**Title**

Essays in Financial Economics



**Author**

Jones, Collin

**Publication Date**

2024

Essays in Financial Economics

By

Collin Jones

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Romer, Chair
Professor Yuriy Gorodnichenko
Professor David Sraer
Professor Jon Steinsson

Spring 2024

Abstract

Essays in Financial Economics

by

Collin Jones

Doctor of Philosophy in Economics

University of California, Berkeley

Professor David Romer, Chair

In chapter one, *New Evidence on Convenient Asset Demand*, I study aggregate demand for short-term convenient assets. I estimate the slope of the aggregate demand curve for these assets, which governs how a given change in convenient assets outstanding changes their convenience yield. I innovate relative to the existing literature by using a new instrument, which is a direct measure of T-bill issuance *surprises* relative to the projections of a well-informed market newsletter, Wrightson ICAP. I argue that Wrightson surprises are plausibly uncorrelated with changes in convenience demand, and are a methodological improvement over the literature's previous approaches. Using local projection methods, I find that the demand curve for short-term convenient assets is meaningfully steep only in the very short-run. A $100 billion increase in the supply of T-bills depresses T-bill convenience yields by 10.4 basis points, on average, in the week of the increase. However, the long-run effect is much more modest, with a $100 billion higher *stock* of T-bills only depressing convenience yields by 1.1 basis points.

In chapter two, *Empirical Network Contagion for US Financial Institutions*, coauthored with Fernando Duarte, we construct an empirical measure of expected network spillovers that arise through default cascades for the US financial system for the period 2002-2016. Compared to existing studies, we include a much larger cross-section of US financial firms that comprise all bank holding companies, all broker-dealers and all insurance companies, and consider their entire empirical balance sheet exposures instead of relying on simulations or on exposures arising just through one specific market (like the Fed Funds market) or one specific financial instrument (like credit default swaps). We find negligible expected spillovers from 2002 to 2007 and from 2013 to 2016. However, between 2008 and 2012, we find that default spillovers can amplify expected losses by up to 25%, a significantly higher estimate than previously found in the literature.

In chapter three, *Money Fund Demand and Regulatory Reform*, coauthored with Abhi Gupta, we introduce an empirical framework for estimating a complete asset demand system in US money markets. The novel approach uses end-of-quarter window dressing by certain

financial firms as a supply shock, to estimate the yield sensitivity of different money market investors. This framework can be used to investor-level demand parameters and compute pricing counterfactuals, to ask whether post-2016 regulatory reforms have led to more or less elastic market demand. Our framework is specially catered to be feasible to estimate with existing data on US money markets.

# Contents

For Dad

# Introduction

The 2008-2009 Great Financial Crisis (GFC) highlighted the importance to macroeconomics and finance of understanding the realities of different actors in the US financial system. The issues encountered during that period generated new theoretical work studying the importance of financial institutions and intermediaries – such as intermediary asset pricing. The GFC's aftermath has also generate new empirical opportunities for studying these questions, as regulators' efforts to improve market transparency have created new datasets on the holdings and actions of intermediaries.

In the three chapters of this dissertation, I contribute to the literature's understanding of the US financial system and work towards three empirical objectives. First, I aim to estimate and understand the *demand curves* of financial actors for different financial assets or different qualities of assets. In traditional, frictionless asset pricing models, demand curves for financial assets are nearly perfectly horizontal, so that changes in outstanding supplies of different securities have essentially no impact on their price or return. However, recent empirical work has shown that supply shocks do *indeed* shift asset prices, in substantial and possibly-persistent ways. In a world where asset demand curves are meaningfully downward-sloping, understanding the foundations of different actors' demand curves becomes critical for understanding asset prices.

My second empirical objective is to better understand the financial stability vulnerabilities of the US financial system. Which sorts of economic or financial shocks will generate large financial losses for intermediaries, adversely affect firms' or household's costs of financing, or or induce risk-taking in the financial system that may lead to the same?

My third objective is to go beyond aggregate studies of the first two objectives, and characterize the *cross-section* of different agents' heterogeneities, and how those heterogeneities contribute to asset prices or financial stability. For instance, when studying asset demand curves, I wish to understand both the degree of marketwide elasticity and which intermediaries or investor sectors contribute most to that elasticity. In studying financial stability, this amounts to learning *which* firms' balance sheets, trades, or debt issuances pose the greatest threat to system-wide risks. As Brunnermeier et al. (2021) write, the goal is that "Instead of abstractly referring to 'arbitrageurs,' 'intermediaries,' and 'noise traders' in our theories, we actually know who they are, what their asset demand curves look like, how large they are, and what their contribution is to fluctuations in asset prices." This sort of work has become more feasible only lately, as regulators have collected (and in many cases, released publicly) new cross-sectional data on financial institutions in the wake of the GFC.

In chapter one, *New Evidence on Convenient Asset Demand*, I study the closely-watched demand curve for short-term convenience. It is now well-established in the economics and finance literatures that fixed income assets that are convenient for investors to hold will trade at a lower yield than a less-convenient alternative. This yield spread is the convenience yield, and measures the yield that investors are willing to forego for a marginal unit of asset convenience. Chapter one studies how this convenience yield shrinks in response to exogenous increases in the outstanding supply of convenient assets. In other words, it estimates the slope of the short-term convenient asset demand curve.

This question has been empirically studied before, but chapter one innovates by introducing a new instrument for short-term convenient asset supply. In order to estimate the slope of this demand curve, I require a component of convenient asset supply that is uncorrelated (at the relevant frequencies) with shifters of short-term asset demand – such as liquidity preferences, risk preferences, and the like. I argue that high-frequency *surprises* in the quantities of Treasury bills (T-bills) sold by the US Treasury at auction satisfy these requirements, and I compile a direct measurement of these surprises, relative to the T-bill quantity forecasts of Wrightson ICAP, one well-informed money market newsletter.

My results transparently show that it is critical to define the horizon of interest in answering this question, which previous studies have not done. After a positive T-bill issuance surprise, convenience yields fall more sharply than previous estimates in this literature would suggest. However, these effects are more short-lived than those previous estimates allow. I estimate that a \$100bn increase in the supply of T-bills only leads to a 1 basis point decrease in T-bill convenience yields on a permanent basis.

Chapter one most directly contributes to this dissertation's first objective, by studying a particularly important asset demand curve. The slope of this particular demand curve is an important inputs into macroeconomic models that feature special investor demand for convenient assets. One strain of this literature studies the sustainability of the US federal debt. My results tend to make certain debt levels appear *more* sustainable, relative to earlier estimates in the literature: a given increase in Treasury borrowing has a smaller permanent effect on convenience yields (i.e. smaller increases in government interest rates).

Chapter one also contributes to the second objective – to better understand threats to US financial stability, because of the stability implications of this particular demand curve. In this literature, short-term borrowing by the US federal government can affect financial stability by crowding out the issuance of money-like alternatives that are more susceptible to runs in crisis – such as financial commercial paper or prime money market mutual fund shares. This mechanism operates through US debt supply's impact on convenience yields. As such, my results suggest that these worries are only acute at short horizons, when convenience yields tend to be most depressed after a T-bill issuance increase.

The heterogeneity in my estimates at the short and medium-run horizons raise the question of why these two slopes should differ so substantially. In its explanation, chapter one also speaks to my third objective, to provide cross-sectional information about which firms or sectors contribute most to market-wide phenomena. In slow-moving capital models such as Duffie (2010), large price changes after a supply shock can dissipate over time if the universe of investors who are able or willing to adjust their holdings after supply surprises widens over time over time after a shock. A key implication of this explanation is that the investors who absorb a surprise increase in T-bill supply should change over time after the surprise. I demonstrate this using end-of-month holdings data from money market mutual funds, a key investor group in this market. I demonstrate that a particular subsector in this market – Treasury-only money funds – appear to play the part of the "fast" investors in this setting, by holding an outsized share of increased supplies.

In chapter two, *Empirical Network Contagion in the US Financial System*, I focus more directly on my objective to study financial stability, in co-authored work on financial network contagion in the US financial system. A lesson of the GFC is that linkages between financial firms can amplify financial losses after some initial economic shock, in a way that affects the health and solvency of many firms – even those without an obvious, direct exposure to the initial shock. The theoretical literature on financial network contagion often suggests empirical measures of how susceptible a current network structure is to certain contagion episodes. However, a drawback of these measures is that calculating them typically requires data on the linkages between individual firms (e.g. how much firm $A$ owes to firm $B$). This data is challenging to use in practice: either because it does not exist, or because it is protected by supervisory confidentially.

In chapter two, my coauthor Fernando Duarte and I construct a measure of US financial network vulnerability that accounts for the entire US financial system and is feasible to construct using publicly-available data alone. The measure implements an object developed theoretically in Glasserman and Young (2015). It is an upper-bound on how the structure of a financial network amplifies expected financial sector losses result from shocks originating in the nonfinancial sector. Amplification in the model takes a simple form: defaults on financial sector debt held by other financial agents cause *additional* within-network losses to the holders of the defaulted debt. [1] The measure compares expected losses in the true, interconnected network to those of a hypothetical network where intra-network linkages have been removed and replaced with other claims.

We estimate this measure for the US financial system using publicly-available data. Necessary inputs include firms' probabilities of default (for which we use market-based measures), firm leverage, and firms total connectedness – the latter being the share of the firms' liabilities held by other financial firms. Computing the last category requires the most judgement on behalf of the researcher, and represents this chapter's largest methodological contribution. We estimate this interconnectedness for US dealers and bank holding companies using publicly-available balance sheet information. This necessitates making decisions on whether certain types of reported liabilities are likeliest to be held by other financial firms or by non-financial actors.

This chapter's primary output is a quarterly series computing our vulnerability measure. Previous work suggests that the simple sort of contagion assessed by this measure is likely small. However, our measure supports this interpretation only in certain quarters. For instance, in 2016 the measure suggests that expected losses in the US financial system rose by a maximum of 3% because of these linkages. During the GFC, however, the measure rises considerably, so that the measure cannot rule out expected losses rising by as much as 35%. Chapter two also contributes to my goal of demystifying the cross-section of US financial system mechanisms. Since the network vulnerability measure is calculated as a sum across firms, each firms' contribution to the measure can be interpreted as a measure of the network vulnerability arising from that firm's balance sheet.

In chapter three, *Money Fund Demand and Regulatory Reform*, my coauthor Abhi Gupta and I work towards all three of this dissertation's objectives, in developing a framework to estimate demand parameters for investors in US money markets. This project belongs to a new strain of literature in asset pricing that estimates full demand systems in different asset classes. This literature uses detailed holdings data – preferably at the investor or fund level – to estimate fund-level demand parameters governing how a fund's portfolio shares in different assets depend on asset characteristics. This is done via a logit demand system, familiar to economists in the industrial organization literature. The assumptions inherent

[1] That is, if Bank of America experiences a loss on its claims to the nonfinancial sector, and thus defaults on its own debt held by Morgan Stanley, then that generates additional losses for Morgan Stanley.

in this approach allow researchers to use the parameter estimates to conduct asset pricing counterfactuals.

In this paper, we adapt the microfoundation and estimation strategies from the asset demand system literature, to make them more applicable to the institutional setting and data availabilities in US money markets. Detailed end of month holdings data is available for money market mutual funds in the United States, and certain volumes and yields data is commercially available at the trade level from the Depository Trust Clearing Corporation – a major clearinghouse in many asset classes. First, we adapt a microfoundation for the logit demand setup of Koijen and Yogo (2019), to better fit the realities of a money market mutual fund's investment problem. The problem features transaction costs, fund outflows, and fixed income assets whose risk qualities change as they age, then predictably mature into cash.

A critical step in estimating our demand system involves estimating how a money market asset's *yield* affects investors' portfolio shares in that security (i.e. estimating the slope of investors' demand curves). This involves resolving an endogeneity problem: a money market asset's yield in equilibrium will surely be correlated with *latent demand* for that security, where latent demand in this setting includes security characteristics that matter for investors but are unobserved by the econometrician. We detail an estimation strategy that estimates this parameter using a well-documented *supply shock* in this market, whereby certain financial firms reduce their issuance of overnight liabilities (such as repurchase agreements or commercial paper) at the end of each month, to window dress leverage statistics reported to their regulators. Our identification strategy first assumes that these end-of-month effects are fundamentally *supply* shocks, and are uncorrelated with any changes in the relative desirability to investors of window dressers and non window dressers' commercial paper.

By providing an empirical methodology to estimate a demand system in US money markets, this chapter very directly works towards the first and third goal. Researchers can use this framework to estimate the predicted price effects of an increase in issuance of many assets in US money markets, thus assessing marketwide demand elasticities at different points in the sample. This approach provides a direct way to assess which money market investors contribute to that elasticity.

Critically for this market, available data allows for estimating demand parameters for both money market mutual funds (which report their holdings) and for the average *residual* (non-MMF) investor in markets such as commercial paper. This is especially enlightening in US money markets, where a slate of reforms in 2016 caused large holdings of US commercial paper to exit the money market fund industry, and enter into some less regulated – but currently unidentifiable – alternate sector. The framework introduced in this chapter will allow for, to our knowledge, the first assessment of how asset demand from this "residual" sector differs from that of the more-regulated money market fund industry.

The empirical methodology introduced in chapter three also allows analyses that are directly relevant to financial stability. US money markets are some of the most crisis-vulnerable asset classes in the financial system – as is evidenced by the Federal Reserve's need to bolster them with liquidity facilities in many of the latest crisis periods. The framework of chapter three allows researchers to study market elasticity in one of the most-watched markets for financial stability questions. Studying how market demand has evolved around 2016 money market reforms will also speak directly to the efficacy of using portfolio restrictions and other regulations to bolster market stability. For instance, if the "residual" sector has asset demand qualities that contribute to greater market volatility, then this shows how new portfolio restrictions on a well-regulated industry (like the money fund sector) could have the counterproductive result of migrating holdings to less preferred investor sectors.

# Acknowledgements

Thank you to my committee members Yuriy Gorodnichenko, Jon Steinsson, and David Sraer, and to my oral exam committee chair Emi Nakamura. Extraordinary teachers and advisors like you all are a huge part of why I pursued my PhD in the first place. Having you all help shape my development as a researcher has been a tremendous privilege. Thank you to Christina Romer, for offering excellent advice (to both Pearl and me) and for being just about the best boss that a teaching assistant could ask for.

Thank you to my chair David Romer, who has been a constant source of support and generosity as he shepherded this dissertation to its completion. I will always be deeply grateful!

Thank you to my coauthors Abhi Gupta and Fernando Duarte, for excellent conversations and collaboration as we developed this dissertation's research. I am also grateful to Rui Yu and Rita Zhang for excellent research assistance.

I am thankful to Michael Fleming of the Federal Reserve Bank of New York, for invaluable mentorship during the first steps of my research career. I am also grateful to all of the economists of the Money Market Analysis group at the Board of Governors of the Federal Reserve System, for their time and feedback as I developed this research as a summer dissertation fellow.

Thank you to the UC Berkeley Clausen Center and to Brian Barish for generous financial support towards this dissertation's research.

I surely could not have completed this dissertation without the endless support of my family. To my parents, especially – you have brought me to, and through, this degree in more ways than I can possibly mention. I am thrilled that the next chapter of my career will come with more time together, including with the many nieces and nephews who were born as I finished this degree!

Most importantly, I am grateful that I was able complete this degree alongside my partner, Pearl. Sharing this all with you – the highs and the lows – has brought me more happiness and support than anything.

# Chapter 1

# New Evidence on Convenient Asset Demand

## 1.1 Introduction

Investor demand for convenient, short term assets has come to occupy an important place in macroeconomic models across a variety of subdisciplines.[1] In models of optimal fiscal policy and government financing, a desire to sate this demand means that governments should issue more debt that satisfies this special investor need. In business cycle models, fluctuations in the strength of this demand can drive business cycles, by shifting the desirability of savings in a way that is not reversed by central banks.

The quantitative predictions and policy prescriptions of these models are linked to the price sensitivity of this convenient asset demand. This is the *slope* of the convenient asset demand curve. In this context, the price of convenience is the difference in yield between some less convenient reference asset and a convenient asset. This difference is the convenience yield, and measures the yield that an investor will *forgo* by holding a convenient asset. A flat demand curve means that large increases in the outstanding supply of short-term convenient assets cause only modest decreases in this convenience yield. When this demand curve is flatter, the US government's debt issuance decisions are less likely to drive business cycles, as in Kekre and Lenel (2023), and a given fiscal deficit is more likely to be fiscally sustainable, as in Mian et al. (2022). In normative terms, a flatter convenient asset demand curve suggests to policymakers that maintaining a larger supply of these assets is good for welfare, all else equal.[2]

Several recent studies have estimated the slope of this demand curve, by observing how convenience yields have historically changed after changes in the outstanding quantity of convenient assets. Krishnamurthy and Vissing-Jorgensen (2012) does this with a long sample and low frequency variation in the outstanding quantity of Treasury securities. Greenwood et al. (2015b) does this using higher frequency variation in the outstanding quantity of

[1] This paper features data scraped with permission from the Wrightson website for research purposes. The content of this paper reflects my own views, and not those of Wrightson ICAP. Any errors committed in this data scraping are solely my own.

[2] Vissing-Jorgensen (2023) discusses this in the monetary policy context. Angeletos et al. (2023) discusses it in the fiscal context.

short-term Treasury bills (T-bills).

As in any setting where the goal is to estimate the slope of a demand curve, identification relies on isolating a component of convenient asset supply that is uncorrelated with any unobservables that independently shift demand. In high frequency studies like Greenwood et al. (2015b), the most-discussed endogeneity concern is *opportunistic issuance*: that the Treasury may respond to movements in the price of convenience by changing their issuance of convenient debt.[3] Those authors address this using a seasonality instrument, arguing that the seasonality in outstanding T-bill supplies is driven by predetermined deadlines in the US tax calendar – *not* by a Treasury response to demand. This seasonality approach is still standard in this literature, used in studies as recent as D'Avernas and Vandeweyer (2023) and Infante (2020).

In this paper, I estimate the slope of the short-term, convenient asset demand curve in the post-crisis period using a fundamentally different instrument. My instrument measures surprises in the quantity of T-bills sold by the Treasury at auction. I measure the size of these surprises directly, using high frequency projections of T-bill auction quantities from Wrightson, a highly-informed and well-respected money market newsletter. Using these surprises as an instrument for T-bill supply, I demonstrate that convenient asset demand appears rate insensitive (steep) in the very short run, but price sensitive (flat) at only slightly longer horizons. My results suggest that a \$100 billion larger *stock* of T-bills depresses T-bill convenience yields by only 1.1 basis points. However, a \$100 billion *increase* in the supply of T-bill will depress convenience yields by 10.4 basis points in the week of the change.

I argue that Wrightson surprises are plausibly uncorrelated with high frequency shifts in convenient asset demand, permitting their use as instruments to estimate the slope of the demand curve. The most important component of this argument is an institutional quirk in the Treasury's T-bill issuance strategy. To avoid being seen as an opportunistic issuer, the Treasury *does not* alter their issuance decisions in response to short-term changes in demand. Treasury statements to this effect are unambigious and speak directly to my exclusion restriction. In one slide deck displayed prominently on the US Treasury's website, the Office of Debt Management writes "Treasury doesn't react to current rate levels or short-term fluctuations in demand".[4] This suggest that when the Treasury surprises Wrightson by issuing more or fewer T-bills than expected, they are not doing so as a response to convenient asset demand. That is, high frequency surprises do not reflect opportunistic issuance.[5]

An inherent complication in studying my research question relates to the supply of T-bill *substitutes*. If other assets provide convenience qualities that are substitutable with the convenience provided by T-bills, then the available quantity of those other assets will affect T-bill convenience yields in equilibrium. Even if the US Treasury does not alter its issuance

[3]To continue the analogy between this setting and a classic demand curve estimation problem: Opportunistic issuance would imply that the US Treasury's supply curve is not vertical. This induces a simultaneous equation estimation problem.

[4]This presentation, which I will discuss further in the main text of this paper, is available as of October 2023 at https://home.treasury.gov/system/files/276/Debt-Management-Overview.pdf.

[5]Importantly, these statements do not necessarily mean that T-bill issuance at lower frequencies is not opportunistic. In response to some persistent change in investor demand for convenient assets, the Treasury may gradually change its issuance strategy. My identification argument must be paired with a high frequency approach, to accord most closely with this stated Treasury policy.

in response to market rates, issuers of substitutable assets may. We might expect those issuers to decrease their issuance after a positive T-bill quantity surprise. This induces a sort of measurement problem. While Wrightson surprises are still a valid supply instrument under my exclusion restriction, observable changes in T-bill supply after a shock may not reflect the total change in relevant short-term convenient asset quantities.

I address this issue chiefly via a sample restriction, motivated by recent insights from D'Avernas and Vandeweyer (2023). Using a simple stylized framework, I clarify that this issue is most problematic in my post-crisis sample when the Federal Reserve's Overnight Reverse Repurchase Agreement (ONRRP) facility is active. In those weeks, the Fed's monetary policy tools make a portion of plausibly-substitutable convenient asset supply (i.e. repurchase agreements issued by the Fed) perfectly rate elastic, dampening the effect of T-bill supply shocks on rates. As such, in the analyses of convenience yields to follow, I exclude from my sample those weeks when the ONRRP facility is active.

Employing this sample restriction, I use Wrightson surprises as instruments to estimate the dynamic response of T-bill convenience yields to a T-bill supply shock, with a local projection instrumental variables (LP-IV) approach. These results show that defining the horizon of interest is hugely important for interpreting the magnitude of the convenience yield response. In the first weeks after a supply shock, T-bill convenience yields drop sharply. However, they recover quickly, with the point estimates returning close to zero approximately three weeks after each shock.

To use these results to estimate the slope of the convenient asset demand curve, I must also estimate how the quantity of T-bills changes over my impulse response horizon. I show that, while Wrightson surprises do not predict a permanent change in T-bill supply, they do predict higher T-bill supplies for several weeks. Supplies rise steadily in the three weeks following a surprise, before beginning their decline at week four. Supplies only return to their pre-surprise level eight weeks after a surprise.

I show that these two estimated impulse responses are consistent with the simple interpretation that *changes* in T-bill supplies have a larger effect on convenience yields than the *stock* of T-bills.[6] I do this by estimating separate stock and flow effects of supply on convenience yields, via a multiple equation GMM approach. This procedure fits the same empirical moments as the LP-IV approach described above, but using fewer (two) estimated parameters. The identifying variation that separates flow and stock effects comes from the dynamics in the T-bill supply response over the impulse horizon. For instance, at horizons of four weeks or greater after a Wrightson surprise, the stock of T-bills is elevated relative to its pre-shock value, while the flow (change) in T-bill supply is sharply negative.

I also propose an alternate specification that trades a somewhat more restrictive, although similar, exclusion restriction for substantially more statistical power. In my baseline estimation, the flow effect is highly statistically significant, but the stock effect is not statistically significantly different from zero. The alternate specification leverages the high-frequency nature of the Wrightson projections, by including Wrightson's own *updates* in its projections of future T-bill supply, in the week of each surprise, as additional instruments. These results imply somewhat larger convenience yield effects, and tighter confidence intervals. Under this

[6] This concept should be familiar to market participants and policymakers, who already frequently discuss Treasury borrowing using the identical concept of "net issuance"

specification, both the stock and flow effects are statistically significant. Yet these results exhibit the same qualitative pattern as the baseline, that flow effects are large, and stock effects are small.

I show that estimated impulse responses for T-bill convenience yields and supplies of T-bill substitute assets are consistent with the stylized framework that led me to restrict the sample to those periods when the Fed's ONRRP facility is inactive. That is, I show how the response of convenience yields to T-bill supply shocks is muted in the subsample where the ONRRP facility is active. In those periods, convenience yields fall on impact, but return to zero more quickly. Instead, outstanding volumes of repurchase agreements appear to be the margin of adjustment: in ONRRP active weeks, repurchase agreement volumes move in the *opposite* direction of T-bill supplies, essentially one-for-one.

My results also indicate that, in ONRRP-inactive periods, outstanding volumes of the other likeliest substitute assets for T-bill convenience do not fall enough to seriously complicate my empirical estimates from those periods. Outstanding discount notes issued by the Federal Home Loan Bank system decrease, consistent with the FHLB issuing opportunistically. However, they decline only \$1.5 billion in response to a Wrightson surprise that moves T-bill supplies by approximately \$25 billion. Outstanding volumes of privately-issued, general collateral repurchase agreements show little reaction to T-bill supply shocks when the ONRRP is inactive.

This paper's separate estimates of flow and stock effects of T-bill supply are relevant for two, possibly-overlapping groups. For policymakers tasked with stabilizing money market rates, a powerful weeklong effect on yields is relevant to their policy objectives. To macroeconomists studying lower-frequency questions like the sustainability of the US fiscal position, the small permanent stock effect is likely most relevant.

My estimated flow effects are larger than convenience yield effects estimated by recent studies, while the estimated stock effect is smaller. In a recent study with a similar sample period, D'Avernas and Vandeweyer (2023) estimate that a \$100 billion increase in T-bill supplies depresses convenience yields by 4 basis points. My results share one important quality with theirs, that convenience yield responses to T-bill supply changes appear smaller in post-crisis data than in pre-crisis data. In the original, pre-crisis estimates of Greenwood et al. (2015b), a \$100 billion increase in T-bill supply depresses convenience yields by 8.27bp.[7]

To put my results in perspective, I demonstrate how they alter one of the positive conclusions from Mian et al. (2022), a leading study in the $R < G$ fiscal sustainability literature. In that paper, the sensitivity of government debt convenience yields to increases in debt issuance helps dictate what range of fiscal deficits appear sustainable in the United States, in the sense that they will lead to stable $\frac{Debt}{GDP}$ ratios in the long-run. In a calibration exercise corresponding to the US fiscal position as of 2019, the original estimates of Greenwood et al. (2015b) suggest a maximum sustainable fiscal deficit of 2.0% of GDP. Recent estimates from D'Avernas and Vandeweyer (2023) imply a maximum sustainable deficit of 2.4% of GDP. My point estimates, with other calibrations from Mian et al. (2022) unchanged, suggests a max-

[7]The estimates of D'Avernas and Vandeweyer (2023) are most similar to mine in sample and methodology, and thus are the most-natural point of comparison to demonstrate that my estimated stock effects are *small* (and flow effects are *big*). However, the specification and research question of D'Avernas and Vandeweyer (2023) are sufficiently different from mine that one would not necessarily expect our estimates to be identical, as I discuss in Section 1.2.

imum deficit of 4% of GDP. The *largest* convenience yield stock impact that my estimates fail to reject at the 90% level implies a maximum deficit of 3% of GDP. My estimates suggest that the slope of convenience demand alone appears insufficient to meaningfully constrain the positive implications of Blanchard (2019), that large levels of US government debt and large deficits appear sustainable when nominal interest rates are low.

The rest of this paper proceeds in eight sections. In Section 2, I discuss the related literature. In Section 3, I clarify my structural parameters of interest with a simple model of convenient asset demand that is very similar to setups from Krishnamurthy and Vissing-Jorgensen (2012) and others. In Section 4, I provide important institutional details about T-bill issuance, such as the Treasury's commitment to avoid short-term opportunistic issuance. In Section 5, I use a stylized framework for understanding convenient asset supplies *other* than T-bills to explain how a post-crisis monetary policy tool effectively limits the post-crisis data sample for me methodology. In Section 6, I introduce Wrightson's T-bill issuance projections and their associated surprises (i.e. projection errors). I discuss the likeliest drivers of these surprises, and the ways in which these surprises improve upon the literature's standard approach of using seasonality instruments. In Section 7, I present my core empirical results regarding convenience yields and T-bill quantities. In Section 8, I present several estimates of interest regarding substitute asset supply, which are consistent with the story underpinning the core empirical results. Before concluding, I demonstrate in Section 9 how my estimates affect the positive conclusions of one strand of the fiscal sustainability literature, as in Mian et al. (2022).

## 1.2 Related Literature

This paper contributes to several literatures at the intersection of macroeconomics and finance. First, is the literature studying the convenience yields on safe, liquid debt securities. Gorton (2017) places the importance of safe assets in the historical context, and summarizes certain theories of their special value of investors. Diamond (2020) shows how a financial sector that issues safe assets to households, financing a risky portfolio of nonfinancial sector loans, is the equilibrium outcome in an economy where households value safety and financial frictions in firm borrowing.

This paper most clearly belongs to the subset of this literature studying these convenience yields' empirical properties. Longstaff (2004) showed that the rates on Treasury securities are *lower* than any traditional asset pricing framework would suggest, given how they compare with equally risk-free bonds issued by the Resolution Funding Corporation following the S&L Crisis. Krishnamurthy and Vissing-Jorgensen (2012) brought the concept of convenience yields to the academic forefront, and introduced the concept of measuring key yield spreads to study the apparent marginal value of long and short-term safety and liquidity. That paper used its low-frequency empirical analysis as a way to test its central notion that convenience yields reflect a special quality of Treasuries, with marginal values that fall in the outstanding quantity. Krishnamurthy and Vissing-Jorgensen (2015) extends that analysis by studying aggregate volumes of safe, private sector bank liabilities, and shows how these quantities tend to fall when Treasury supply rises. Of course, it is worth noting that the large literature on convenience yields, and special investor demand for safety or liquidity attributes, is identical in many ways to the preferred habitat literature of Vayanos and Vila (2021). Indeed, a "preferred habitat" for short-term, convenient assets is the precise focus of this literature.

The literature that most resembles this paper's specific questions and empirical approach is that which studies variation in convenience yields and convenient asset supplies at high frequencies. Greenwood et al. (2015b), discussed at length above, showed how convenience yields in short-term T-bills vary with the quantity outstanding of those bills. Sunderam (2015) showed that private sector issuance of asset-backed commercial paper – a possible private sector substitute for T-bills – rises at high frequencies as the volume of T-bills rises. They show that their findings are consistent with a model featuring substitutability of private and public sector convenient assets. Infante (2020) studies how convenience demand interacts with the issuance of private *repurchase agreements* – themselves likely seen as convenient by investors – to finance holdings of long-term Treasuries. That study uses T-bill supply as a plausibly-exogenous shifter of demand for privately issued convenient assets, and features a high-frequency seasonality instrument in the spirit of Greenwood et al. (2015b). Klingler and Sundaresan (2023) shows how a proxy for T-bill demand from US primary dealers predicts movements in T-bill convenience yields since the financial crisis, suggesting that the demand of those agents (possibly driven by regulatory considerations) has become an important driving factor.

In a contemporaneous paper that is complementary to this one, Phillot (2023) also studies the financial market effects of Treasury supply announcement surprises using a local projection instrumental variable approach. They measure Treasury supply surprises using daily changes in Treasury futures prices on Treasury auction announcement dates. The present paper differs from Phillot (2023) in several ways. Two are particularly important. First, my focus is on supply surprises in short-term, convenient Treasury bills, whereas Phillot (2023) studies supply surprises in longer-term notes and bonds. My analysis of issuance projections from Money Market Observer suggests that supply surprises are more frequent in T-bill issuance than they are in Treasury coupons and notes, so that my setting provides a more regular and frequent source of surprise variation  and in a market that arguably operates very differently from the market for notes and bonds. Second, my instrument is a direct measure of the surprise quantity component in Treasury quantity announcements, rather than a surprise inferred from movements in prices.

Recently, D'Avernas and Vandeweyer (2023) convincingly argued that the relationship between T-bill supply, private sector liquid asset issuance, and money market rates depends critically on the segmented nature of these markets, and on the availability of the Federal Reserve's ONRRP facility. That paper's focus was on testing the predictions of a structural model microfounding the interplay between private sector safe asset issuance, money market segmentation, and liquidity demand from a subset of money market investors. Their empirical specifications studying the effect of T-bill supply on convenience yields uses a seasonality instrument estimated in four-week differences, much like Greenwood et al. (2015b). That paper is transparent about their model *not* reflecting T-bill demand from government money market mutual funds – a large holder of T-bills in aggregate. As such, that paper does not directly study the question of interest in this paper, which is to estimate the *aggregate* elasticity of short-term convenient asset demand. At several points in this paper, I will reference the estimates of D'Avernas and Vandeweyer (2023), to help understand whether certain of my empirical estimates are large or small. This comes with the necessary caveat that that paper's empirical estimates are for a closely-related, but not identical, empirical question.

Other studies have discussed other implications of the ONRRP facility for convenient asset supply and demand. Ahnert and Macchiavelli (2021) show another way that the Fed's ONRRP facility has changed the nature of money markets, but providing money market mutual funds another source for liquid assets, to hold as insurance against investor outflows. Carlson et al. (2014) also addresses the potential for a close substitutability between ONRRP volume and T-bills, for market participants.

While not identical, a related literature makes measuring the substitutability of convenience from different asset classes its main focus. The first paper to make measuring convenience substitutability a primary focus was Nagel (2016), which stressed the need to consider the stance of monetary policy when studying the impacts of Treasury supply on convenience yields in low-frequency studies. Krishnamurthy and Li (2023) use low-frequency variation in rates and outstanding quantities of convenient assets to estimate elasticities of substitution between asset classes. Kacperczyk et al. (2021) show that short-term CDs in the Euro area appear to share some convenience substitutability with T-bills – although in times of market stress that substitutability appears to disappear. The subset of my results that are most relevant to this literature are those that measure movements in rates and volumes of assets *other* than T-bills after a T-bill supply shock. A fruitful direction for future research is to use this paper's notion of a T-bill supply shock to directly estimate that literature's parameters of interest, which are typically elasiticities of substitution between different safe asset classes.

This paper also belongs to a wider and more general literature that uses high frequency or event study methods to understand supply shocks in asset markets. Most relevant are those studies that study supply shocks in safe asset markets. Krishnamurthy and Vissing-Jorgensen (2011) study the impact of Quantiative Easing on bond yields in a classic event study. Gorodnichenko and Ray (2018) attempt to understand quantitative easing using a high frequency approach, focusing on unexpected *demand* revealed at Treasury auctions. Lou et al. show that supply changes in longer-term Treasury coupon security markets need not be *surprises* to have an impact on prices.

Given that my empirical results suggest that T-bill supply shocks have a *flow* effect that appears much larger than their stock (i.e. permanent) effect, my results provide evidence for theories that allow for overreaction in prices after an asset supply shock. Duffie (2010) shows how this may be caused by a subset of sluggish investors in markets, with potentially-elastic demand but only infrequent rebalancing practices. Particularly relevant for my setting, with shocks that occur amidst the highly-segmented US money market, Greenwood et al. (2018) document how price dynamics after a supply shock in one, partially-segmented market take time to affect other markets that are linked only by slow-moving, generalist arbitrageurs. In their model, prices in the market with a supply shock initially overreact to changes in supply, and prices of substitutes may underreact. D'Amico and King (2013), in another study on large-scale asset purchases by the US Federal Reserve, highlight seemingly-different stock and flow effects of Fed purchases, in much the same was as I do for T-bill supply shocks.

While the work described thus far shares a setting, method, or research question with this paper, there is a much wider literature in macroeconomics for which my identified parameter estimates will be relevant. The $R < G$ fiscal sustainability literature is immediately relevant, given the importance of the slope of convenient sovereign debt demand for that literature's quantitative conclusions. Relevant studies include Mian et al. (2022), Reis (2021), Blanchard (2019), Mehrotra and Sergeyev (2021), and Angeletos et al. (2023). Convenient asset demand curves likewise feature in a number of business cycle models, such as Kekre and Lenel (2023), Drechsler et al. (2018), and Bayer et al. (2023).

## 1.3 Model

In models with special investor demand for convenient assets, a yield spread between two appropriately chosen assets will measure the value to investors of the marginal unit of convenient assets. In a common, reduced form approach that is agnostic about the exact use source of this extrapecuniary value, I define the structural parameters that this paper estimates.

Most of the components of the model are standard, and replicate similar setups in Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood et al. (2015b). Specifically, the representative investor has expected utility function

$$\mathbb{E}_0\left(\sum_{t=0}^{\infty}\beta^t u(C_t + v(\frac{B_t}{GDP_t}))\right)$$

where $C_t$ is real consumption and $B_t$ is the nominal volume of convenient assets. Like in Krishnamurthy and Vissing-Jorgensen (2012), the normalization $\frac{B_t}{GDP_t}$ internalizes the notion that greater income and spending should come with greater desire for the sort of liquidity benefits that convenient assets provide. The object $v(\cdot)$ is the convenience function, which captures total extrapecuniary benefits. The investor is subject to a standard flow budget constraint,

$$B_{t-1}(1+i^b_{t-1}) + A_{t-1}(1+i_{t-1}) + I_t = P_tC_t + B_t + A_t + T_t$$

This budget constraint introduces quantity notation for $A_t$, the investor's nominal investment in risk-free (but inconvenient) assets; $I_t$ the investor's non-investment income; and $T_t$ the investor's transfer income. It introduces interest rate notation for $i^b_t$ the nominal net interest rate on bills; and $i_t$ the nominal net interest rate on risk-free illiquid assets.

The standard first order conditions in this model say that the marginal convenience value should equal the current convenience yield. Call $\frac{B^*_t}{GDP_t}$ the quantity of convenient asset holdings at which this is true, so that

$$v'_t(\frac{B^*_t}{GDP_t}) = \frac{i_t - i^B_t}{1+i_t} \geq 0 \tag{1.1}$$

In other words, at this quantity $B^*_t$, the marginal interest earnings that an investor is willing to *forgo* by holding the marginal convenient asset should be exactly offset by the extrapecuniary value of marginal convenience.

As with other papers in the literature, I will assume a tractable parameterization for the marginal convenience function $v'$. I say that the marginal benefit of an additional unit of convenience declines linearly in the quantity of convenient assets held.

$$v'(\frac{B_t^*}{GDP_t}) = \alpha + \beta\frac{B_t}{GDP_t}^{*} + \xi_t = \frac{i_t - i_t^B}{1+i_t} \tag{1.2}$$

If the investor's actual holdings of convenient assets satisfy these first order conditions in each moment, then there is no need to define separate notions of short-run and long-run demand curves. A 1-unit, permanent increase in the outstanding quantity of convenient assets should lower convenience yields by $\beta$, instantly and permanently.

In the high frequency setting of this paper, there is reason to believe that the very short-run price impacts of a change in supply may differ from its long-run effects. The theoretical models of Duffie (2010) and Greenwood et al. (2018) have just such a feature. D'Amico and King (2013) document separate flow and stock effects of the Federal Reserves Large-Scale Asset Purchases program – which also featured Treasury securities. Lou et al. (2013) show that this phenomenon is not restricted to *surprise* supply shocks. Well-anticipated auctions sizes of longer-term Treasury coupon securities show a similar rebound effect. This paper's own empirical analyses to follow suggest that allowing for such differences is necessary to account for impulse responses to T-bill supply shocks.

I allow for this possibility in a simple way via an *inertia* equation for the investor's actual convenient asset holdings. Borrowing a setup from Gabaix and Koijen (2021), I write

$$\frac{B_t}{GDP_t} - \frac{B_{t-1}}{GDP_{t-1}} = \mu\left(\frac{B_t^*}{GDP_t} - \frac{B_{t-1}}{GDP_{t-1}}\right)$$

where $0 < \mu \leq 1$. This parameterization implies that actual convenient asset holdings respond only sluggishly to changes in the long-run optimal level of convenience holdings $B_t^*$. This particularly tractable setup suggests that this sluggishness lasts exactly one period, so that convenience yields settle to their new long-run level one period after a permanent supply shock. To see this, note that

$$\begin{aligned}
\Delta\frac{i_t - i_t^B}{1+i_t} &= \beta\Delta\frac{B_t^*}{GDP_t} + \Delta\xi_t \\
&= \beta\left(\frac{1}{\mu}\Delta\frac{B_t}{GDP_t} - \frac{B_{t-1}^*}{GDP_{t-1}} + \frac{B_{t-1}}{GDP_{t-1}}\right) + \Delta\xi_t \\
&= \frac{\beta}{\mu}\Delta\frac{B_t}{GDP_t} + \beta\frac{B_{t-1}}{GDP_{t-1}} + \alpha + \xi_{t-1} - \frac{i_{t-1} - i_{t-1}^B}{1+i_{t-1}} + \Delta\xi_t \\
\frac{i_t - i_t^B}{1+i_t} &= \alpha + \frac{\beta}{\mu}\Delta\frac{B_t}{GDP_t} + \beta\frac{B_{t-1}}{GDP_{t-1}} + \xi_t \\
&= \alpha + \beta(\frac{1}{\mu} - 1)\Delta\frac{B_t}{GDP_t} + \beta\frac{B_t}{GDP_t} + \xi_t
\end{aligned} \tag{1.3}$$

I call equation (1.3) the *short-run convenient asset demand* curve, and it will serve as the primary estimating equation in the analyses to follow. I call equation (1.2) the *long-run convenient asset demand curve*. Equation (1.3) nests the possibility that $\mu = 1$, in which case long and short-run demand curves are identical. This paper aims to estimate both $\beta$ and $\beta(\frac{1}{\mu}-1)$.[8]

For most of this paper, my focus is on convincingly estimating the long-run slope of the convenient asset demand curve $\beta$, which most obviously relates to my estimates' applications in the macroeconomics literature. Those studies are primarily concerned with the long-run impacts of a one-time, permanent increase in the outstanding quantity of convenient assets (typically government debt). With this framework, I have clarified how the quantitative answer to that thought experiment might differ from the short-term responses that I estimate with high frequency data.

