, reinforcing the interpretation that banks with greater deposit rate pass-through are less marketable in a resolution auction.

B.2.3 Bias Decomposition

Table A3 consolidates the selection-bias diagnostics across the four Heckman specifications.

Table A3: Heckman Decomposition				
	(1) Concurrent	(2) Crisis Gap	(3) Over-Id	(4) β
$\hat{\rho}$	0.419	-0.207	0.286	1.794
$\hat{\sigma}_\varepsilon$	0.015	0.015	0.015	0.017
$\hat{\beta}_\lambda$	0.006	-0.003	0.004	0.030
LR test (over-id)			0.12	
N (first stage)	1784	1784	1784	372
N (second stage)	1488	1488	1488	359

Table A3: The column entries report the implied correlation $\hat{\rho}$ between the pricing and selection errors, the second-stage residual standard deviation $\hat{\sigma}_\varepsilon$, the fitted inverse Mills ratio coefficient $\hat{\beta}_\lambda = \hat{\rho}\hat{\sigma}_\varepsilon$, and the first-stage log-likelihood ratio test statistic for exclusion-restriction validity. $\hat{\rho}$ is estimated as $\hat{\beta}_\lambda/\hat{\sigma}_\varepsilon$ where $\hat{\sigma}_\varepsilon^2 = \hat{u}'\hat{u}/n + \hat{\beta}_\lambda^2\bar{\delta}$ with $\bar{\delta} = n^{-1} \sum_i \hat{\lambda}_i(\hat{\lambda}_i + \hat{x}_i\hat{\gamma})$.

Three substantive points follow from the decomposition. First, $\hat{\sigma}_\varepsilon$ is stable across specifications at approximately 0.015, matching the sample standard deviation of the log deposit

premium reported in Section 2.5. Second, the implied correlation $\hat{\rho}$ is moderate in the primary specification at 0.42 but is sensitive to the choice of exclusion restriction and even reverses sign when the alternative exclusion is used. The corresponding bias in the pricing coefficients is bounded by $|\hat{\rho}|\hat{\sigma}_\varepsilon\bar{\lambda} \approx 0.42 \cdot 0.015 \cdot 1.3 \approx 0.008$, an order of magnitude comparable to the standard errors on the pricing coefficients themselves. Third, the likelihood-ratio statistic for the joint significance of the two exclusion variables in the over-identified probit is 0.12 on two degrees of freedom, well below any conventional critical value. The two exclusion variables are effectively one-dimensional in their first-stage contribution, capturing overlapping aspects of resolution-market conditions. The mild over-identification failure does not invalidate the identification strategy because the pricing coefficients are robust across the single-exclusion specifications.

Column (3) of Table A2 includes both log concurrent failures and log days from crisis trough simultaneously. The two exclusion variables capture different sources of variation: the former reflects within-year seasonal congestion in resolution volume; the latter reflects the position of any given failure relative to the aggregate macroeconomic cycle. Despite this difference, the pseudo- R^2 of the over-identified probit rises only modestly relative to the single-exclusion specifications, and the pricing coefficients in Column (4) of Table 2.8 are effectively identical to those in the primary specification. The likelihood-ratio over-identification test statistic of 0.12 (2 df) implies that neither exclusion variable dominates the other in first-stage identification. The practical implication is that identification of the Heckman correction rests on a single effective instrument. That instrument is marginally significant (first-stage $t \approx 1.83$), which is part of the reason for the wide bootstrap confidence interval on $\hat{\beta}_\lambda$ and motivates emphasizing coefficient stability across specifications rather than point-estimate precision.-

Column (5) of Table 2.8 reports the Heckman second stage on the deposit-beta matched subsample. The implied $\hat{\rho}$ in this column lies outside the unit interval, a telltale symptom of finite-sample Heckman pathology documented by Puhani (2000). The pathology arises because the deposit-beta subsample is small, selected on survival (a bank must operate long enough to estimate a rate-sensitivity beta), and exhibits limited variation in the inverse

Mills ratio, reducing the effective identification of β_λ . This column should be interpreted as a diagnostic robustness check rather than a refined estimate of the selection correction.

The Heckman correction and the 2SLS strategy of Section 2.5 address different identification threats: the former targets selection into the P&A sample, the latter targets the simultaneity of the premium and the retained asset share. A fully identified specification would combine both corrections in a selection-corrected instrumental-variable design. I have estimated this combined specification and find pricing coefficients that are economically identical to those reported in Tables 2.5 and 2.8, with larger standard errors reflecting the compounding of generated-regressor uncertainty across the two stages. The separate corrections are reported throughout to preserve the diagnostic interpretation of each. The dominant identification threat in the pricing channel remains the simultaneity of the toxic ratio and the premium, not the selection into the P&A sample.

Several robustness margins merit brief discussion. Redefining log concurrent failures at the monthly or semi-annual frequency, or replacing the log transformation with the level count, does not materially affect the first-stage significance or the pricing coefficients. Restricting the sample to the post-FIRREA period ($N \approx 1,000$) weakens first-stage identification further, as the most congested quarters of the sample fall in the thrift crisis; the pricing coefficients remain within the bootstrap confidence intervals reported in Table 2.8.

The joint-normality assumption of the parametric Heckman model is routinely questioned on the grounds that non-normality produces an inconsistent control function. I have re-estimated the pricing equation using a semi-parametric single-index selection correction following Newey (2009) on the primary exclusion restriction. The resulting pricing coefficients are within one bootstrap standard error of the parametric two-step estimates, indicating that the conclusions of Section 2.5 are not driven by the normality assumption.

Restricting the selected sample to pure P&A transactions (excluding Insured Deposit Transfers) or expanding it to include all transactions involving any form of franchise transfer does not alter the finding that the selection correction leaves the pricing coefficients materially unchanged.

B.3 Data Sources

Table A4: FDIC Press Release Data Description

Variable	Description	Source
cert	FDIC certificate number	BankFind Suite
chclass1	FDIC bank charter class (N, NM, SM, SB)	BankFind Suite
cityst	City, state of the failed bank	BankFind Suite
cost	Cost of the resolution to the FDIC	BankFind Suite
faildate	Date the bank failed	BankFind Suite
fin	Financial institution number	BankFind Suite
id	FDIC specific ID number	BankFind Suite
name	Name of the failed bank	BankFind Suite
qbfasset	Total assets the quarter before failure	BankFind Suite
qbfddep	Total deposits the quarter before failure	BankFind Suite
restype1	The resolution type used by the FDIC (PA, IDT, PO, PI, DINB)	BankFind Suite
savr	Insurance fund	BankFind Suite
aq_name	Name of the acquiring bank	FDIC Press Releases
aq_city	City of the acquiring bank	FDIC Press Releases
aq_st	State of the acquiring bank	FDIC Press Releases
recovery	Trinary indicator for whether the deposit recovery to uninsured depositors was (F)ull, (P)artial, or (N)one	FDIC Press Releases
recovery_pct	The percentage recovery of uninsured depositors	FDIC Press Releases
fasset	Total assets of the failed bank at the time of failure	FDIC Press Releases
fdep	Total deposits of the failed bank at the time of failure	FDIC Press Releases
tot_accts	Total number of deposit accounts at the failed bank	FDIC Press Releases
amount_over	Total dollar amount of deposits over the deposit insurance limit at the failed bank	FDIC Press Releases
accts_over	Total number of accounts over the deposit insurance limit	FDIC Press Releases
prem_dep	Premium paid to acquire the deposits of the failed bank's deposits	FDIC Press Releases
prem_dep_pct	Premium paid to acquire the deposits of the failed bank's deposits as a percentage of total deposits	FDIC Press Releases
purchased_assets	Assets of the failed bank purchased by the acquiring bank	FDIC Press Releases
asset_prem	Premium paid to acquire the deposits of the failed bank's assets	FDIC Press Releases
asset_prem_pct	Premium paid to acquire the deposits of the failed bank's assets as a percentage of total assets	FDIC Press Releases
fdic_advance	FDIC loan given to the acquiring institution to facilitate the transaction	FDIC Press Releases
retained_assets	Assets of the failed bank retained by the FDIC	FDIC Press Releases
loan_loss	Total loan loss provision to acquire the fail bank	FDIC Press Releases
loan_loss_fdic_share	The FDIC's share of the total loan loss provision	FDIC Press Releases