The difficulty in estimating the parameters of equation (1.3) is a familiar difficulty with estimating the slope of any demand curve: the objects $B_t$ and $\Delta B_t$ may be correlated with the structural residual $\xi_t$. This represents other factors affecting the marginal value of extrapecuniary convenience. This object likely varies at high frequencies, for instance during much-discussed flight to safety episodes. For this reason, the VIX volatility index is sometimes considered as an observable empirical proxy for $\xi$. But $\xi_t$ also likely varies at lower frequencies. Post-GFC changes to financial intermediary regulations could have made convenience either more or less valuable, depending on the intermediary.

## 1.4 Institutional Setting

Treasury bills are the most classic example of an asset that offers convenience attributes. The outstanding supply of T-bills is driven by the cash needs of the federal government. T-bill supplies exhibit substantially more high frequency variation than those of longer-term Treasury bonds, both because they are issued more frequently and because the Treasury prioritizes predictability in their supply less. To construct convenience yields, we compare T-bill rates to the fixed leg of an Overnight Indexed Swap (OIS) contract of similar maturity, as has become standard.

T-bills are a natural example of an asset that is convenient for investors to hold. As Feldhütter and Lando (2008) describe, Treasury securities have a number of qualities that might contribute to this convenience. Their accepted status as a nominally riskless instrument allows them to be used as high-quality collateral for short-term borrowing arrange-

[8] One could imagine characterizing the investor's sluggishness with more parameters. For instance, one might instead write

$$B_t - B_{t-1} = \mu\left(\frac{B_t^*}{GDP_t} - B_{t-1}^*\right) + \phi\left(B_{t-1}^* - B_{t-1}\right)$$

with $0 \leq \phi \leq 1$. In that world, convenience yield responses to permanent supply shocks may take more than one period to settle to their long-run level. In Appendix A.2, I show that an alternate estimation procedure that allows for inertia of this form produces qualitatively similar estimates for the long-run convenience yield impact, but with larger standard errors. With a simulation-based bootstrap exercise, I also show that assuming my baseline, single parameter inertia equation when the true process is a two parameter version will tend to produce estimates of $\beta$ that are biased *upward* in absolute value, but have substantially lower variance. As the central conclusion of this paper is that my estimated $\beta$ is *small* in absolute value, such a bias, if present, would not change any of this paper's qualitative conclusions. In other words, estimation strategies based on Equation (1.3) perform favorably in the bias-variance tradeoff with that more-involved alternative.

ments. A liquid secondary market means that T-bills can be sold with little price impact. T-bills also contribute favorably to the various regulatory requirements imposed on certain financial institutions.

Likewise, the fixed leg of an OIS contract is a sensible proxy for the level of short rates, without any embedded risk or convenience premia. An overnight indexed swap is a fixed-length contract wherein one side agrees to pay a fixed interest payment at contract end, and the other pays the geometric average of the federal funds rate over the same period. No principal payments are exchange at maturity, nor are any additional payments made at contract start. My proxy for the short rate, without any embedded convenience attributes, is the fixed length on this contract. As Greenwood et al. (2015a) and Sunderam (2015) argue, the fact that OIS contracts are not a way to invest principal (and are thus not a store of value) makes it unlikely that this rate would include a convenience premium. It has been shown that these rates are also unlikely to reflect any counterparty risk, as in Feldhütter and Lando (2008).[9]

Consistent with the notion that T-bills enjoy a convenience yield, T-bill yields have been lower than OIS rates of the same maturity over my sample, mostly consistently. Figure 1.1 shows the difference between the quoted, fixed leg OIS rate and the yield on a like-maturity Treasury bill, at the 4, 13, and 26 week frequencies. With the occasional exception of the 26-week maturity point late in the sample, sustained negative realizations of the convenience yields have been rare, suggesting that T-bills offer an extrapecuniary convenience benefit to investors. [10]

[9]He et al. (2022) make the interesting point that, because OIS contracts are not a store of value, they also carry a smaller capital change for large financial institutions than Treasury securities. In that sense, the yield spread between OIS and Treasury rates may include a component of *inconvenience* value as well. Nevertheless, we might still expect marginal inconvenience to rise as the stock of T-bills rises, so that this mechanism still contributes to smaller convenience yields after increases in T-bill supply. Indeed, this is what He et al. (2022) assume. I do not view the difference between these two interpretations – declining marginal convenience versus increase marginal inconvenience – as critical to this paper's results.

[10]For the 4-Week convenience yield, which is most important to this paper's results, some short-lived negative realizations are clearly tied to debt ceiling impasses. Convenience yields in those unusual periods are the subject of Cashin (2023).

Figure 1.1: T-bill Convenience Yields

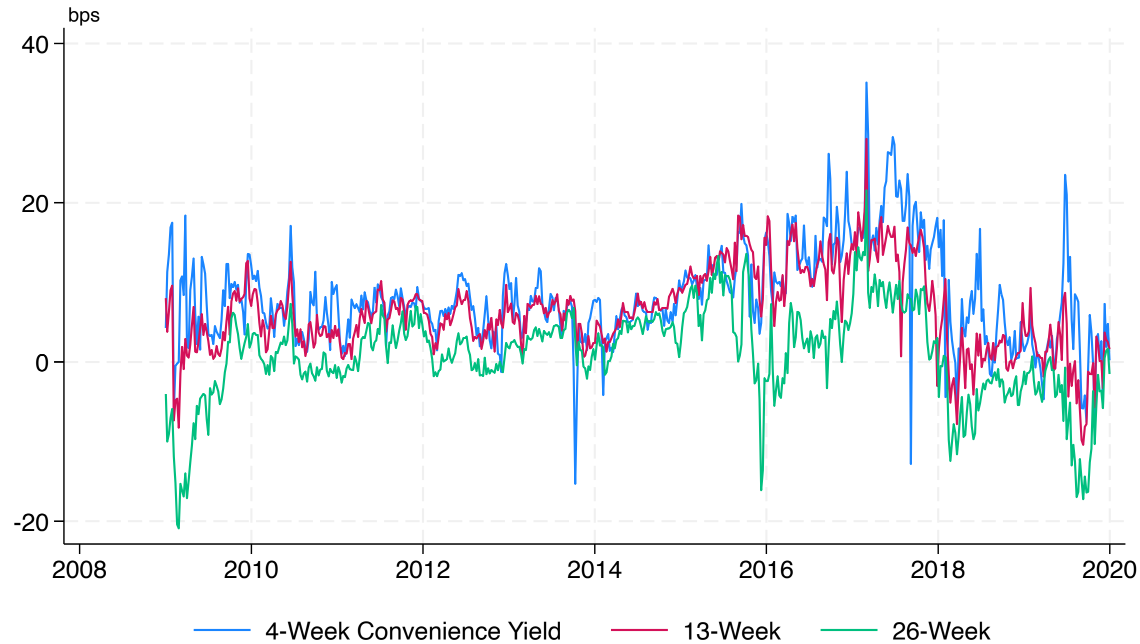


*Note:* T-bill convenience yields, at weekly frequency, from 2009-2019. Weeks are delineated by the 4-Week auction date, so that each realization is on such an auction date. Single-week spikes downward in the 4-Week maturity correspond to debt ceiling impasses. Convenience yields defined as $\text{OIS}_{m,t} - \text{T-bill}_{m,t}$ for each maturity $m$. Source: Federal Reserve Board of Governors H15 release, US Treasury via treasurydirect, Bloomberg, and Author's Calculations.

Week-to-week changes in the outstanding supply of T-bills are driven by fluctuations in the underlying financing needs of the US federal government. Fiscal outflows contribute positively to this cash need. These are government expenditures, including government consumption expenditures like payroll and military procurement, as well as transfer payments like Medicare and Social Security. Fiscal receipts contribute negatively to this cash need. These are almost entirely tax receipts (less tax refunds), including income taxes paid by individuals and corporate taxes paid by firms. The approximate timing of many of these flows, both outflows and receipts, is often predictable to an informed observer.

These financing needs can be met by one of two sources: net borrowing, or decreases in the balance of the Treasury's cash account. Since the 2008 Financial Crisis, functionally all of the Treasury's cash has been held in a reserve account at the Federal Reserve.[11] This reserve account operates in much the same way as the reserve account of a commercial bank. To finance net outflows arising from a week's expenditures exceeding a week's receipts, the Treasury can either permit the balance in this cash account to shrink, or increase debt issuance to limit the impact on the cash balance.

The Treasury has a well-established goal to keep their debt issuance quantities "regular and predictable", but enforcing perfect predictability is infeasible[12]. In the face of sometimes-

[11]See Santoro (2012) for an informative summary of changes in Treasury cash management policy around this time.

[12]Garbade (2007) is an excellent historical primer

unexpected inflows and outflows, perfectly avoiding surprises to market participants by keeping debt issuance plans unchanged would carry an unacceptably-volatile balance in the cash account.

In practice, the component of debt issuance that the Treasury permits to sometimes fluctuate unpredictably is the supply of short-term T-bills. In principle, the Treasury could also allow its issuance of longer-term *coupon* securities to fluctuate in the face of funding need surprises. However, the Treasury appears reluctant to do so. There are also major differences in issuance variability *within* the T-bill asset class. Over my post-crisis sample, T-bills are regularly issued in maturities of 4-Weeks, 13-Weeks, 26-Weeks and 52-weeks.[13] Figure 1.2 shows gross weekly issuance sizes, in billions of dollars of principal, for each of these maturities in 2016. Issuance sizes for 4-Week bills are substantially more variable than those for either of the longer maturities.[14]

Figure 1.2: Issuance Sizes of 4-Week, 13-Week, and 26-Week Bills in 2016

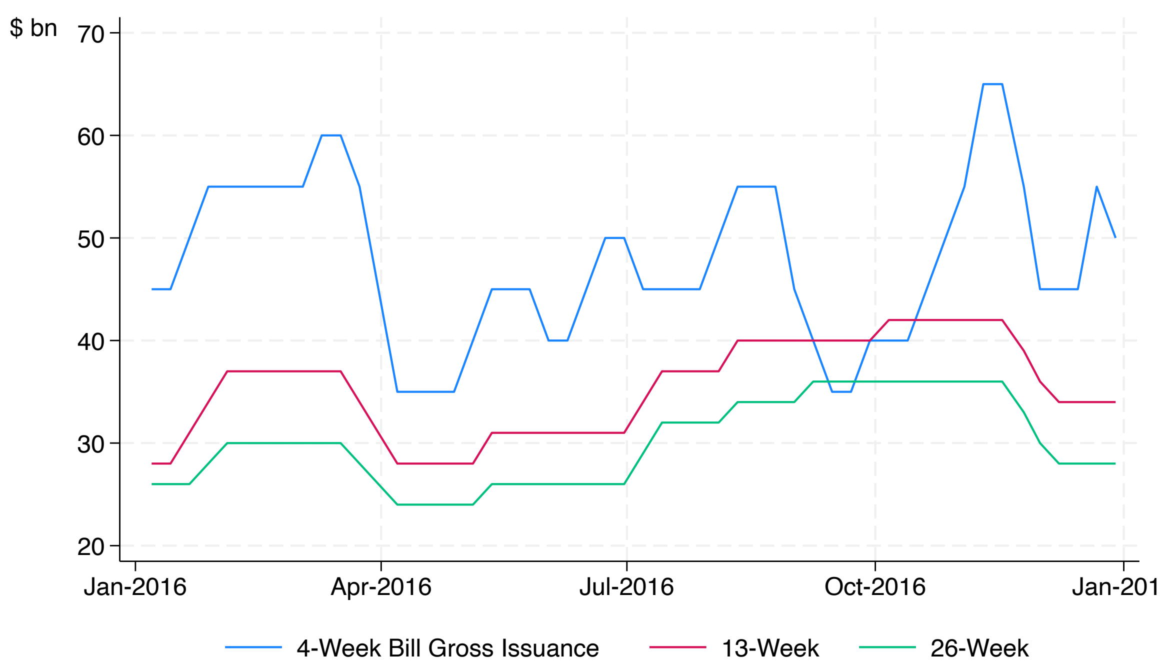


*Note:* T-bill weekly issuance sizes for regularly-issued bill maturities, in an representative example year, 2016. Source: US Treasury via treasurydirect, and Author's Calculations.

Treasury debt managers do not vary their issuance strategy to take advantage of short-term movements in demand for their debt. Figure 1.3 shows a presentation slide, taken directly from a debt management slide deck displayed prominently on the US Treasury's website. The phrasing on this slide is unambiguous: Treasury debt managers write "Treasury

[13]In 2018, the Treasury also began issuing 8-Week bills.

[14]This pattern is born out over the rest of the sample: the standard deviation of week-to-week changes in issuance for the 4-Week bill is 0.03% of then-current nominal GDP. Standard deviations for the 13-Week and 26-Week bill issuance changes are both approximately 0.013%of nominal GDP.

doesnt react to current rate levels or short-term fluctuations in demand". In Section 1.6.2, I discuss the role that this policy plays in my identification strategy.

Figure 1.3: A Summary of Treasury's Debt Issuance Philosophy


*Note:* A slide from a US Treasury presentation titled "An Overview of Treasury's Office of Debt Management". As of October 2023, this presentation is publicly available at Here. This presentation is displayed somewhat prominently alongside the Treasury's Quarterly Refinancing documents, here. Source: US Treasury via https://home.treasury.gov/system/files/276/Debt-Management-Overview.pdf.

In discussions, economists often find this practice difficult to reconcile with the "Objective" on this slide, to "Fund the government at the least cost to the taxpayer over time." It seems conceivable that timing debt issuance to coincide with higher demand for Treasury debt might lead to lower financing costs. The Treasury does not appear to share that interpretation. Rather, debt managers appear believe that keeping debt issuance uncertainty low for investors keeps the Treasury's borrowing costs lower, in the long-run. In 1982, the Treasury deputy assistant secretary summarized their logic, when saying "regularity of debt management removes a major source of market uncertainty, and assures that Treasury debt can be sold at the lowest possible interest rate consistent with market conditions at the time of sale."[15] Other portions of the slide deck featured in Figure 1.3 suggest that the justifica-

[15]Glasserman et al. (2017) features an informative and insightful discussion of the history, costs, and

tion for this strategy has not fundamentally changed in the intervening years.[16] While this justification does seem plausible, in reality it is not important for my identification strategy that this practice be *optimal* for the Treasury. It is only important that the Treasury *perceives* it as optimal, which they appear to do.

## 1.5 Estimation Challenges with T-bill Substitutes

This paper will use a notion of T-bill supply shocks to estimate the slope of convenient asset demand. Doing this is complicated by the possibility that supplies of T-bill convenience *substitutes* may vary after the shock. As D'Avernas and Vandeweyer (2023) argue with a slightly different model, this issue is likely most-severe when the Fed's overnight reverse repurchase agreement (ONRRP) facility is active. I demonstrate this intuition using a simple, stylized, linear supply and demand framework. Given this issue, I will limit my estimation sample for convenience yield effects to those periods when the ONRRP facility is *not* active.

### 1.5.1 Pre-ONRRP Estimation Problem

First, I discuss how an endogenous supply of convenience substitute assets would complicate my estimation, in the basic setting without the Federal Reserve's ONRRP facility.

While the discussion of Section 2.2 focused on T-bills as the only source of short-term convenience, there are almost surely other assets that satisfy a similar desire from investors for convenience. Two likely candidates are Federal Agency securities and certain overnight repurchase agreements that are used by money market participants to invest principle.[17] While not as liquid as T-bills, short-term agency securities are still backed by the full faith and credit of the US federal government. As such, to the extent that T-bill convenience is directly tied to the virtually-zero default risk of T-bills, then agency securities likely share that quality. Similarly, safe overnight repurchase agreements are overnight lending agreements that are *collateralized* with safe assets like US Treasury or agency securitites. Given that repurchase agreements are usually for overnight lending, default risk is low and liquidity is high (in the sense that repurchase agreements are *daily* convertible to cash via redemption). Finally, if convenience is driven by the regulatory desirability of convenient assets, then T-bills, Agencies, and repurchase agreements collateralized by Treasuries or Agencies are likely similarly convenient.[18]

If investors view the convenience properties of these alternate investments as substitutable with the convenience of T-bills, then the demand curve in equation (1.3) is misspecified. Instead, we might write

benefits of this strategy.

[16] One such slide is featured in the Appendix A.3.

[17] A repurchase agreement (repo) is a financial contract in which one party sells a security to the other, and pledges to repurchase the security at a later date. Most repos are for overnight maturities, so that repurchase occurs on the following day. One side in the transaction is fundamentally a cash investor. The other is fundamentally a cash borrower, offering the exchanged security as collateral.

[18] For instance, government-only money market mutual funds, which have a restricted set of allowable investments, can typically invest in any all three of these asset classes.

$$\frac{i_t - i_t^B}{1+i_t} = \alpha + \beta(\frac{1}{\mu} - 1)\Delta\frac{Q_t}{GDP_t} + \beta\frac{Q_t}{GDP_t} + \xi_t$$
$$Q_t = B_t + RP_t + \text{Agency}_t$$

where $RP_t$ are outstanding general collateral repurchase agreements (repo), and $\text{Agency}_t$ are outstanding Agency securities.[19]

This corresponds to the case where convenience from these three sources are *perfect* substitutes for one another. Using Wrightson T-bill issuance surprises to *estimate* the degree of substitutability between these asset types is an interesting direction for future research, but is outside the scope of this paper.[20]

In this case, using changes in T-bill supply as proxies for changes in $Q$ may not yield consistent estimates of $\beta$. Suppose that a researcher has isolated a component of T-bill supply $B_t$ that is uncorrelated with the unobservable components of convenient asset demand, $\xi_t$ (I discuss the manner in which I have done so shortly, in Section 1.6.2). For expositional purposes, suppose that $\mu = 1$, so that a 1-unit change in $\frac{Q_t}{GDP_t}$ will tend to depress convenience yields by $\beta$ immediately, all else equal. In this world, we will have

$$\frac{d}{d\text{B}_t}\left(\frac{i_t^{\text{Ref}} - i_t^{\text{B}}}{1+i_t^{\text{Ref}}}\right) = \beta\frac{dQ_t}{dB_t} = \beta\left(1 + \frac{dRP_t}{d\text{B}_t} + \frac{d\text{Agency}_t}{d\text{B}_t}\right) \leq \beta \tag{1.4}$$

This convenience yield response will depend in part on $\frac{dRP_t}{d\text{Bills}_t}$ and $\frac{d\text{Agency}_t}{d\text{Bills}_t}$. In general, we would expect both of these objects to be weakly *negative*. When T-bill supplies rise, convenience yields on both T-bills and their convenience substitutes will tend to fall. Because non-Treasury issuers of convenience substitutes may well issue opportunistically, they could respond to this price change by decreasing issuance quantities. The quantitative importance of this mechanism will depend on the size of $\frac{dRP_t}{d\text{B}_t}$ and $\frac{d\text{Agency}_t}{d\text{B}_t}$.

The upper-left panel of Figure 1.4 visualizes a stylized, linear supply and demand framework is sufficient to discussing the problem. My econometric task is to use variation in $B_t$ to identify the slope of the depicted demand curve, $\beta < 0$. In time 0, total safe asset supply equals $B_0 + SUB_0$ where $SUB_t = \text{Agency}_t + RP_t$ ("$SUB$" for "substitutes"). The supply curve is shown in red, the demand curve is in blue and equilibrium convenience yields and quantities are determined by the intersection. The other three panels of Figure 1.4 will be discussed in the following subsection.

[19] The market for all repo borrowing in the United States is extremely large. I wish to focus attention on the segment of this market that is frequented by investors looking for short-term investments of their cash principle. These are the investors that might view a repo as substitutable for a T-bill. This does not include, for instance, those market segments where the primary motivation for trading is securing a particular security as collateral. It also does not include securities lending transactions, even though the mechanics of those contracts are functionally identical to repos.

[20] For instance, we might parameterize a more general case of possibly-imperfect substitutability as

$$Q(B_t, RP_t, \text{Agency}_t) = (B_t^\rho + RP_t^\rho + \text{Agency}_t^\rho)^{\frac{1}{\rho}}$$

The discussion in this section assumes the case where $\rho = 1$. In principle, one can use this relationship as the basis for estimating $\rho$. That is the focus of Krishnamurthy and Li (2023).

Figure 1.4: Stylized Convenience Yield Responses, by ONRRP State

Before ONRRP

ONRRP $> 0$

ONRRP $= 0$

ONRRP $= 0$

*Note:* Expected influence of the Fed's Overnight Reverse Repurchase Agreement (ONRRP) facility on the convenience yield response to T-bill supply shocks. Uses a simple, linear supply and demand framework described in the text. Source: Author's calculation.

As argued above, it is unlikely that the US Treasury's T-bill issuance decisions are opportunistic. In other words, it is unlikely that T-bill supply depends directly on convenience yields. In the diagram, this is reflected in a vertical region of the supply curve

Each supply curve has a region where total quantities rise in convenience yields. When issuers of substitute assets are issuing strictly positive quantities of those assets, then the total convenient asset supply curve may be upward-sloping. This is driven by the possibly-opportunistic (i.e. rate sensitive) issuance of repurchase agreement issuers and, possibly, US government agencies like Fannie Mae, Freddie Mac, and the Federal Home Loan Banks.[21]

The upper-left panel allows us to consider the effect of an increase in T-bill supplies $B$ that are not accompanied by any shift in the demand curve. I depict this via a shift in the supply curve from $\text{Supply}_0$ to $\text{Supply}_1$.

In this setting, a T-bill supply shock can be used to identify the slope of the demand curve, $\beta$. However, the demand curve's slope is identified as

$$\beta = \frac{i_1^{Ref} - i_1^{Conv} - i_0^{Ref} - i_0^{Conv}}{Q_1 - Q_0}$$

.

In the depicted example, this is not equal to $\frac{i_1^{Ref} - i_1^{Conv} - i_0^{Ref} - i_0^{Conv}}{B_1 - B_0}$. That is – assuming that the change in T-bill supply equals the total change in convenient asset supply will give an incorrect estimate of $\beta$.

One could imagine resolving this issuing by measuring $\frac{dQ_t}{dB_t}$ directly, using high frequency data on $RP_t$ and $\text{Agency}_t$. This is straightforward in principle, but difficult in practice. High frequency data on outstanding repurchase agreement volumes has only recently become

[21] Some readers may question whether a private actor issuing a repurchase agreement reflects a *net* creation of convenient assets, in aggregate. If the asset used to collateralize the repo is itself a short-term, convenient asset, then it may not, given that the cash borrower must relinquish that collateral to the lender for the duration of the contract. However, if the repo is collateralized with a default-free, but longer-term security like a Treasury coupon note, then the repo may create convenient assets, on net. This paper is about the special investor demand for short-term convenient assets, which Krishnamurthy and Vissing-Jorgensen (2012) convincingly argue are largely separate from the convenience attributes of longer-term safe securities. A repo collateralized by a long-term Treasury, in that sense, creates a greater supply of short-term convenient assets.

publicly available. Similarly, it is challenging to directly measure the outstanding volume of US Agency discount notes at high frequencies.[22]

With this caveat on data availability in mind, I will demonstrate in Section 1.8 that publicly available data suggests $\frac{dQ_t}{dB_t} \approx 1$ for my T-bill supply shocks, in the sample period where the upper-left panel of Figure 1.4 is a sensible depiction of market supply. The next subsection outlines why the Fed's ONRRP facility makes that panel a poor approximation of reality for the remainder of the sample.

### 1.5.2 The Estimation Problem, with the Fed's ONRRP Facility

In a world where T-bills and repos have perfectly substitutable convenience, The Fed's ONRRP facility makes a portion of the market supply curve for convenient assets perfectly elastic. I demonstrate why this is problematic for my estimation, motivating an important sample restriction for my empirical analyses of the effect of T-bill supply shocks on convenience yields.

Whatever relationship exists between convenience yields and substitute safe asset issuance was undoubtedly altered by the Federal Reserve's introduction of the standing ONRRP facility. The ONRRP facility is a monetary policy implementation tool introduced in 2013 that allows cash investors in US money markets to "purchase" repurchase agreements from the Fed. The interest rate on these repo agreements is predetermined by the Federal Open Market Committee, and does not fluctuate in response to short-run market movements. The monetary policy goal of this facility is to enforce a Fed-determined *floor* on a wider range of money market rates.

Because the Federal Reserve predetermines the interest rate at which they are willing to issue repurchase agreements, the ONRRP limits the ability of repo convenience yields to move in response to shocks.[23] In a world where repo and T-bill convenience are perfect substitutes, this will also limit the responsiveness of T-bill convenience yields to shocks.[24]

The upper-right panel of Figure 1.4 shows how this process operates in our stylized setting. In this depiction, the Federal Reserve commits to providing a perfectly-elastic supply of repurchase agreements to money market participants, so that equilibrium convenience yields equal the level consistent with the Fed's choice of ONRRP rate and overall short-term rates. This practice creates a *flat* portion in the convenient asset supply curve, at the convenience yield consistent with that level. The upper-right panel also allows us to see the equilibrium division of convenient assets between T-bills $B_0$, Private Repurchase Agreements

[22] This is because, unlike for T-bills, many US agency securities are issued both via a regular, well-document auction system, and via a daily issuance "window" that is often less transparently documented.

[23] For instance, one definition of the repurchase agreement convenience yield might be $\text{IOER}_t - \text{Repo Rate}_t$, where *IOER* is the interest on excess reserves determined by the FOMC. Under this definition, by issuing repurchase agreements at an FOMC-determined rate, the Fed directly chooses *both* the convenient asset yield and the reference asset yield in this relationship.

[24] A less extreme story that is consistent with the results shown in this paper is that the marginal, "fast" T-bill investor that is able to quickly respond to T-bill supply shocks views T-bill and repo convenience as only imperfect substitutes. This is consistent with a results show in Appendix A.8, that Treasury-only money market funds provide most of the market elasticity immediately following a T-bill issuance surprise. Treasury-only funds cannot hold repurchase agreements – only Treasuries. In that world, we might expect T-bill supply shocks to have some convenience yield impact in the short-run.

and Agency Securities $Sub_0$, and ONRRP volume $ONRRP_0$.

The movement from period 0 to period 1 in the upper-right panel shows the expected effect of a T-bill supply shock on convenience yields and safe asset supplies, again without any accompanying movement in the demand curve. In a stylized world of perfect substitutability, the increase in T-bill supplies from $B_0$ to $B_1$ has no effect on the convenience yield. However, the increase in T-bill supplies decreases the issuance of repurchase agreements at the ONRRP, from $ONRRP_0$ to $ONRRP_1$.

An important practical detail of the ONRRP is that its volume cannot be strictly negative. The Federal Reserve uses the facility to *issue* repurchase agreements to money market investors (i.e. borrow cash). The Federal Reserve does not respond to low convenient asset demand by purchasing repurchase agreement (i.e. lending cash).

A scenario where takeup at the ONRRP is zero is depicted in the bottom-left panel of Figure 1.4. In this setting T-bill supply shocks once again have convenience yield effects. The bottom-right panel is functionally the same as the upper-left panel, in which the ONRRP did not exist. An ONRRP facility with zero takeup is identical, from the perspective of identifying $\beta$, to no ONRRP facility at all.

The bottom-right panel of Figure 1.4 makes the point that the bias induced by using $\frac{i_1^{Ref}-i_1^{Conv}-i_0^{Ref}-i_0^{Conv}}{B_1-B_0}$ to estimate $\beta$ instead of $\frac{i_1^{Ref}-i_1^{Conv}-i_0^{Ref}-i_0^{Conv}}{Q_1-Q_0}$ depends crucially on the *slope* of the private substitute asset supply curve. A supply curve closer to vertical implies less bias, because $Q_1 - Q_0 \approx B_1 - B_0$ when the supply curve is more vertical (i.e. as the substitute asset supply response approaches zero). But this is unhelpful in a world where $ONRRP > 0$. When ONRRP takeup is strictly positive, that necessarily introduces a perfectly elastic portion to the convenient asset supply curve, regardless of the rate sensitivity of non-Fed issuers.

### 1.5.3 Empirical Takeaways

Ignoring the endogenous supply response of T-bill substitutes to T-bill supply shocks will tend to bias our estimate of $\beta$ towards zero. This bias becomes less quantitatively important when substitute asset supplies are *less* responsive to changes in convenience yields. When the Fed's ONRRP facility is not active, this depends on the convenient asset issuance behavior of private repo issuers and Federal Agencies like the FHLBs. I will show evidence that those suppply responses do not appear large enough to meaningfully alter my estimates of $\beta$. Estimation when the ONRRP facility is active are likely more problematic, because quantity at the ONRRP is perfectly rate elastic by construction. I will avoid using data in those periods to draw any conclusions about $\beta$.

Figure 1.5 shows takeup volumes at the ONRRP facility over my sample. This is exactly equal to zero before the facility was first implemented. It is meaningfully greater than zero for a period from 2014 to 2018. It is close to zero from 2018-2020. The figure also shows the "ONRRP=0" period that I use as the ONRRP inactive period. This includes the 2018-2020 period, the ONRRP volumes are often strictly positive, but fairly close to zero.

Figure 1.5: Cash Investor Repurchase Agreement Volumes

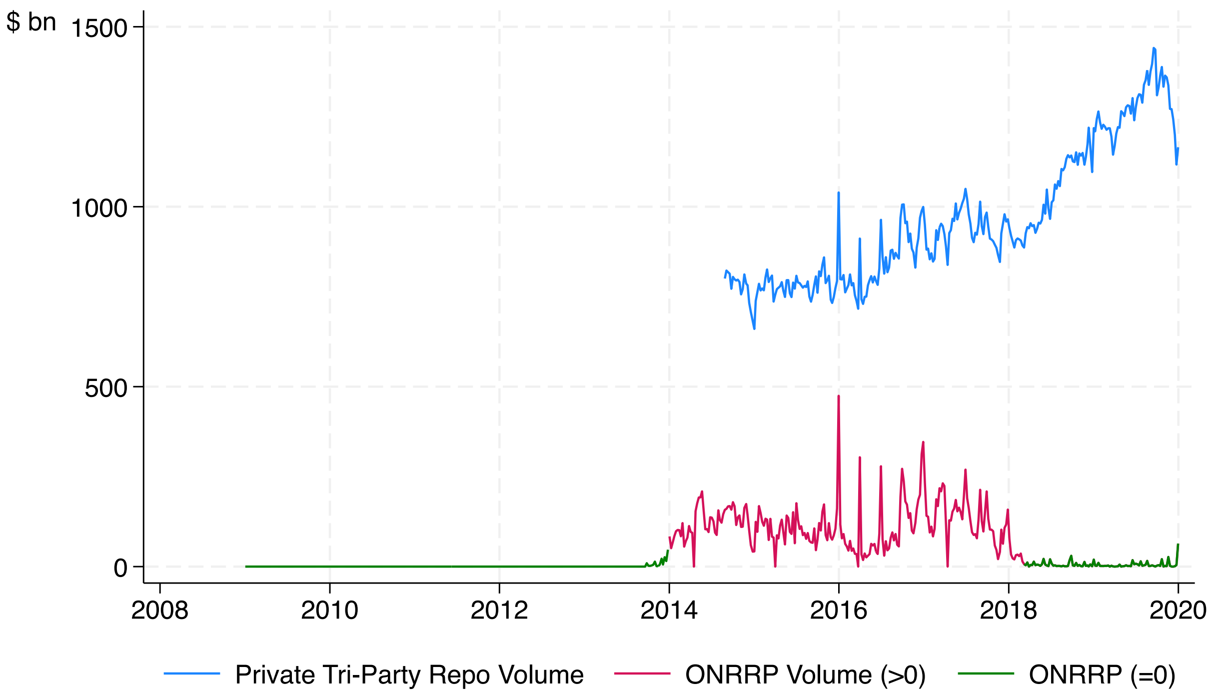


*Note:* Shows volume in private tri-party repo and at the Fed's ONRRP repo facility. Private tri-party repo volume date became publicly available via the US Treasury OFR's Repo Markets Data Release starting in August 2014. The ONRRP facility was first implemented in September 2013. Takeup became consistently large beginning in 2014, before returning again to small values for 2018-2019. Weekly realizations plotted are *daily* realizations on the day of a given week's 4-Week T-bill auction, consistent with the convention in Figure 1.1. Sources: US Treasury Office of Financial Research via its Repo Market Data Release; Federal Reserve Bank of New York via its ONRRP volume release.

## 1.6 A New Measure of T-Bill Supply Shocks

Identification in this setting requires isolating a component of week-to-week convenient asset supply that is uncorrelated with week-to-week variation in convenient asset demand. I introduce a notion of T-bill issuance *surprises*, based on projections from Wrightson's Money Market Observer. I argue that these surprises plausibly satisfy my exclusion restriction, and have several advantages over identification via *seasonal* T-bill supply variation.

### 1.6.1 Identification

I define the exclusion restriction necessary for identifying the impulse response function of T-bill convenience yields to a T-bill supply shock.

In this paper, my central empirical analyses will be estimated via a local projection, instrumental variables (LP-IV) approach. That is, I will be estimating equations of the following general form via 2SLS

$$
\begin{aligned}
\text{2nd: } \frac{i_{t+h} - i^B_{t+h}}{1+i_{t+h}} &= \alpha_h + \beta_h \frac{B_t}{GDP_t} + \gamma_h' X_t + e_{t+h} \\
\text{1st: } \frac{B_t}{GDP_t} &= \psi_1 + \zeta z_t + \phi_1' X_t + w_{t+h} \\
h \in \{0, 1, 2, \cdots, H\}
\end{aligned}
$$

where $X_t$ are control variables and $z_t$ is an instrument for T-bill supply $B_t$. This procedures describes estimating a separate second stage equation for each horizon $h$ of interest. The coefficients of interest in this estimation are the sequence of $\{\beta_h\}_{h=0}^H$.

Our impulse response function of interest is the impulse response for an increase in T-bill supply at time $t = 0$. As such, our objective is for $\beta_h$ to be a consistent estimator of that object, so that

$$
\beta_h \to \mathbb{E}\left(\frac{i_{t+h} - i^B_{t+h}}{1+i_{t+h}} | B_t = \bar{B} + 1, X_t\right) - \mathbb{E}\left(\frac{i_{t+h} - i^B_{t+h}}{1+i_{t+h}} | B_t = \bar{B}, X_t\right)
$$

That is, we want $\beta_h$ to measure the effect of a 1-unit larger realization of $B_t$ on projections of the convenience yield $h$ periods hence.

Recalling the structure of equation (1.3), where $\xi_t$ denotes unobservable factors shifting the convenient asset demand curve, consistency will following from the conditions

$$
\begin{aligned}
\mathbb{E}(B_t^\perp z_t^\perp \neq 0) \qquad (1.5) \\
\mathbb{E}(\xi_t^\perp z_t^\perp = 0) \qquad (1.6) \\
\mathbb{E}(\xi_{t+h}^\perp z_t^\perp = 0), \qquad (1.7) \\
h \in \{0, 1, \cdots, H\} \qquad (1.8)
\end{aligned}
$$

where $x^\perp_{t+j} = x_{t+j} - \text{Proj}(x_{t+j} | X_t)$.

### 1.6.2 Surprises As Instruments

In this paper, I will use T-bill issuance surprises from a highly-informed market newsletter as an instrument for T-bill supply. That is, letting $\mathbb{E}_{Priv,t-\delta}\left(\frac{B_t}{GDP_t}\right)$ be a measurement of private-sector expectations of time-$t$ T-bill supply, formed at time $t - \delta$, I will say

$$
z_t = B_t - \mathbb{E}_{Priv,t-\delta}\left(\frac{B_t}{GDP_t}\right)
$$

Researchers in this literature most-commonly discuss an endogeneity worry driven by opportunistic issuance. If the Treasury tends to respond to fluctuation in demand by issuing more or fewer T-bills, then that will tend to induce a correlation between $B_t$ and $\xi_t$.

My instrument addresses this concern by invoking the Treasury's policy from Figure 1.3 that it "doesn't react to current rate levels or short-term fluctuations in demand." I assume what seems like a natural consequence, which is that, when the Treasury surprises private actors with their T-bill issuance, they are not doing so *in response* to some fluctuation in demand for Treasury debt.

The dual facts that my analysis is at high frequencies and uses very-recent projections make this logic possible. While the Treasury does not react to "short-term" fluctuations in demand, but there is no policy prohibiting them from reacting to a lower-frequency, longer-term demand fluctuation.[25] As such, this argument for identification will tend to become less valid in lower-frequency analyses. At the yearly or quarterly frequency, the Treasury may or may not view an issuance reaction to demand conditions as "reacting to short-term fluctuations in demand". In this paper I only argue that over horizons of several weeks they certainly would.[26]

The standard approach in this literature is to use as an instrument a measure of *seasonally-expected*. This alleviates concerns about opportunistic issuance because this dimension of variation comes from some *other*, well-understood sources. Seasonal variation in T-bill supply is undeniably driven by the predictable timing of tax receipts and certain payments like Social Security, Medicare, and military payroll. Naturally, because this variation in T-bill supply is expected, that source of variation is incompatible with the impulse response logic and framework introduced above. Instead, this literature typically estimates the effect of T-bill supply changes in a single-equation approach, based on Equation (1.2).[27]

Focusing on T-bill issuance *surprises* as an alternative to resolve the opportunistic issuance problem is empirically attractive for two reasons. First, it permits the use of the impulse response estimation methods outlined above. This allows me to transparently show dynamics in the convenience yield response after a T-bill supply shock. As I will show below, these dynamics are empirically meaningful, and affect the qualitative nature of one's conclusions about whether the convenient asset demand curve is flat or steep.

Second, my instrument is more robust to the possibility that *demand* for short-term convenient assets is seasonal. If convenient asset demand has a seasonal component, then the validity of a seasonal supply instrument relies on the seasonal components of demand and supply being uncorrelated. That sort of correlation might exist by happenstance. For instance, empirical work in finance has convincingly shown that financial institutions manage their balance sheets in different ways at the ends of months or quarters, because of the predictable seasonality in when post-crisis leverage constraints are most binding[28]. This could plausibly alter those institutions' demand for T-bills are month or quarter-end.

A correlation between seasonal T-bill supply and seasonal T-bill demand could also be

[25]Indeed, in a different presentation slide included in the Appendix A.3, the Treasury notes that they will "slowly adjust to shifts in expected cost." This seems to open the door to a policy of lower-frequency opportunistic issuance.

[26]To provide a concrete example: It is the impression of some market participants that the Treasury responded to money market fund reforms in 2016 by increasing their issuance of T-bill. This may or may not have been the case. However, even if true, that plausibly would *not* be interpreted as a "short-term" fluctuation in demand, and thus may not directly contradict Treasury policy.

[27]Most common is to estimate equation (1.2) in four-week differences, with some notion of seasonal T-bill supply as an instrument.

[28]See Du et al. (2018), Anbil and Senyuz (2018), and Munyan (2015)

driven by the same calendar events that drive the Treasury's seasonal cash needs. If individuals or corporations tend to move their expected tax payments into convenient assets (which are highly liquid) prior to a major tax deadline, then liquidate those assets upon tax payment, then that will tend to induce a positive correlation between between seasonal T-bill supply and demand. T-bill supply *rises* in the lead-up to major tax deadlines, then falls afterwards. In Appendix A.1.3, I show that the structural demand residuals implied by my estimates of $\beta$ and $\mu$ in equation (1.3) imply a positive correlation between seasonal Treasury supply and demand.[29]

If T-bill issuance surprises are not driven by the Treasury responding to demand shocks, then what are they driven by? I argue below that they are driven by differences in private actors' and the Treasury's projections of near-future tax receipts and expenditures. The Treasury surprises market participants with larger issuance (i.e. more borrowing) when they expect greater near-future net *outflows* than private actors do.

In Appendix A.5, I address a remaining challenge to identification: That high frequency T-bill surprises, through their informativeness about future government cash flows, carry some additional macroeconomic information that is independently relevant to convenient asset demand. For instance, because government tax receipts are undeniably cyclical, a positive T-bill surprise today, which predicts lower government inflows in the near-future, could conceivably suggest that the economy is headed towards a recession.