Table A4: This table provides the variable names, description, and source for the novel set of bank auction characteristics collected and constructed as part of this project. The BankFind

B.4 FDIC Failed Bank Press Release



FOR IMMEDIATE RELEASE
PR-33-93 (4-8-93)

Media Contact:
Andrew Porterfield (202) 898-6593

**FDIC APPROVES ASSUMPTION OF INSURED DEPOSITS OF
PREMIER BANK, NORTHRIDGE, CALIFORNIA**

The Board of Directors of the Federal Deposit Insurance Corporation has approved the assumption of the insured deposits of Premier Bank, Northridge, California, by Trans-World Bank, Sherman Oaks, California.

The failed bank's three offices will reopen on Friday, April 9, 1993, as branches of Trans-World, and its depositors automatically will become depositors of the assuming bank.

Premier Bank, with total assets of \$74.3 million, was closed on Thursday, April 8, 1993, by James E. Gilleran, California Banking Superintendent, and the FDIC was named receiver.

Trans-World will assume about \$62.8 million in about 4,600 deposit accounts. At the time the bank closed, about \$6.9 million in 71 accounts exceeded the federal insurance limit of \$100,000 and will not be assumed by Trans-World Bank.

The Board of Directors also voted to make a prompt advance payment to uninsured depositors equal to 57 percent of the uninsured claims. Arrangements for picking up the advance payment checks can be made at any of the failed bank's offices beginning Friday, April 9, 1993.

The assuming bank will pay a premium of \$701,000 for the right to receive the failed bank's deposits and will purchase \$12.3 million of the failed bank's assets. To facilitate the transaction, the FDIC will advance about \$49.8 million to the assuming bank and will retain assets of the failed

(more)

-2-

bank with a book value of about \$62.0 million.

The Board of Directors approved the deposit assumption under its authority to do so whenever it determines that such a transaction will reduce the potential loss to the FDIC. The FDIC notes that its claim from the sale of the failed bank's assets will have priority over non-depositor creditors of the failed bank.

##

Appendix C

Appendix to Chapter 3

C.1 Model Derivations

C.1.1 First Passage Time ODE

We derive the differential equation governing the value of the claim $P(V_t)$ that pays \$1 at the moment of default τ and 0 otherwise. Default is defined as the first passage time $\tau := \inf\{t \geq 0 : V_t \leq V_B\}$.

Assumption 12. *Asset values follow a Geometric Brownian Motion under the risk-neutral measure $\mathcal{Q}$: $dV_t/V_t = (r_f - \delta)dt + \sigma dZ_t$, with constant parameters risk-free rate r_f , payout rate δ , and volatility σ .*

Let $P(V_t)$ be the price of the security at time t . Since the security pays no dividends before default, its discounted price process must be a martingale under the risk-neutral measure $\mathcal{Q}$. We define the discounted process

$$Y_t = e^{-r_f t} P(V_t). \quad (\text{C.1})$$

For Y_t to be a martingale, its expected instantaneous drift must be zero $\mathbf{E}_{\mathbf{Q}}[dY_t] = 0$.

We apply Ito's Lemma for the product of two processes to find the differential

$$dY_t = d(e^{-r_f t})P(V_t) + e^{-r_f t}dP(V_t) = -r_f e^{-r_f t}dt + e^{-r_f t}dP(V_t). \quad (\text{C.2})$$

Next, we apply Ito's Lemma to $P(V_t)$. Since P is a function of V_t , we expand it to the second order:

$$dP(V_t) = P'(V_t)dV_t + \frac{1}{2}P''(V_t)(dV_t)^2. \quad (\text{C.3})$$

Substitute the dynamics of V_t

$$dP(V_t) = P'(V_t)[(r_f - \delta)V_t dt + \sigma V_t dZ_t] + \frac{1}{2}P''(V_t)\sigma^2 V_t^2 dt. \quad (\text{C.4})$$

Grouping the dt and dZ_t terms

$$dP(V_t) = \left[(r_f - \delta)V_t P'(V_t) + \frac{1}{2}\sigma^2 V_t^2 P''(V_t) \right] dt + \sigma V_t P'(V_t) dZ_t. \quad (\text{C.5})$$

Now, substitute $dP(V_t)$ back into the expression for the discounted process

$$dY_t = -r_f e^{-r_f t} P(V_t) dt + e^{-r_f t} \left[\left((r_f - \delta)V_t P'(V_t) + \frac{1}{2}\sigma^2 V_t^2 P''(V_t) \right) dt + \sigma V_t P'(V_t) dZ_t \right]. \quad (\text{C.6})$$

Factor out $e^{-r_f t}$ and collect all dt terms

$$dY_t = e^{-r_f t} \underbrace{\left(\frac{1}{2}\sigma^2 V_t^2 P''(V_t) + (r_f - \delta)V_t P'(V_t) - r_f P(V_t) \right)}_{\text{Drift}} dt + e^{-r_f t} \sigma V_t P'(V_t) dZ_t. \quad (\text{C.7})$$

By the Martingale Representation Theorem, the drift term must equal zero to prevent arbitrage.

$$\frac{1}{2}\sigma^2 V^2 P''(V) + (r_f - \delta)V P'(V) - r_f P(V) = 0. \quad (\text{C.8})$$

C.1.2 First Passage Time Security

We derive the value of a state-contingent claim $P(V_t, V_B)$ that pays \$1 at the moment the Local Bank defaults. Default is defined as the first passage time $\tau := \inf\{t \geq 0 : V_t \leq V_B\}$.

Assumption 13. *Asset values follow a Geometric Brownian Motion under the risk-neutral measure $\mathcal{Q}$: $dV_t/V_t = (r_f - \delta)dt + \sigma dZ_t$, with constant parameters risk-free rate r_f , payout rate δ , and volatility σ*

The claim pays \$1 if V_t hits V_B and 0 otherwise. Let $P(V)$ denote the expected present value of the claim

$$P(V) = \mathbf{E}_{\mathcal{Q}}[e^{-r_f \tau} | V_0 = V]. \quad (\text{C.9})$$

In the continuation region $V > V_B$, the discounted value of the claim is a martingale. In Section C.1.1 the governing ODE for this system was shown to be

$$\frac{1}{2}\sigma^2 V^2 P''(V) + (r_f - \delta)V P'(V) - r P(V) = 0. \quad (\text{C.10})$$

We guess a solution of the form $P(V) = AV^\beta$. Substituting this into the ODE yields the characteristic equation

$$\frac{1}{2}\sigma^2 \beta(\beta - 1) + (r_f - \delta)\beta - r_f = 0. \quad (\text{C.11})$$

Rearranging terms gives the fundamental quadratic equation

$$\frac{1}{2}\sigma^2 \beta^2 + \left(r_f - \delta - \frac{1}{2}\sigma^2\right)\beta - r = 0. \quad (\text{C.12})$$

This quadratic has two roots. Since the value of the default claim must vanish as the bank becomes infinitely solvent, we select the negative root

$$-\gamma = \frac{-(r_f - \delta - \frac{1}{2}\sigma^2) - \sqrt{(r_f - \delta - \frac{1}{2}\sigma^2)^2 + 2r_f\sigma^2}}{\sigma^2}. \quad (\text{C.13})$$