I offer two results that argue this sort of mechanism should be small. First I show that, even at quarterly frequencies, surprises in government receipts and expenditures appear much more disconnected from the state of the macroeconomy than surprises in other, more-closely watched macroeconomic indicators like housing starts, unemployment, and industrial production. I do this using newly-digitized data on projections (and thus surprises) of these variables from the FOMC's Tealbooks. These results do not suggest that government expenditures and receipts are acyclical. Rather, they suggest that the cyclical component of fiscal flows is mostly *expected*, at high frequencies.

Second, I show that the T-bill issuance surprises that are likely to be most informative about the business cycle have convenience yield impulse responses that are nearly identical to those that are less informative. Because tax receipts are more-closely tied to output and income than expenditures, T-bill issuance surprises near major tax deadlines are likely most-informative for macreconomic fundamentals. The fact that convenience yield responses to T-bill surprises in those weeks is extremely similar to responses to surprises at other times suggests that any modest informativeness of T-bill surprises for the macreconomic is unlikely to come with any sizable convenience yield response over the horizons that I study. As such, it is unlikely to threaten my exclusion restriction

[29]In Appendix A.1.2, I also demonstrate a curious quality of seasonality instruments in post-crisis data: They do not share the same relationship with OLS estimates that the original logic (and original estimates, given their pre-crisis sample) of Greenwood et al. (2015b) suggests. Estimates via seasonality instruments imply smaller convenience yield changes after T-bill supply movements than OLS estimates. This conflicts with the interpretation that seasonality instruments *remove* an underlying, positive correlation between demand the supply. This suggests that the seasonality of T-bill supply and demand could very well have *changed* in some important way since 2008.

### 1.6.3 Wrightson Supply Projections

Wrightson's projection errors of 4-Week T-bill issuance will serve as our proxy for T-bill issuance surprises. Wrightson projections for 4-Week T-bill issuance offer an unusually rich dataset of private projections of our independent variable of interest. They are released soon before each 4-Week T-bill auction, and are reevaluated frequently with the most up-to-date market information.

In this paper, I use T-bill issuance projections from Wrightson ICAP as a proxy for private market expectations of the Treasury's issuance decisions. Wrightson ICAP is an independent research firm, founded in 1978, with data available by subscription. Wrightson's research team delivers high frequency projections of a number of data releases relevant to money markets and to the macroeconomy, including overnight funding rates, FOMC interest rate decisions, and employment data releases, among many others. Wrightson's chief economist is frequently quoted in articles by other business news publications, typically discussing issues of Treasury finance and Fed policy.

Wrightson's T-bill issuance projections are very granular. Each set of projections delivers a prediction about the quantity of bills to be auctioned in every one of the T-bill maturities auctioned on a regular basis. Each set of projections delivers a projection for every such auction in the following eight to nine weeks.

Wrightson's T-bill projections are updated regularly and frequently, using the most up-to-date available information relevant to Treasury financing. Wrightson typically releases its weekly Money Market Observer publication in the late evening on Sunday or very early morning on Monday. Each Money Market Observer publication includes a table of T-bill projections, for the following eight to nine weeks of auctions.

Through the text of the Money Market Observer newsletter, as well as the shorter but-more focused daily Treasury Commentary newsletter, one obtains an unusually rich lens into the thought process of the forecaster. Through this commentary, it is clear that Wrightson's projections are based on a wide range of data sources, including daily Treasury cash flow releases from the Daily Treasury Statement; the latest debt ceiling negotiations or fiscal spending debates; expertise on often-obscure seasonalities in cash flows; and the Treasury's perceived intentions, partially revealed through their recent issuance patterns.

I define my measure of a T-bill issuance surprise to be Wrightson's same-week projection error for the quantity sold at the Treasury's auction for the 4-Week T-bill. I do this for several reasons. First, the 4-Week bill is the regularly-auctioned T-bill whose issuance quantity exhibits meaningful high-frequency variation. Thus, the 4-Week bill's size is the least predictable of the regular bills that are auctioned over the entirety of my sample, leaving a shock measure with ample variance. Second, the size of the 4-Week T-bill is announced by the Treasury very soon after Wrightson publishes their issuance projections for the week. For most of my sample, 4-week bill sizes are announced on Monday afternoon, approximately half a day after Wrightson publishes their projections. This increases the chance that my shock does indeed surprise markets, when it is realized. Third, focusing on supply surprises of a regularly-auctioned bill increases my certainty that surprises in my measure are not driven by any opportunistic issuance. Regularly-scheduled bill auctions are bound by the Treasury's goal of regular and predictable issuance, which is the justification for avoiding opportunistic issuance. While I have not seen any statements by the Treasury that explicitly lead me to believe that opportunistic issuance is a greater problem for the more irregularly-auctioned cash management bills, the fact that those bills are clearly less bound by the regular and predictable principle naturally decreases one's certainty.

### 1.6.4 Time Series Properties

Consistent with their interpretation as rational market expectations, Wrightson T-bill supply projections have substantial forecasting power, at all relevant horizons. Modest positive serial correlation in the surprises, as well as a negative correlation with the projection size, will inform our choice of controls in the analyses to follow.

Wrightson's 4-Week T-bill issuance projections exhibit substantial in-sample predictive power, meaningfully outperforming simple time series techniques for projecting future issuance. In Table 1.1 below, I regress current 4-Week T-bill issuance sizes on different combinations of predictors. In the first, I regress current issuance on Wrightson's same-week issuance projection. Next, I estimate an AR(4) in 4-Week issuance sizes via OLS. Last, I include both the Wrigthtson projection and four lags of issuance. The results show three important qualities. First, issuance sizes are fairly persistent, so that a simple AR(4) process has a within-sample $R^2$ measure of 79%. Second, even given that, the Wrightson projections substantially outperform the simple AR(4), reaching an $R^2$ measure of 87%. In order words, the Wrightson projections are able to explain 38% of the deviation from the mean issuance size that the AR(4) left unexplained. Finally, lags of issuance offer essentially no predictive power when added to the regression alongside the Wrightson projections.

Table 1.1: Wrightson In-Sample Predictive Power, Same-Week

| | 4W Issuance | 4W Issuance | 4W Issuance |
|---|---|---|---|
| 4W Proj | $0.94^{***}$ | | $0.78^{***}$ |
| | 0.02 | | 0.08 |
| 4W Issuance$_{t-1}$ | | $1.10^{***}$ | $0.26^{**}$ |
| | | 0.05 | 0.11 |
| 4W Issuance$_{t-2}$ | | $-0.24^{***}$ | $-0.07^{*}$ |
| | | 0.07 | 0.04 |
| 4W Issuance$_{t-3}$ | | $-0.06$ | $-0.07^{*}$ |
| | | 0.05 | 0.04 |
| 4W Issuance$_{t-4}$ | | 0.05 | 0.04 |
| | | 0.04 | 0.03 |
| Constant | $2.33^{***}$ | $5.43^{***}$ | $2.41^{***}$ |
| | 0.68 | 1.25 | 0.80 |
| N | 575 | 575 | 575 |
| $R^2$ | 0.87 | 0.79 | 0.87 |

*Note:* Shows OLS regression estimates, with the Treasury's weekly realized 4-Week T-bill issuance size as the dependent variable. Independent variables are some combination of Wrightson's same-week projection of the same, and weekly lagged realizations of 4-Week T-bill issuance. Standard errors are heteroskedasticity-robust. Sample is weekly, 2009-2019. Sources: Wrightson ICAP, US Treasury via treasurydirect, and author's calculations. *** 0.01, **0.05, *0.10

Wrightson's projections' superior forecasting power is also evident in out-of-sample forecasting exercises. In Table 1.2 below, I perform such an exercise over the post-crisis sample. In each exercise, at each date in the sample, I form an out-of-sample projection of the following week's 4-Week T-bill issuance via estimating a model using only data available from the previous week. In RW, I assume a random walk in 4-Week T-bill issuances, with the projection for next week's issuance always equaling the previous week's. In AR(4), I estimated an AR(4) in issuance, using data from the previous week and earlier, and use the fitted model to project next week's issuance. In AR(4)+FE, I augment the AR(4) model with week-of-year fixed effects, to capture a degree of seasonality. In Wrightson, the projection of next week's issuance equals Wrightson's same-week forecast. In Wrightson+, I fit an OLS model predicting next week's issuance using Wrightson's projection, four lags of issuance, and four lags of Wrightson's projection errors. The results show that all of the specifications with Wrightson projections as an input have RMSE values substantially lower than the other specifications. Additionally, the out-of-sample forecasting power from taking Wrighton's projections as given (as suggested by theory, if they are indeed rational expectations) is nearly the same as that from using Wrightson's projections as inputs to a linear predictive model.

Table 1.2: Wrightson Out-of-Sample Predictive Power, Same-Week

| Maturity (Weeks) | RW | AR(4) | AR(4)+FE | Wrightson | Wrightson+ |
|---|---|---|---|---|---|
| 4 | 5.46 | 5.35 | 5.21 | 4.10 | 4.08 |
| 13 | 1.38 | 1.46 | 1.43 | 1.37 | 1.39 |
| 26 | 1.45 | 1.61 | 1.58 | 1.31 | 1.48 |

*Note:* Shows root mean squared errors of several out-of-sample forecasting exercises for T-bill issuance, of most regularly-issued maturities. Units on T-bill issuance are in billions of dollars in principal. "RW" projects this week's issuance to be equal to last week's issuance of the same. "AR(4)" and "AR(4)+FE" perform out-of-sample forecasting via an AR(4) in issuance, with or without week-of-year fixed effects to capture a degree of seasonality. "Wrightson" takes Wrightson's same-week projection of issuance as the projection. "Wrightson+" includes Wrightson's same-week projection as one of several independent variables in a regression-based out-of-sample forecast (see text for details). Sample is weekly, from 2009-2019. Sources: Wrightson ICAP via its Money Market Observer publication, US Treasury via treasurydirect, and author's calculations.

There are some modest ways in which the Wrightson projection errors, in my finite sample, behave differently than the rational expectations of theory. Chiefly Table 1.3 below shows OLS estimates from regressing same-week 4-Week T-bill projection errors on the size of the Wrightson projection itself; lags of projection errors; and lags of 4-Week issuance. In theory, and in a large sample, perfect rational expectation errors should be uncorrelated with any information that should have been known to Wrightson at the time they made their projections. In practice, we find statistically significant positive regression coefficients on the size of the projection, and negative coefficients on the previous week's issuance. This is consistent with there being a modest average tendency, in this sample, for Wrightson to overestimate increases in issuance size from this week to the next.

In light of this evidence, I will conduct robustness checks of the analyses to follow that control for lagged issuance surprises and the size of Wrightson's projection. In practice, I will find that all of the results of this paper are robust to the inclusion of those controls. The $R^2$ measures from these OLS regressions indicate why. Even though some of the regression coefficients on time-$t$ information are statistically significant, those measures collectively explain only a very small fraction of the variance in the 4-Week issuance surprises themselves. As such, their inclusion as controls has little bearing on the results.

Table 1.3: Wrightson Autocorrelations

| | 4W Error | 4W Error | 4W Error | 4W Error |
|---|---|---|---|---|
| 4W Error$_{t-1}$ | 0.09** | 0.11*** | | 0.03 |
| | 0.04 | 0.04 | | 0.07 |
| 4W Error$_{t-2}$ | | | | −0.03 |
| | | | | 0.06 |
| 4W Error$_{t-3}$ | | | | 0.02 |
| | | | | 0.05 |
| 4W Error$_{t-4}$ | | | | −0.02 |
| | | | | 0.04 |
| 4W Projection | | −0.07*** | −0.22*** | −0.22** |
| | | 0.02 | 0.08 | 0.09 |
| 4W Issuance$_{t-1}$ | | | 0.26** | 0.23* |
| | | | 0.11 | 0.13 |
| 4W Issuance$_{t-2}$ | | | −0.07* | −0.03 |
| | | | 0.04 | 0.07 |
| 4W Issuance$_{t-3}$ | | | −0.07* | −0.10* |
| | | | 0.04 | 0.05 |
| 4W Issuance$_{t-4}$ | | | 0.04 | 0.05 |
| | | | 0.03 | 0.04 |
| N | 575 | 575 | 575 | 575 |
| $R^2$ | 0.01 | 0.04 | 0.07 | 0.07 |

*Note:* Shows OLS regression estimates with "Wrightson surprises" in 4-Week T-bill issuance as the dependent variable. Independent variables include lags of the same, the Wrightson *projection* of that week's 4-Week T-bill issuance, and lags of 4-Week issuance. Standard errors are heteroskedasticity-robust. Sample is weekly, from 2009-2019. Sources: Wrightson ICAP, via its Money Market Observer newletter, US Treasury via treasurydirect, author's calculations. *** 0.01, **0.05, *0.10

Consistent with the exclusion restriction requiring that these surprises be uncorrelated with changes in convenient asset demand conditions, I find that the 4-Week T-bill surprises are uncorrelated with a host of weekly variables that might serve as proxies for convenient asset demand. Those include current and lagged realizations of the VIX index; current and lagged changes in the VIX index; the same for the MOVE interest rate volatility index; and current and lagged log differences in the S&P 500 index. There is some modest serial correlation in the errors themselves, as the previous results suggested.

Table 1.4: Wrightson Correlations

| | |
|---|---|
| VIX | 0.02 |
| MOVE | 0.04 |
| L.4w Error | 0.09 |
| S4.VIX | −0.02 |
| L.VIX | 0.04 |
| L.S.VIX | 0.07 |
| L.MOVE | 0.05 |
| S.MOVE | −0.04 |
| S.Log SP 500 | 0.05 |
| L.S.Log SP 500 | −0.08 |

*Note:* Shows simple correlation coefficients between Wrightson 4-Week T-bill surprises and other variables of interest. Sample is weekly, 2009-2019. Sources: Wrightson ICAP via its Money Market Observer newsletter, US Treasury via treasurydirect, Federal Reserve Bank of St Louis via FRED.

### 1.6.5 Understanding the Shocks

For most of this paper's post-crisis sample, Wrightson projection errors in 4-Week T-bill issuance can be explained as reflecting differences in net cash flow projections between the Treasury and Wrightson. Many of the largest projection errors occur in the months surrounding post-crisis debt ceiling showdowns. During these periods, some surprises are better explained by differences in technical debt management decisions in these periods.

Figure 1.6 shows my sample of Wrightson issuance surprises in two ways. In the left panel, I show the time series of surprises, split into two subsamples. A "Debt Ceiling Period" is defined as a week within 40 days (either before or after) the final resolution of a debt ceiling episode. The "Not Debt Ceiling Period" subsample includes all other weeks. In the right panel, I show the same data via a histogram, with widths equal to \$1 billion. In both figures, and everywhere in this paper, a surprise is defined as

$$\varepsilon_t^B = \text{Actual 4-Week Bill Issuance}_t - \mathbb{E}_{\text{Wrightson},t-\delta}\left(\text{4-Week Bill Issuance}_t\right)$$

That is, a positive value indicates T-bill issuance that is *greater* (i.e. more borrowing) than Wrightson expected.

Figure 1.6: Wrightson Surprises, 4-Week T-bill

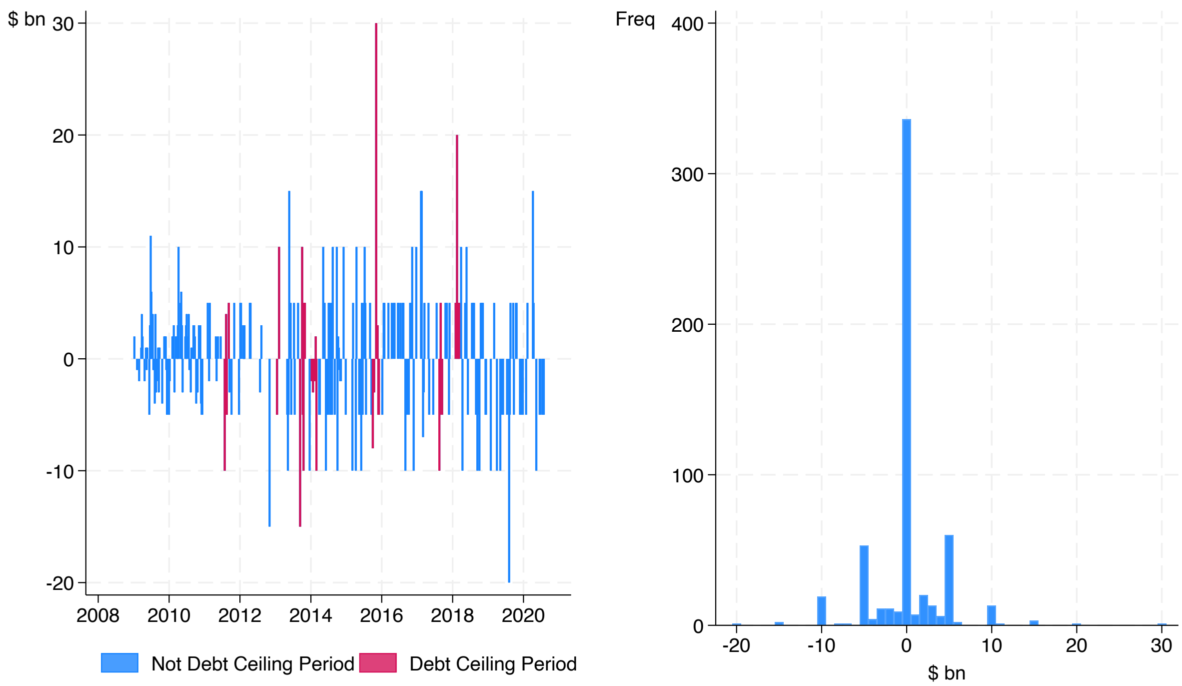


*Note:* Time series and histogram of Wrightson surprises, defined in the text as same-week Wrightson projection errors for the 4-Week T-bill issuance size. The left panel shows weekly realizations of this object. The right panel shows a histogram, with bin width of $1 billion. The left panel highlights surprises that occur around debt ceiling reset periods. As described in the text, the underlying *source* of surprises in these periods can differ from those in other weeks. Sources: Wrightson ICAP via its Money Market Observer newsletter, author's calculations.

While the most common outcome is *no* surprise in 4-Week bill issuance, there is still substantial surprise issuance variation over the post-crisis sample. My sample from January 2009 to December 2019 includes 575 weeks of 4-Week bill issuance. Of these, 336 weeks had zero surprise in the 4-Week bill issuance size of that week and 239 weeks had a nonzero surprise. Of 239 surprise weeks, 194 surprises had an absolute value of $5 billion or smaller.

Outside of debt ceiling episodes, Wrightson's discussion of surprises in Treasury's 4-Week T-bill issuance is consistent with the notion that surprises arise from differences between their net cash flow projections and the Treasury's. These discussions are visible in the daily "Treasury Commentary" that Wrightson publishes each morning, in which they sometimes discuss the previous day's developments in fiscal financing. For instance, after October 4th, 2010's +$4 billion surprise, Wrightson writes

> We were surprised by the lack of any increase in the 4-week bills this week, but we were also surprised by a relatively strong Treasury cash balance for Friday...The Treasury's own cash flow forecasts are probably slightly more optimistic than ours, as we doubt that the Treasury would cut things quite as close as our current projections imply.

This sort of discussion is not limited to smaller issuance surprises like October 4th 2010's.

For instance, after February 13th 2017's +$15 billion surprise (one of the eight largest surprises in the sample), Wrightson writes

> The Treasury continues to surprise us by not cutting its bill offering sizes more aggressively...The smaller cutbacks in the 4-week bills thus far clearly indicate that the Treasury has a very different near-term cash balance forecast than we do.

In Appendix A.7, I show empirical evidence consistent with this understanding of the T-bill issuance surprises. Namely, I show that Wrightson T-bill issuance surprises today predict almost-exactly offsetting changes in T-bill supply and net Treasury cash flows over the following weeks. This is consistent with a simple interpretation that the Treasury borrows $1 billion more over the next several than private agents expect when net outflows are approximately $1 billion larger than those agents expect.

Wrightson's Treasury commentaries also show that, while not necessarily reflective of the *average* surprise, some of the largest surprises are driven by other, more technical debt management factors. This most commonly occurs around (either before or after) resolutions of debt ceiling episodes, when the Treasury has the widest array of technical concerns to consider, such as the amount of "extraordinary accounting measures" to utilize in order to permit more marketable debt issuance. For instance, after the January 14, 2013 issuance surprise (which was several weeks before the resolution of a debt ceiling episode on February 4, 2013) they write

> The Treasury surprised us by cutting the size of the 4-week bill by $5 billion, to $35 billion this week...The unexpected cutback in the 4-week bill presumably reflects an effort to preserve flexibility as debt limit problems approach....There is a lot of uncertainty on two different levels in the outlook. The Treasury's cash flow over the coming weeks will be more difficult to predict than usual due to all of this year's tax changes, and we don't know how much of the Treasury's remaining accounting flexibility is being absorbed by nonmarketable debt issuance to government trust funds.

Indeed, the largest surprise in my sample is a $+\$30 billion$ surprise announced on November 2, 2015. This surprise comes *after* the conclusion of the fall 2015 debt ceiling episode, which effectively resolved when the US House and Senate passed HR1314 on October 30th. Wrightson writes of the surprise

> We had speculated in this week's issue of the Money Market Observer that the Treasury might allow its cash balances to run higher than usual in the months ahead in order to accommodate a more rapid rebound in the supply bills...Of course, the fact that bill supplies would be rebuilt faster than originally expected was old news by the time the Treasury's new quarterly borrowing projections were announced yesterday afternoon, as the Treasury had already announced a startling $45 billion increase in the size of today's 4-week bill auction.

These two debt ceiling-adjacent surprises, which are emblematic of much of the Wrightson discussions around these episodes, show two occassions when a T-bill issuance surprise was driven by technical, debt ceiling-related factors. In the first, Wrightson's and the Treasury's understanding of debt ceiling headroom likely differed – possibly as a result of nonmarketable debt issuance, about which the Treasury could plausibly have superior, private information. In the second, the Treasury decided to increase T-bill supplies more quickly than expected, after a debt ceiling-related low point.

Importantly, in *neither* case does Wrightson give an indication that the reason Treasury's technical decisionmaking differed from expectations was driven by something obviously relevant for *demand*. Surely, new information about the default probability of the US government is revealed during debt ceiling episodes, and that information shifts private demand for T-bills. However, these two episodes (and others around debt ceiling episodes) seem to reflect different debt management decisions, *given* the current state of government default probability. In the second case, for example, all uncertainty about government default was effectively resolved the previous week, when the House and Senate passed the corresponding debt limit increase bill.[30]

## 1.7 Core Empirical Results

My local projection estimates suggest that Wrightson T-bill supply surprises have a powerful, but transitory, effect on convenience yields. The convenience yield effect is more transitory than the T-bill supply effect, suggesting that flow effects are stronger than stock effects in this setting. I use the same empirical moments as the local projection approach, with a multi-equation GMM estimation, to estimate these separate flow and stock effects, arriving at this paper's estimate of the long and short-run convenience yield impacts of a T-bill supply change.

### 1.7.1 Methodology

I first provide additional details about how the LP-IV setup introduced in Section 1.6.1 will be used in this setting, to estimate impulse responses to T-bill supply shocks. I then detail how I use the simple structure of Equation (1.3) and the shape of those impulse responses to estimate separate stock and flow effects.

In this section, I estimate impulse response functions to a T-bill supply surprise. I do this via a 2SLS, LP-IV specifications of the form

$$\text{2nd: } Y_{t+h} = \alpha_{2,h} + \gamma_h \frac{B_t}{GDP_t} + \phi'_{2,h} X_{t-\delta} + e_{t+h} \tag{1.9}$$

$$\text{1st: } \frac{B_t}{GDP_t} = \alpha_1 + \chi \varepsilon_t^B + \phi'_{1,h} X_{t-\delta} + w_t \tag{1.10}$$

[30]It is also worth noting that, in the first of these examples, differences in future projected cash flows may indeed have driven the surprise, as Wrightson discusses "all of this year's tax changes" creating cash flow uncertainty.

where $\frac{B_t}{GDP_t}$ is T-bill supply as a percentage of GDP; $\varepsilon_t^B$ is the Wrightson T-bill supply surprise variable; and $X_{t-\delta}$ is a list of control variables. The lefthand side variable $Y$ indicates some dependent variable of interest. In this section, $Y$ is either a measure of the T-bill convenience yield, or a measure of T-bill supply. The time indices in equations (1.9) and (1.10) indicate weeks.

The vector of control variables $X_{t-\delta}$ will vary across specifications, but will only ever include objects which should be publicly known and observable prior to the realization of the shock $\varepsilon_t^B$. The time subscripting $t-\delta$ is meant to signify this. In all specifications, $X_{t-\delta}$ will include lags of the dependent variable of interest $Y$, so that our impulse response estimates are estimated with high frequency variation in $Y$ alone. When estimating the response of a variable that varies at the daily frequency, $X_{t-\delta}$ will include lagged daily observations for $Y$ in the days before the T-bill auction announcement when $\varepsilon_t^B$ becomes observed. While $X_{t-\delta}$ does not include lagged realizations of the T-bill surprise in the baseline specification, robustness results in Appendix A.10 do. Their inclusion has no material impact on the results. The vector $X_{t-\delta}$ also includes Wrightson's expected future T-bill supply in each of the following 6 weeks. These are expectations formed before the realization of the surprise in time $t$.

After presenting these results, I will use the same empirical moments to estimate the parameters $\beta$ and $\beta(\frac{1}{\mu}-1)$ via GMM, by imposing the simple dynamics implied by equation (1.3). To do this, I construct moment conditions of the form

$$\mathbb{E}\left(e_{t+h}X_{t-\delta}\right) = 0 \tag{1.11}$$

$$\mathbb{E}\left(e_{t+h}\varepsilon_t^B\right) = 0 \tag{1.12}$$

$$\tag{1.13}$$

where

$$e_{t+h} = \frac{i_{t+h} - i_{t+h}^B}{1+i_{t+h}} - (\alpha_{2,h} + \phi_{2,h}' X_{t-\delta} - \frac{\beta}{\mu}\Delta\frac{B_{t+h}}{GDP_{t+h}} - \beta\frac{B_{t+h}}{GDP_{t+h}})$$

for $h \in \{0, 1, 2, \cdots, H\}$. In my baseline application, $H = 9$. These moment conditions are constructed so that the structural GMM estimation and the reduced form LP-IV approach are estimated using the same moments in the data. The difference between the two approaches is that the GMM estimation will require the shapes of the estimated responses of T-bill supply $B_{t+h}$ and convenience yields to be related via the estimated parameters $\beta$ and $\frac{\beta}{\mu}$. As I will show below, this imposed structure allows a good fit.

When $h \geq 1$, this estimation is overidentified. In the GMM procedure, I perform a two-step, optimal GMM estimation using a Newey-West, autocorrelation-consistent approach to estimate the variance of the estimated moment conditions. Results are similar when I instead use a one-step GMM approach, with moment conditions weighted by the identity matrix.

While this approach seems like a natural way to estimate separate flow and stock effects using these shocks and data, I note that the LP-IV estimates themselves are more robust to model misspecification than the structural GMM approach that follows. LP-IV estimates in this setting will be consistent estimates of the true underlying impulse response function, provided that the Wrightson surprises are uncorrelated with other shocks to convenient asset demand over the impulse horizons. To interpret the GMM estimates as structural, we must also assume that the convenience yield response operate strictly through the channel that Equation (1.3) allows – namely, the contemporaneous stock and flow of T-bill supply.

### 1.7.2 LP-IV Results: Future Quantity

In order to interpret the convenience yield impulse response estimates to follow, we must first understand how a Wrightson surprise predicts both levels and changes of future T-bill supply over the response horizon. These results show that the response of T-bill supply rises for several weeks, before falling back to its pre-surprise level.

To reach these estimates, I estimate equations (1.9) and (1.10) by 2SLS, replacing $Y$ with $\frac{\text{Bills}}{\text{GDP}}$. The results are shown in Figure 1.7 below. I depict 90% confidence intervals along with the point estimates. Given recent results from Olea and Plagborg-Møller (2021) that simple, heteroskedasticity-robust standard errors are sufficient in local projection settings such as this, the standard errors shown here are heteroskedasticity robust. I have found that including Newey-West, autocorrelation-consistent standard errors instead tends to *shrink* the estimated confidence intervals, making this choice conservative.

Figure 1.7 shows a hypothetical "Benchmark" impulse response, which is the response we would see if a 4-Week Wrightson T-bill surprise in week 0 predicted *no* other changes in Treasury issuance over the response horizon. A 4-Week bill contributes positively to total T-bill supply for four weeks. After that point, the bill matures and leaves the outstanding stock.

An impulse response function like this "Simple Benchmark" is resoundingly rejected by the data. Instead, this response suggests that a positive \$10 billion surprise in T-bill issuance today predicts additional increases in T-bill supply for the next several weeks. Like in the "Simple Benchmark", supply tends to begin falling at week 4, when 4-Week bills issued in week 0 mature. At its peak in week 3, the T-bill supply response is \$25 billion.[31]

We can understand this result by recalling the Treasury's stated objective of "regular and predictable" issuance. In a world where the Treasury wishes to keep T-bill issuance predictable, but where unexpected cash flow needs necessitate occasionally surprising market participants, the Treasury should attempt to *smooth* the effect of unexpected cash flows on T-bill issuance. Doing so would tend to lead to smaller issuance surprises, week-to-week.

[31] In this impulse response and all of those to follow, the initial shock is normalized to raise T-bill supply by \$10 billion this week. As is standard in this literature, the underlying estimation normalizes T-bill supplies by then-current nominal GDP. I report values in billions of dollars by converting the empirical estimates to dollars, using nominal GDP for 2017Q1.

Figure 1.7: Impulse Response of Future T-bill Supply to T-bill Issuance Surprise

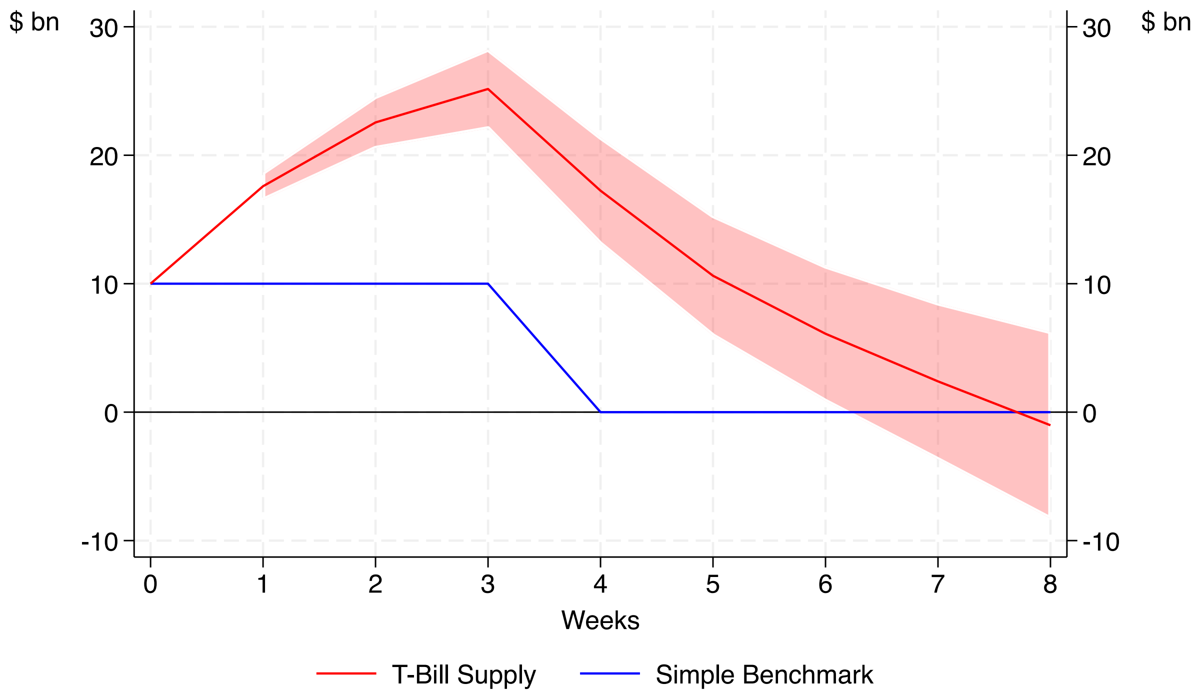


*Note:* Shows estimate impulse response function of future T-bill supply to a T-bill supply shock in period 0, defined as a Wrightson surprise that elevates period 0 T-bill supplies by $10 billion. Estimates and standard errors computed via LP-IV, as described in the text. Confidence bands are 90%, computed via heteroskedasticity-robust standard errors. "Simple Benchmark" line depicts the hypothetical response one would expect, if a Wrightson surprise in period 0 predicted no other changes in T-bill issuance over the horizon. Sources: Wrightson ICAP via its Money Market Observer Newsletter, US Treasury via treasurydirect, author's calculations.

### 1.7.3 LP-IV Results: T-bill Convenience Yields

I now turn to the principle empirical result of this paper, and estimate the convenience yield response to a T-bill supply shock. The response of T-bill convenience yields to a surprise T-bill issuance is large and statistically significant in the first weeks following the surprise. However, the convenience yield response decays to zero much sooner than T-bill supply itself returns to zero. Indeed, convenience yields four weeks after the initial shock are modestly *larger* four weeks after the initial surprise, even though T-bill supply is still substantially elevated at this time.

A T-bill issuance surprise creates a large and statistically significant response in the convenience yield in the weeks after the shock. These results are shown in Figure 1.8. At the time of issuance, a $10 billion T-bill surprise comes with a 0.82bp decline in the size of the convenience yield.

In the most-recent estimate in the literature with a similar sample period, D'Avernas and Vandeweyer (2023) report that a $100 billion increase in T-bill supply depresses T-bill convenience yields by 4bp. That study, which does not use an impulse response framework, does not differentiate between the short and long-run convenience yield response. For ease of comparing my estimates to the literature, I include in Figure 1.8 a "Benchmark" line, which applies the most-natural interpretation of those earlier results to construct a hypothetical T-bill convenience yield impulse response. This green line is equal to 0.04bp, multiplied by the then-current realization of the impulse response for T-bill supply to the shock. As Figure 1.8 shows, my point estimate for the on-impact response of convenience yields is over twice as large as that benchmark result would suggest.

The shape of the rest of the impulse response function supports the notion that the short-run impact of a change in T-bill supply is much larger than the long-run impact. In addition to giving a sense of the magnitude of my results, the "Benchmark" line also shows the shape of the impulse response that we would expect, if the then-current outstanding stock of T-bills is the most important determinant of the convenience yield. That is not what the estimates suggest. Rather, we see the peak response at week 1 – not at week three, when the stock of T-bills is at its highest point.

Indeed, the point estimate for the convenience yield response is positive for several weeks, beginning in week 4. Note also from the Figure that, at week 4, T-bill supplies are still well-elevated, relative to their level before the surprise. The assumed structural equation (1.3) has no difficulty explaining this response. In week 4, T-bill supply begins to fall, even though the stock is still elevated. If flow effects are much stronger than stock effects, then a positive convenience yield response at those horizons could well result.[32]

Figure 1.8 summarizes my principal empirical result, which the rest of this paper will work to better understand - that T-bill supply shocks appear to have a transitory effect that is substantially greater than the the long-term effect that these impulse responses allow. The GMM estimation that follows formalizes this intuition, and delivers conclusions about the permissible sizes of any long-run supply effect, to be consistent with these impulse responses.

[32]Nonetheless, it is worth noting what sorts of market frictions might cause this type of response. This sort of effect sees most-consistent with a model where T-bill investors view a rollover of their maturing T-bill holdings into the newest issues as less costly than moving their capital into some alternate investment.

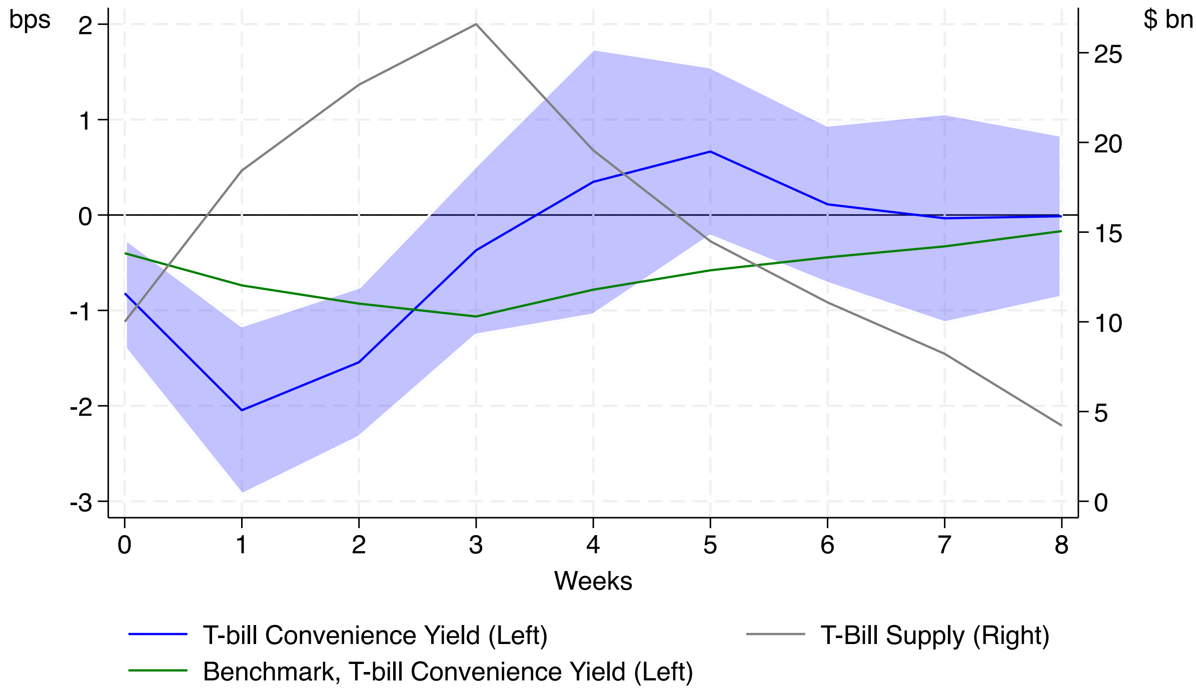


Figure 1.8: Impulse Response of T-bill Convenience Yield

*Note:* Show estimated impulse response function of 4-Week T-bill convenience yields to a T-bill supply shock, defined as a week 0 Wrightson surprise. Results are estimated via LP-IV. Confidence bands are 90% and standard errors are heteroskedasticity-robust. "T-bill supply" line depicts the same impulse response as Figure 1.7, for reference. "Benchmark" shows impulse response one would expect, if T-bill supply changes move convenience yields immediately and permanently by 0.4bp / $10bn, in line with an estimate from D'Avernas and Vandeweyer (2023). Sources: Wrightson ICAP via its Money Market Observer newsletter, Bloomberg, Federal Reserve Board of Governors via its H15 release, author's calculations.

### 1.7.4 GMM Results

I use the moments from the LP-IV setup above to estimate separate flow and stock effects via GMM, as described in Section 1.7.1. I estimate that the same-week effect of an increase in T-bill supplies of $100 billion depresses convenience yields by 10.4bps. I estimate a long-run effect of only 1.1bps/$100 billion, with a tight enough confidence band to reject long-term effects larger than 3.0bps/$100 billion.

Figure 1.9 below shows the fit of the structural estimates in describing the shape of the LP-IV impulse response. The blue line of the figure reproduces the convenience yield impulse response in Figure 1.8. This is an *unrestricted* estimated fit for the impulse response, in the sense that a separate estimate of $\hat{\beta}_h$ is estimated, to fit each horizon $h$. The red line in Figure 1.9 uses the estimated impulse response for T-bill supply; my estimate of the flow effect; and my estimate of stock effect to describe the impulse response. This is a *restricted* estimate for the impulse response.

The close similarity between the blue line of Figure 1.9 and the red line representing the GMM fit show that the fit is quite good. In other words, imposing the structure of Equation (1.3) sacrifices little information, compared to the unrestricted LP-IV estimate, so that my two estimates of the flow and stock effect are an effective way of summarizing the dynamics suggested by the response. As another assessment of model fit, I am unable to reject the null hypotheses that all of the moment conditions are true. That chi-squared hypothesis test carries a p-value of 0.24. This test is possible because the GMM model is overidentified.