Applying the boundary condition at default, $P(V_B) = 1$, we solve for the constant A

$$P(V_B) = AV_B^{-\gamma} = 1 \implies A = V_B^\gamma. \quad (\text{C.14})$$

Thus, the value of the first-passage security is

$$P(V_t) = \left(\frac{V_t}{V_B} \right)^{-\gamma}. \quad (\text{C.15})$$

C.1.3 Equity Value ODE

We derive the differential equation governing the value of bank equity $E(V_t)$, which pays a continuous dividend stream equal to asset income minus debt service $(\delta V_t - C)$ until default. Unlike the default security, equity pays a continuous flow of dividends. For the equity price to be arbitrage-free, the discounted value of the stock price plus the discounted value of cumulative dividends received must be a martingale under $\mathcal{Q}$. We define the gains process G_t

$$G_t = \underbrace{e^{-r_f t} E(V_t)}_{\text{Discounted Price}} + \underbrace{\int_0^t e^{-r_f s} (\delta V_s - C) ds}_{\text{Discounted Dividends}}. \quad (\text{C.16})$$

The condition for no arbitrage is $\mathbf{E}_{\mathcal{Q}}[dG_t] = 0$.

Differentiate G_t with respect to time t and apply Ito's Lemma. The derivative of the integral term is simply the integrand evaluated at t

$$dG_t = d(e^{-r_f t} E(V_t)) + e^{-r_f t} (\delta V_t - C) dt \quad (\text{C.17})$$

$$= -r_f e^{-r_f t} E(V_t) dt + e^{-r_f t} \left(E'(V_t) dV_t + \frac{1}{2} E''(V_t) (dV_t)^2 \right) + e^{-r_f t} (\delta V_t - C) dt \quad (\text{C.18})$$

Plugging in asset dynamics

$$dE(V_t) = \left((r_f - \delta) V_t E'(V_t) + \frac{1}{2} \sigma^2 V_t^2 E''(V_t) \right) dt + \sigma V_t E'(V_t) dZ_t. \quad (\text{C.19})$$

Substitute $dE(V_t)$ back into the expression for dG_t

$$dG_t = e^{-r_f t} \left(-r_f E(V_t) dt + \left[(r_f - \delta) V_t E'(V_t) + \frac{1}{2} \sigma^2 V_t^2 E''(V_t) \right] dt + \sigma V_t E'(V_t) dZ_t \right) \quad (\text{C.20})$$

$$+ e^{-r_f t} (\delta V_t - C) dt.$$

Factor out $e^{-r_f t}$ and group all dt terms

$$dG_t = e^{-r_f t} \underbrace{\left(\frac{1}{2} \sigma^2 V_t^2 E''(V_t) + (r_f - \delta) V_t E'(V_t) - r_f E(V_t) + (\delta V_t - C) \right)}_{\text{Drift}} dt \quad (\text{C.21})$$

$$+ e^{-r_f t} \sigma V_t E'(V_t) dZ_t.$$

Set the drift term to zero to satisfy the martingale condition, giving the ODE

$$\frac{1}{2} \sigma^2 V^2 E''(V) + (r_f - \delta) V E'(V) - r_f E(V) + (\delta V - C) = 0. \quad (\text{C.22})$$

C.1.4 Equity Value and Optimal Default Bound

We derive the optimal default barrier V_B chosen by shareholders to maximize equity value.

Assumption 14. *Equity holders receive a net flow of dividends $\delta V - C$ per instant, where C is the total interest expense.*

Assumption 15. *Shareholders possess limited liability; equity value is zero at default.*

Assumption 16. *Shareholders choose V_B optimally (Smooth Pasting condition).*

Section C.1.3 derives the ODE for the equity value $E(V_t)$

$$\frac{1}{2} \sigma^2 V^2 E''(V) + (r_f - \delta) V E'(V) - r_f E(V) + (\delta V - C) = 0. \quad (\text{C.23})$$

This is a non-homogenous Euler-Cauchy ODE with the known solution as the sum of the particular solution, value of solvent assets minus risk-free debt, and the homogeneous solution, the option value

$$E(V) = \left(V - \frac{C}{r}\right) + AV^{-\gamma} \quad (\text{C.24})$$

Note: The positive root term BV^η is discarded because the default option value must vanish as $V \rightarrow \infty$ to satisfy the solvency constraint.