Figure 1.9: GMM Model Fit to LP-IV Moments

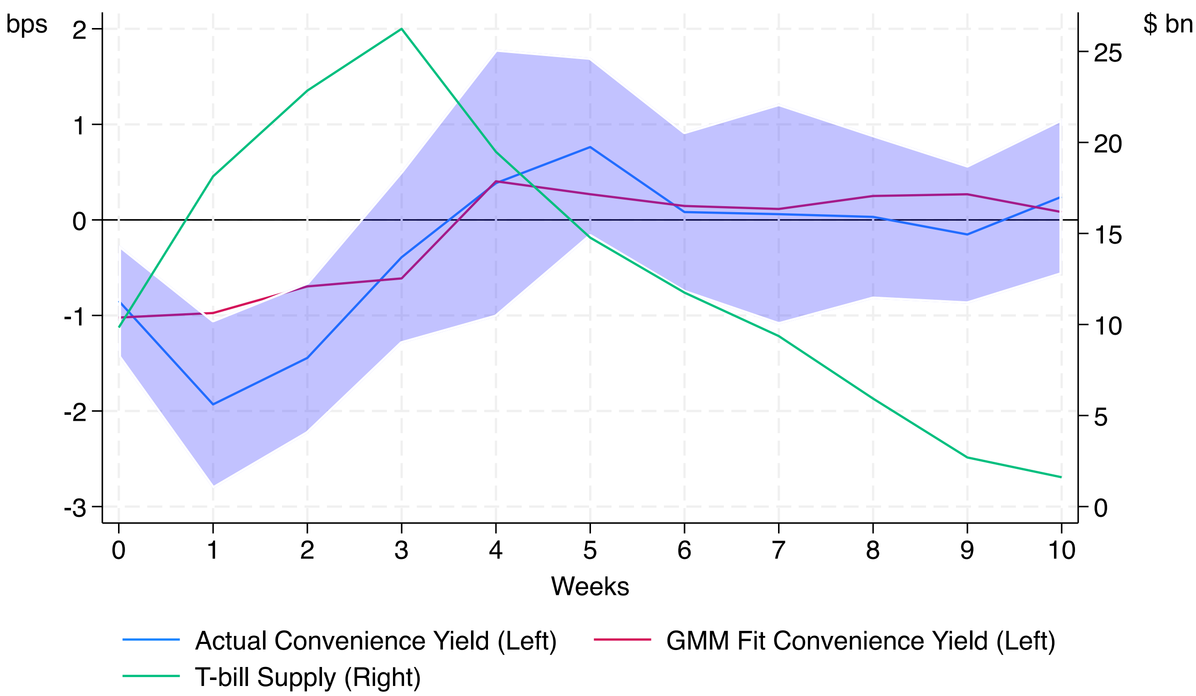


*Note:* The red line shows the two-parameter GMM fit of the empirical moments from the LP-IV estimates depicted in Figure 1.8, as described in the text. The blue and green lines shown are the same as in that figure. Sources: Wrightson ICAP via its Money Market Observer newsletter, Bloomberg, Federal Reserve Board of Governors via its H15 release, author's calculations.

The GMM procedure produces estimates of the stock effect that are small and estimates of a flow effect that are large. The point estimate for $\beta$, the stock effect in the notation of equation (1.3), is -1.14bp / \$100 billion. The point estimate for the flow effect is -9.25bp / \$100 billion.

In Figure 1.10, the point estimate and 90% confidence interval for the stock effect is shown as "Long Baseline". The point estimate and associated confidence interval of the flow effect – $\beta(\frac{1}{\mu-1})$ in Equation (1.3) – is shown as "Short Baseline". The D'Avernas and Vandeweyer (2023) estimate of -4bp/\$100 billion is show as a black "x", under the confidence interval of the "Long" effect. The means that my point estimates can reject a stock effect of -4bp / \$100 billion.

Figure 1.10: Point Estimates for "Short" and "Long" Convenience Yield Effect


*Note:* "Baseline" estimates shows the point estimates and confidence intervals for the flow and stock effects underlying the fit in Figure 1.9. The "Exp Augmented" results show flow and stock effects for the power-preserving specification described in section 1.7.5. Confidence bands are 90%. Standard errors are calculated via standard, GMM delta-method approaches, with a Newey-West, autocorrelation-consistent kernel for the moment covariance matrix. Sources: Wrightson ICAP via its Money Market Observer newsletter, Bloomberg, Federal Reserve Board of Governors via its H15 release, author's calculations.

Some readers may be skeptical of the fit shown in Figure 1.9, because the *fitted* impulse response at week 1 lies just outside of the 90% confidence interval of the LP-IV estimate. That point in the impulse response is indeed the most challenging for the GMM procedure to fit. Given that the *change* in T-bill supply from week 0 to week 1 is approximately the same as the change from week 1 to week 2, estimating large flow effects and small stock effects suggests that the convenience yield response at those two horizons should be nearly the same. Instead, the LP-IV estimates suggest that the convenience yield response at week 1 is larger.

There are modest ways to alter the GMM procedure, which produce better fits at those shorter horizons, such that the fitted response fits within the LP-IV confidence intervals. Doing so produces only modestly different estimates of the stock the flow effects. In Appendix A.4, I show results for a one-step GMM procedure, with a GMM weight matrix equal to the identity matrix. Estimates with pre-specified weight matrices are asymptotically inefficient, but still consistent. That weight matrix only modestly increases my estimate of the "stock" effect, to 1.85 bp/$100 billion.

### 1.7.5 An Alternate, Power-Preserving GMM Procedure

Next, I propose and estimate an alternate version of the GMM estimation above, which leverages the rich, high-frequency nature of Wrightson's projections data to obtain more power in the estimation. This procedure has a stronger exclusion restriction than the estimates above, but yields substantially smaller standard errors. Under this power-preserving estimation, the "stock" effect estimated above becomes statistically significantly different from zero, at the 10% level.

While the estimation procedure above uses fairly standard local projection techniques, in some ways it does not fully take advantage of the richness of the Wrightson T-bill projections. The LP-IV approach above (and the GMM estimation that uses the same moments) estimates the average convenience yield response after a Wrightson supply shock, and relates that to the average T-bill supply response after those same shocks. Figure 1.7 shows that this average response is a semi-persistent, but not permanent, T-bill supply change that rises for several weeks, then falls.

However, my setting is unusual in that I can measure how projections of future T-bill supply have changed fairly soon after each measured surprise. Using this information, I can supplement the GMM estimation above with another source of variation, that accounts for Wrightson's own expectations of how persistent the supply change from each supply will be. That is, I can supplement the moment conditions in equations (1.11) and (1.12) with an additional set of moments, corresponding to a new set of instruments:

$$\mathbb{E}\left(e_{t+h}\text{Update}_{t,t+h}\right) = 0$$
$$\mathbb{E}\left(e_{t+h}\text{Update}_{t,t+h-1}\right) = 0$$

$$\text{Update}_{t,t+h} = \left(\mathbb{E}_{Wrightson,t+1-\delta}B_{t+h} - \mathbb{E}_{Wrightson,t-\delta}B_{t+h}\right) \times \mathbf{1}\left(\varepsilon_t^B \neq 0\right)$$

That is, I use Wrightson's projection *updates* for future T-bill supplies at each future horizon as additional instruments for T-bill supply levels and changes at those horizons. These new instruments are the change in projections that Wrightson reports between their pre-surprise projections and their first post-surprise projection (i.e. the following Monday morning). I interact this update variable with dummy that equals 1 when week t had a nonzero Wrightson T-bill surprise. This limits the update variable to those information updates that could plausibly be driven by information revealed in a 4-Week T-bill surprise.

The exclusion restriction for using this extra dimension of variation is stronger than that from earlier specifications. First, these new instruments inherently compare surprises that are expected to be persistent to those that are expected to be more transitory. To use this variation for identification, it must be the case that convenient asset *demand* is not systematically different between states where a surprises' effects are perceived to be permanent. Second, this approach sacrifices some of the desirable ways in which identification strategy resembles high frequency identification methods in the empirical macroeconomics literature. Because Wrightson does not publish updated projections *immediately* after a 4-Week T-bill supply surprise, this measure of projection updates is necessarily taken as of several days after the associated surprise. It is conceivable that there is additional information, beyond the surprise directly measured via $\varepsilon_t^B$ that is included in my measurement of the projection update.

That said, there are some ways in which these additional assumptions are not overly restrictive. The identification logic of Section 1.6.2 is that the Treasury does not change its issuance strategy in response to short-term fluctuations in demand. It seems plausible that this means *updates* to Wrightson's future T-bill supply projections in week 0 should not be directly affected by short-term fluctuations in demand. Arguing that *surprises* in week 0 are unrelated to demand is conceptually similar to arguing that updates to T-bill projections in week 0 are unrelated to demand.

GMM estimates using this alternate set of instruments are also presented in Figure 1.10. Point estimates for the stock effect are listed as "Long, Exp Augment". Point estimates for flow effect are listed as "Short, Exp Augmented". Estimates for both the flow effect and stock effect are somewhat larger, moving to 13.53bps/\$100 billion and 1.73bp/\$100 billion, respectively. Consistent with this new approach using additional T-bill variation ignored by the earlier estimates, the confidence bands are substantially smaller. Under this alternate estimation, the estimated stock effects are statistically significantly different from zero, at all conventional confidence levels.

## 1.8 Empirical Results, Convenient Asset Substitutes

I show several estimates in support of the assumptions about convenient asset substitutes that were made in Section 1.5. First, I show that issuance volumes of FHLB discount notes do fall after a T-bill supply shock, but not by amounts that are large enough to materially impact my estimates for $\beta$. Next I show that, as expected, the convenience yield effect of T-bill supply shocks in periods when the ONRRP is *active* is smaller than in the inactive period, used in the results above. I find that repurchase agreement volumes falls substantially in the period when ONRRP is active, as the stylized framework in Section 1.5 describes.

I first assess whether rate sensitivity in the issuance of short-term Federal Government Agency notes is quantitatively large enough to meaningfully affect my estimates of $\beta$, in the manner discussed in Section 1.5.2. To do this, I estimate the impulse response of four-week discount note issuance by the Federal Home Loan Bank system, via LP-IV with

$$\text{2nd: } \sum_{\ell=0}^{3} \text{4W FHLB Issuance}_{t+h-\ell} = \alpha_{2,h} + \gamma_h B_t + \phi'_{2,h} X_{t-\delta} + e_{t+h} \tag{1.14}$$

$$\text{1st: } B_t = \alpha_1 + \chi \varepsilon_t^B + \phi'_{1,h} X_{t-\delta} + w_t \tag{1.15}$$

The 4-week moving sum on the lefthand side of the second stage regression reflects the fact that I am most interested in the impulse response of then-current FHLB discount note *supply* (i.e. stock). FHLB issuance data is a flow that requires some empirical assumptions to be used in this manner.

Figure 1.11 shows these results in the left panel. The direction of the response of FHLB discount note volumes moves in the expected direction, so that a surprise increase in T-bill supplies decreases FHLB discount note supplies over time. The response is statistically significant after the first week, for several weeks thereafter.

The results of Figure 1.11 show that the likely magnitudes of the response of FHLB issuance is not large enough to substantively alter my estimates of $\beta$. That is, in the notation of Section 1.5.2, $\frac{d\text{Agency}_t}{d\text{Bills}_t}$ appears small. A T-bill supply surprise that causes \$23 billion of additional T-bill supply after 3 weeks causes a peak decrease in FHLB supply of only approximately \$1.5 billion.[33]

Figure 1.11: Impulse Response of FHLB Note Convenience Yields and Issuance Volumes

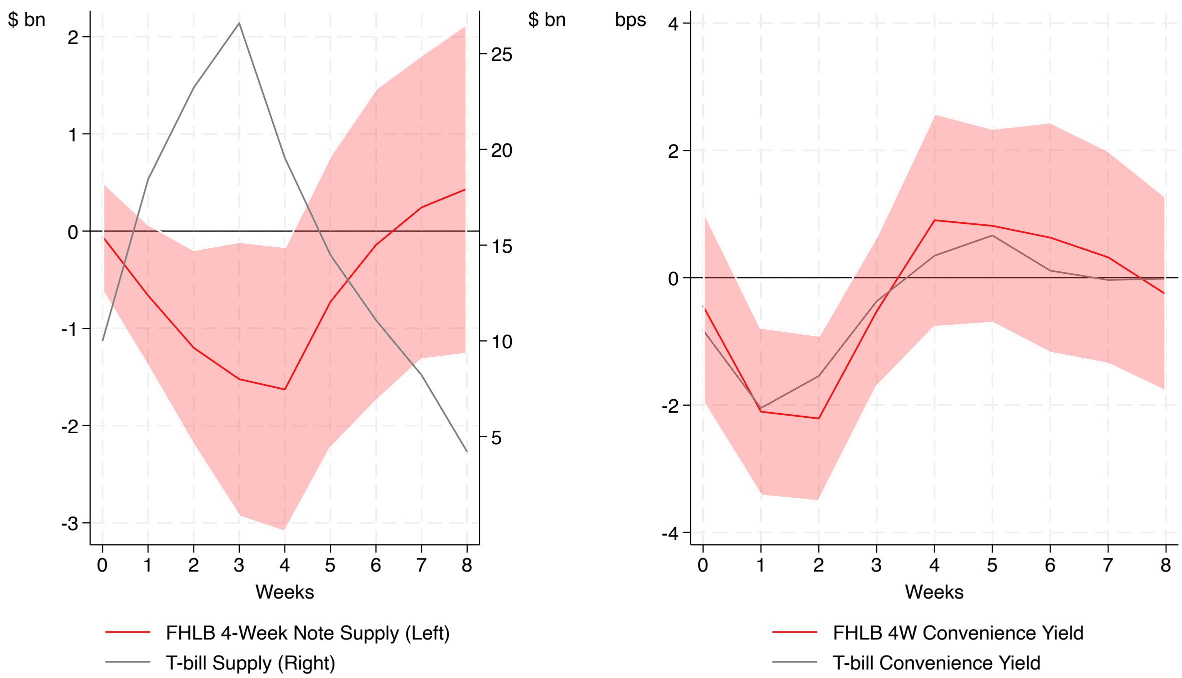


*Note:* Shows estimated impulse responses for convenience yields and issuance volumes of 4-Week discount notes issued by the FHLB – a commonly-discussed substitute for T-bills, for many market participants. Confidence intervals are 90% and standard errors are heteroskedasticity-robust. Sources: Wrightson ICAP via its Money Market Observer newsletter, Bloomberg, Federal Reserve Board of Governors via its H15 release, Federal Home Loan Banks system via its released auctions data, author's calculations.

The right panel of Figure 1.11 also shows the impulse response function of FHLB convenience yields to a Wrightson T-bill surprise. Reassuringly, this impulse response function looks nearly identical to the T-bill convenience yield response featured above. This shows that that the results above also affect the rates of the likely closest T-bill substitutes.

Next, I assess the extent to which movements in repo volumes might affect our estimates of $\beta$. In principle, we might conduct the same analyses as in Figure 1.11, for repo convenience yields and repo volumes. In practice, this is complicated by a lack of appropriate, high frequency repo volume data for my entire post-crisis sample. The US repo market is

[33]The fact that this analysis is restricted to FHLB discount notes is not very restrictive. The FHLB system issues the vast majority of short-term discount notes issued by any non-Treasury US government agency. In August 2023, Fannie Mae had \$15.23 billion of discount notes outstanding. Freddie Mac had \$5.6 billion of discount notes outstanding. In June 2023, the FHLB system had \$321 billion of discount notes outstanding.

composed of many subsectors, which are used by their participants for different reasons. I am most concerned with the repo market subsector that is used by cash investors (like money market mutual funds) to invest principle. Cash investors in these subsectors are the most likely investors to view T-bills and repurchase agreements as substitutes. The largest repo subsectors that fit this criteria are the Federal Reserve's ONRRP facility and the tri-party repo market. High frequency data on tri-party repo volumes became publicly available in September 2014.

I proceed by estimating impulse response functions for repo volumes and repo convenience yields by subsample, estimating separately for the ONRRP$= 0$ and ONRRP $> 0$ subsamples.[34] To do this, I define the repo convenience yield as $\text{IOER}_t - \text{Repo Rate}_t$, where IOER is the interest on excess reserves set by the Federal Reserve. Repo Rate is the triparty repurchase agreement repo rate reported by the Federal Reserve Bank of New York.[35]

These estimated impulse responses are shown in Figure 1.12. These results are largely consistent with the stylized story of Section 1.5. When volume at the ONRRP is strictly positive, repo volume is the margin of adjustment to a T-bill supply shock. The red line in the right panel of Figure 1.12 shows that repo volumes decline after a T-bill issuance shock, in a statistically significant way. Moreover, the volume of the decline is similar to the increase in T-bill volume after a shock. This suggests that repo volumes decrease after a T-bill supply shock in this subperiod nearly 1-for-1. Also consistent with this story, the impulse response of repo convenience yields is close to flat over this period. When ONRRP $= 0$, the repo convenience yield adjust substantially (and in the theoretically expected direction) after a T-bill supply shock.

An interesting and potentially surprising result is that the blue line in the right panel of Figure 1.12 is relatively flat and statistically indistinguishable from zero. This suggests that while repurchase agreement volumes respond *substantially* to T-bill supply shocks in the ONRRP $> 0$ period, they do not appear to respond nearly as strongly in the ONRRP $= 0$ subperiod.[36] Returning to the stylized examples of Figure 1.4, this suggests that the bottom-right panel is a reasonable representation of reality in the ONRRP $= 0$ period. That is, private repo supply curves are sufficiently inelastic that $\frac{dRP_t}{d\text{Bills}_t} \approx 0$.

This suggests that the structural estimates of flow and stock effects from the ONRRP$= 0$ subsample should reflect the T-bill convenience yield response to T-bill supply shocks with only little response in the supply of non-Treasury substitutes. That is, the impulse response of T-bill supplies in this subsample is a good proxy for the impulse response of *total* safe asset supplies. This is the assumption that motivated focusing my attention in the preceding analyses to changes in T-bill supply.

[34] These estimating equations are the natural analog of equations (1.14) and (1.15), adapted to these new lefthand side variables.

[35] In the early subsample, this rate is a survey rate reported by the largest dealers participating in this market. Later in the subsample, I use the Treasury Triparty general collateral rate reported by FRBNY.

[36] Indeed, the point estimate is positive, which is of a theoretically unexpected sign.

Figure 1.12: Impulse Response of Repo Convenience Yields and Repo Volumes

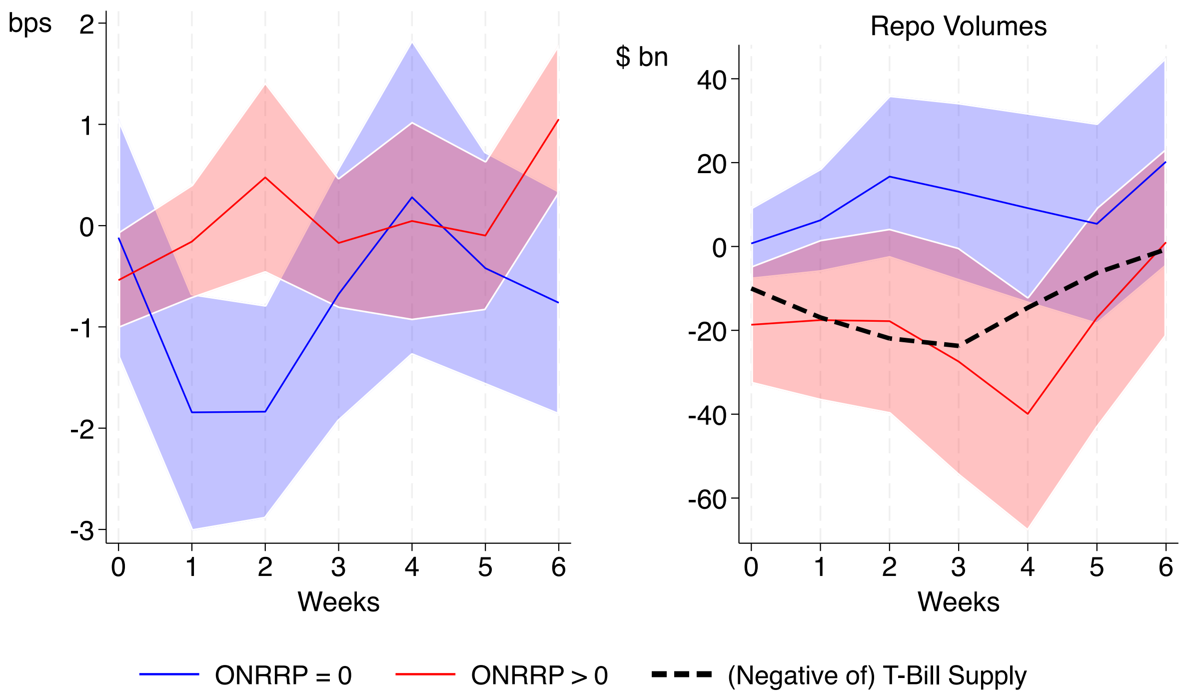


*Note:* Shows estimated impulse response functions for repo volumes and repo "convenience yields" to a T-bill supply shock, defined as a week 0 Wrightson surprise. Repo volumes are defined as the sum of private tri-party general collateral repo volumes as reported by the OFR, and ONRRP takeup as reported by FRBNY. Repo convenience yields are the difference between the interest on excess reserves and then-current, representative repo rates. When available, the repo rate used is the broad general collateral repo rate reported by FRBNY, When not available, it is the general collateral repo survey rate reported by the same. Sources: Wrightson ICAP via its Money Market Observer newsletter, Federal Reserve Board of Governors via its H15 release, FRBNY via its published repo rates and volumes, Treasury OFR via its repo market data release, author's calculations.

Lastly, Figure 1.13 shows that T-bill convenience yields do indeed less in the ONRRP$>0$ period. While point estimates at week 0 are nearly the same as in the ONRRP$\approx 0$ period, estimates in the following weeks are smaller. These indicate a smaller, and less persistent, convenience yield response when takeup volume at the ONRRP is an available margin for market adjustment.

Figure 1.13: Impulse Response of T-bill Convenience Yields by Subsample

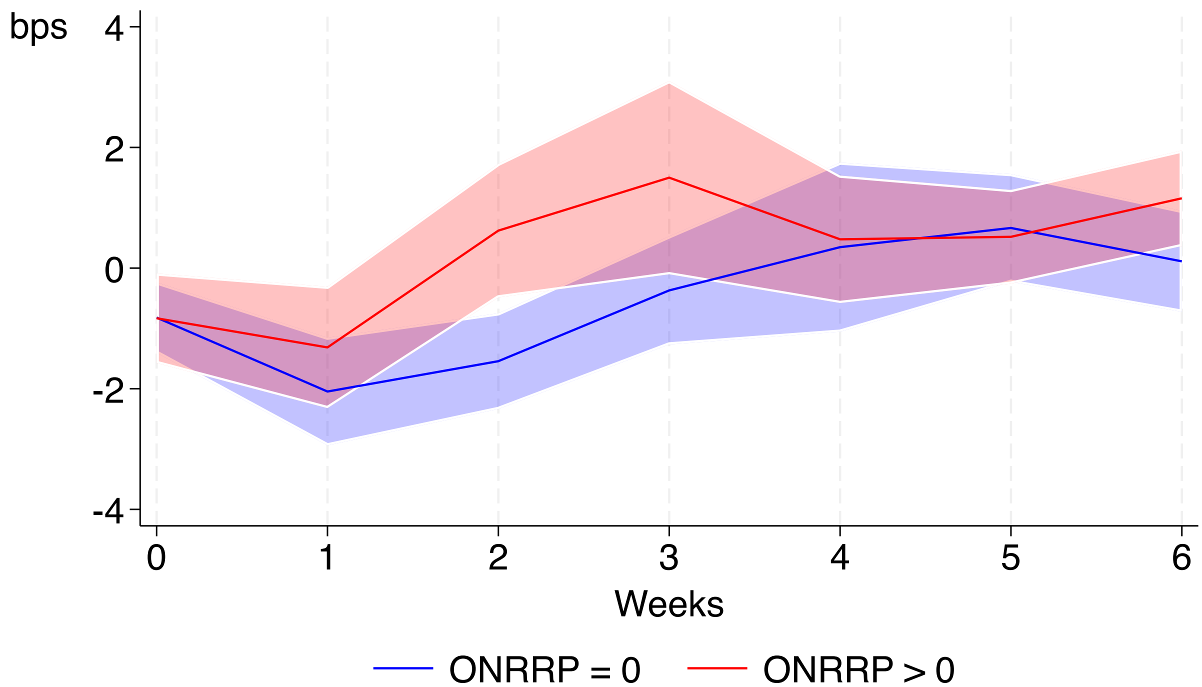


*Note:* Compares estimated T-bill convenience yield impulse responses, to a T-bill supply shock, in subperiods where ONRRP takeup is close to zero or substantially positive. Confidence intervals are 90% and standard errors are heteroskedasticity-robust. Sources: Wrightson ICAP via its Money Market Observer newsletter, Federal Reserve Board of Governors via its H15 report, author's calculations.

## 1.9 Application: $R < G$ and Debt Sustainability

To place my core empirical results in context, and demonstrate how they will tend to impact the quantitative conclusions of macroeconomic models featuring public debt convenience yields, I replicate a calibration exercise from Mian et al. (2022) on fiscal sustainability when $R < G$, using my estimates of a substantially flatter long-run convenient asset demand curve. My estimates suggest that endogeneity of convenience yields is not a strong enough force to meaningfully constrain the implications of Blanchard (2019), that large fiscal deficits appear sustainable when $R < G$.

The model of Mian et al. (2022) offers an excellent setting to understand the positive, quantitative impact of my results for fiscal sustainability in a transparent, minimalist setting. With those goals in mind, I present the simplest, most-stylized form of their model, without considering complications created by the zero lower bound, or aggregate risk. In this section, I summarize the components of that model which are most important for understanding its conclusions about the importance of the slope of the convenient asset demand curve.

The model is deterministic and exists in continuous time. It features households, which are separated into populations of savers and spenders. Savers have the ability to save via holdings of government debt, for which they enjoy flow convenience benefits in addition to the security's interest payments, as in Section 2.2. The model features a central bank which, in this simplest form of the model, is able to perfectly maintain its inflation target by keeping interest rates at the natural rate.

A government in the model sets fiscal policy via its choice of government spending $x$, borrowing $b_t$, and lump-sum taxes on savers $\tau_t$. Its choices must satisfy the flow budget constraint

$$\frac{db_t}{dt} = z_t + R_t b_t$$

where $z_t$ is the primary fiscal deficit, such that $z_t = x - \tau_t$. $R_t$ is the nominal interest rate on government debt, and $\dot{b}_t$ is net government borrowing in moment $t$. $R_t$ is the nominal interest rate on government debt.

This is a model about maintaining a stable *ratio* of government debt to GDP $y_t$. As such, it is convenient to redefine $\tilde{h}_t = \frac{h_t}{y_t}$, and rewrite this flow budget constraint as

$$\frac{d\tilde{b}_t}{dt} = \tilde{z}_t + \tilde{b}_t \left(R_t - G\right)$$

where $G$ is the growth rate of nominal GDP. In the version of the model without a ZLB, where the central bank is always able to keep output at potential, this is the same as normalizing by potential GDP, which the authors do in their own presentation.

The Euler equation of the saver households in the model is

$$\frac{d\log(\tilde{c}_t)}{dt} = R_t - G - \rho + v'(\tilde{b}_t)\tilde{c}_t \tag{1.16}$$

where $\rho$ is the (continuous time) discount rate, and $v'(\cdot)$ is the marginal extrapecuniary value of convenience from public debt.

Considering a steady state in a detrended model where potential output is constant, the natural rate $R^*$ is the rate that supports constant consumption by savers in Equation (1.16). When government debt $b$ is constant in steady state, this gives

$$R^*(\tilde{b}) = \rho + G - v'(\tilde{b})\tilde{c}_s$$

where $\tilde{c}_s$ is the steady state consumption of savers. In the model, this equals $1 - \mu - x$, where $1 - \mu$ is the labor endowment of savers.

Continuing to consider a steady state where consumptions and government debt are constant, the government's flow budget constraint suggests

$$z(\tilde{b}) = (G - R^*(\tilde{b}))\tilde{b} \tag{1.17}$$

where $z(b)$ is the primary deficit (as a share of GDP) that supports the steady-state level of debt-to-GDP $\tilde{b}$.

Equation (1.17) characterizes the model's useful deficit-debt diagram. This diagram characterizes the relationship between steady state levels of the primary deficit $\tilde{z}$ and level of government debt $\tilde{b}$. The left-hand panel of Figure 1.14 shows the deficit-debt diagram for the baseline parameters in the calibration exercise of Mian et al. (2022).

Figure 1.14: Deficit-Debt Diagrams, Under Alternate Measurements of Slope of Convenient Asset Demand Curve

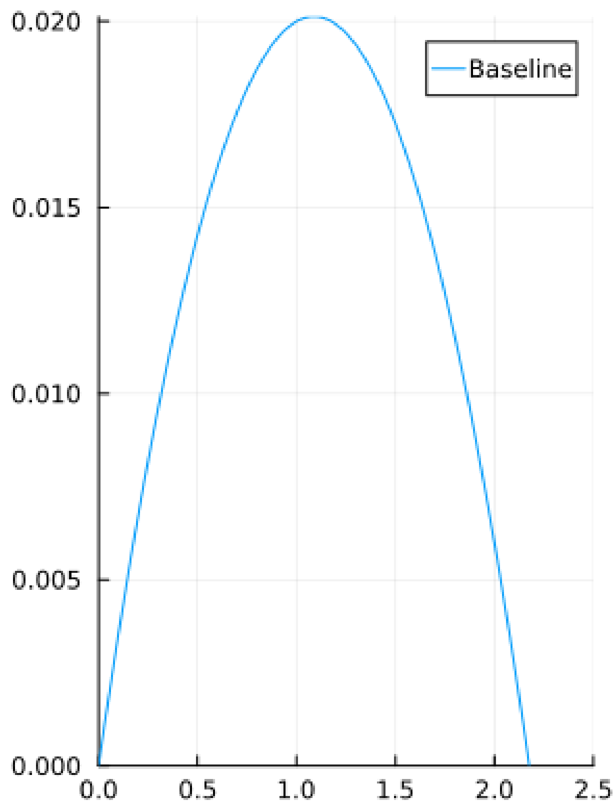


*Note:* Recreates deficit-debt diagram implied by the Mian et al. (2022) model, with different assumptions for the slope of the convenient asset demand curve. Other model calibrations are identical to those in Mian et al. (2022), and reflect the US macroeconomic state in late 2019. Sources: Author's calculations, D'Avernas and Vandeweyer (2023), Greenwood et al. (2015b), and Mian et al. (2022).

This curve shows combinations of deficit and debt that support a steady-state. It can also be used to understand transition dynamics away from a steady state. Suppose an economy at steady state $A$ increases deficits, immediately and permanently, from $\tilde{z}_1$ to $\tilde{z}_2$. Because debt, deficit pairs *above* the locus (i.e. deficits that are larger than those for which debt is constant) correspond to increasing levels of $\tilde{b}$, this suggests that $\tilde{b}$ will increase until contacting the locus again, at a higher level of debt $\tilde{b}_2$.

The deficit-debt diagram in this model can show the *maximum* sustainable quantity of government debt, and the deficit that supports it. These are the highest levels of $\tilde{z}$ and $\tilde{b}$ on the locus – $\tilde{b}^*$ and $\tilde{z}^*$. In the baseline parameters assumed by Mian et al. (2022), these occur at a deficit equal to 2% of GDP, and government debt level equal to 109% of GDP.

Mian et al. (2022) is open about how important the assumed slope of the convenient asset demand curve is for these conclusions. Flatter convenient asset demand means that a given level of government debt $\tilde{b}$ is associated with a *larger* primary deficit $\tilde{z}$, via smaller interest expenses $R\tilde{b}$. The figure in the left panel of Figure 1.14 assumes $\tilde{b}\frac{\partial(\rho+G-R)}{\partial\tilde{b}} = -1.7\%$.

With a series of arithmetic calculations described in Appendix E of their paper, Mian et al. (2022) note that the estimates of Greenwood et al. (2015b) suggest a value of $-1.4\%$. Performing the same arithmetic transformation for the estimates in the published version of D'Avernas and Vandeweyer (2023) yields an estimate of $-0.78\%$. The implied deficit-debt diagrams are shown in the right panel of Figure 1.14.

The same transformation for the estimates suggested by my results suggest a value of $-0.3\%$. The value corresponding to the 90% confidence interval *lower-bond* for my stock effect is $-0.5\%$.[37].

The flatter demand curves suggested by my estimates have very large implications for the quantitative conclusions of this simple model – suggested by the visual differences in the deficit-debt diagrams. The values of Greenwood et al. (2015a) suggest a maximum debt, deficit combination of 114% and 2.03% of GDP, respectively. The flatter demand curve estimates of D'Avernas and Vandeweyer (2023) correspond to a maximum debt, deficit of 177% and 2.4% of nominal GDP, respectively. My point estimates of the convenience yield effect of T-bill supply shocks suggest a pair of 345% and 4.0% of GDP for debt and deficit. The largest convenience yield responses not rejected by my estimates at the 90% level suggest 240% and 3% of GDP.

In discussing this calibration exercise, I do not suggest that these values should be taken as conclusive measurements of fiscal sustainability. This is a simple, stylized model that omits many realistic qualities of the fiscal sustainability problem, like private capital and aggregate – many of which are discussed in later sections of Mian et al. (2022). It surely does not suggest any normative conclusions that these levels of debt and deficit are *optimal*.

That said, this exercise shows how steepness in the convenient asset demand curve acts as much *less of a constraint* on the sustainability of different deficit, debt combination, given my estimates. This mechanism will still feature in richer, more realistic models. My estimates suggest that this steepness alone does not appear to the general fiscal policy implication of Blanchard (2019), that $R < G$ makes many combinations of debt and deficits appear sustainable.

## 1.10 Conclusion

In this paper, I have introduced a new short-term instrument for convenient asset supply. This instrument is based on T-bill issuance surprises, relative to the projections of Wrightson's Money Market Observer, a prominent money market newsletter. These surprises avoid concerns about opportunistic issuance by the Treasury, which might otherwise identification. Unlike the literature's previous approach, Wrightson surprises are more robust to the possibility of seasonality in convenient asset *demand*

With this new instrument and a local projection approach that differs from preceding empirical frameworks in this literature, I show that short-run effects of T-bill supply surprises are substantially larger than long-run effects. Imposing a simple, structural restriction that

[37] This is the lower-bound suggested by the power-preserving GMM estimation of Section 1.7.5

the T-bill convenience yield depends on the then-current *stock* and *flow* of T-bill supplies, I use my estimated impulse response to estimate these two effects separately. The estimates suggest that a \$100 billion increase in the supply of T-bills depresses convenience yields by 10.4bps – a much larger effect than previous studies have reported. However, this large effect is short-lived, leading to a point estimate of the stock (i.e. long-run) effect of only 1.1 basis points. This stock effect is not statistically significant under my baseline GMM approach. It is statistically significant, and somewhat larger, under an alternate power preserving GMM approach that uses *updates* in Wrightson's projections of future T-bill supply in surprise weeks as additional instruments for future T-bill supply. This alternate approach features a more restrictive exclusion restriction, but retains valuable variation about which surprises are expected to produce transitory or persistent changes in future supply.

My estimates, which suggest a steeper short-run convenient asset demand curve but a *flatter* long-run demand curve, will tend to suggest more fiscal sustainability in models where $R < G$. I demonstrate this by repeating a simple but powerful calibration exercise from Mian et al. (2022), based on US data as of the end of 2019. My estimates, used in their simplest framework, suggest that larger long-run deficits can still support a steady state with constant fractions of debt to GDP.

# Chapter 2

# Empirical Network Contagion for US Financial Institutions

## with Fernando Duarte

## 2.1 Introduction

The financial crisis of 2007-2008 renewed economic interest in the network structure of the financial system and the interactions between the financial sector and the real economy.[1] Since then, academic research on financial networks has grown substantially, vastly improving our understanding of how interconnectedness among economic agents arises, evolves and ultimately affects the economy. Because detailed empirical data on financial networks is almost always insufficient to perform consequential analyses, the literature has predominantly focused on theoretical aspects, while acknowledging the limitations that a lack of empirical results imposes on our understanding. Although there is a decent understanding of certain empirical aspects of the structure of the financial network, like its core-periphery topology or its increasing complexity, there is not as much knowledge about welfare-relevant attributes of networks, such as the size of network spillovers, the degree of propagation and amplification of shocks through network effects, or how network vulnerability varies as a function of the shock size. In general, the literature approaches welfare-relevant questions by constructing top-down measures of systemic risk or interconnectedness that rely on more-readily available data (such as stock market returns) instead of actual network-specific data. To tie non-network specific data to welfare-relevant network variables requires a model, or at least some auxiliary assumptions that are difficult to test without network-specific data.

In this paper, using node-specific data, we empirically estimate a measure of expected network default spillovers for the US financial system for the period 2002-2016. Although default spillovers are only one dimension of potential network effects, they have been repeatedly cited as a major factor during the financial crisis, motivating existing regulation and studies by the theoretical network literature. We build our measure of spillovers by using the general and elegant framework of Eisenberg and Noe. The nodes of the network are

[1] The views expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Bank of New York, or the Federal Reserve System.

US financial institutions, including bank holding companies, broker-dealers, and insurance companies. The connections between nodes are defined by the bilateral payment obligations between them. We refer to these claims as *inside* assets or liabilities. In addition, each node has assets from and liabilities to the *outside* sector, composed of non-financial firms, households, governments, and financial firms outside the US. Our primary measure of network vulnerability will quantify expected network spillovers in the face of an exogenous shock to these outside assets. The primary method for contagion in the model is a "default cascade", whereby a shock to outside assets can cause some institutions to default on their in-network counterparties, which could in turn cause those counterparties to default on their own inside obligations, and so on. This domino effect can propagate through the financial system, creating network spillovers.

Empirically estimating the expected value of these spillovers requires knowledge of the bilateral claims between each pair of nodes. Such granularity of data is not publicly available[2]. However, Glasserman and Young (2015) show that, for a large family of exogenous shock distributions, a meaningful upper bound on the expected value of default spillovers can be constructed with knowledge of node-specific information only (i.e. without a precise breakdown of the nodes' counterparties or the magnitudes of obligations to them). In particular, the bound is based on each node's probability of default, its total outside assets and its ratio of inside liabilities to total liabilities. The empirical measure of spillovers that we estimate is this upper bound on the expected value of default spillovers proposed by Glasserman and Young (2015). Thus, at the cost of estimating an upper bound on spillovers instead of their actual values, the data requirements are greatly reduced. For a significant portion of all US financial institutions (constituting 21% of total assets in the network), we use detailed balance sheet data from the FR-Y9C reporting form to construct all relevant variables. When detailed line-item balance sheets are unavailable (as is the case for most insurance companies and broker-dealers in our sample), we combine remaining firms into more-aggregated, sector-level nodes, and use the Federal Reserve's Financial Accounts of the United States (formerly known as Flow of Funds) to estimate outside assets at the sector level[3]. We obtain the last ingredient, firm-specific probabilities of default, from Moodys Analytics' (formerly KMV) Expected Default Frequency series.

We find that between 2002 and 2007 the upper bound on default spillovers is rather small, which means that the financial network is robust to contagion arising from counterparty risk. However, between 2008 and 2012, the upper bound on spillovers is meaningfully above zero. Our results suggest that the financial network is most fragile in the first quarter of 2009, when we estimate that network default spillovers can amplify initial losses by up to 25 percent. After 2012, the upper bound on default spillovers starts to decline and reverts to pre-crisis levels by 2015. In 2016 and 2017, the last two years of our sample, our measure of spillovers starts to increase again — slowly but consistently.

One way to understand our results is to decompose the upper bound on default spillovers

[2]In the network simulaion literature, a common method to compute exact payments in the response to a shock is to assume a "maximum entropy" form to inside obligations. In the presence of any uncertainty about the structure of bilateral claims, this essentially spreads a node's inside obligations as evenly as possible across potential counterparties.

[3]Including these additional assets brings the quantity of assets in the domestic financial sector accounted for in our sample up to 35%

into two factors: A weighted average default probability for the sample and a *connectivity multiplier* that captures how the initial losses in outside assets could be transmitted and amplified by a default cascade. We find that both factors are important in explaining the overall dynamics of spillovers. Between 2002 and 2007, default likelihoods were negligible and the connectivity multiplier declined by around 10 percent. In 2008, default probabilities spiked, but financial connectivity declined sharply as financial institutions reduced exposures among each other amidst stressed financial and economic conditions. Even though our estimates for expected spillovers in 2008 increased, the reduction in financial connectivity was an important mitigant. In 2009, both default likelihoods and financial connectivity increased, leading to a large jump in our spillover estimates.