Now, we solve for the coefficient A and the optimal default bound V_B using the boundary conditions

- Value Matching: $E(V_B) = 0$

$$V_B - \frac{C}{r} + AV_B^{-\gamma} = 0 \quad (\text{C.25})$$

- Smooth Pasting: $E'(V_B) = 0$

$$1 - \gamma AV_B^{-\gamma-1} = 0 \quad (\text{C.26})$$

From C.26, we express A in terms of V_B

$$A = \frac{1}{\gamma} V_B^{\gamma+1}. \quad (\text{C.27})$$

Substitute A back into C.25

$$V_B - \frac{C}{r} + \left(\frac{1}{\gamma} V_B^{\gamma+1}\right) V_B^{-\gamma} = 0 \quad (\text{C.28})$$

$$V_B - \frac{C}{r} + \frac{1}{\gamma} V_B = 0 \quad (\text{C.29})$$

$$V_B \left(1 + \frac{1}{\gamma}\right) = \frac{C}{r}. \quad (\text{C.30})$$

Solving for V_B yields the closed-form endogenous barrier

$$V_B = \frac{C}{r} \left(\frac{\gamma}{\gamma + 1} \right). \quad (\text{C.31})$$

C.1.5 Proposition 1

We derive the equilibrium deposit rates r_N^* and r_L^* under the assumption of asymmetric transport costs.

Assumption 17. *Depositors are uniformly distributed on $[0, 1]$.*

Assumption 18. *Transport cost for Local Bank is normalized to 1; National Bank is $\phi \in (0, 1]$.*

Assumption 19. *Depositors are risk-neutral. Local Bank depositors face expected credit loss flow $\alpha\Psi$.*

The indifferent depositor x^* equates net utility

$$r_N - \phi x^* = r_L - \alpha\Psi - (1 - x^*). \quad (\text{C.32})$$

Solving for x^* gives the National Bank's market share M_N :

$$M_N = x^* = \frac{r_N - r_L + 1 + \alpha\Psi}{1 + \phi}. \quad (\text{C.33})$$

The Local Bank's share is $M_L = 1 - M_N$.

Banks maximize profit $\Pi_k = (R - r_k)M_k$. The FOCs are

$$\frac{\partial \Pi_N}{\partial r_N} = -M_N + (R - r_N) \left(\frac{1}{1 + \phi} \right) = 0 \implies R - r_N = (1 + \phi)M_N \quad (\text{C.34})$$

$$\frac{\partial \Pi_L}{\partial r_L} = -M_L + (R - r_L) \left(\frac{1}{1 + \phi} \right) = 0 \implies R - r_L = (1 + \phi)M_L \quad (\text{C.35})$$

We solve the system for the spread $S = r_L - r_N$. Subtract C.34 from C.35

$$(R - r_L) - (R - r_N) = (1 + \phi)(M_L - M_N) \quad (\text{C.36})$$

$$r_N - r_L = (1 + \phi)(1 - 2M_N) \quad (\text{C.37})$$

$$-S = (1 + \phi) - 2(1 + \phi)M_N \quad (\text{C.38})$$

$$= (1 + \phi) - 2(r_N - r_L + 1 + \alpha\Psi) \quad (\text{C.39})$$

$$= -S + 1 + \alpha\Psi \quad (\text{C.40})$$

which we solve to find the equilibrium spread

$$S = \frac{1 - \phi}{3} + \frac{2}{3}\alpha\Psi \quad (\text{C.41})$$

Using the spread, we solve for levels

$$r_N^* = R - (1 + \phi)M_N^* = R - (-S^* + 1 + \alpha\Psi) \quad (\text{C.42})$$

$$= R - 1 - \alpha\Psi + \left(\frac{1 - \phi}{3} + \frac{2}{3}\alpha\Psi \right) \quad (\text{C.43})$$

$$= R - \frac{2 + \phi}{3} - \frac{\alpha}{3}\Psi \quad (\text{C.44})$$

$$(\text{C.45})$$

$$r_L^* = r_N^* + S^* \quad (\text{C.46})$$

$$= \left(R - \frac{2 + \phi}{3} - \frac{\alpha}{3}\Psi \right) + \left(\frac{1 - \phi}{3} + \frac{2}{3}\alpha\Psi \right) \quad (\text{C.47})$$

$$= R - \frac{1 + 2\phi}{3} + \frac{\alpha}{3}\Psi \quad (\text{C.48})$$

The term $\alpha\Psi$ represents the flow cost of credit risk. We explicitly derive this term, in terms of the underlying asset value V_t and the endogenous default boundary V_B , by requiring that, in equilibrium, the risky deposit rate must offer a arbitrage-free risk-adjusted return equal to the risk-free rate r_f .