Another way to analyze our results is by constructing a node-specific "contagion index", which quantifies the ability of a node to transmit and amplify losses. We find that JP Morgan Chase, Bank of America, Citigroup and Wells Fargo have the largest contagion indices. Similarly to the network-wide measure, we can decompose the contagion index into two sub-components, the node-specific financial connectivity and the size of each node's outside assets relative to its equity capital. While the contagion index of JP Morgan Chase, Bank of America and Citigroup are generally driven by their large financial connectivity, the contagion index of Wells Fargo is mainly driven by its outside assets being large relative to equity capital.

To the best of our knowledge, our paper is the first to empirically assess network spillovers across many years and for a wide cross-section of US financial institutions. Having a panel has several advantages. First, it allows us to better identify the drivers of spillovers. Second, it places tighter restrictions on theoretical models that seek to model default spillovers. Third, it provides information that is potentially useful to policymakers and regulators, such as the quantitative contribution of spillovers to systemic risk.

**Related Literature.** Some of the first studies into financial network topologies and the relative vulnerabilities of different networks were the seminal theoretical models of Allen and Gale (2000) and Freixas et al. (2000), both of which model liquidity crises at depository institutions. In both, the central takeaways were that different configurations of networks (all of which, in their studies, were purely hypothetical) could either alleviate the risk of contagion or exacerbate it. Since then, many papers have used simulations to estimate the severity of losses in interconnected banking networks in the face of a variety of shocks. Examples include Upper and Worms (2004) for the German banking system, Elsinger et al. (2006a) and Elsinger et al. (2006b) for the Austrian Banking System, and van Lelyved and Liedorp (2006) for Dutch Banks[4]. Drehmann and Tarashev (2013) develop and test bank-specific, simulation-based measures of systemic risk for 20 large financial institutions.

Another strand of the literature focuses on characterizing the topology of financial networks, using either degree distributions as in Boss et al. (2004) or searches for a core-periphery structure in banking systems as in Craig and von Peter (2014)[5]. For these studies, the lack of usable data on bilateral claims in the financial system has necesitated assumptions to fill gaps in balance sheet data or, in the case of van Lelyved and Liedorp (2006),

[4]See Upper (2011) for a useful survey of simulation-based contagion risk estimations

[5]See Glasserman and Young (2016) for a general survey of the networks literature, including a dedicated discussion of networks specific-measures such as degree distributions, core-periphery structures, and the related concept of node depth.

incomplete bilateral claims data restricted to a subset of the balance sheet of a subset of all financial institutions. Studies quantifying losses through simulation require explicit and potentially-stringent assumptions about shock distributions.

A separate subset of papers sheds the analysis of counterfactual shocks entirely, in favor of strictly empirical analysis. Gropp et al. (2009), for instance, find evidence of comovement in market-based estimates of probability of default for large European financial institutions which, they argue, are reasonably attributable to contagion effects. Similarly, Hawkesby et al. (2007) analyze comovements between asset prices of several large multinational financial institutions. As Hawkesby et al. (2007) note, while they may give interesting insight into market perceptions of interconnectedness or potential co-exposure to common factors, these studies do not attempt to 'capture the degree of contagion that may occur during periods of financial stress'. Another related strategy has been to use market data, such as stock returns, to construct "top down" measures of systemic risk that indirectly relate to the actual network structure, as in Adrian and Brunnermeier (2016), Acharya et al. (2017), or Brownlees and Engle (2016).

The framework of Eisenberg and Noe (2001) is a common thread through much of the recent financial networks literature. Their model presents an intuitive and general system for intra-network defaults and payment shortfalls in the presence of fully general shocks to assets outside the network, bound by simple rules such as limited liability and debt seniority[6]. Building on their findings, Glasserman and Young (2015) derive useful bounds on contagion losses without additional assumptions regarding bilateral claims, and for a broad family of shock distributions. A substantial portion of our paper can be viewed as an empirical estimation of these bounds.

After deriving their theoretical upper bounds, Glasserman and Young (2015) employ data from the European Banking Authority's (EBA) 2011 stress test and simplifying assumptions to argue that default spillovers, on their own, are likely to be small unless further frictions, such as bankrupcy costs or bank runs, are also present. Much of the literature supports this view (e.g. Upper and Worms (2004)). We find that even for pure default spillovers that do not interact with any other frictions, we cannot dismiss the possibility of sizable spillover effects (for the first quarter of 2009, we find that default spillovers can amplify initial exogenous losses by up to 25%). The main reason why other studies find negligible spillovers while we do not is that the default probabilities for the financial institutions we analyze are substantially larger than the probabilities of default of the institutions used in other studies (which are almost always European banks). In addition, the share of in-network liabilities to total liabilities (a measure of network connectivity) that we empirically estimate are also somewhat larger than those in Glasserman and Young (2015) and the rest of the literature.

[6]Particularly, many simulation-based studies of contagion risk rely on the Eisenberg and Noe (2001) algorithm to find a sequence of network 'clearing' payments after a shock to assets outside the network.

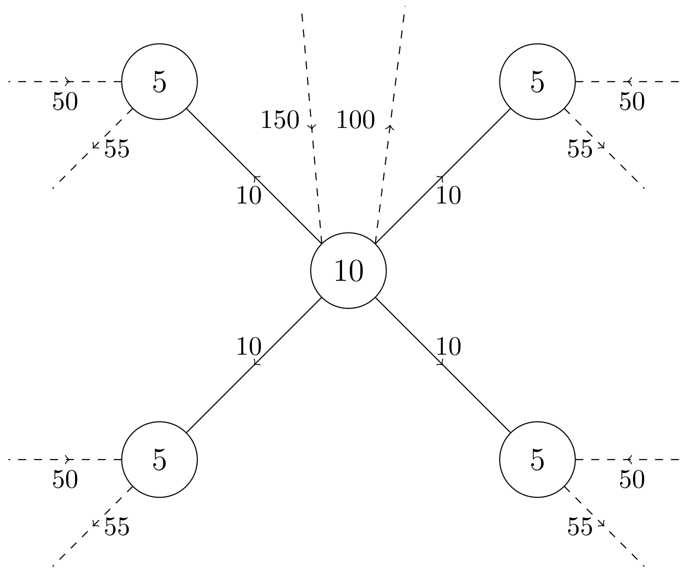


Figure 2.1: **A simple network**. Nodes are financial institutions. There is a connection from node $i$ to node $j$ if $i$ is a net borrower from $j$. The dashed lines show connections to the outside sector.

## 2.2 Network Model

### 2.2.1 Overview

The network model we use is exactly as in Eisenberg and Noe (2001) and Glasserman and Young (2015). The nodes of the network are all US domestic financial institutions. The connections between nodes are defined by institutions borrowing from and lending to one another. There is a link from node $i$ to node $j$ if $i$ has any payment obligations towards node $j$. In addition to lending to one another, nodes can borrow and lend to the rest of the domestic and global economy. These assets and liabilities are termed as *outside* the financial system. In our application, the *outside* sector is comprised of domestic and foreign non-financial institutions, governments, and households, as well as foreign financial institutions.

Figure 2.1 shows an example of a simple network, taken from Glasserman and Young (2015). The four arrows originating in the central node and pointing to the four peripheral nodes show that the central node owes 10 to each of the peripheral nodes. The four peripheral nodes have no borrowing or lending among themselves. For this network, we say that the central node has *inside liabilities* of 40, while each of the peripheral nodes has *inside assets* of 10. In practice, we find that inside assets and liabilities for US financial institutions are primarily composed of deposits, loans and securities lending transactions.

In addition to its claims inside the network, the central node has lent 150 and has borrowed 100 from the outside sector, depicted by the dashed lines with arrows going into and out of the central node. We refer to positive claims with respect to the outside sector as *outside assets* and to negative claims as *outside liabilities*. Outside assets typically consist of securities, loans to firms and households (including mortgages), and public debt. Outside liabilities mostly involve deposits and lines of credit.

The difference between all assets and all liabilities gives each node's net worth. The central node has a net worth of 10, shown inside the circle that represents the node. Each of the peripheral nodes has outside assets of 50, outside liabilities of 55 and an inside asset of 10 with respect to the central node, for a net worth of 5.

### 2.2.2 Shocks and Propagation

The shocks we consider are exogenous reductions in the value of outside assets. Therefore, all initial losses always originate outside the network. One example of such a shock is an increase in defaults for residential mortgages held by financial institutions.

For sufficiently high initial losses in outside assets, some nodes in the network will be unable to pay their creditors in full. When this happens, all debts for the defaulting node (including those outside the network) are written down pro rata and creditors receive only a fraction of their promised payments. Note that under a pro rata allocation, a node defaults on either all of its creditors or none of them. When creditors for some node are not paid in full, they may themselves be unable to pay their own creditors, and so on. Initial losses thus get transmitted inside the network through this "domino" effect. We do not include in our analysis any liquidity or equity injections, and only net claims between two nodes are assumed to be of relevance (as opposed to gross positions). In addition, nodes do not renegotiate claims, even if it may be mutually beneficial to do so.

As a numerical example, consider what happens when the outside assets of the central node in Figure 2.1 receive a shock of size 80. Outside assets for the central node decrease from 150 to 70. Total liabilities are initially 140. After the shock, under a pro rata allocation, only 50 percent of each liability is repaid as the central node only has 70 remaining in assets. Each of the peripheral nodes receives 5 from the central node, just enough to balance their assets and liabilities. A shock to the outside assets of the central node of magnitude greater than 80 would reduce the value of assets for peripheral nodes below the value of their liabilities. In this case, the peripheral nodes would default on their creditors. In this case, the central node has created contagion to the peripheral nodes through network contagion. The peripheral nodes default even though none of their outside assets were affected by the initial shock.

### 2.2.3 The Disconnected Network

To quantify the amplification of losses stemming from the network structure –as opposed to the initial losses from exogenous shock to outside assets– we compare expected losses for the system (the network plus the outside sector) to the losses in a hypothetical system in which all connections inside the network have been severed. Both networks are subject to the same distribution of exogenous shocks to outside assets, and to no other shocks. We create this hypothetical *disconnected* system by removing all connections between nodes inside the original network but keeping the links with the outside sector intact. We also assume the net worth at each node remains unchanged by creating, for each node, a fictitious claim to the outside sector equal in value to the net value of all the connections that were removed. Depending on the sign of the net value of removed connections, the new fictitious claim can

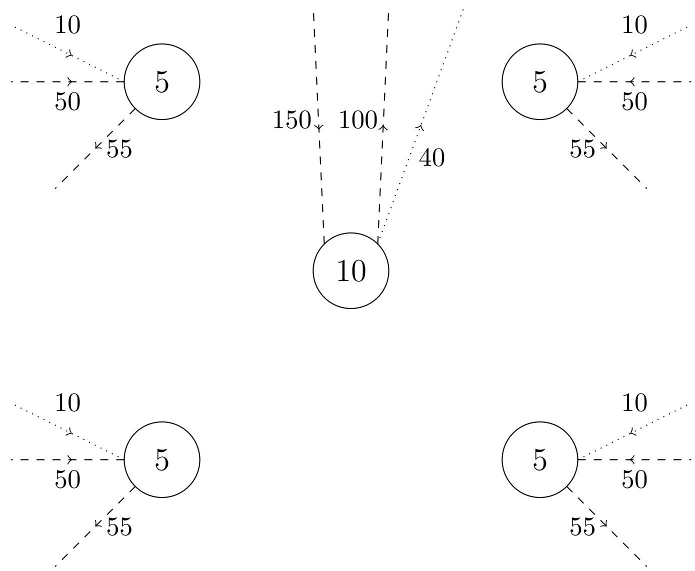


Figure 2.2: **A Simple Disconnected Network**. The disconnected version of the networks in Figure 2.1 are obtained by removing all connections between nodes inside the original network but keeping the links with the outside sector intact. Net worth remains unchanged by creating fictitious outside assets or liabilities. Dashed lines indicate actual balance sheet assets and liabilities, and dotted lines indicate fictitious assets from or liabilities to the outside sector.

be an asset or a liability. If it is an asset, we assume it is not subject to the shocks to outside assets to keep the set of assets initially shocked identical to that of the original network. If the new fictitious claim is a liability, we assume it has the same priority as all other liabilities. In case of default, the new fictitious liability gets haircut pro rata just like all other non-fictitious liabilities, and any "losses" imposed on that obligation are counted towards the value of total system losses. Figure 2.2 shows the disconnected version of network displayed in Figure 2.1.

### 2.2.4 An Upper Bound on Network Spillovers

We are interested in whether the expected system losses in our real-world, interconnected system are substantially greater than those in the hypothetical *disconnected* system, where node connections have been excised. We define $R$ to be the ratio of expected losses for the actual network to the expected losses in the disconnected network. That is, if $L$ denotes total system losses,

$$R = \frac{E(L_{\text{Actual}})}{E(L_{\text{Disconnected}})} \tag{2.1}$$

The value of $R$ gives the relative magnitude of additional losses imposed on the system because of the interconnected structure of the network - to wit, network effect losses. With perfect information on the bilateral claims in the system, this ratio could be calculated exactly in response to a variety of shocks by using the Eisenberg and Noe (2001) algorithm to compute the set of node payments that 'clear' the system (i.e. follow the system's rules of limited liability and pro rata allocation). In the United States financial system, detailed and publicly-available data on bilateral obligations between financial firms does not exist.

The main result in Glasserman and Young (2015) is that a useful upper-bound on $R$ can be derived without any information on the makeup of each node's bilateral claims. We call this upper-bound $B$. If the tails of the distribution of exogenous shocks to outside assets are not too fat-tailed, then $B$ can be calculated using node-specific information only[7]. Glasserman and Young (2015) show that $B$ depends only on each node's total outside assets $c$, each firm's probability of default due to direct shocks to outside assets $\delta$, and the maximum *liability connectivity* among nodes in the system $\beta^+$. Each node's liability connectivity is defined as its ratio of inside liabilities to total liabilities.

Glasserman and Young (2015) show that

$$B = 1 + \frac{1}{(1-\beta^+)} \frac{\sum_{i \in S} \delta_i c_i}{\sum_{i \in S} c_i}, \tag{2.2}$$

$\delta_i$ : probability of default from outside shocks for node $i$,

$c_i$ : the dollar value of outside assets for node $i$,

$\beta^+$ : maximum liability connectivity, i.e., $\beta^+ = \max_{i \in S} \beta_i$, with $\beta_i =$ the fraction of firm i's liabilities held by other nodes in the networks,

$S$ : Set of financial institution nodes within the network.

The upper bound $B$ for network spillovers is increasing in the maximum financial connectivity of the system, $\beta^+$, and in the quantity $\sum \delta_i c_i / \sum c_i$, most-easily interpretable as a weighted average probability of default for the system (with each firm's weight given by its share of total outside assets). When $\beta^+$ is close to 1, aggregate financial connectivity is high and any initial shock to outside assets has the potential to be transmitted broadly across the network. In contrast, when $\beta^+$ is close to zero, any initial shock dissipates quickly and expected losses should be similar to those in a truly disconnected network.

For most systems calibrated to real-world data, previous studies have found that the upper bound $B$ is small. For example, picking $\beta^+ = 0.8$ and $\delta_i = 1$ percent for all nodes $i$, we get $B = 1 + 0.01/(1-0.8) = 1.05$. This means that the connected system has expected losses that are at most 5 percent larger than those in the system of isolated nodes. In their example exercise, Glasserman and Young (2015) find an even smaller upper bound of 1.0175 for European banks using data from the the 2011 European Banking Authority stress test.

[7]More technically, we consider shocks that have an "increasing failure rate" (IFR). A random variable with distribution function $G(x)$ and density $g(x)$ is said to have an IFR if $g(x)/(1-G(x))$ is an increasing function of $x$ . This family encompasses the normal, exponential, and uniform distributions. There are no restrictions on the correlation structure of shocks.

In addition, the joint distribution of potential shocks is assumed to be invariant to scale (homogeneous in assets). For example, if total assets of a node double, expected losses are assumed to also double.

### 2.2.5 The Network Vulnerability Index

We define the *Network Vulnerability Index* (NVI) to be the upper bound on the magnitude of additional expected losses created in the system by network spillovers, expressed as a share of expected disconnected system:

$$NVI = (B - 1) = \frac{1}{(1-\beta^{+})} \frac{\sum \delta_i c_i}{\sum c_i}. \tag{2.3}$$

Being an upper bound, the $NVI$ is most useful when its value is small, since the model then clearly indicates low vulnerability to potential network spillovers. When the index is large it is less informative. In this case, the true value of potential network spillovers could be as large as the upper bound or as low as zero, as dictated by the bilateral claims between nodes. The model does not produce any additional information that can help pinpoint the true value of network spillovers within that the range $[0, NVI]$. As an extreme, when the $NVI$ is equal to infinity, it provides no information[8].

### 2.2.6 A Firm-Specific Risk Measure: The 'Contagion Index'

Glasserman and Young (2015) also presents a firm-specific measure of the potential to cause contagion, which they term a firm's 'contagion index'. For a wide family of shocks, the index is defined as

$$\text{contagion index} = w_i \beta_i \lambda_i$$

where $w_i$ is a firm's net worth, $\beta_i$ is liability connectivity as in equation 2.3, and $\lambda_i = \frac{c_i}{w_i}$ is the leverage of firm $i$'s outside assets.

Given that the magnitude of exogenous shocks to outside assets in the model is bounded by each firm's actul quantity of outside assets, the contagion index calculates the total payment shortfall that a firm could potentially pass on to other nodes following a shock to its own outside assets. Glasserman and Young (2015) show that an outside asset shock to node $i$ cannot possibly cause default to node $j$ if node $j$'s net worth is greater than node $i$'s contagion index. They also show that the *probability* of node $j$ defaulting solely because of a shock to node $i$'s assets *must* be less than the probability of node $j$ defaulting from a shock to its own assets if $i$'s contagion index is less than $j$'s quantity of outside assets, $c_j$[9].

## 2.3 Data and Empirical Methodology

xThis paper combines a number of different data sources to estimate the fields in equation 2.3. What follows is a description of those data sources and any decisions made in how

[8]In this section, we have used the words "small" and "large" to characterize different levels of the $NVI$ without being explicit about their meaning. This was a deliberate choice, since the model provides no welfare analysis and no other indication on how to evaluate the overall magnitude of the $NVI$. In short, the burden of interpreting what constitutes small or large values for the $NVI$ is the policymaker's.

[9]The bounds derived by Glasserman and Young (2015) are actually stronger than this - applying to the probability of node $i$ causing default through contagion to a given *group* of firms.

to best utilize them. The resulting datasets yields a quarterly series for the NVI spanning 2002:Q1 to 2016:Q4.

### 2.3.1 Assets and Liabilites of Bank Holding Companies

Line-item balance sheet information for bank holding companies comes from quarterly filings of the Federal Reserve's FR-Y9C reporting form[10]. The public nature of this data, as well as the level of granularity in reported asset and liability classes, make this form particularly well-suited to our analysis.

Our objective in using FR-Y9C data is to estimate the outside assets and liability connectivity of each firm in the FR-Y9C's sample. This involves classifying each of the form's asset and liability line items as *inside* or *outside* the financial system. We produce this classification for each of the line items in the current FR-Y9C balance sheet, and apply those classifications across each firm in the sample[11]. In cases where this binary classification seems inappropriate, we split the value of the field, classifying fifty percent of its magnitude as inside the system and fifty percent as outside the system[12]. The final two columns of Tables B.7 and B.8 provide these classification breakdowns for current variables (or groups of variables) in the form.

In past versions of the form, line-items were often less granular. To apply our inside-vs-outside classifications (made based on the current form's line-items) backward to previous form versions, we find the variables in each past form that include the same assets or liabilities as a given group of variables in the current form (the latter group of variables is typically larger, reflecting a movement towards. increasing form granularity over time). We then compute a firm-specific percentage of the total value of the current-form variable group that is attributable to each individual variable in the group during the first year that the variable group was reported[13]. By applying this estimated share back through time to 2002, we create a series for each FR-Y9C variable that is roughly consistent over time[14]. For a detailed view of the current-form variable groups identified, and the method used to extend them back to 2002, see Tables B.7 and B.8.

[10]An FR-Y9C filing is required by each domestic bank holding company (BHCs), savings and loan holding company, US intermediate holding company, and securities holding company with total assets exceeding one billion dollars

[11]The 'current' iteration of the form used in this paper is that from December 2016. For brevity, we will continue to refer to this as the 'current' form.

[12]Section 2.5 shows that our estimates are not very sensitive to alternative assumptions about the share of inside and outside assets and liabilities in these more ambiguous categories.

[13]To consider a simple but illustrative example - say we have determined that the asset categories contained in variables Y1 and Y2 of the current form are the same as those in variable X from some earlier version of the form. We then define $P_{Y1,i}$ and $P_{Y2,i}$ for firm i as the average of $\frac{Y1_i}{Y1_i+Y2_i}$ and $\frac{Y2_i}{Y1_i+Y2_i}$ in the first year that both Y1 and Y2 are reported. Firm i's imputed values for Y1 and Y2 in the early sample then becomes $P_{Y1,i}X$ and $P_{Y2,i}X$.

[14]In practice, these breakdowns are only important when the group of current-form variables includes two or more different in-vs-out classifications. Otherwise, the total sum of variables is directed into the same categorization, and any variable-by-variable divisions within the total sum become irrelevant.

### 2.3.2 FDIC-Insured Deposits of BHCs

We wish to avoid classifying any FDIC-insured deposits from the commercial bank subsidiaries of BHCs as inside the financial system, since those deposits are likely not held by financial firms and are ultimately government liabilities (which reside outside the network). To separate FDIC-insured deposits from a BHC's total deposits, we use quarterly data from the FFIEC 041 (also known as the Call Report), as collected by the Federal Financial Institutions Examination Council[15]. After matching each commercial bank to its BHC parent, we subtract the estimated quantity of FDIC-insured deposits, as reported in the Call Report, from the BHC's total deposits. Only this final "uninsured" value of deposits is considered inside the financial system, for the purposes of further analysis[16].

The process of matching Call Report data to balance sheet data on its BHC parent can become complicated, particularly around BHC mergers, acquisitions, or legal classification changes. To find the final BHC parent of each commercial bank, we use a bank-parent matching hierarchy maintained by the Federal Reserve Bank of New York. We then match commercial banks to their parent BHCs based on the BHC's RSSID identifier code. When this process does not lead to any commercial bank matches for a given BHC, we then match that BHC to all the commercial banks owned by the BHC's parent organization. This can occassionally lead to overestimates of the FDIC-insured deposits of these BHCs, but we have found that this process almost always yields sensible-looking series for these BHCs' financial connectivities. The alternative approach, where we do not conduct the second matching procedure and list FDIC-insured deposits as '0' for these firms, often yields impractically-large financial connectivities.

### 2.3.3 Probabilities of Default

The probabilities of default $\delta$ for each firm in equation 2.3 are the true - or physical - probabilities of default. As such, any risk-neutral estimate of a firm's default probability (such as those commonly extracted from credit default swaps or corporate bond spreads) would be inappropriate for calculating our NVI.

We consider Moody's Analytics' (formerly KMV's) Expected Default Frequency (EDF) series to be suitable for our analysis. The EDF measure uses typical lognormal assumptions and an options-pricing approach to equities to determine which variables should theoretically be important for determining a given firm's probability of default. They then use them to

[15]More specifically, our primary variables of interest from the form are RCON2200 (total domestic deposits) and RCON5597 (estimate of uninsured domestic deposits). Our estimate of insured deposits becomes the difference between these two fields. It is worth noting that our final estimate of uninsured deposits (which is then used in our index) is the difference between this estimate of insured deposits and the FR-Y9C form's value for firm domestic deposits (*not* the Call Report's domestic deposits variable). It is our understanding that the FR-Y9C form, as it pertains to entire BHCs instead of just commercial bank subsidiaries, includes a better estimate of total deposits for our purposes.

[16]In our benchmark setup for the NVI, 100% of uninsured domestic deposits are counted as *inside* the system. While this is likely close to accurate for the custodian banks (banks whose deposits are primarily safeguards of the assets of other banks) in our sample – namely State Street and Bank of New York Mellon – this is certainly unrealistic for many of the other BHCs in our panel. Section 2.5 includes robustness exercises on different configurations, including one allowing for more firm-specific allocation percentages. In short, this decision makes little difference in our final NVI series.

fit an empirical model of default probabilities using Moody's extensive database of historical defaults, as explained in Nazeran and Dwyer (2015)[17].

Equation 2.3 calls for the probability of default due to shocks to outside assets, although Moody's Analytics' EDF makes no distinction between the actual sources of default losses. Rather than attempt to back out the theoretically-appropriate default probability from these EDFs, we simply include the EDF value itself as $\delta_i$ for each firm in equation 2.3, with the understanding that this probability is in fact an upper-bound on the direct default probability. Relying on the fact that the NVI is itself an upper-bound, the bounds obtained from an NVI calculated this way will still be valid.

Moody's EDF model produces a daily series of physical expected default frequencies at one-year horizons. We define a firm's quarterly EDF measure to be the average of its daily measures over a given quarter.

### 2.3.4 Non-BHC Financial Firms

For financial firm subsectors whose firms do not file FR-Y9C forms, we include nodes into the NVI using less granular firm-level balance sheet information and subsector-level data on assets and liabilities from the Financial Accounts of the United States, maintained by the Board of Governors of the Federal Reserve System.

To incorporate a new firm (or, in this case, group of firms) into our NVI measure requires us to know each firm's individual default probability $\delta_i$, its outside assets $c_i$, and that firm's liability connectivity $\beta$ *if* that firm's $\beta$ becomes the new $\beta^+$ for the system. We first make the necessary simplifying assumption that the $\beta^+$ selected from the firms in our FR-Y9C sample correctly identifies the $\beta^+$ for the entire network[18].

Left to determine is how the inclusion of other financial subsectors affects the other component of the NVI, the weighted average default probability $\frac{\sum \delta_i c_i}{\sum c_i}$. We approximate the value of this componenet for the subsectors not covered by the FR-Y9C by first constructing an estimate of the total outside assets of each of those subsectors from the Financial Accounts of the United States and then computing an average default probability weighted by assets for each new subsector using total firm asset values and Moody's EDF measures. Total quarterly assets for each firm, compiled from that firm's financial releases and filings, are also available in the Moody's EDF dataset[19].

More explicitly, let $Y$ denote the set of firms in our FR-Y9C sample, $S = \{S_1, S_2, \cdots\}$

[17]Moody's historical defaults dataset considers government rescues as default events, if the rescue specifically saved the firm from default. So, in that sense, the EDF series can be considered as a probability of default without government intervention. As the model of Glasserman and Young (2015) does not include the possibility of government rescue, this empirically estimated probability closely matches its model counterpart.

[18]In fact, the $\beta^+$ we select from the FR-Y9C sample for the NVI is the highest financial connectivity found in the top 20 BHCs by assets. See 2.5 for a discussion of this decision, and an analysis of robustness to different selections.

[19]This is done using a method similar to that for the FR-Y9C, categorizing different Financial Accounts asset classes as inside or outside the system. See Table B.10 for the precise 'inside' vs 'outside' classification used for different variables in the release. Line-items from the Financial Accounts are much coarser than those in the FR-Y9C, making this an admittedly cruder method of classification. As Section 2.5 shows, however, this breakdown has very little effect on the final NVI measure.

denote a set of sets, with each individual element $S_j$ being the set of firms belonging to some new financial subsector, and let $A$ denote the entire financial network $Y \cup S$. Then,

$$\frac{\Sigma_A \delta_i c_i}{\Sigma_A c_i} = \frac{\Sigma_Y \delta_i c_i + \Sigma_{S_j \in S}(\Sigma_{i \in S_j} \delta_i c_i)}{\Sigma_Y c_i + \Sigma_{S_j \in S}(\Sigma_{i \in S_j} c_i)} \approx \frac{\Sigma_Y \delta_i c_i + \Sigma_{S_j \in S}(\bar{\delta_{S_j}} \Sigma_{i \in S_j} c_i)}{\Sigma_Y c_i + \Sigma_{S_j \in S}(\Sigma_{i \in S_j} c_i)} \tag{2.4}$$

$$\text{where } \bar{\delta_{S_j}} = \frac{\Sigma_{i \in S_j} \delta_i a_i}{\Sigma_{i \in S_j} a_i}, \text{with } a = \text{total assets of firm i} \tag{2.5}$$

This computation is only an approximation of the true $\frac{\Sigma_A \delta_i c_i}{\Sigma_A c_i}$ for two reasons. First, the weighted average probability of default per sector $\bar{\delta_{S_j}}$ is weighted here by *total assets* per firm, where a more precise measure for the NVI would be weighted by *outside assets* per firm. Second, and more significantly, our sample of average default probability is limited to those firms for which we have a Moody's EDF measure - namely, to publicly-traded firms. Provided that our computed averages are good representations of the entire sector, then the NVI constructed using this approximation should remain a useful upper-bound on network spillovers that allows us to include a much larger portion of the US financial system than the FR-Y9C sample alone would allow.

### 2.3.5 Subsector EDF Samples

Per Section 2.3.4, we wish to include subsector-wide averages in our NVI for security broker dealers, insurance companies, real estate investment trusts, and an 'other' category for several other types of financial firms. To determine which firms should comprise each subsector's sample in equation 2.5, we use Moody's Analytics' own internal sectoral classification system, pairing their categorizations with the subsector definitions given in the Financial Accounts of the United States. For the purposes of calculating outside assets $c$ for the 'other' sector, we sum across the Financial Accounts subsectors for credit unions, finance companies, funding corporations, and issuers of asset-backed securities[20].

To show the relative magnitude of assets assigned to these difference subsectors by the Financial Accounts of the United States, we plot the percentage of total network assets (defined as the sum of total financial assets in each subsector described above, plus the Financial Accounts' total financial assets for BHCs) attributable to each of these subsectors in Figure 2.3. As Figure 2.3 shows, BHCs are by far the largest financial subsector by assets, meaning that the weights given to BHCs' default probabilitys will, in aggregate, be larger than the weights assigned to any of the included subsectors.

To assess whether the samples used to compute each of the default probabilities within equation 2.4, in Figure 2.4 we plot the percentage of total network assets (with the 'network' defined as the sum of total financial assets in the subsectors of Figure 2.3) that are accounted for by the total assets of firms whose default probabilities directly enter into the NVI - either individually if that firm files an FR-Y9C, or as part of the sample for computed a subsector's average default probability. The red line in Figure 2.4 plots the same value, if we consider

[20]The 'other' category is the only one for which finding an appropriate subsample within the EDF dataset it not straightforward. We choose to include any firms with sectoral tags of 'Finance Companies', 'Investment Management', or 'Finance Not Elsewhere Classified' in this node's average probability calculation.

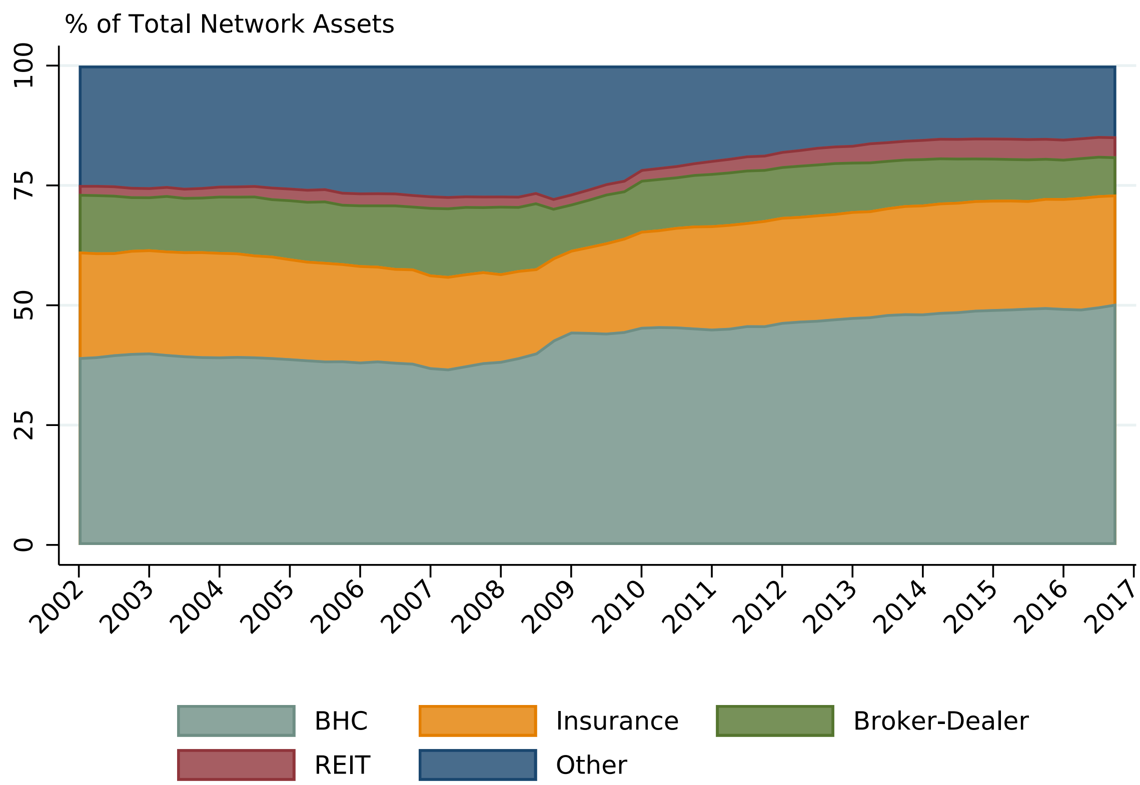


Figure 2.3: **Total Financial Assets for each Network Subsector, as a Percentage of Total Network Assets.** According to the Financial Accounts of the United States, BHCs comprise by far the largest percentage of network assets

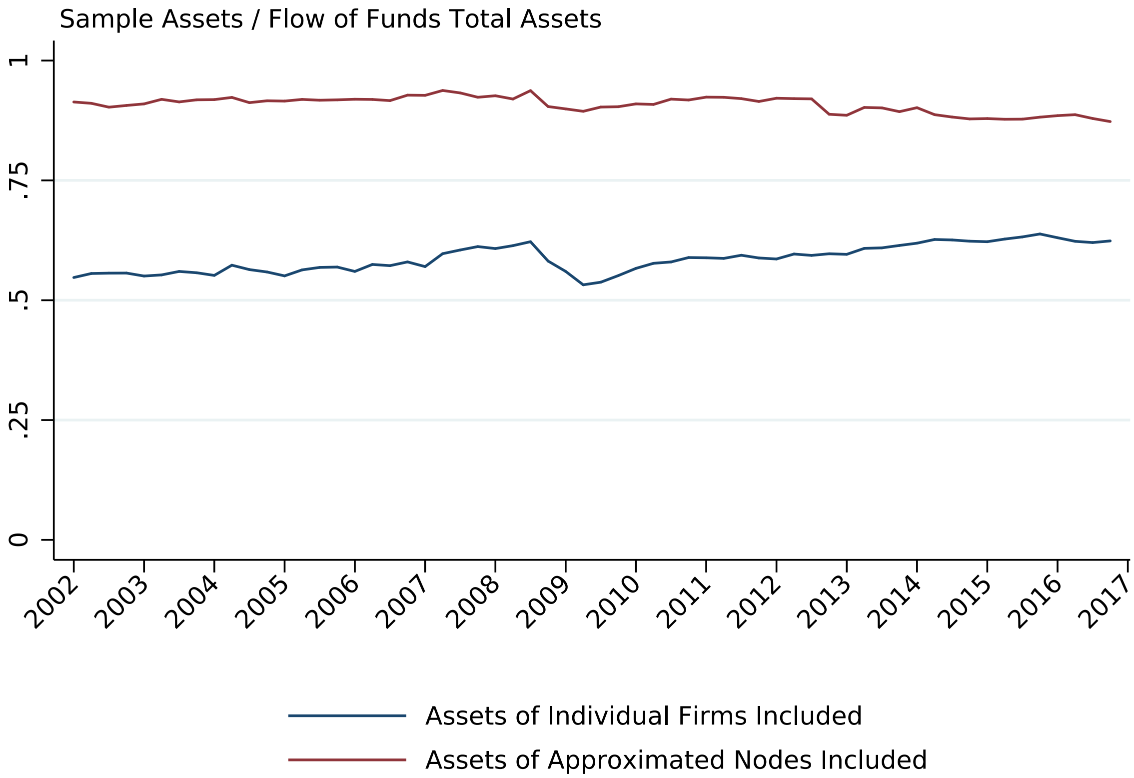


Figure 2.4: **Data Coverage of EDF Sample for NVI, Against Total Network Assets.** Only a subsample of firms in each subsector have their default probabilities directly enter the NVI. The first line sums the total assets of firms whose EDF measures enter the NVI in some form - either individually if that firm filed an FR-Y9C, or as part of a subsector average default probability sample - divided by total network assets from the Financial Accounts of the United States. The second line plots the same, if we instead consider each approximated subsector node per equation 2.4 to cover the entire sum of that subsector's assets.

each approximated node in equation 2.4 to cover all of the assets from that subsector's Financial Accounts entries. The blue line shows that even the most conservative measure covers consistently more than 50% of the assets of the entire U.S. financial system. If we consider our network to cover the sectors of each approximated subsector node, that coverage becomes even higher, as shown in the red line of Figure 2.4[21].

### 2.3.6 Defaulting firms

The model of Glasserman and Young (2015) includes an explicit assumption that no nodes included in the system are initially in default (defined as having book liabilities greater than book assets). To avoid including any such firms in our estimates of equation 2.4, we use Moody's Analytics' Default and Recovery Database to identify dates of bankruptcy filing. If a firm files for bankruptcy at any point during our sample period (2002-Q1 to 2016-Q4), then no expected default frequency data is used for that firm after the date of filing.

## 2.4 Results

### 2.4.1 Network Vulnerability Index Estimates

Figure 2.5 plots the NVI — the upper bound on expected network default spillovers. The figure shows the main result of our paper: When estimated empirically, vulnerability to network spillovers can range from negligible to large.

The NVI was essentially zero from 2002-Q1 to 2007-Q4, with only a slight increase in 2007-Q3 and 2007-Q4, which immediately implies that expected vulnerability to network default spillovers were negligible for this period. Contrary to some narratives of the crisis, we do not observe any substantial buildup of network fragility of the kind we study in the years leading up to the crisis. To understand this result, we decompose our spillover measure into two factors: the weighted average of probabilities of default ($\frac{\Sigma \delta_i c_i}{\Sigma c_i}$) and a 'connectivity multiplier' ($\frac{1}{1-\beta^+}$) that captures the magnitude with which initial losses in outside assets can be transmitted and amplified through network connections. The final NVI measure is the product of these two components.

As Figure 2.6 shows, both factors contribute to the low spillover measure in the period 2002-Q1 to 2007-Q4. Because probabilities of default were miniscule in this period, the weighted average default probabilities were close to zero. Since Moody's EDF probabilities are physical, they are adjusted for risk and thus unlikely to arise because of any low risk premium observed during this period[22]. Over the same period, the connectivity multiplier

[21]Note that the actual assets attributed to each subsector for the purposes of calculating these coverage statistics are total subsector assets after any deductions from FR-Y9C sample overlap. This adjustment, as well as the fact that FR-Y9C coverage of BHC assets in the Financial Accounts is not 100%, are why the second line in Figure 2.4 is not mechanically 100%.

[22]In addition, version 9 of KMV generally adjusts probabilities of default (upwards) for this period taking into account the ex-post defaults observed during the crisis that were not expected before it, minimizing the concern that our results are driven by any potential underestimation of default probabilities before the crisis.

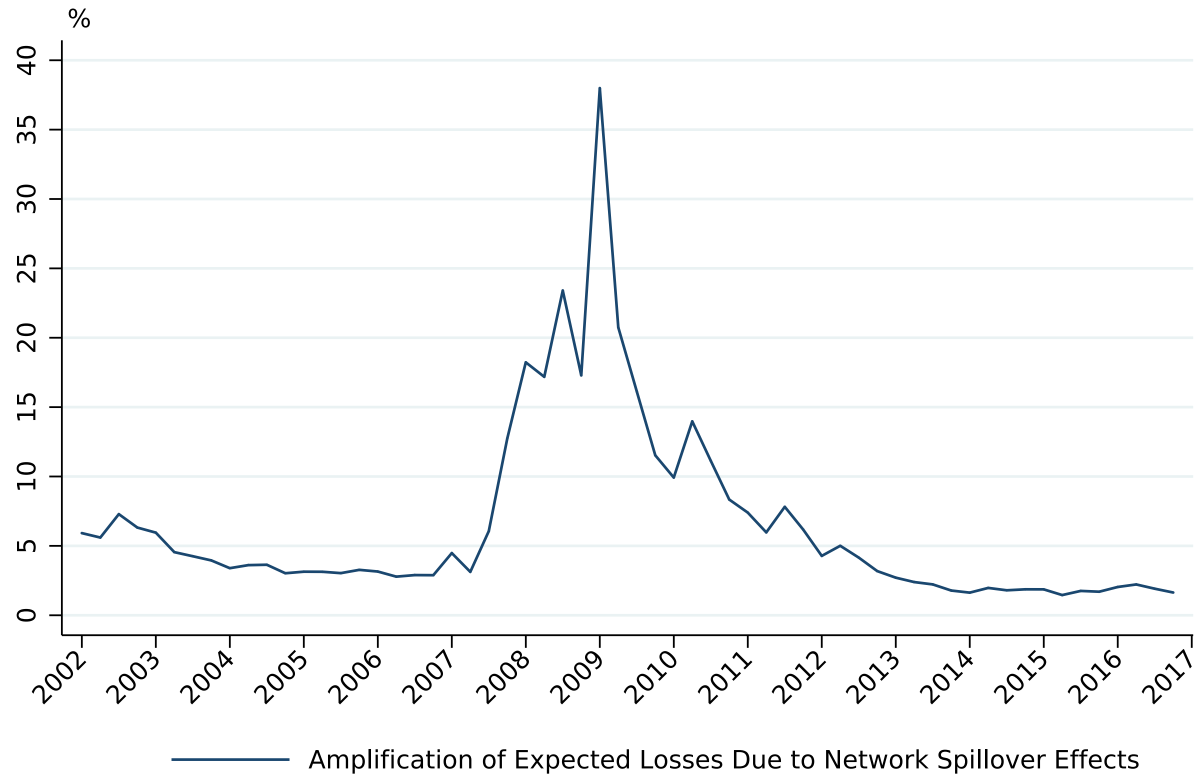


Figure 2.5: **Network Vulnerability Index (NVI).** The NVI is an upper bound on expected losses due to default network spillovers in the U.S. financial system, expressed as a share of initial (exogenous) losses to assets outside the network. Between 2008 and 2012, network default network spillovers amplified expected losses by between 5 and 25 percent.

declined by 10 percent. Thus, neither the vulernability of firms to default nor the inner topology of the financial network signaled any increased vulnerability.

During the height of the crisis, between 2008-Q1 and 2008-Q4, outside assets (especially real estate) experienced sharp declines in realized and future expected values, pushing up our measure of spillovers. The connectivity multiplier, in contrast, was a mitigating factor, as it noticeably declined, reflecting financial institutions desire to reduce their counterparty exposure to each other in times of stress. In fact, the decline in the connectivity multiplier over 2008 was as large as the decline observed over the six preceding years 2002-2007. Figure 2.7 shows $\beta^{+}$, the maximum liability connectivity selected at each point of the sample, which drives this dynamic. Overall, in late 2008 the increase in default probabilities outweighed any mitigation from lower connectivity. Our estimates indicate that expected network default spillovers over this period could amplify total initial losses by at most 11.4 percent. Whether 11.4 percent should be considered a small or large number is in the eye of the beholder.

In 2009-Q1, after the failure of several financial institutions and with the crisis now

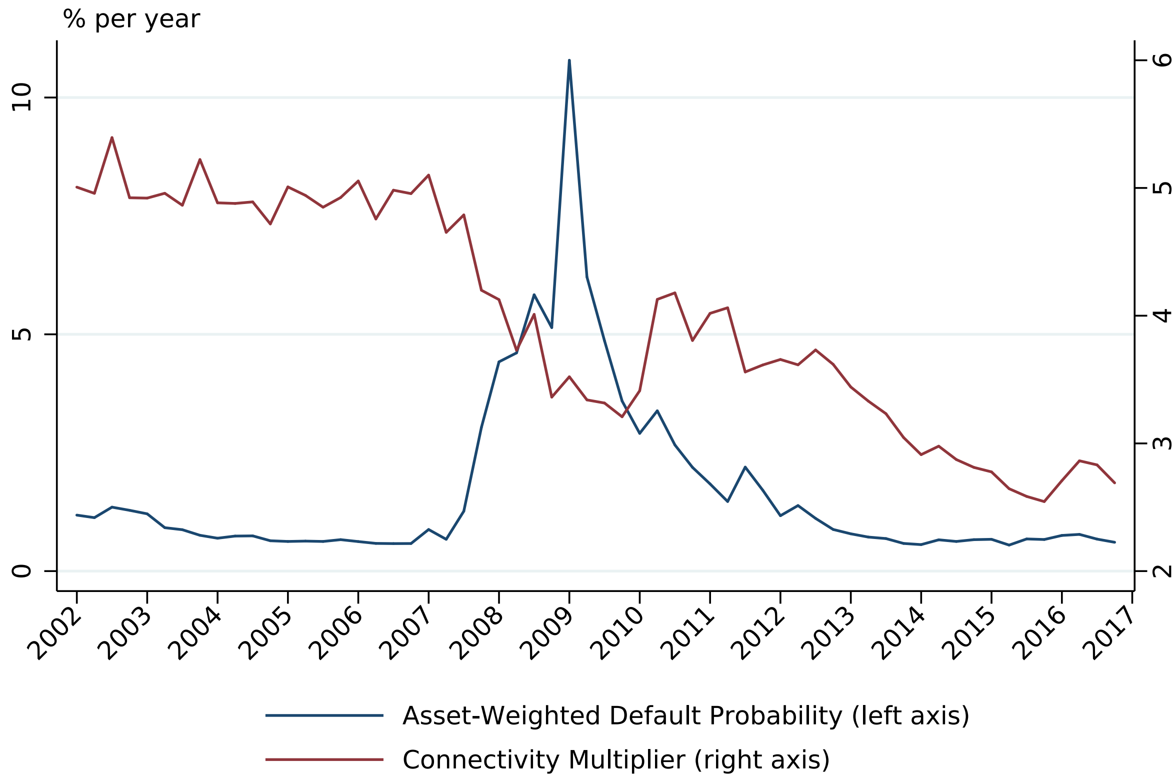


Figure 2.6: **NVI Decomposition.** The NVI is the product of the asset-weighted probability of default of firms inside the network and a "connectivity multiplier" that captures the degree of amplification and transmission created by defaults inside the network..

global in scope, the spillover measure jumped markedly, with our estimates indicating that expected network default spillovers over this period could amplify total initial losses by up to 25 percent, the largest value observed in our sample. The large increase from 2008-Q4 to 2009-Q1 was driven by both default probabilites and financial connectivity. Expected losses increased not only because real estate kept deteriorating, but also because the slowdown in real economic activity induced an increase in expected losses for almost all categories of outside assets, including commercial, industrial and consumer loans. Financial connectivity also increased, driven by the failure of some network nodes and the merger and consolidation of various other nodes. Keeping in mind that our estimates are always upper bounds and not point forecasts, to the best of our knowledge, a 25 percent amplification is the largest empirical estimate for network spillovers in the literature. Estimates that exceed 25 percent in the literature usually rely on additional amplification mechanisms (like bankruptcy costs – see Section 2.5 – or the interaction of default cascades with other phenomena, such as runs and fire-sales. The NVI remains highly elevated in 2009-Q2 at 23 percent before dropping to 18.6 percent and 12.5 percent in 2009-Q3 and 2009-Q4, respectively.

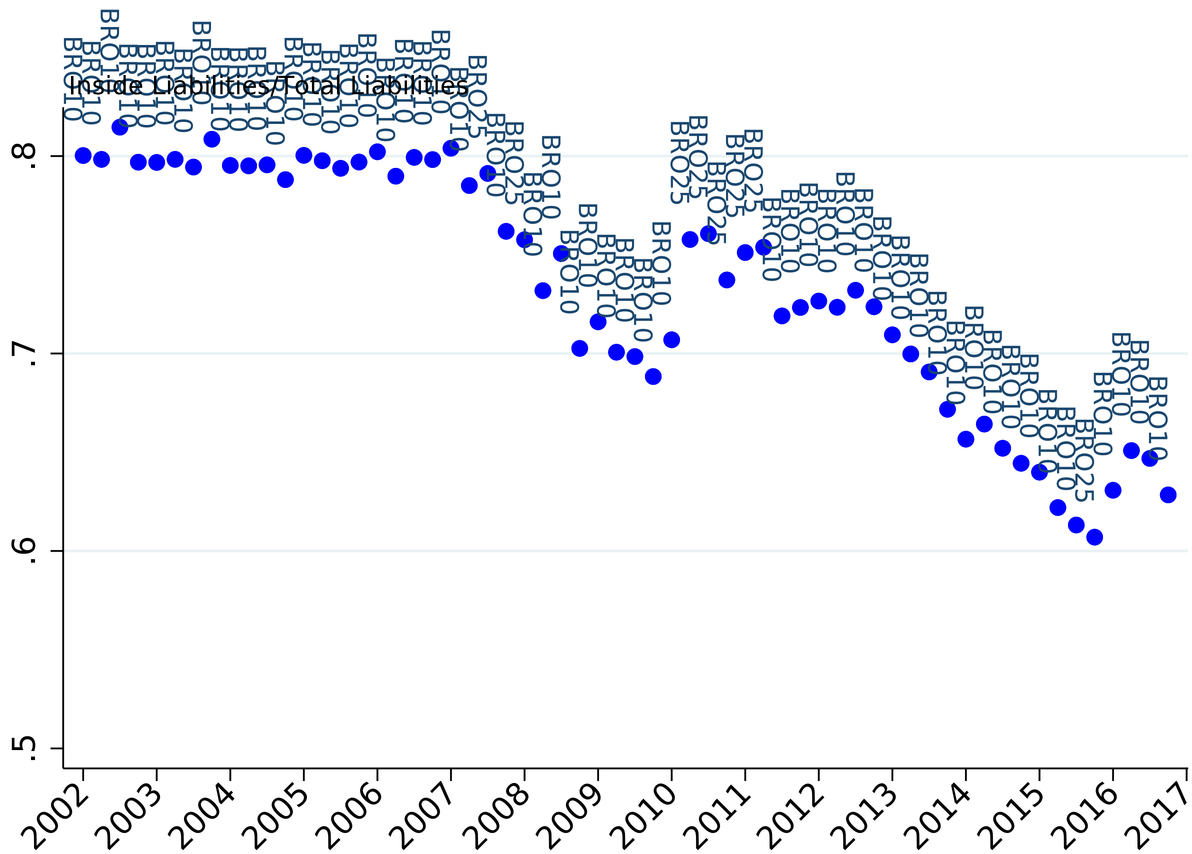


Figure 2.7: **Maximum Liability Connectivity Among Large BHCs** ($\beta^+$). The most interconnected BHC before the financial crisis was JP Morgan & Chase. After Goldman Sachs and Morgan Stanley join the BHC sample in 2009, they become the most interconnected firms for the remainder of the sample.

After 2009, our spillover measure hovered between 5 and 10 percent until 2012, when the European crisis started to recede. From 2013 onward, the measure steadily decreased and reached pre-crisis levels by 2015. The most important contributor to this decrease was the reduction in the probability of default of financial institutions, particularly bank holding companies that strengthened their equity capital positions substantially over this period. Financial connectivity remained elevated until 2014 and has been declining ever since. Starting in 2015, average default probabilities have slightly increased. However, because financial connectivity has continued to decline between 2015 and 2017, our overall measure of spillovers is little changed.

### 2.4.2 Results of FR-Y9C Asset and Liability Line-Item Classifications

Tables 2.1 and 2.2 show what percentage of BHC inside and outside assets and liabilities are attributable to each of several broad categories of balance sheet items, based on our classifications in Tables B.7 and B.8.

These being BHCs, it is unsurprising that deposits comprise a large portion of both inside and outside liabilities. Inside assets are mostly comprised of repurchase agreements, federal funds, and deposits, while outside assets are mostly loans and mortgage-backed securities.

| BHC Assets Inside Financial System (%) | | BHC Assets Outside Financial System(%) | |
|---|---|---|---|
| Repos and Fed Funds | 31.93 | Loans | 60.25 |
| Interest Bearing Deposits | 28.53 | Agency MBS | 13.82 |
| Private Label ABS | 6.60 | State, Treasury, and Agency Debt | 7.37 |
| Goodwill | 5.70 | Other Securities | 4.61 |
| Other Trading Assets | 4.83 | Interest Bearing Deposits | 3.62 |
| Derivatives | 3.67 | Noninterest Bearing Deposits | 1.48 |
| Private Label MBS | 1.92 | Goodwill | 1.34 |
| Other MBS | 1.05 | Other Trading Assets | 1.14 |
| Other | 15.77 | Other | 6.37 |
| % of BHC Assets | 19.06 | % of BHC Assets | 80.94 |

Table 2.1: **Shares of BHC Assets Inside and Outside the Financial System By Category, 2016-Q4.**

| BHC Liabilities Inside Financial System (%) | | BHC Liabilities Outside Financial System (%) | |
|---|---|---|---|
| Uninsured Domestic Deposits | 61.42 | Insured Domestic Deposits | 62.78 |
| Repos and Fed Funds | 10.73 | Foreign Deposits | 17.38 |
| Longer Term Debt | 9.67 | Longer Term Debt | 8.08 |
| Trading Liabilities | 4.38 | Short Term Debt | 3.27 |
| Short Term Debt | 3.92 | Subordinated Debt | 2.62 |
| Derivatives | 2.96 | Other | 5.87 |
| Other | 6.92 | | |
| % of BHC Liabilities | 45.51 | % of BHC Liabilities | 54.49 |

Table 2.2: **Shares of BHC Liabilities Inside and Outside Financial System by Category, 2016-Q4.**

### 2.4.3 Sector-Specific Average Default Probabilities

As we discuss in Section 2.3.4, we calculate an asset-weighted average default probability for several different groupings of firms as proxies for the actual average EDF measure of their entire respective financial subsectors, so that the assets of firms in those subsectors can be incorporated into the final NVI. Figure 2.8 shows the final series for these average default frequencies at each point in our quarterly sample. For ease of comparison, we also plot an analagous average default probability (again asset-weighted) for the portion of our sample included in the FR-Y9C report. Figure 2.8 shows that the default probabilities for each subsector exhibit similar movements. All sectors show greatly heightened default probabilities during the financial crisis, with the average default probabilities of broker dealers and the 'other' category elevating earliest in the crisis and remaining heightened for the longest. The largest default probability magnitudes come from the 'other' category, and from real estate investment trusts, which experienced a number of defaults around this time[23].

### 2.4.4 Firm-Specific Contagion Indices

A useful understanding of our NVI measure can come from investigating the NVI's variables of interest for some of the largest firms in our sample. Figure 2.9 plots several important firm-specific variables that contribute to the NVI for four large BHCs - JP Morgan & Chase, Wells Fargo, Bank of America, and Citigroup. Measures that feed directly into the NVI - outside assets $c$ and connectivity $\beta$ - are plotted alongside the contagion index measure defined in Section 2.2.

Figure 2.9 shows how the general, system-wide dynamics described above play out for a few important financial institutions. The path of financial connectivity for these firms differs

[23]To see the firms whose default probabilities are included in each sector subsample at a snapshot of our data sample (2016-Q4), as well as the asset-weights assigned to them, see Tables B.6, B.2, B.3, and B.4.

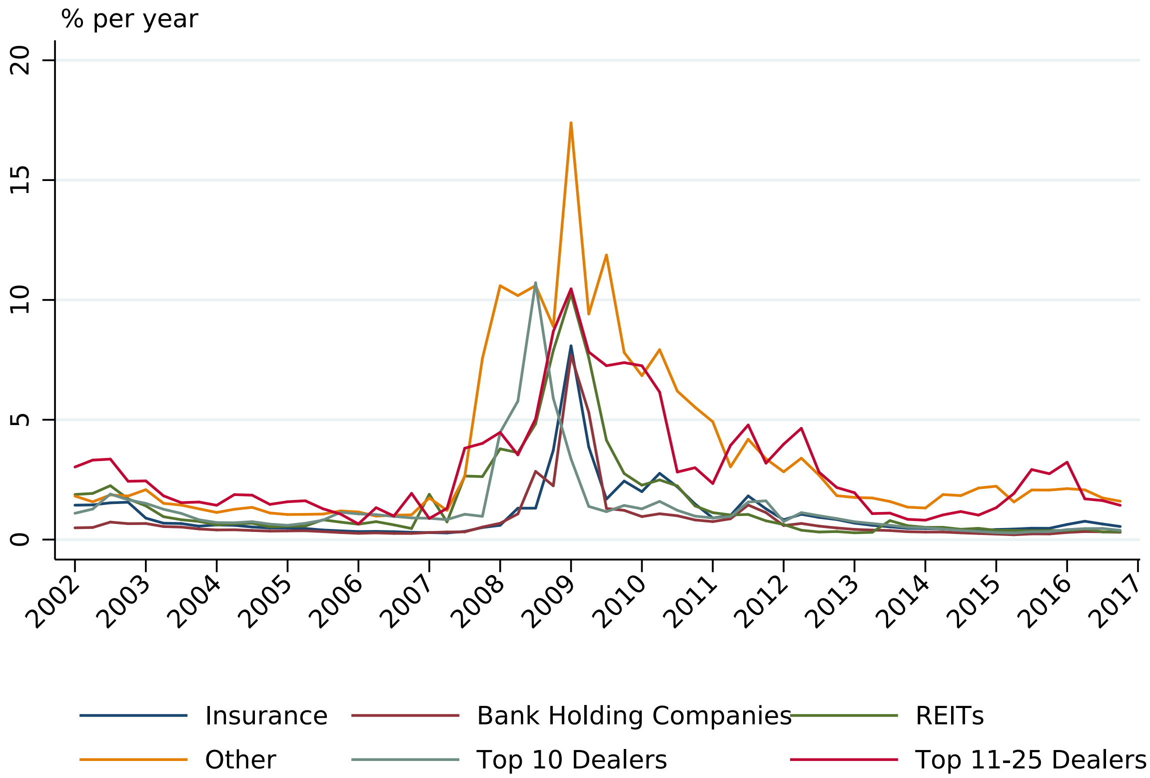


Figure 2.8: **Sector-Wide Asset-Weighted Average Probabilities of Default.** Using firm-specific EDF measures and total firm assets, we calculate asset-weighted average expected default frequencies for each financial subsector in the estimated network at each quarter in our sample. These probabilities are used in our final network vulnerability measure to fill in portions of the financial sector not covered by the FR-Y9C. All sectors show greatly heightened default probabilities during the financial crisis.

from 2002-2008, but falls or remains steady for each of them either during the financial crisis or soon thereafter. Financial connectivity for each of these large firms had risen back to pre-crisis levels by the end of 2016.

Three of these four firms were parties to large-scale acquisitions of other financial firms at the time of the crisis (Bear Stears for JPM, Merrill Lynch for BAC, and Wachovia for WFC), which caused their outside assets (and assets generally) to increase around that time. Naturally, this causes increases in the ‘contagion index’, which is linked to the probability of a failure by that firm causing subsequent contagion defaults. As the smallest of these four firms’ contagion index is larger than any one of their net worths, we cannot rule out the possbility that a large exogenous shock to one firms’ assets could cause a contagion failure in another of these four firms. However, given the large size of each of these firms’ outside assets compared to each other’s contagion indices, we know that the probability of a shock to one

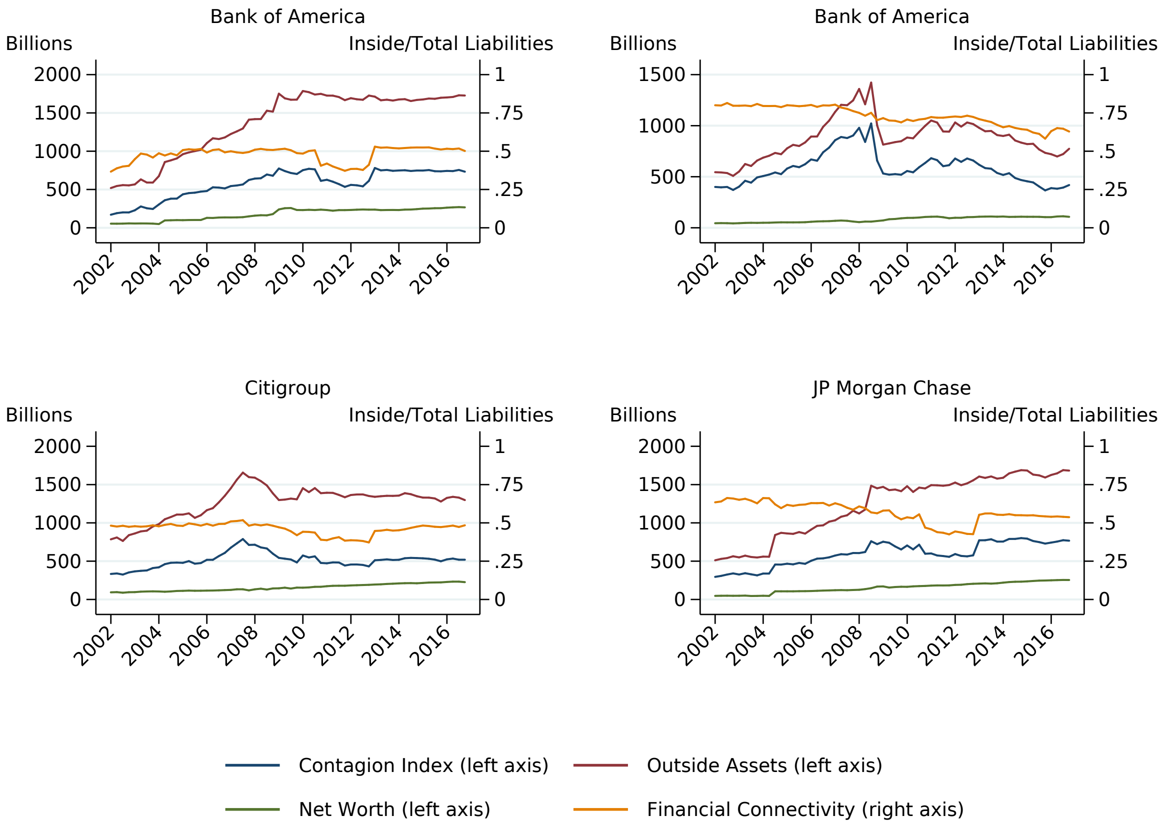


Figure 2.9: **Firm-Specific Variables for Select Large BHCs.** The figures show the net worth (the difference between total liabilities and total assets), total assets outside the financial system, the ratio of inside liabilities to total liabilites (connectivity), and the contagion index for several large BHCs.

firm's assets causing default in any other of these firms must be lower than the probability of that firm defaulting because of its own exogenous shock (see Glasserman and Young (2015). Table 2.3 shows the same field values for 19 of the largest BHCs in the sample for the final period of our data, 2016-Q4.

| | Contagion Index | Financial Connectivity | Outside Assets | Net Worth |
|---|---|---|---|---|
| JP Morgan Chase Co | 767.62 | 0.54 | 1683.01 | 254.40 |
| Bank of America Corp | 732.66 | 0.50 | 1727.24 | 266.84 |
| Wells Fargo Co | 624.17 | 0.45 | 1601.25 | 200.50 |
| Citigroup | 519.09 | 0.48 | 1297.92 | 226.14 |
| Top 10 Dealers | 418.97 | 0.63 | 774.12 | 107.39 |
| U S Bank | 164.01 | 0.45 | 415.33 | 47.93 |
| Top 11-25 Dealers | 121.74 | 0.52 | 292.91 | 56.78 |
| Pnc Financial Services Group | 104.62 | 0.39 | 316.84 | 46.85 |
| Bank of NY Mellon Corp | 91.46 | 0.52 | 216.21 | 39.58 |
| Capital One Finance Company | 80.22 | 0.28 | 332.25 | 47.51 |
| BBT Corp | 75.87 | 0.44 | 203.87 | 29.93 |
| Suntrust Bank | 73.89 | 0.44 | 193.30 | 23.62 |
| Fifth Third Bank | 56.52 | 0.49 | 131.46 | 16.23 |
| State Street Corp | 54.03 | 0.49 | 131.97 | 21.22 |
| Keycorp | 53.33 | 0.46 | 130.03 | 15.24 |
| American Express Co | 51.94 | 0.43 | 140.27 | 20.50 |
| Citizens Financial Group | 50.01 | 0.41 | 141.88 | 19.75 |
| Ally Financial | 38.53 | 0.28 | 148.69 | 13.32 |
| Regions Finance Company | 37.28 | 0.37 | 117.42 | 16.66 |

Table 2.3: **Select Firm-Specific Variables for Large BHCs, 2016-Q4**. The table shows the net worth (the difference between total liabilities and total assets), total assets outside the financial system, the ratio of inside liabilities to total liabilites (connectivity), and the contagion index for BHCs in the last period of our sample.

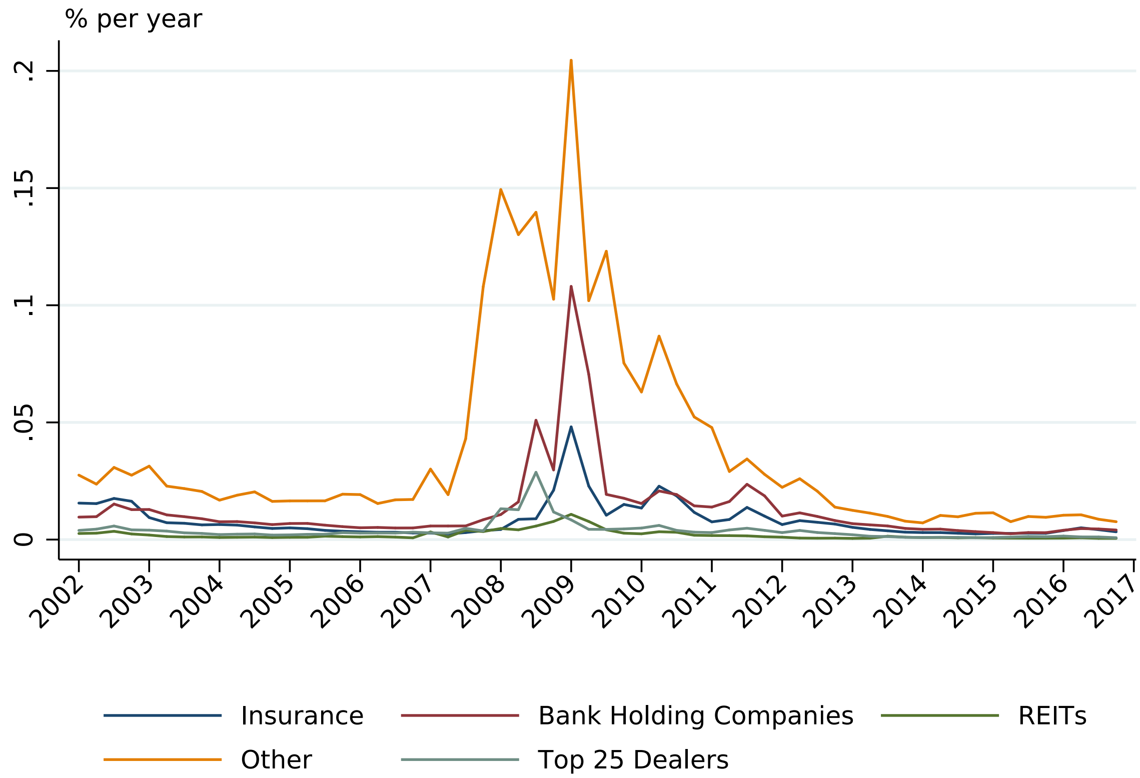


Figure 2.10: **Additive Contributions to NVI by Sector.** The contribution of each sector to the NVI is that sector's average default probability, weighted by the portion of system-wide outside assets belonging to the individual sector, and multiplied by the connectivity component of Figure 2.6. The 'Other' category consistently contributes the largest magnitude of any sector, followed by bank holding companies. Real estate investment trusts, due to their relatively low quantity of outside assets, contribute the least. The final NVI measure is the sum of each sector's contribution.

## 2.5 Robustness

We next test the robustness of our results to changes in a number of data treatment procedures and model assumptions.

### 2.5.1 Bankruptcy Costs

A common choice in the financial contagion literature is to impose additional costs of bankruptcy on firms that default. These additional costs are frequently cited as a potential factor for contagion risk. A necessarily incomplete list of the reasons for such costs includes:

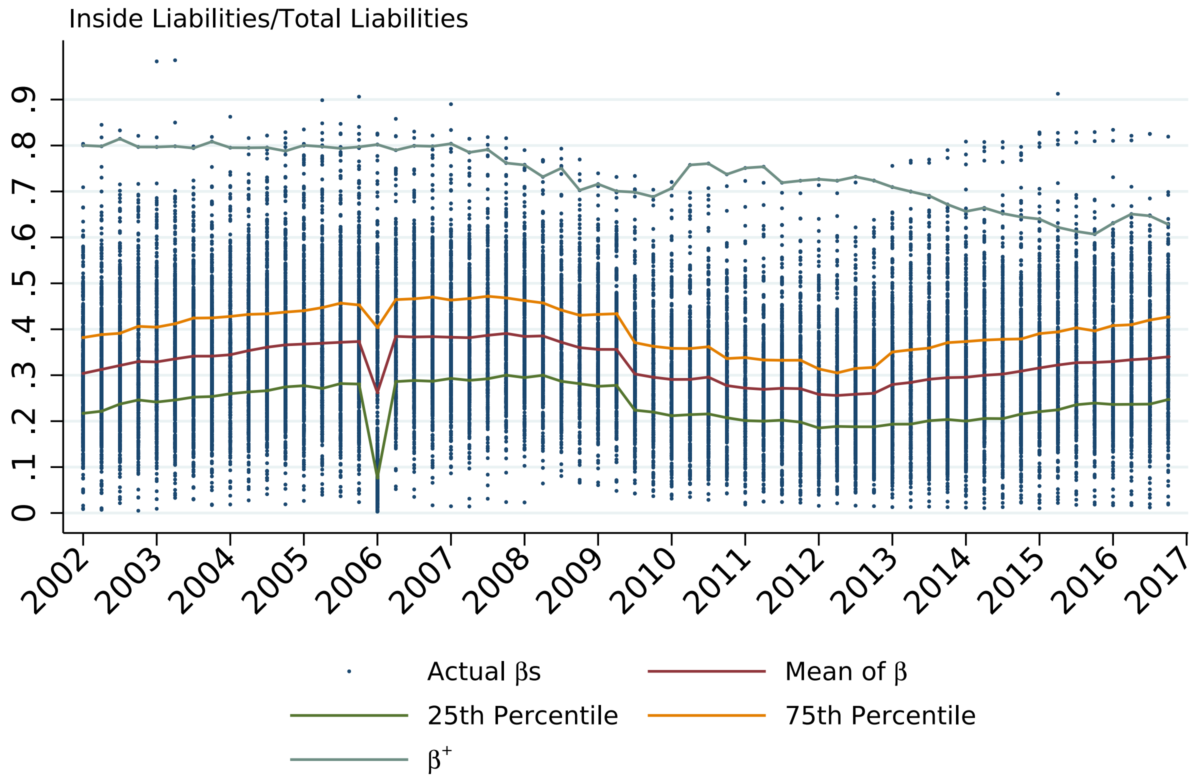


Figure 2.11: **Distribution of Liability Connectivities.** The figure shows distribution statistics for $\beta$, the balance sheet estimated portion of firm liabilities held by other financial firms, for bank holding companies in each quarter of our sample. The value $\beta^+$ refers to the high $\beta$ chosen for ther purposes of calculating the financial connectivity multiplier in the benchmark NVI.

Delay of payments, inefficient liquidations, penalties, funding shortages, downgrades on debt instruments, runs, legal fees, administrative expenses and, more generally, disruptions to the provision of financial intermediation services necessary to the real economy[24]. The Eisenberg and Noe (2001) framework can be easily modified to include these sorts of costs, and Glasserman and Young (2015) find a new upper bound on relative network spillovers in their presence. This new upper bound is

$$B = 1 + \frac{1}{(1-(1+\gamma)\,\beta^{+})} \frac{\sum \delta_i c_i}{\sum c_i}, \tag{2.6}$$

where $\gamma \in [0,1]$ are imposed when a firm defaults by reducing asset values by a share $\gamma$ of payment shortfalls..

[24]For a good example on these and other costs of failure, see the study of Lehman Brothers' case by Fleming and Sarkar (2014).

As Glasserman and Young (2016) note, estimating $\gamma$ empirically can be quite challenging. To test whether different bankruptcy costs change the central story of our NVI, Figure 2.12 plots the new upper bound in the presence bankruptcy costs of different magnitudes.

The dynamics of the NVI are the same under reasonable levels of bankruptcy costs. The level of the NVI under different $\gamma$ specifications also remains similar for every $\gamma$ except for the largest bankruptcy cost we consider, $\gamma = 30\%$. Even in the case of $\gamma = 30\%$, the conclusions to be drawn from the measure are much the same as those from our benchmark setting. That is, when the upper-bound is small enough to draw meaningful conclusions, it is small in both setups. In times when the benchmark NVI is too large to make definitive statements about the relative magnitude of contagion losses, it is similarly too-large in both configurations.

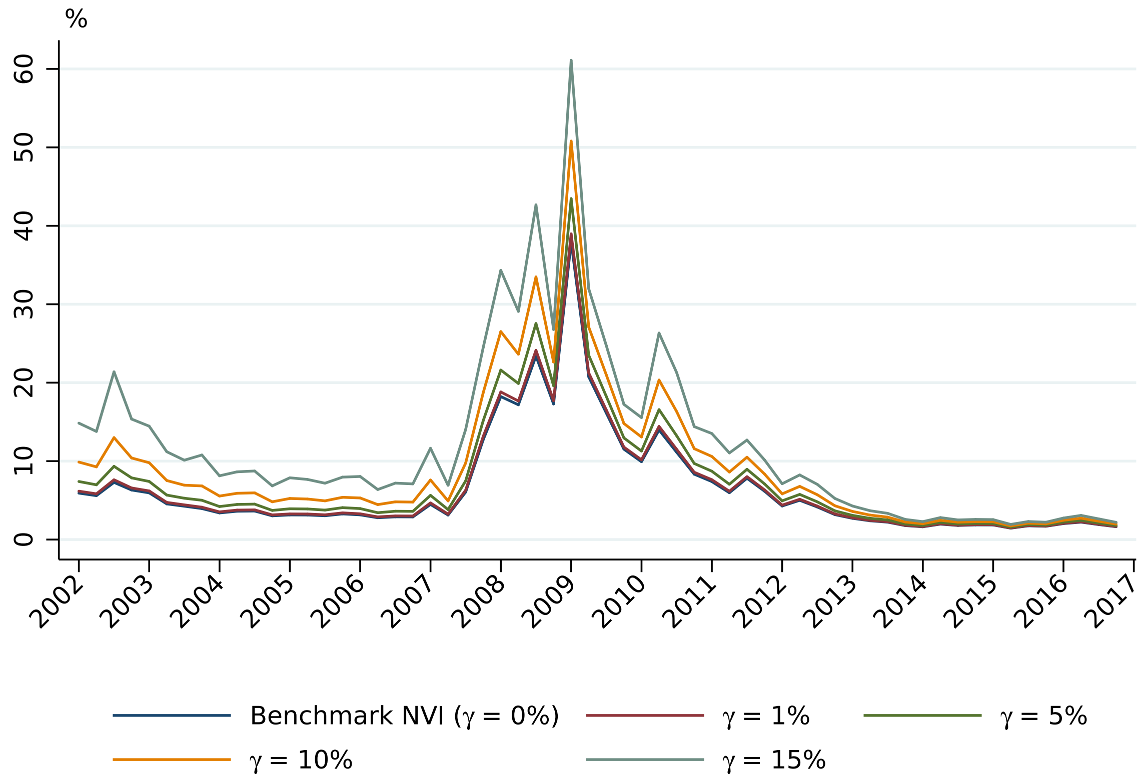


Figure 2.12: **Network Vulnerability with Additional Costs of Bankruptcy.** Adding bankruptcy costs to the model increases the vulnerability of the system to network spillovers, but does not change the qualitative nature of our results.

### 2.5.2 FR-Y9C Balance Sheet Classifications

Whenever a more absolute classification seemed inappropriate for a particular line item, we allocated 50% of the line item as inside and 50% as outside the financial system. Figure

2.13 shows how different allocations of these more uncertain balance sheet items as inside or outside the financial system change the NVI. The allocation of these fields can have a material effect on the magnitude of the series, particularly around the financial crisis. However, much as before, the conclusions to be drawn from the NVI, considering its nature as an upper-bound, are qualitatively the same across different allocation schemes of these assets and liabilities.

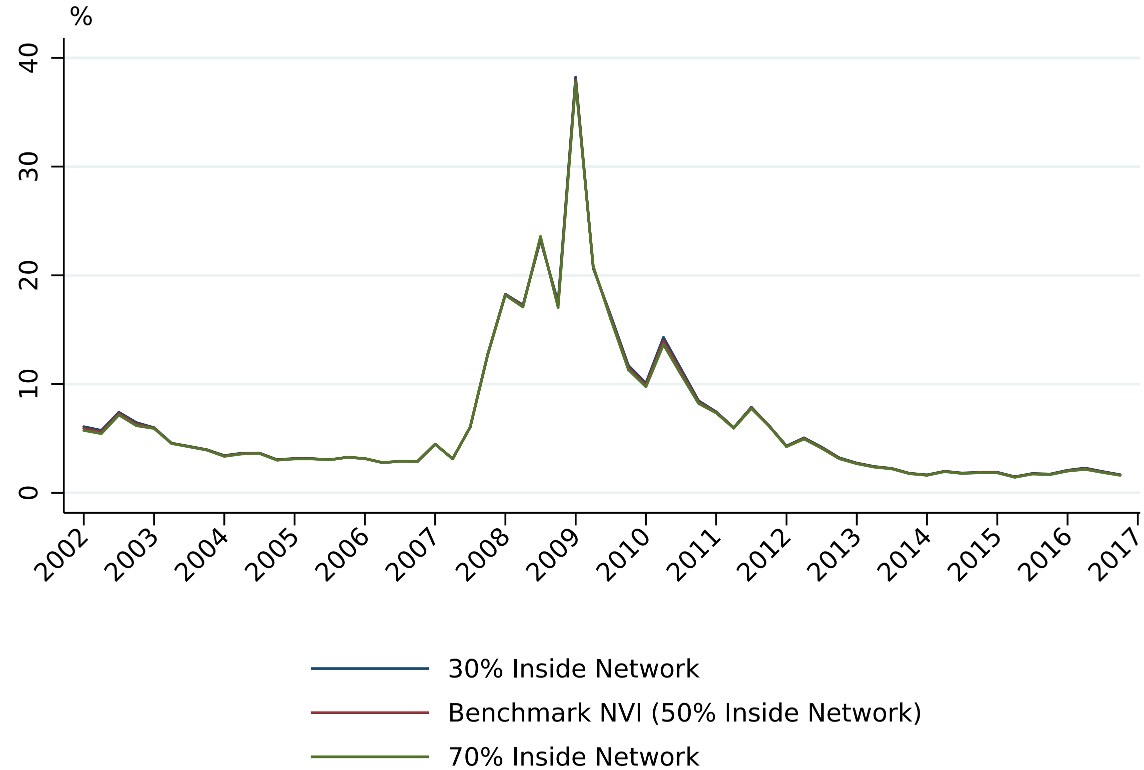


Figure 2.13: **Network Vulnerability Under Different Classifications of Hard-To-Classify and Liabilities.** The lines of this figure show the upper bound on expected network spillovers when we make different classification decisions for balance sheet items that are neither clearly inside nor outside the network. While different schemes can alter the magnitude of the measure, especially in the financial crisis, the measure remains qualitatively similar.

### 2.5.3 $\beta^+$ Selection Sample

We are also interested in learning how sensitive our NVI measure is to different selections of firm liability connectivity $\beta$ to use as the maximum connectivity $\beta^+$ in the NVI. In the benchmark setup, $\beta^+$ is chosen as the largest $\beta^+$ among the top 20 BHCs by assets at any point in the sample.

Figure 2.14 assesses the importance of this selection by applying different criteria for choosing a specific firm's $\beta$ as $\beta^+$ for the quarter. Panel (a) shows the NVI under these different criteria and Panel (b) shows the maximum liability connectivity $\beta^+$ chosen under each scheme.

In two of the test cases, we simply select $\beta^+$ as the second or third highest connectvitiy among large BHCs, instead of the highest. This can have a large impact on the size of the connectivity value $\beta^+$ in the measure, but ultimately any magnitude shifts are insufficient to cause any notable differences in the NVI itself.