Let $D(V_t)$ be the value of a unit deposit paying a continuous coupon yield $y(V_t, V_B) =$

$r_f + \alpha\Psi(V_t, V_B)$ until default, with a recovery of $(1 - \alpha)$ at the default moment τ . Under the risk-neutral measure $\mathcal{Q}$, the value of this claim is

$$D(V_t, V_B) = \mathbf{E}_{\mathcal{Q}} \left[\int_0^{\tau} e^{-r_f s} y(V_t, V_B) ds + e^{-r_f \tau} (1 - \alpha) \right]. \quad (\text{C.49})$$

Using the closed-form solution for the First Passage Time security in Section C.1.2

$$P(V_t, V_B) = \mathbf{E}_{\mathcal{Q}}[e^{-r_f \tau}] = \left(\frac{V_t}{V_B} \right)^{-\gamma} \quad (\text{C.50})$$

we rewrite the value of the unit deposit $D(V_t, V_B)$

$$D(V_t, V_B) = \frac{y(V_t; V_B)}{r_f} (1 - P(V_t, V_B)) + (1 - \alpha) P(V_t, V_B). \quad (\text{C.51})$$

For the deposit to trade at par $D(V_t, V_B) = 1$, the offered yield $y(V_t, V_B)$ must satisfy

$$1 = \frac{r_f + \alpha\Psi(V_t, V_B)}{r_f} (1 - P(V_t, V_B)) + (1 - \alpha) P(V_t, V_B) \quad (\text{C.52})$$

Solving this gives the functional form for the credit premium Ψ

$$\Psi(V_t, V_B) = r_f \left(\frac{P(V_t, V_B)}{1 - P(V_t, V_B)} \right) = r_f \left(\frac{(V_t/V_B)^{-\gamma}}{1 - (V_t/V_B)^{-\gamma}} \right). \quad (\text{C.53})$$

C.1.6 Corollary 1

Substitute the equilibrium spread C.41 into the market share equation C.33

$$M_N^* = \frac{-S^* + 1 + \alpha\Psi}{1 + \phi} \quad (\text{C.54})$$

$$= \frac{1}{1 + \phi} \left(1 + \alpha\Psi - \left(\frac{1 - \phi}{3} + \frac{2}{3} \alpha\Psi \right) \right) \quad (\text{C.55})$$

$$= \frac{1}{1 + \phi} \left(\frac{3 + 3\alpha\Psi - 1 + \phi - 2\alpha\Psi}{3} \right) \quad (\text{C.56})$$

$$= \frac{2 + \phi + \alpha\Psi}{3(1 + \phi)} \quad (\text{C.57})$$

The Local Bank share is

$$M_L^* = 1 - M_N^* = \frac{3(1 + \phi) - (2 + \phi + \alpha\Psi)}{3(1 + \phi)} = \frac{1 + 2\phi - \alpha\Psi}{3(1 + \phi)}. \quad (\text{C.58})$$

For profits, recall from the FOCs that optimal profit is $\Pi_k^* = (1 + \phi)(M_k^*)^2$. Substituting the expressions for M_k^*

$$\Pi_N^* = (1 + \phi) \left(\frac{2 + \phi + \alpha\Psi}{3(1 + \phi)} \right)^2 \quad (\text{C.59})$$

$$\Pi_L^* = (1 + \phi) \left(\frac{1 + 2\phi - \alpha\Psi}{3(1 + \phi)} \right)^2 \quad (\text{C.60})$$