Another potential selection method would be to select $\beta^+$ as the largest financial connectivity among all firms in a given quarter, regardless of the size of that firm. This setup could lend the measure more theoretical validity, as $\beta^+$ in the model of Glasserman and Young (2015) is in fact the largest of *any* node in the system. Figure 2.14 shows that this can have a dramatic impact on both $\beta^+$ and the NVI. Panel (b) of Figure 2.14 shows why this is the case. Early in the sample, the Investors Financial Services Corporation (ticker IFIN) has a very high $\beta$, consistently higher than 0.95. Following IFIN's acquisition by State Street in 2007, the full-sample $\beta^+$ drops substantially, coming much close to that chosen from the largest BHCs. Panel (a) of Figure 2.14 shows that, at this same time, the NVI calculated from this unrestricted $\beta^+$ becomes more similar in magnitude to that from the benchmark setup. Even though it is more theoretically appealing to use the largest $\beta$ across all firms in our sample, we judge the NVI to be more useful as an empirical gauge of network spillovers when it is not driven by the balance sheet composition of a single small firm[25].

### 2.5.4 Comparison with FR Y-15 Data

Beginning in 2012, the Board of Governors of the Federal Reserve System began requiring large US BHCs to file an FR-Y15 Systemic Risk Report, which reports (among other indicators) certain variables relating to total intrafinancial assets and liabilities[26]. The low yearly frequency, short sample, and narrower panel of firms available with this data make this form less appealing as a main source of data. However, information from the form is still useful as a cross-check, especially given that the form line items more closely correlate to the model's variables. The following figures incorporate these fields into our NVI measure in a variety of different ways.

First, Figure 2.15 uses FR-Y15 data by the most direct method, substituting applicable FR-Y15 fields into the NVI equation 2.3. One setup in Figure 2.15 uses the FR-Y15 value for intrafinancial liabilities to construct liability connectivity $\beta$ for firms who file an FR-Y15, then subsequently chooses the maximum of those newly-generated $\beta$ values as the maximum connectivity $\beta^+$ for the NVI. Another of Figure 2.15's configurations directly uses FR-Y15 data for all balance-sheet items in the NVI's computation. To do this, we first limit our

[25]However, the fact that IFIN has the largest $\beta$ in the sample is unsurprising, and serves as a reassuring check on the validity of our inside/outside liability classifications. IFIN specifically provided asset management services to US financial services industry, making it the perfect candidate for large financial connectivity.

[26]Available at time of publication at https://www.ffiec.gov/nicpubweb/nicweb/Y15SnapShot.aspx.

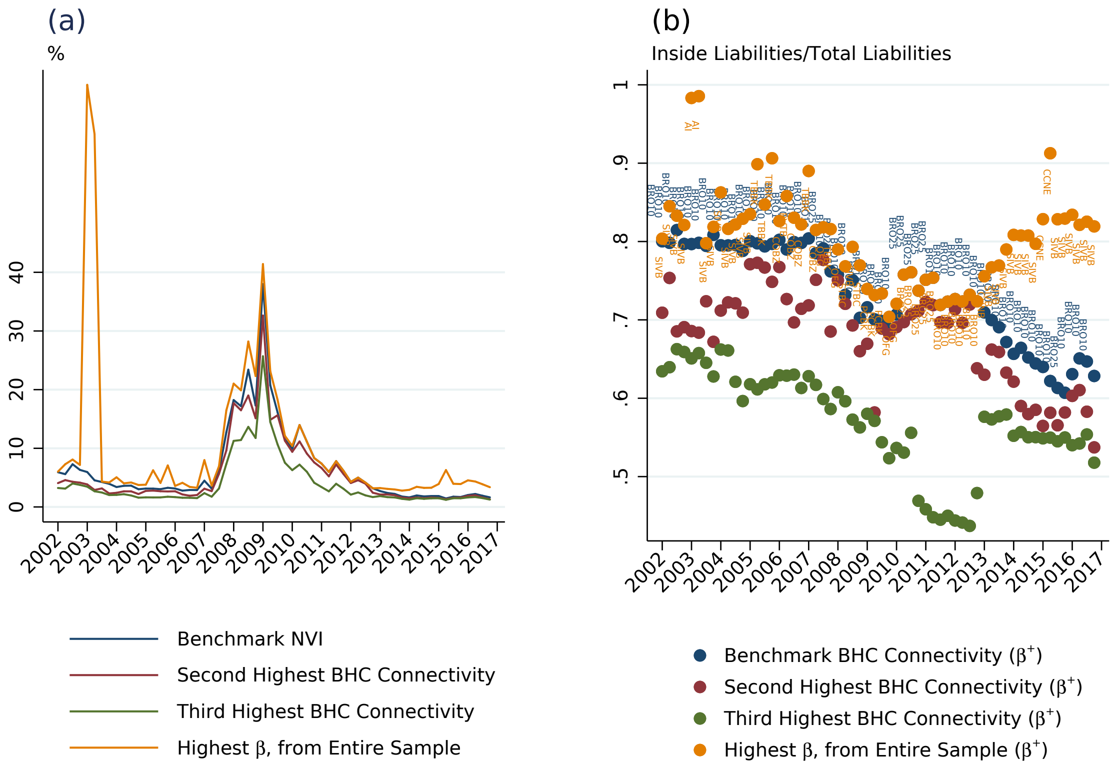


Figure 2.14: **Network Vulnerability and Maximum Connectivity, Different Selection Criteria.** Panel (a) shows the upper bound on network spillover effects when we make different selection decisions for which BHC to have its liabilitity connectivity counted as the largest liability connectivity for the system. When we limit ourselves to the top BHCs by assets, it does not materially affect the NVI if we select the highest connectivity, second highest, or third highest. If we relax our restriction of high-asset BHCs, however, and allow any firm in the FR-Y9C to have the highest connectivity, then the movements of the NVI become unpredictable. Panel (b) shows the level of connectivity chosen as the highest under each of those four selection criteria. Selecting the highest connectivity from any FR-Y9C firm allows several small, highly connected firms to have an undue influence on this value.

sample to those firms who have filed an FR-Y15 at some point from 2012-2016 (this limits our BHC sample, and completely removes our non-BHC appromixated subsector nodes). With our panel reduced in this way, we are able to both use the FR-Y15 for maximum connectivity $\beta^+$ as before, and use intrafinancial assets to calculate $c_i$ for each firm in the sample. To differentiate the effects of the FR-Y15 data from those of a panel size reduction, we also plot a version of the NVI computed with our standard data sources, but that limits its sample to those same firms.

While these new data sources cause changes that are moderately-sized in relative terms, they only serve to shift downward our upper bound measure, in a period when that upper bound was already quite low. To that extent, the FR-Y15 data does not change the conclusions of the NVI as an upper bound in these periods - namely, that the potential for network spillover losses from direct counterparty exposures is very small from 2013-2016.

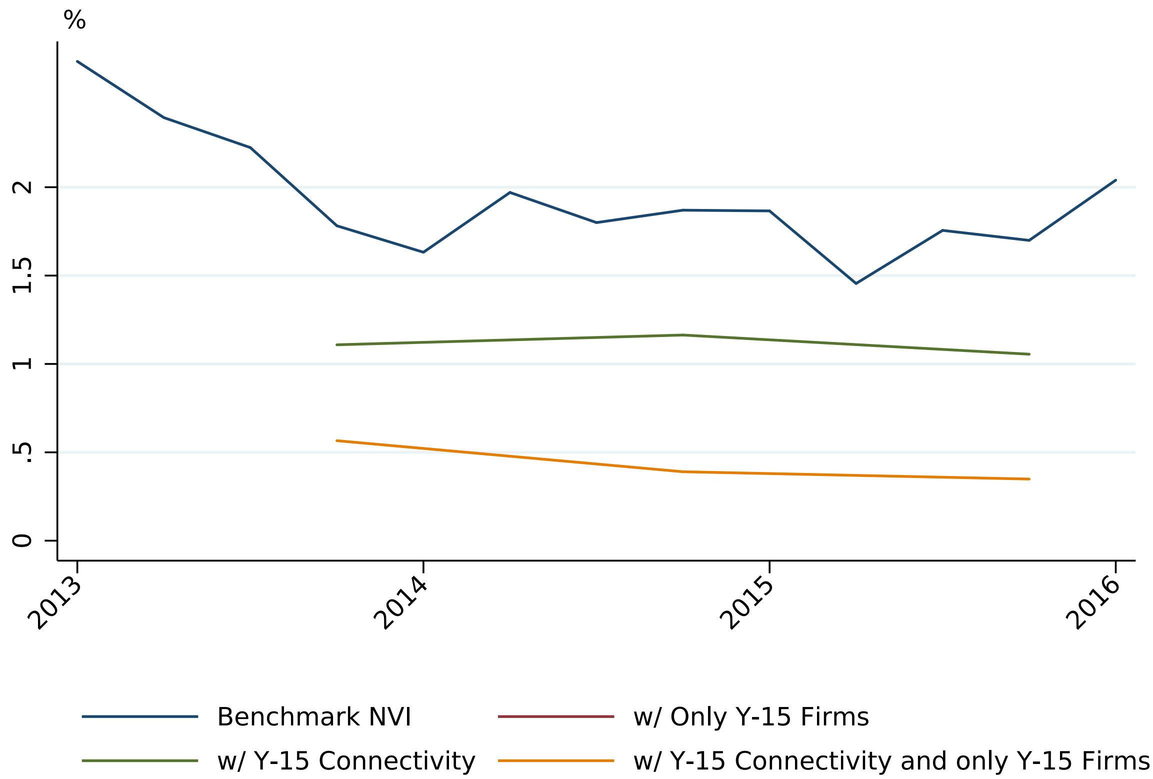


Figure 2.15: **Network Vulnerability Calculated with FR-Y15 Data.** The green line uses FR-Y15 data, when available, to calculate the maximum liability connectivity $\beta^+$ that enters the NVI. The red line limits the panel of firms in the NVI to the FR-Y15 sample, then uses the FR-Y15 for all relevant balance sheet fields. Finally, the orange line calculates the NVI using our standard data sources, but only for the panel or firms with FR-Y15 data. This different data source yields fairly different NVI results in magnitude, but the conclusions to be drawn from that measure in this period remain unchanged.

A second way to use FR-Y15 data is to use the reported fields for intrafinancial deposit liabilities to help inform our classification of deposit liabilities. Figure 2.16 shows the NVI with this change. Specifically, we find a firm-specific average percentage of non-insured deposits inside the financial system from the FR-Y15 sample, then assume that that same percentage of non-insured deposits are inside the system for the entire sample. This allows us to recalculate $\beta$ for firms that filed an FR-Y15 from 2013-2016 (which roughly includes the same firms from which we select $\beta^+$ in the benchmark setup), and select a new $\beta^+$ for the

NVI. Figure 2.16 also includes a much-coarser robustness check that changes the quantity of non-insured domestic deposits classified as inside the system, from 100% in the benchmark, to 20%.

While these adjustments have some impact on the measure - particularly in mid-to-late 2008 - they certainly do not change the nature of any of our conclusions. We view this as reassuring that our 100% inside-system assignment for non-insured deposits, while unrealistic for most firms, has little impact on our actual upper-bound.

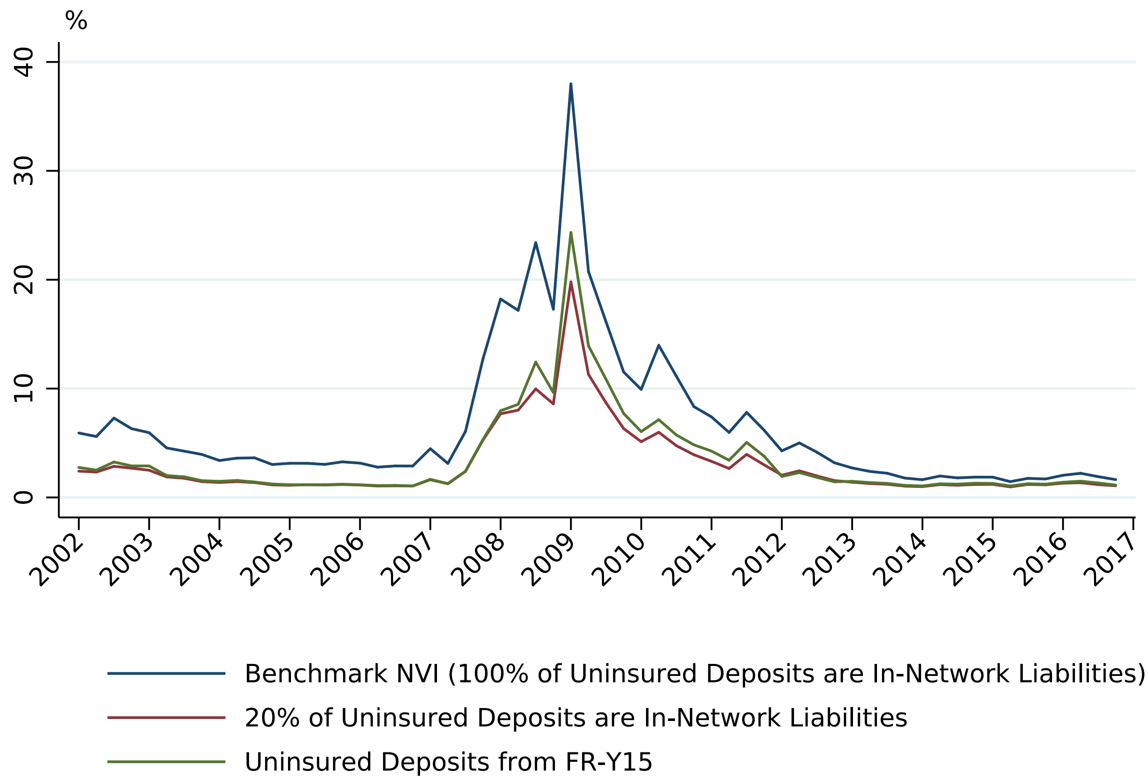


Figure 2.16: **Network Vulnerability with Alternate Percentages of Uninsured Deposits Inside the Network.** The red line above shows the NVI when onyl 20% of uninsured domestic deposits are classified as 'inside' the financial system, for the purpose of calculating liability connectivity (as opposed to 100% in the Benchmark setup). The green line uses FR-Y15 data to construct an average percentage of non-insured deposits inside the system for those firms who file the FR-Y15. That average percentage from the FR-Y15 sample period is then applied for that firm uniformly across each quarter. Neither of these configurations alter the NVI in any substantial way.

Lastly, we wish to see whether any of the off-balance sheet fields on the FR-Y15 can, when combined with our FR-Y9C classifications, alter our NVI in any way. Particularly, the fields allowing firms to record the magnitudes of any undrawn lines of credit with financial institutions and the magnitude of any potential future exposure on over-the-counter derivatives are potentially practically-meaningful assets or liabilities for a firm, but would not appear on the FR-Y9C balance sheet. Figure 2.17 incorporates information from the FR-Y15 on these quantities into the measure in a similar way to Figure 2.16. We limit our NVI sample to FR-Y15 filing firms, then estimate a firm-specific average percentage of total firm assets or liabilities added by including these fields on the balance sheet. Finally, we add extra inside assets or liabilities to that firm in each quarter using those percentages and the firm's total assets or liabilities at the time. Figure 2.17 shows that, while these values can change the NVI (primarily by increasing liability connectivity) the connectivity increases implied by their magnitudes in 2013-2016 are not sufficient to impact the NVI in any meaningful way.

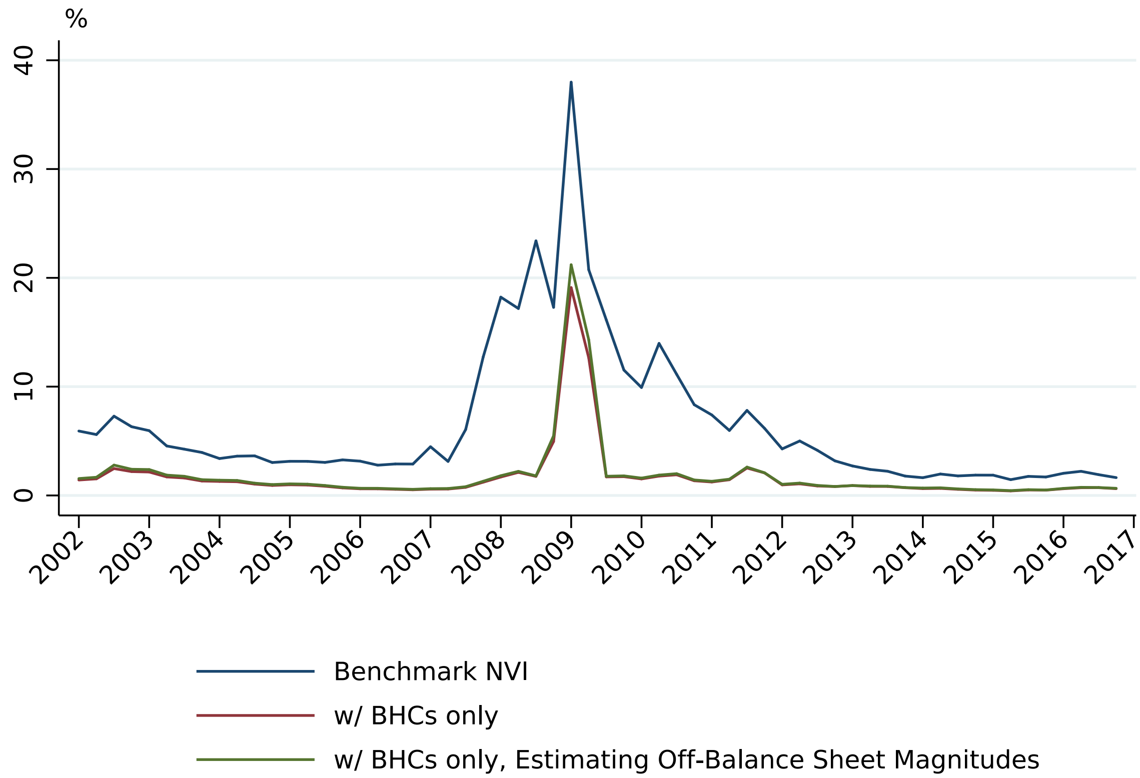


Figure 2.17: **Network Vulnerability with Extrapolated Quantities for FR-Y15 Off-Balance Sheet Items.** The green line shows the value of the NVI when limited to FR-Y15 filing firms, and when certain off-balance sheet items from the FR-Y15 are applied to earlier quarters in the sample. This is done by constructing an average percentage of total firm assets or liabilities attributable to those fields, and assuming that same percentage of assets or liabilities should be added as 'inside' the financial system throughout the sample. The NVI constructed only from FR-Y15 filing firms (i.e. removing all approximated subsector firms) is included in the red line, for reference. These changes have no discernible impact on the NVI, when compared to the benchmark setup over the same panel of firms.

## 2.6 Conclusion

By using detailed data on balance sheet exposures for US financial firms, we have constructed a measure of network spillovers that arise through default cascades in the period 2002-2016. We find that default spillovers, on their own, can amplify expected losses by up to 25% during the financial crisis, but are close to zero before 2008 and after 2012. Default spillovers can be large when nodes inside the network are more exposed to losses outside the network or when the topology of the network implies a higher degree of connectivity among nodes. We find that both elements are important contributors to the time-series dynamics of spillovers and that they can move together or in opposite directions depending the time period examined. In contrast to some narratives of the crisis, we find that neither the exposure to the outside sector nor the connectivity of the financial network increased before the financial crisis. Instead, we find that the events *during* the crisis made the network fragile. After the crisis, our measure of spillovers returned to its low pre-crisis levels, although in the last two years of our sample (2015 and 2016) it has shown a slight increase that may provide a useful signal for policymakers. Considering further amplification mechanisms, such as bankruptcy costs, exacerbates the magnitude of default spillover losses but does not change the conclusion that spillovers were important in 2008-2012 and negligible in the rest of our sample.

# Chapter 3

# Money Fund Demand and Regulatory Reform

## with Abhi Gupta

## 3.1 Introduction

Armed with steadily-improving data on investor-level portfolio holdings, a growing strand of literature in asset pricing has considered prices and quantities as the market clearing outcome of an *asset demand system*. This estimation approach has been applied to a variety of settings including equities, households' portfolios, and foreign exchange markets. [1]. At its most ambitious, this literature estimates *investor-level* demand curves as in Koijen and Yogo (2019). With those estimates, researchers can answer important questions about the role of investor heterogeneity, the evolution of markets over time, and the sensitivity of prices to supply or demand shocks that are hard to analyze with traditional portfolio choice models.

No study has yet applied the asset demand system approach to US money markets, which are not well modeled by existing approaches. As the large literature on short-term asset convenience yields has shown, the prices of very short-term, fixed income assets differ in meaningful ways from the predictions of standard models. [2]. This unexplained variation suggests important aspects of these markets are currently unmodeled. For example, participants and the business press often emphasize the importance of *flows* to investors and borrowers, a set of forces that play a minimal role in standard models. The more-flexible, asset demand system approach is ideally suited to empirically assessing the impact of these flows and other relevant market features on prices and allocations.

US money markets also provide an ideal setting to explore the interaction of regulation and investor demand, which the asset demand literature has identified as crucial. Brunnermeier et al. (2021) notes that one of the primary goals of the asset demand system literature is to assess how investor-characteristics like "regulatory environment" impact those investors' asset demands. By historical happenstance, money markets also present an ideal laboratory

[1] See Koijen and Yogo (2019), Koijen and Yogo (2020)

[2] See the growing literature on convenience yields of short-term assets, such as Krishnamurthy and Vissing-Jorgensen (2012), Nagel (2016), and Krishnamurthy and Li (2023)

for studying this question. The explosion of publicly-available *holdings* data for this market (the key ingredient in the asset demand system approach) coincides with a period of major regulatory changes for money market mutual funds (MMFs) – the most important demand-side investor sector in this market. As policymakers consider new changes to MMF regulation in the wake of Covid-era money market turbulence, understanding the impact of that regulation on the money market asset demand system is more important than ever.

Existing asset pricing models are also ill-suited to assessing the aggregate impacts of 2016's MMF reforms. A key feature of these SEC reforms was that funds that invested in private-sector assets were subject to additional rules compared to government-focused funds, generating substantial flows from the former into the latter. The reforms' rationale was the stabilization of short-term funding markets after the run on MMFs in 2008, but the recurrence of instability in 2020 has left their effectiveness in doubt. Money markets are characterized by significant and important heterogeneity across investors. Our asset demand system approach lets us jointly match prices and quantities in the market and construct credible counterfactuals. Fortuitously, the explosion of publicly-available *holdings* data for this market (the key ingredient in the asset demand system approach) coincides with this regulatory change.

In this paper, we introduce a model and empirical approach for estimating an asset demand system in US money markets over the period 2014-2018. This framework makes two primary contributions. First, it adapts the typical microfoundations in this literature to better suit the setting of a US money market investor. Second, the framework proposes a strategy for identifying a critical parameter, which governs how a money market investor's portfolio shares in assets vary with those asset's yield (i.e. the slope of their demand curves). This strategy focuses on end-of-month window dressing in US money markets, which causes some financial firms to severely reduce their issuance of short-term debt at month-end. Identification relies on these end-of-month effects being fundamentally *supply shocks*.

The model in this paper features money market investors with several salient, real-world features that are missing from the asset demand system's previous setups. The investors of our model invest exclusively in fixed income securities, to maximize an objective function with a dynamic mean-variance form with transaction costs, as in Gârleanu and Pedersen (2016). The fixed income securities of the model predictably mature into cash at the end of their term, allowing the investor to reallocate the matured security's portfolio share to other assets with fewer transaction costs. These features are necessary to capture the institutional details of US money market mutual funds, which invest in assets with relatively short maturities, and appear averse to selling those assets before their maturity (i.e. they are "buy and hold" investors).

This setup retains an important feature of this paper's predecessor's in the asset demand system literature: an investor's portfolio shares, across assets in a particular moment in time, exhibit a empirically-feasible form (i.e. the "portfolio equation"). These shares depend linearly on the money market asset's *yield* (observable both to the econometrician and to the investor) and the asset's *characteristics*, which capture how each dollar invested in that asset contribute to portfolio-wide holding costs (e.g. risk). Some, but not all, of these characteristics will be observable to the econometrician. A central assumption of the asset demand system, which this paper continues to make, is that the parameters of these portfolio equations are *structural*, such that they can be used to compute market clearing yields of money market assets in different counterfactuals.

This paper then outlines how existing data, combined with the structure of the model, can be used to estimate the parameters of each investor's linear portfolio equation. This involves an endogeneity problem: demand-relevant asset characteristics that are unobserved to the econometrician (the "latent demand" of the model) are likely correlated with yields, in the cross-section of assets. This complicates the estimation of the asset yield demand parameter, in a similar way that estimating negative utilities over *price* are complicated in microeconomic models of product choice.

The identification strategy outlined in this paper utilizes a well-documented phenomena in US money markets, whereby certain financial firms (i.e. securities dealers) window dress their regulatory ratios at the end of each calendar month. To improve end-of-month snapshots of leverage ratios, certain financial firms limit their borrowing via short-term debt (such as repurchase agreements or commercial paper) in each month's final trading days. These firms are *suppliers* in our market of interest, so end-of-month reductions in their outstanding quantity of commercial paper are interpretable as supply shocks.

To formalize this notion, we introduce a simple model of the supply side for overnight financial commercial paper (FCP). Motivated by market commentary describing the prevalence of so-called "Fed arbitrage" trade in overnight funding markets, we model issuers of overnight FCP as maximizing their profits of borrowing in overnight FCP markets and using the lent funds to earn the Federal Reserve's interest on excess reserves (IOER), which is typically *above* overnight FCP rates for these firms. Each firm has a "cost of leverage", which is modeled in reduced form as an increasing marginal cost of overnight FCP borrowing that is separate from interest expense. We interpret end-of-month window dressing as a predictable end-of-month increase in the marginal cost of leverage for *some* FCP issuers, but not others. These firms are most likely the US-subsidiaries of foreign financial firms, which have had a greater incentive to window-dress their end-of-month and end-of-quarter leverage ratios since these ratios were adopted [3].

We next detail how this intuition – that end-of-month effects are fundamentally supply shocks – can be implemented to estimate three important model parameters, across two nonoverlapping sample periods: before and after the full implementation of 2016 money market reforms. These parameters are the demand-side yield-sensitivity of money market funds' portfolio shares across assets; the yield-sensitivity of money market investors that are *not* money market funds; and an FCP supply-side parameter governing how quickly "costs of leverage" rise in the quantity of FCP issued. These parameters can be estimated via GMM with three moment conditions. This first condition is that the likeliest window dressing firms do not, on average, experience different *demand shocks* for their FCP at the end of each month than do the less likely window dressing firms. The second condition is that FCP issuers are satisfying their first order conditions suggested by the supply-side of the model, which implies their current marginal "cost of leverage" as a function of their outstanding quantities, current yield, and model parameters. The third condition is that heterogeneity in these marginal costs of leverage represent *supply* shifters, that are relevant to money market investor's asset demand only through their effect on equilibrium *yields*. [4]

[3]See discussions in Munyan (2015) and Anbil and Senyuz (2018)

[4]Put another way: In the model, the same parameters should make sense of *panel* variation, whereby

A major advantage of this novel framework is that it will be feasible to estimate using the existing universe of publicly or commercially-available data on US money markets. This data include publicly-available, end-of-month holdings data for US money market mutual funds from the SEC; publicly-available, daily data on investor-level volume at the Federal Reserve's overnight reverse repurchase agreement (ONRRP) facility from the Federal Reserve Bank of New York; and commercially-available, trade-level data on primary market transactions in the US commercial paper market from the Depository Trust and Clearing Corporation. Even though window dressing is a high frequency (daily) phenomenon, the estimation does *not* require daily data on the holdings of US money market participants.

The estimates from our empirical approach can be directly used to answer several academic and policy-relevant questions. First, has total market elasticity changed substantially, after 2016 money market reforms? This involves calculating the results of a counterfactual, which is straightforward to do with these estimates: how do equilibrium money market yields change, as a result of a supply-driven increase in the quantity of some short-term term instrument or instruments?

Second, these estimates can be used to ask *why* market elasticity has (or has not) changed after 2016. Is it because demand elasticities from money market mutual funds have changed, or because demand elasticities from other, non-money fund investors have changed? It is well-documented that total money fund holdings of assets like commercial paper have fallen since the 2016 reforms were implemented. To our knowledge, there is no consensus (neither among academics nor among market participants) about *which* investor sectors have increased their commercial paper holdings as a result. While this estimation framework may not provide the identity of these investors, it will characterize their apparent elasticity and show how (or whether) it differs from that of the more-regulated money fund sector. These sorts of estimates, made possible with our estimation framework, are highly relevant for understanding whether additional portfolio restrictions on well-regulated investor sectors can be self-defeating, by directing funds to alternate sectors will less-desirable demand properties.

This paper relates to several literatures in asset pricing and post-GFC financial market empirics. Naturally, this paper most-closely belongs to the asset demand system literature of Koijen and Yogo (2020) and Koijen and Yogo (2019). More broadly, the asset demand system literature belongs to a recent collection of papers in asset pricing that document how empirical facts in asset pricing differ from the implications of frictionless representative agent models. Particularly relevant are those papers that document how empirical asset demand curve appear much less elastic that traditional models suggest, as in Gabaix and Koijen (2021). In focusing on how asset demand from a key financial intermediary (money market funds) influences asset prices, this paper also shares much in common with the intermediary asset pricing literature of Adrian et al. (2014), He and Krishnamurthy (2013), and He et al. (2017) – albeit with an admittedly-narrower empirical focus, on money markets. The model employed in this paper also shares much in common with the framework of Bacchetta et al. (2023), which also used a Gârleanu and Pedersen (2016) inspired setup to estimate key parameters portfolio equations – there for international equity mutual funds.

window dressing firms see larger yield changes at the end of each month, and *cross sectional* variation, whereby assets with a greater exogenous supply component are held in greater quantity at the end of the month.

This rest of this paper proceeds in 5 sections. In Section 2, we discuss our data and the institutional setting, introducing the notion that balance sheet costs of financial firms vary both in the cross-section section and predictable across time, at month-ends. In Section 3, we introduce a model of FCP supply and demonstrate how a simplified version can be estimated successfully using public data and variation arising from window dressing. In Section 4, we introduce our asset demand system model for Prime MMFs and discuss how it could be estimated with more comprehensive data and a similar identification strategy. Lastly, we conclude and outline future avenues for research along the lines developed here in Section 5.

## 3.2 Institutional Setting and Data

### 3.2.1 Background on Money Market Funds and Recent Reforms

US money market mutual funds are major intermediaries in short-term funding markets. They link borrowers such as financial firms, non-financial firms, and governments to a variety of individual and institutional investors. MMFs typically invest in highly-rated, short-term, US-dollar denominated assets issued by these borrowers including commercial paper, certificates of deposit, short-term treasuries, and asset-backed securities such as repurchase agreements. In return, investors in money market funds are promised some measure of principal protection, ample liquidity, and low but consistent returns.

Not all assets held by MMFs match the safety promised by these funds to their investors, however, which has caused issues for the sector in previous financial crises. For example, the failure of Lehman Brothers in 2008 caused the value of commercial paper it issued to collapse, damaging the balance sheets of major money market funds. The subsequent surge in redemptions by MMF investors and contraction of the commercial paper market shared many traits with a classic bank run and prompted major reforms. As then SEC Commissioner Mary Jo White described in 2013[5]

> "[O]ur goal is to implement an effective reform that decreases the susceptibility of money market funds to run risk and prevents money market fund events similar to those that occurred in 2008 from repeating themselves."

These new rules, codified in 2014 and implemented in 2016, aimed to reduce the risk of runs and shore up MMF balance sheets by requiring some funds to float their shares – explicitly breaking the guarantee of perfect nominal safety – and allowing for fund managers to impose redemption gates and fees in times of stress to forestall a potential run by investors. Importantly, these regulatory changes fell more strongly on prime funds, which invest in both private and government securities, than on government-only funds.

Unfortunately, these reforms did not prevent a recurrence of instability in money markets in early 2020. As such, there has been renewed interest in another round of reform with a greater focus on the interaction of regulation and the behavior of market participants. SEC commissioner Gary Gensler wrote in 2021[6],

[5] See https://www.sec.gov/news/statement/2013-06-05-open-meeting-statement-mjw for the original speech.

[6] See https://www.sec.gov/news/statement/gensler-mmf-20211215 for more infomation.

”So, were at it again  grappling with how to build greater resiliency into money market funds...[t]he economics will be critical here.”

Existing research has already demonstrated that the 2016 reforms did affect the behavior of investors in money market funds. For example, research from Li et al. (2021) suggests that these the new redemption gates and fees may have actually increased the likelihood of a run on prime funds by lowering the liquidity benefits of investments in prime money market funds in times of financial stress. Consistent with this story, we show in figure 3.1 that MMF investors moved strongly from prime funds to government only funds when these new regulations came into effect at the beginning of 2016. Our focus in this paper is the related topic of how reform affected the behavior of money market funds themselves.

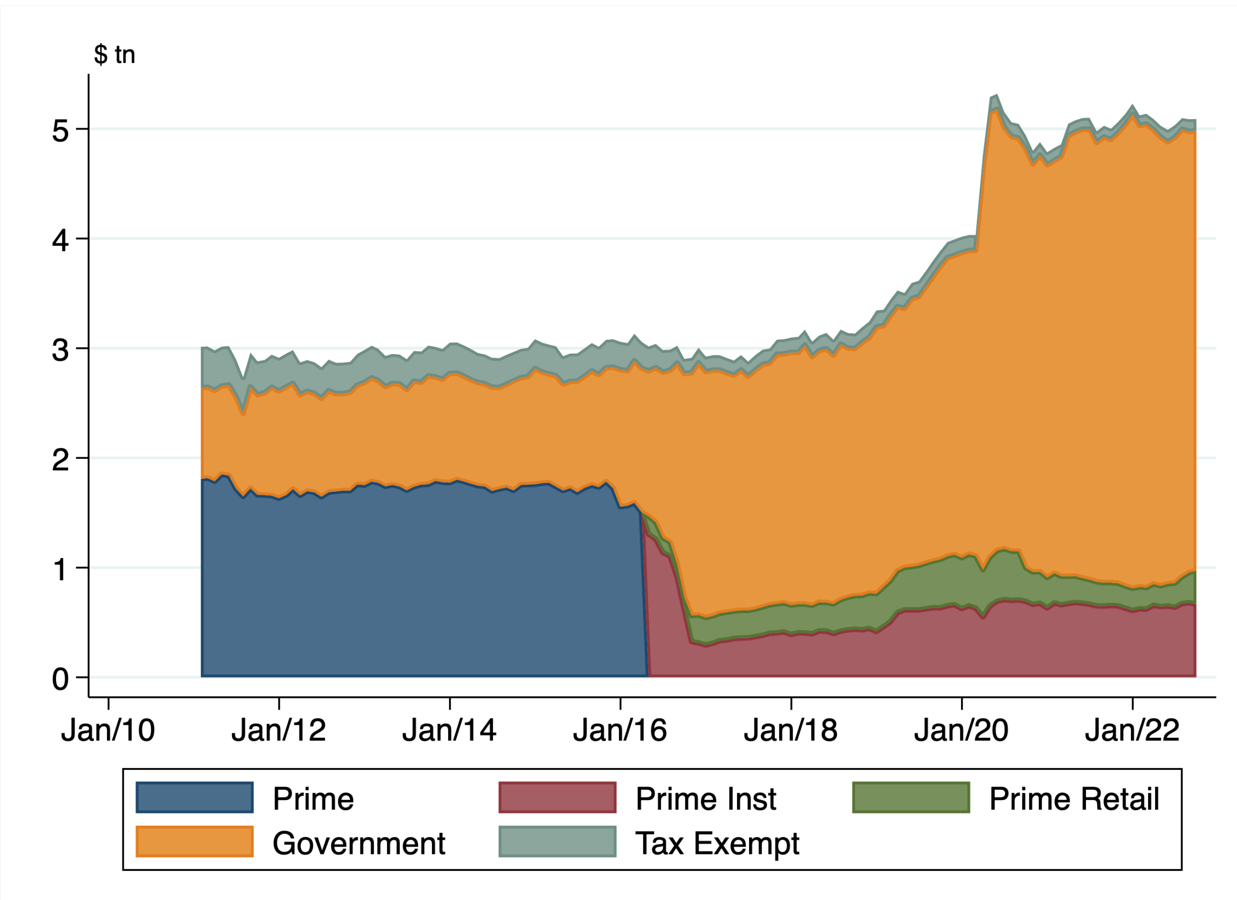


Figure 3.1: Assets under management at MMFs by fund typutilitiese as reported to the SEC. The “Prime” fund category was split into “Prime Retail” and “Prime Institutional” after reforms were implemented at the beginning of 2016. Both new categories have similar investment universes but operate under slightly different sets of regulation.

### 3.2.2 Data

To study the behavior of money market funds we use detailed data on each of their portfolio holdings at the end of each month as reported to the SEC in form N-MFP. This data is high quality, standardized, and publicly-available. These data give both the quantity and the yield of each security held by US money market funds and allow us to consider fund behavior both before and after the 2016 reforms. We focus our attention on the period of January 2014 through February 2018 and observe the universe of prime and government-only money market funds. In future work we plan on making use of more comprehensive data on FCP and related asset transactions from DTCC, the financial clearing house that handles the universe of these trades. This data provides daily coverage of these transactions, allowing us to see MMFs' FCP holdings between month-ends and the FCP holdings of non-MMF entities.

### 3.2.3 Financial Firms, Window Dressing, and IEOR Arbitrage

An additional set of actors relevant to our study are the large financial firms that borrow extensively from money market funds. These institutions interact with MMFs via the commercial paper market; prime MMFs are important sources of funding for these institutions and financial commercial paper is a large and important component of prime funds' balance sheets. As in 2008, in 2020 issues with money market funds happened alongside major disturbances in financial commercial paper markets.

These financial firms have also experienced large regulatory changes since 2008. One of these changes – the leverage ratio requirements of the Basel III Capital Accords – was implemented differently in the United States and in the Eurozone. Importantly, US firms report month an average of their daily leverage ratios while European ones report leverage only as of the end of the month or quarter. This incentivized European firms and their American subsidiaries to *window-dress* their borrowing at the end of each month to appear less leveraged to regulators and the public. The leverage ratio requirements began being reported to regulators at the beginning of 2013 and so are in effect for the duration of our sample.

Following the existing literature such as Du et al. (2018), we conceptualize these regulations as imposing "balance sheet costs" to these financial firms for taking on leverage. While US firms face constant regulatory costs throughout the quarter, European firms' balance sheet spike at the end of month and are especially severe at the end of each quarter.

In our model of overnight borrowing by financial firms that follows, we will also assume that financial firms are borrowing in short-term unsecured markets (like the CP market) primarily to engage in "Fed arbitrage". In this simple trade, private actors borrow in overnight markets to invest the proceeds in a Federal Reserve account and earn the (higher) IOER rate. This trade is common amongst private CP issuers over our period and is qualitatively similar to the low-risk CIP arbitrage trades discussed in Du et al. (2018). This trade and the difference between within and end of month balance sheet costs will be key for our empirical strategy and will be discussed in more detail in the following section.

## 3.3 A Model of Financial CP Supply

We first present a general model of FCP supply and the "Fed arbitrage" trade. We then make some simplifying assumptions and examine market-level data to build intuition for our identification strategy. Lastly we estimate this simplified FCP demand curve using window dressing as an instrumental variable to demonstrate this the potential for this approach to identify parameters in a more generalized demand system.

### 3.3.1 General Setup

In our model of FCP supply, issuers borrow in overnight markets until the marginal cost of additional borrowing equals the marginal revenue of earning the Fed's IOER rate. The marginal cost of borrowing is comprised of both the interest expenses associated with increased borrowing and implicit balance sheet costs reflecting regulatory constraints. Issuers are heterogeneous both in demand for their CP and in the implicit marginal cost of borrowing, arising from heterogeneity in these balance sheet costs.

Consider an issuer (or "dealer") $d$ of overnight Financial commercial paper (CP). Suppose that this issuers perceives two costs two borrowing in overnight CP markets. The first, interest expense, is an *explicit* cost, which is observable to an econometrician. A second, balance sheet cost, is an *implicit* cost $BC$, which the econometrician must infer using information on issuance rates and quantities.

We say that this issuer's Total $\text{Cost}_{d,t}$ of overnight CP issuance is

$$\text{Total Cost}_{d,t} = y_{d,t}Q_{d,t} + BC_d(Q_{d,t}) \tag{3.1}$$

in this notation, $y_{d,t}$ refers to issuer $d$'s interest rate paid in overnight markets. $Q_{d,t}$ is quantity *borrowed* in overnight CP markets, denominated in dollars. The object $BC_d(\cdot)$ measures aggregate balance sheet costs to issuer $d$ for issuance. While we consider $BC_d(\cdot)$ as primarily coming from regulatory costs, some component of this implicit cost may be imposed by investors or private markets as well.

We assume that overnight Financial CP borrowers are engaging in an IOER arbitrage trade. That is, we assume that Total $\text{Revenue}_d$ from the issuer's overnight borrowing position is

$$\text{Total Revenue}_{d,t} = \text{IOER}_t \times Q_{d,t} \tag{3.2}$$

If we assume that the issuer chooses $Q_{d,t}$ to maximize Total $\text{Profits}_{d,t}$ = Total $\text{Revenue}_{d,t}-$ Total $\text{Costs}_{d,t}$ then the issuer's first order condition will give

$$\text{IOER}_t = y_{d,t} + Q_{d,t}\frac{dy_{d,t}}{dQ_{d,t}} + \frac{dBC_d}{dQ_{d,t}} \tag{3.3}$$

The response of yields $\frac{dy}{dQ}$ and of balance sheet costs $\frac{dBC}{dQ}$ can potentially depend on the level of borrowing, the issuer, and the time period.

### 3.3.2 Exposition with a Simplified Demand Curve

To aid in intuition building and demonstrate that window dressing provides variation that can successfully identify demand curve parameters consider a simplified version of this model. First, suppose a tractable form for the issuer's demand curve such that

$$y_{d,t} = \beta_D \sigma_d Q_{d,t} + \varepsilon^D_{i,t} \implies \frac{dy}{dQ} = \beta_D \sigma_d$$

Here, $\sigma_d$ is a *scale* parameter. In the current application, we will suppose that $\sigma_d$ is observable as $\sigma_d = \bar{Q}_d^{-1}$. This is meant to capture the notion that the effect of a changes in borrowing quantities on that issuer's yields will depend on the typical size of the investor universe for that issuer's debt. We further assume that each issuer's commercial paper is an imperfect substitute for those of other issuers so that the demand curve has a finite, positive $\beta_D$.

We will also assume a balance sheet cost function with marginal costs that rise in the quantity borrowed and allow for an issuer and time-specific idiosyncratic component $\varepsilon^S_{d,t}$:

$$BC(Q_{d,t}) = \varepsilon^S_{d,t} Q_{d,t} + \frac{1}{2}\beta_S \left(\sigma_d Q_{d,t}\right)^2 \tag{3.4}$$

In this formulation the *marginal* balance sheet cost of additional borrowing is linear in the normalized current quantity borrowed, $\sigma_d Q_{d,t}$, and the idiosyncratic cost disturbance $\varepsilon^S_{d,t}$.

This structure suggest the following, two-equation system of simultaneous equations

$$\text{IOER}_t = y_{d,t} + \hat{Q}_{d,t}(\beta_D + \beta_S) + \varepsilon^S_{d,t} \tag{3.5}$$
$$y_{d,t} = \beta_D \hat{Q}_{d,t} + \varepsilon^D_{d,t} \tag{3.6}$$

where $\hat{Q}_t(d) = \sigma_d Q_{d,t} = Q_{d,t}/\bar{Q}_d$

We represent this system in 3.2 below. The issuer chooses the optimal quantity level $\hat{Q}$ so that $MC = IOER$, where $MC$ includes both marginal explicit interest expenses and marginal implicit balance sheet costs. [7] At that quantity, the yield on issuer $d$'s overnight CP borrowing is $y_0$ – the point corresponding to that quantity of the demand curve.

[7] Note that the MC line over twice as steep as the demand curve. If $\beta_S = 0$ so that marginal balance sheet costs are constant in $Q$, then the slope of the MC curve would be exactly twice as steep as the demand curve.

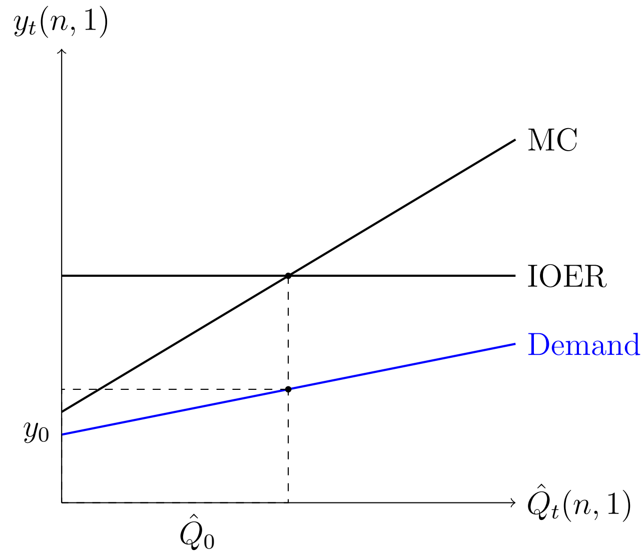


Figure 3.2: Supply and demand curves for ovrnight commercial paper issued by a bank engaging in the IOER arbitrage trade.

The interpretation of the idiosyncratic balance sheet cost disturbances $\varepsilon^S$ that we will pursue in this paper is that they are *supply shocks*. Our goal in this section is then to convincingly estimate cross-sectional variation in $\varepsilon^S_{d,t}$ – each issuer's balance sheet cost intercept. Once estimated, this object will serve as our instrument for *yields* in estimating an asset demand system for Prime money fund investors. That analysis will require an instrument that is correlated with yields through supply alone, and not through demand. This measure of cross-sectional variation in a *supply-side* factor will achieve this.

### 3.3.3 Identifying Demand Parameters with Window Dressing

We estimate the model's parameters using two identifying restrictions. First, we assume that the non-systematic (i.e. idiosyncratic) component of demand is uncorrelated with the non-systematic component of balance sheet costs in the cross-section, i.e. that $\varepsilon^S_{d,t}$ and $\varepsilon^D_{d,t}$ are independent. Second, we assume that the incentive to window dress for European banks is larger at the end of each quarter than at within-quarter month-ends, which generates heterogeneity in $\varepsilon^S_{d,t}$ across issuers and between within-quarter month-ends and ends-of-quarters.

The plausibility of this second assumption can be seen in MMFs' takeup of alternative investments to overnight commercial paper at the end of each month. The firm's optimization problem implies that a rise in balance sheet costs will raise marginal costs and lower commercial paper issuance. In figure 3.3, we show that the volume of ONRRP transactions engaged in by the Federal Reserve rises at the end of each month and to a greater extent at the end of each quarter. Since 2013 MMFs have had access to the ONRRP window and have invested in them as they provide essentially risk-free returns. However, the yields on ONRRP are the bottom of the Federal Funds Rate "corridor" and so are typically lower than those on overnight commercial paper. Given the low risk typically ascribed to overnight CP, MMF funds prefer to hold them over reverse repos for their greater yield. We can also observe the end-of-quarter fall in commercial paper holdings on MMF balance sheets directly as shown in figure 3.4. Note that while we cannot see within-month variation since all holdings are reported as of the end of each month, we do see evidence that financial CP holdings fall at the end of each quarter substantially. Lastly, we can see in figure 3.5 that these end of month shifts have noticeable impacts on the effective Federal Funds rate and that the ONRRP rate is always below the effective Federal Funds rate.

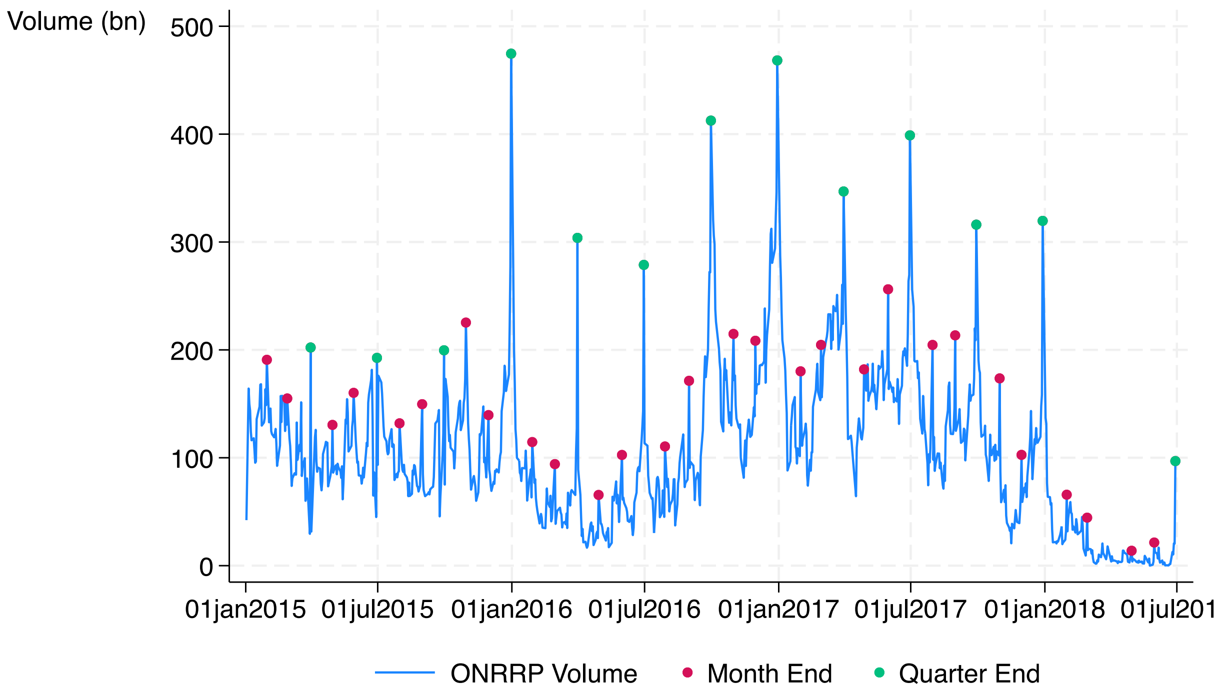


Figure 3.3: Daily value of overnight reverse repurchase (ONRRP) agreements entered into by the Federal Reserve. ONRRP is a close substitute for overnight commercial paper from the perspective of money market funds.

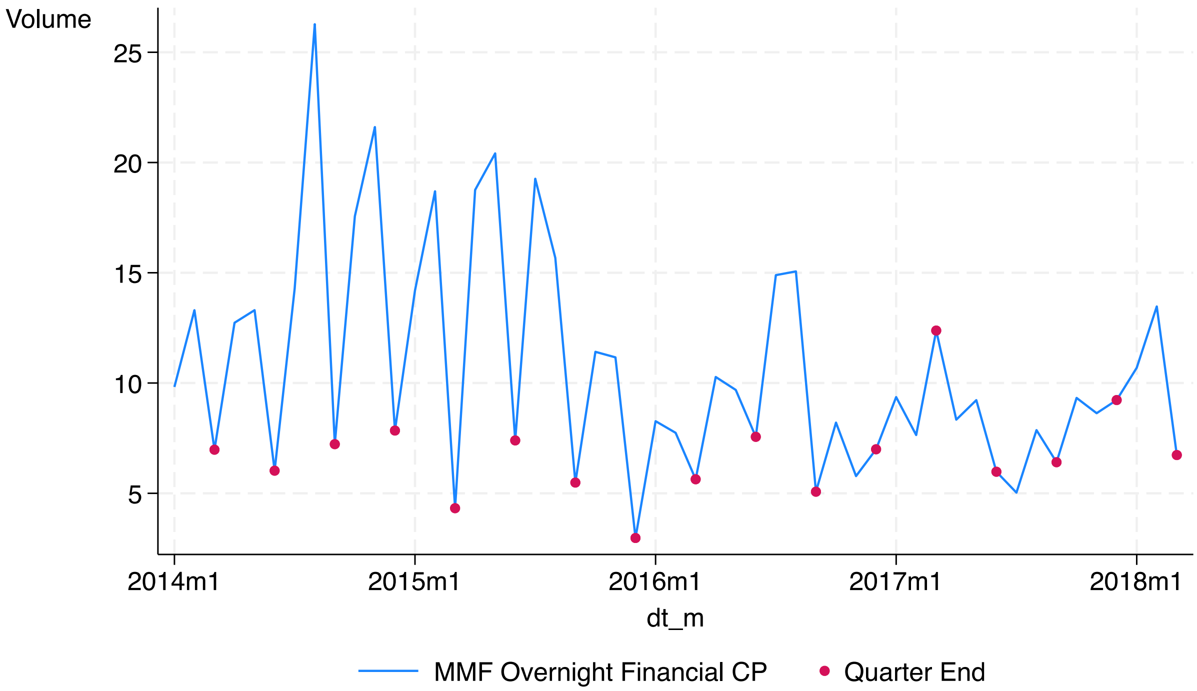


Figure 3.4: MMF holdings of overnight commercial paper as of the end of each month.

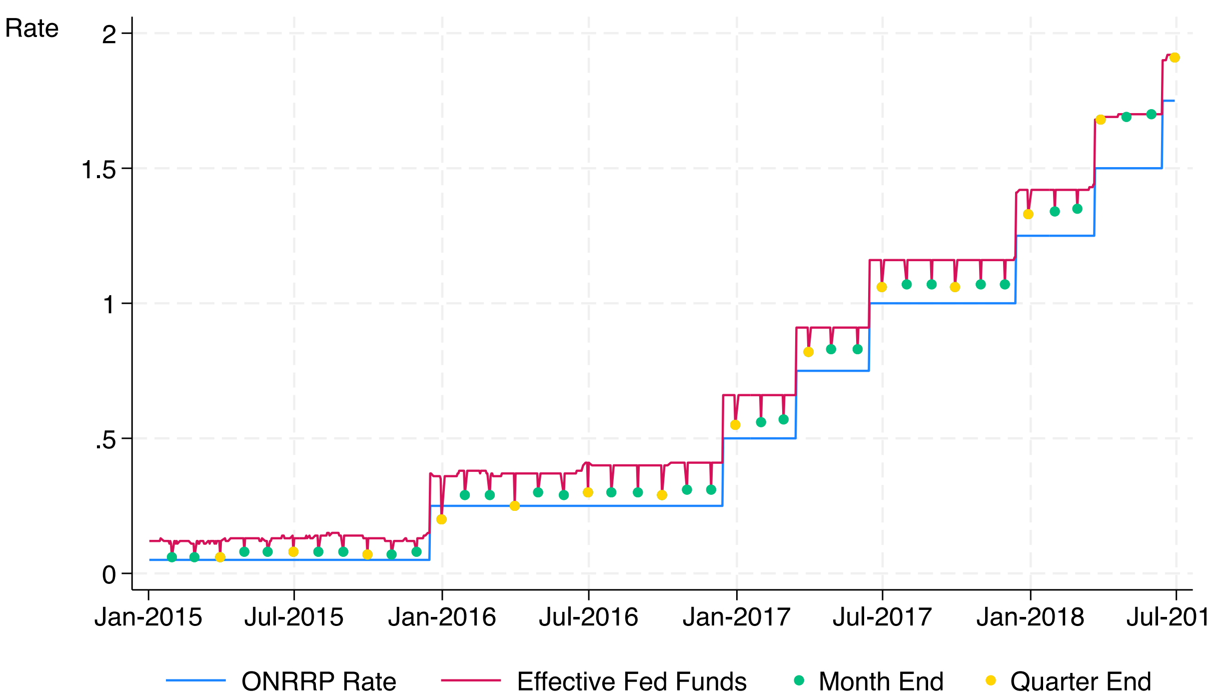


Figure 3.5: Effective Federal Funds rate and ONRRP rate.

As discussed before, heightened quarter-end window dressing by European firms acts as a issuer-specific *supply shock* in this setting. We model this as a increase in $\varepsilon^S_{d,t}$, the intercept of issuer $d$'s total $MC$ curve, at month-end for those firms. This suggest that the marginal balance sheet cost shifts upward equally at all quantities.

However, because month-end window dressing affects a highly nontrivial share of all firms in the market, month-end window dressing also acts as an *aggregate* supply shock in the market. In this simple framework, an *aggregate* supply shock functions as a *demand* shock to individual issuers. A decrease in aggregate supply will lower yields on issuer $d$'s competitors, allowing them to borrow at lower rates for every quantity borrowed.

Figure 3.6 below demonstrates how our identification leverages the panel nature of our CP data to overcome this challenge. Suppose that at time 0, issuers $A$ and $B$ were identical, so that each was issuing quantity $\hat{Q}_0$ and borrowing at rate $y_0$. Suppose that this is a mid-month realization, where neither firm has a window dressing incentive.

Now suppose that at time $t = 1$, firm $A$ has a quarter-end window dressing shock that is shared by some large share of all borrowers in this market. The aggregate nature of the supply shock causes the demand curve for both firms to shift from Demand to "Demand 2". Because each firm's "MC" curve is tied to the interest expense of borrowing (the explicit portion of total costs), any shift in the Demand line comes with an equal shift in the $MC$ curve. For firm $B$, which has no window dressing shock in time 2, this causes a shift down in its $MC$ curve. Firm $B$ will increase their issuance of CP to $A_{B,1}$, and enjoy a lower borrowing rate at $y_{B,1}$.

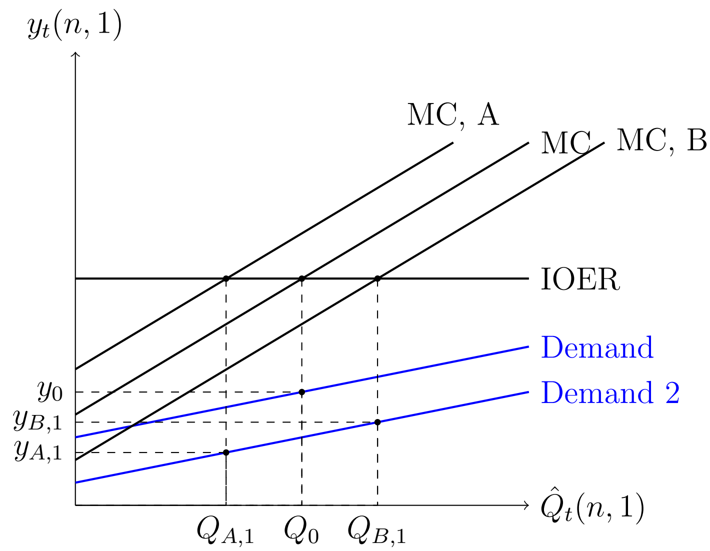


Figure 3.6: Impact of end-of-quarter window dressing on the overnight commercial paper market. At the end of a quarter window-dressing firms, represented here by firm A, face a rise in balance sheet costs that shifts their MC curve up. Since these firms are a large part of the market, their pullback shifts the market demand curve down and lowers marginal costs for non-window-dressing issuers such as firm B.

However, firm $A$ differs from firm $B$ in period 1 because they also experience a shift upward in their marginal balance sheet costs. The $MC$ curve for firm $A$ shifts *upwards*, not downwards. Firm $A$ issues $Q_{A,1} < Q_0 < Q_{B,1}$ in period 1, and borrows at rate $y_{A,1} < y_{B,1} < y_0$.

### 3.3.4 Demand Curve Estimation

In this scenario, we can identify the slope of the demand curve using observables, as

$$\beta_D = \frac{(y_{B,1} - y_{B,0}) - (y_{A,1} - y_{A,0})}{(Q_{B,1} - Q_{B,0}) - (Q_{A,1} - Q_{A,0})}$$

That is, we can identify the slope of the issuer-specific demand curve using a difference-in-differences approach, comparing quantity and yield differences between window dressers and non-window dressers at quarter-end versus a within-quarter month-end. We'll implement this in practice using the following regression specification:

$$\begin{aligned}
\text{2nd}: y_{d,t} &= \alpha_{d,q} + \delta \mathbf{1}(t = \text{Q-end}) + \beta_D \hat{Q}_{d,t} + e_{d,t} \\
\text{1st}: \hat{Q}_{d,t} &= \psi_{d,q} + \chi \mathbf{1}(t = \text{Q-end}) + \sum_{i \in dealers} \phi_i(\mathbf{1}(t = \text{Q-End}) \times \mathbf{1}(d = i) + f_{d,t}
\end{aligned}$$

where $d$ indexes issuers, $t$ the month, $\alpha_{d,q}$ and $\phi_{d,q}$ are dealer-quarter fixed effects, and $\delta$ and $\xi$ are the average end-of-quarter changes in yields and quantities, respectively, at the end of a quarter. $\sum_{i \in dealers} \phi_i(\mathbf{1}(t = \text{Q-End}) \times \mathbf{1}(d = i)$ are issuer by end-of-quarter fixed effects to capture that issuers differ in their average level of quarterly window dressing across the sample.

This is an instrumented difference in difference regression that uses variation in issuers' window dressing behavior to estimate exogenous shifts in quantities and then in turn the response of yields to quantities. This specification combines evidence from all issuers and from comparing both within-quarter month ends to each end-of-quarter. The instruments are the issuer dummies interacted with quarter ends. The number of instruments is equal to the number of issuers. As a robustness check, we also consider an alternative set of instruments that interact an issuer's country of origin with the end-of-quarter indicator. This binned instrument imposes an equal level of window dressing for all French banks, all German banks, and so on.

| | (1) spread | (2) spread | (3) spread |
|---|---|---|---|
| Quantity Outstanding | -0.0126*** | 0.0726*** | 0.0271*** |
| | (0.00446) | (0.0197) | (0.00707) |
| $N$ | 619 | 601 | 618 |
| Specification | OLS | IV | IV |
| Month FE | Yes | No | No |
| Issuer x Quarter FE | No | Yes | No |
| Country x Quarter FE | No | No | Yes |
| Quarter End FE | No | Yes | Yes |

Standard errors in parentheses
$^{*}$ $p < 0.1$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$

Table 3.1: Estimates of the FCP demand curve slope $\beta_D$.

In table 3.1 we report our OLS and IV results. The OLS estimates of $\beta_D$ are negative, matching the aggregate market-level declines in short-term yields and FCP volumes shown in figures 3.5 and 3.4. Both of our instruments, however, give positive and significant demand curve slopes matching what we would expect given the model. The sign of these IV estimates mean that issuers whose end-of-quarter outstanding quantities are lower than their end-of-month quantities have smaller end-of-quarter yields (relative to their end-of-month yields). We wouldn't expect this if these end-of-quarter quantity differences were driven primarily by an endogenous response to *demand* shocks. These results suggest that window dressing does indeed generate sufficient variation to estimate demand curves and that our model of FCP issuance accords well with observed market dynamics.

## 3.4 An Asset Demand System for Money Market Investors

In this section, we depart from the previous section's assumption of homogenous, linear CP demand curves, and model asset demand using a logit-style framework. We detail how end-of-month window dressing, and heterogeneous balance sheet costs across CP issuers, can still motivate a valid identification strategy in this richer setting. The strategy that we describe is feasible using existing public or commercially-available data.

### 3.4.1 Investor Problem

We introduce the objective function of a money market investor in this setting, and show that our setting provides a similarly-tractable portfolio choice equation as others in the asset demand system literature.

Consider a money market investor $f$ who maximizes the following objective function

$$\max_{z_{f,0}} \mathbb{E}_{f,0} \sum_{t=0}^{\infty} \beta^t \left( z'_{f,t} R_{t+1} - \frac{\gamma_i}{2} z'_{f,t} \Omega_f z_{f,t} - (z_{f,t} - z^{bh}_{f,t})' \Lambda (z_{f,t} - z^{bh}_{f,t}) \right) \tag{3.7}$$

$$\text{s.t. } z'_{f,t}\mathbf{1} = 1 \tag{3.8}$$

where $z_{f,t}$ is an $(|N_f| \times 1)$ vector of portfolio shares and $|N_f|$ is the number of available assets in fund $f$'s investment universe. $R_t$ is the vector of returns of the different available assets. $z^{bh}_{f,t}$ represents the buy-and-hold portfolio shares of the $|N_F|$ assets. These are the portfolio shares that the fund begins the period $t$ with, as a result of previously-purchased fixed income assets reaching another period closer to maturity.

This setup, substantially similar to that introduced in Gârleanu and Pedersen (2016), can easily represents a dynamic variant of a mean-variance investor that experiences quadratic costs from trading. In that setting, the $(|N_f| \times |N_f|)$ matrix $\Lambda$ represents the size of the quadratic costs from trades in the assets, given that $z_{f,t} - z^{bh}_{f,t]}$ is the vector of asset *trades* that must be engaged in to support the portfolio share $z_{f,t}$. In the mean-variance interpretations, $\Omega$ represents *risk*, in which case we might say $\Omega = \mathbb{V}_t(R_{t+1})$. However, it is more precise to say that $\Omega$ represents that matrix of quadratic holding costs for different asset positions. This leaves open the possibility that portfolio holding costs may derive from a source other than the risk of returns. For instance, they may derive from a money market fund's regulatory requirements for limiting position exposures to any individual counterparties.

We now represent the fixed-income nature of the $|N_f|$ assets, via a law of motion equation for the buy and hold portfolio of investor $f$. That is, we say that

$$z^{bh}_{ft} = \frac{1}{1 + z'_{f,t-1} R_t} I_{h \to h-1} \text{diag}(\mathbf{1} + R_t) z_{f,t-1} \tag{3.9}$$

where $I_{h \to h-1}$ is an $(|N_f| \times |N_f|)$ matrix, populated with zeros and ones, that transforms a period $t$ holding in a maturity $h$ security into its period $t+1$ counterpart of maturity $h-1$. We will call $I_{h \to h-1}$ the ageing matrix.

The buy-and-hold position in each security is thus a function of the previous period's positions, *aged* into shorter maturity positions today; the appreciation in the value of each position due to its return; and the appreciation of the entire portfolio value.

In our setting, we will replace this exact law of motion with the tractable approximation

$$z^{bh}_{ft} = \tilde{I}_{h \to h-1} z_{f,t-1} \tag{3.10}$$

where $\tilde{I}_{h \to h-1} = \frac{1}{1+\bar{R}_f} I_{h \to h-1} \text{diag}(\mathbf{1} + \bar{R})$, where $\bar{R}$ are average daily returns of each asset, and $\bar{R}_f$ is the average daily appreciation of fund $f$'s portfolio.

The first order conditions of the investor in this setting imply, for each asset $n$,

$$0 = e'_n \mathbb{E}_{f,t} R_{t+1} - \gamma_i e'_n \Omega z_{i,t} - e'_n \Lambda (z_{f,t} - z^{bh}_{f,t}) + \beta e'_n \tilde{I}'_{h \to h-1} \Lambda \mathbb{E}_{f,t} (z_{f,t+1} - z^{bh}_{f,t+1}) - \mu_{f,t} \tag{3.11}$$

where $e_n$ is an $(|N_f| \times 1)$ vector with a one in the index for asset $n$, and zeros elsewhere. The object $\mu_{f,t}$ is the lagrange multiplier for the constraint that the sum of portfolio shares in period $t$ must equal one.

In this project, we follow Koijen and Yogo (2020) in making a series of simplifying assumptions about equation 3.11, to make computing pricing counterfactuals more empirically tractable. The first such assumption is

$$\Omega_f = k_f k + f' + \sigma_i^2 I_{|N_f|} \tag{3.12}$$

where $k$ is an $(|N_f| \times 1)$ vector and $\sigma_i^2$ is a scalar. This assumption implies that total quadratic holding costs in a period depend on the (square of) the size of each individual position, and on a (square of the) single weighted sum of all positions. If we interpret quadratic holding costs as risk, then this suggests that each asset's risk has an idiosyncratic component, and a component that is correlated across securities. This latter piece can be summarized by each asset's factor loading on some single risk factor.

Second, we assume that quadratic transaction costs are constant across securities, and have no spillovers into costs associated with a different security. That is, we say

$$\Lambda_f = \lambda_f I \tag{3.13}$$

with scalar $\lambda$.

Third, we say that the expected return of each asset $n$ is observable as the *yield* of that security. This is exactly true of 1-day ($h = 1$) maturity securities. For $h \geq 2$ securities, this is an approximation motivated by the fact that an average of the next $h$ days' returns for asset $n$ are observable (i.e. because a fixed income security's payment upon maturity is known).

With these simplifications, 3.11 becomes

$$z_{f,n,t} = \lambda\phi z^{bh}_{f,n,t} - \phi\mu_{f,t} + \phi y_{n,t} - \gamma_i\phi(k'z_{f,t})k_n + \beta\lambda\phi\overline{Val}_n\mathbb{E}_{f,t}(z_{f,n-1,t+1} - z^{bh}_{f,n-1,t+1})$$

where $\phi = \frac{1}{\gamma_i\sigma_i^2+\lambda}$ and where $n-1$ is the asset that $n$ *ages* into after one day. $\overline{Val}_n$ is the average expected one-day appreciation of asset $n$, as a share of the average appreciation of the entire portfolio.

This first order condition helps explain the *dispersion* of portfolio shares from fund $f$ across the assets in its investment universe. There are several reasons why a fund may end period $t$ with a greater portfolio share in asset $n$ than in some other asset $m$ in the investment universe. First, it may have started the period with more of asset $n$ than asset $m$. Second, asset $n$ may have a higher yield than asset $m$. Third, A marginal investment in asset $n$ will contribute less to portfolio-wide portfolio costs than asset $m$. Interpreting holding costs as risk, this means that the risk in a marginal dollar in asset $n$ is less correlated with the current risk in the rest of the portfolio. Given assumption 3.12, this means asset $n$ has a smaller factor loading than asset $m$. Finally, a fund may hold more of asset $n$ if it intends to hold more units of asset $n$ in the next period. In this case, greater purchases today require smaller purchases (and transaction costs) tomorrow.

It is convenient to consider this first order condition with respect to a reference asset with 1-period maturity, which does not contribute to the portfolio-wide component of holding costs (e.g. has no "factor exposure" risk). In our application, investments by fund $f$ in the Federal Reserve's overnight reverse repurchase agreement facility (ONRRP) will act as the reference asset.

Assets like the ONRRP that are newly-issued as 1-period investments for money market participants simplify this FOC in two ways. First, if there are effectively no 2-period maturity assets that age into the 1-period investment in question, then $z^{bh}_{f,h,t} = 0$ for those securities. Second, because a 1-period investment matures into cash the next period, $z_{f,n-1,t+1} - z^{bh}_{f,n-1,t+1} = 0$ for those assets $n$.

Together with the assumption that investments in the Fed's ONRRP facility carry no factor exposure (i.e. $k_{ONRRP} = 0$) this suggests

$$z_{f,ONRRP,t} = \phi y_{ONRRP,t} - \phi \mu_{f,t}$$

Motivated by this, and without loss of generality, we let $\hat{\mu}_{f,t} = z_{f,ONRRP,t}$ so that we can reexpress the first order condition with respect to some other asset $n$ as

$$\begin{aligned} z_{f,n,t} = \hat{\mu}_{f,ONRRP,t} + \lambda \phi z^{bh}_{f,n,t} + \phi(y_{n,t} - y_{ONRRP,t}) - \gamma_i \phi(k' z_{f,t}) k_n \\ + \beta \lambda \phi \overline{Val}_n \mathbb{E}_{f,t}(z_{f,n-1,t+1} - z^{bh}_{f,n-1,t+1}) \end{aligned} \tag{3.14}$$

Equation 3.14 will become the basis of a tractable, logit-style demand system below.

### 3.4.2 A Logit-Style Demand Equation

Let $\delta_{f,n,t}$ be the portfolio share of asset $n$, relative to the ONRRP reference asset. We write that

$$\delta_{f,n,t} = \frac{z_{f,n,t}}{z_{f,ONRRP,t}} = 1 + \psi_f z^{bh}_{f,n,t} + \tilde{\phi}_f (y_{n,t} - y_{ONRRP,t}) \tag{3.15}$$

$$-\Pi_{f,t} k_n + \xi_f \overline{Val}_n \mathbb{E}_{f,t}(z_{f,n-1,t+1} - z^{bh}_{f,n-1,t+1}) \tag{3.16}$$

A critical simplifying assumption in this literature, which this paper will also assume, is that the parameters of this equation are *structural*, such that we can construct counterfactuals that hold fixed the objects $\{\psi_f, \tilde{\phi}_f, \Pi_{f,t}, \text{and } \varepsilon_f\}$. As such, we consider estimating the parameters of this equation for $\delta_{f,n,t}$ as the critical objective of this framework.

Note a useful consequence of the $f$'s budget constraint, that

$$z_{f,ONRRP,t} = \frac{1}{1 + \sum_{m \neq ONRRP} \delta_{f,m,t}}$$

That is, fund $f$'s portfolio share in the outside outside (ONRRP) depends on the overall attractiveness of the other assets in fund $f$'s investment universe, as captured by $\delta$.

This setup allows the model to reproduce some critical qualities of any sensible pricing counterfactual. Higher yields and lower risk will lead a fund to invest more in a particular security. It will do so by reducing portfolio shares in other securities, in a way that does not violate its budget constraint. This system shares the familiar logit demand quality that, if the desirability (capture by $\delta$) of two securities is unchanged in a counterfactual, then the *ratio* of portfolio shares of those two securities will stay constant through the counterfactual.

### 3.4.3 Empirical Implementation and Endogeneity

As in Koijen and Yogo (2020), we do not assume that the factor loadings $k_{f,n}$ of securities are observable to the econometrician. Instead, we say that

$$k_{f,n} = \chi_f' x_n + \varepsilon_{f,n,t} \tag{3.17}$$

where $x_n$ are characteristics of security $n$ that are observable to both the econometrician and the fund. In this setup, $\varepsilon_{f,n,t}$ are interpretable as security characteristics that are demand relevant (in the model, via their relevance for holding costs), observable to the fund, but unobservable to the econometrican. These terms serve as the latent demand components in the model.

Substituting into equation 3.16, and considering a newly-issued security of maturity 1-day for convenience, we see that

$$\delta_{f,n,t} = 1 + \tilde{\phi}_f (y_{n,t} - y_{ONRRP,t}) - \Pi_{f,t}(\chi_f')x_n - \Pi_{f,t}\varepsilon_{f,n,t} \tag{3.18}$$

Koijen and Yogo (2020) estimate a portfolio equation like 3.18 using a fund's cross-sectional dispersion in portfolio weights across assets, at a single date $t$. This estimation comes with an endogeneity concern, because unobserved latent demand $\varepsilon$ are likely correlated with security yield. In Koijen and Yogo (2020), the authors construct a cross-sectional instrument for yields, using the measured investment universes of different funds, and assuming that investment universes for funds are exogenous. Koijen and Yogo (2020) focused on the equity market, where "supply shocks" are challenging to find, given that the number of outstanding shares of a company stock is not economically meaningful.

One advantage in focusing on money markets is that these markets regularly experience price movements that are generally attributed to supply shocks. In this paper, we use one particularly important and often-discussed supply shock to estimate the parameter $\phi_f$, which dictates how asset yields affect equilibrium portfolio shares.

In practice, the available data does not provide sufficient variation to estimate the object $\tilde{\phi}_f$ at the investor level. Instead, we will allow for heterogeneity in $\tilde{\phi}$ at the investor class level, meaning for money market mutual funds and other investors. That is, we will estimate separate $\tilde{\phi}_{MMF}$ and $\tilde{\phi}_{Non-MMF}$.

### 3.4.4 End-of-Month Window Dressing, via Market Clearing

We next detail how to implement a similar intuition as from Section 3.3, in a richer environment where market clearing derives from the asset demand equations described above.

Consider now a single issuer of overnight financial commercial paper (FCP) $d$. As before, $d$ is likely a securities dealer, from one of several different regulatory jurisdictions (American, Canadian, European, Japanese, etc.). Market clearing in the market for dealer $d$'s overnight FCP requires that

$$Q_{d,t} = \sum_{f} \delta_{f,d,t} A_{f,t}$$

where $Q_{d,t}$ is the dollar value of dealer $d$ overnight FCP borrowing, and $A_{f,t}$ is money market investor $f$'s volume at the Fed's ONRRP facility (i.e. the reference asset).

Consider time $t$ to be the last trading day of its month, and time $t-1$ to be the peunultimate trading day of its month. We assume that these two periods, separated only by a trading day, are sufficiently close that fund each fund's parameters governing the desirability of different static asset characteristics are nearly unchanged. That is, we assume $\Pi_{f,t} \approx \Pi_{f,t-1}$. And, naturally, the static asset characteristics $x_n$ will likewise not change between dates $t$ and $t-1$.

With these assumptions, is straightforward to show that this market clearing equation suggests the *change* in the dollar value of dealer $d$'s overnight FCP borrowing between the end-of-month $t$ and previous day $t-1$ is

$$\begin{aligned}
\Delta Q_{d,t} = & \sum_{f} (A_{f,t} - A_{f,t-1}) \delta_{f,t} \\
& + \phi_{MMF} \Delta(y_{d,t} - y_{ONRRP,t}) \sum_{f \in MMF} A_{f,t-1} \\
& + \phi_{NonMMF} \Delta(y_{d,t} - y_{ONRRP,t}) \sum_{f \in NonMMF} A_{f,t-1} \\
& + \sum_{f \in MMF} A_{f,t-1} \Pi_f \Delta\varepsilon_{f,d,t} + \sum_{f \in NonMMF} A_{f,t-1} \Pi_f \Delta\varepsilon_{f,d,t}
\end{aligned} \tag{3.19}$$

Note that, when considered *across* dealers $d$, in a single end-of-month period $t$, Equation 3.19 represents a cross-sectional relationship between the change in total outstanding overnight FCP supply $\Delta Q_{d,t}$, and the change in overnight FCP *yield*, via $\Delta(y_{d,t} - y_{ONRRP,t})$. This relationship depends on both the yield sensitivity of MMF investors $\phi_{MMF}$ and the yield sensitivity of non-MMF investors, $\phi_{NonMMF}$.

Equation 3.19 motivates the first moment condition in the Generalized Method of Moments estimation strategy that we propose. That moment condition becomes

$$\begin{aligned}
& \frac{1}{|WD|} \sum_{d \in \{\text{Window Dressers}\}} \left( \sum_{f \in NonMMF} A_{f,t-1} \Pi_f \Delta\varepsilon_{f,d,t} + \sum_{f \in MMF} A_{f,t-1} \Pi_f \Delta\varepsilon_{f,d,t} \right) \\
= & \frac{1}{|NWD|} \sum_{d \in \{\text{Non-Window Dressers}\}} \left( \sum_{f \in NonMMF} A_{f,t-1} \Pi_f \Delta\varepsilon_{f,d,t} + \sum_{f \in MMF} A_{f,t-1} \Pi_f \Delta\varepsilon_{f,d,t} \right)
\end{aligned}$$

where $|WD|$ and $|NWD|$ are the number of likely window-dressing and non-window dressing firms, respectively (identified via their regulatory jurisdiction). That is, the first estimating moment condition is that the average aggregate demand shocks for window-dressing and non window-dressing firms are equal. This is the asset demand system analogue of assuming that end-of-month window dressing by certain financial firms is a *supply-side* phenomenon, and does not reflect any average change in the demand-side desirability of overnight FCP from these two firm types.

Note that this moment condition leaves open the possibility that $\varepsilon_{f,d,t} \neq \varepsilon_{f,d,t-1}$ for some (or even all) funds $f$, for one or more dealers $d$. That is, is is conceivable that some dealer $d$ experiences a *demand shock* for its overnight commercial paper at end-of-month $t$. For identification, we merely require that window dressers and non window dressers experience the same demand shocks on average.

Of course, this single moment condition in insufficient to separately identify both $\phi_{MMF}$ and $\phi_{NonMMF}$. Compiling additional moment conditions to separately identify these objects requires a return to the overnight FCP supply model above.

### 3.4.5 Balance Sheet Cost Heterogeneity as an Instrument

Next, we outline the final set of moment conditions, for separately identifying the crucial yield sensitivities in the asset demand system. These conditions return to the framework of Section 3.3, detailing the supply problem of overnight FCP issuers.

Recall from Section 3.3 that overnight FCP issuers in the model issue FCP such that the following will hold, for dealer (issuer) $d$ in time $t$

$$\varepsilon^S_{d,t} = \text{IOER} - y_{d,t} - Q_{d,t}\left(\frac{dy_{d,t}}{dQ_{d,t}} + \beta_s\right) \tag{3.20}$$

The strategy outlined in this section relies on the identifying assumption that heterogeneity in balance sheet costs $\varepsilon^S_{d,t}$ across dealers $d$ is uncorrelated with "latent demand" – meaning demand-relevant asset characteristics that unobserved by the econometrician. Intuitively, it requires balance sheet costs to be a supply-driven phenomenon, that affects investors' demand for different overnight FCP only through their equilibrium effects on FCP yields.

In section 3.3 above, we described a simpler setup for expositional purposes where $\frac{dy}{dQ}$ is assumed to be constant across dealers $d$. In the asset demand system outlined above, $\frac{dy}{dQ}$ will vary across dealers, in a particular way implied by market clearing. Specifically, one can show that market clearing implies that

$$\frac{dy_{d,t}}{dQ_{d,t}} = \frac{1}{\sum_f \tilde{\phi}_f A_{f,t}(1 - z_{f,d,t})} \tag{3.21}$$

That is, the (nonconstant) "slope" of the demand curve for each dealer will depend on $\phi_{MMF}$, $\phi_{NonMMF}$, investors' portfolio shares in the dealer's FCP, and the size of the outside asset of each fund for which dealer $d$'s FCP appears in their investment universe.

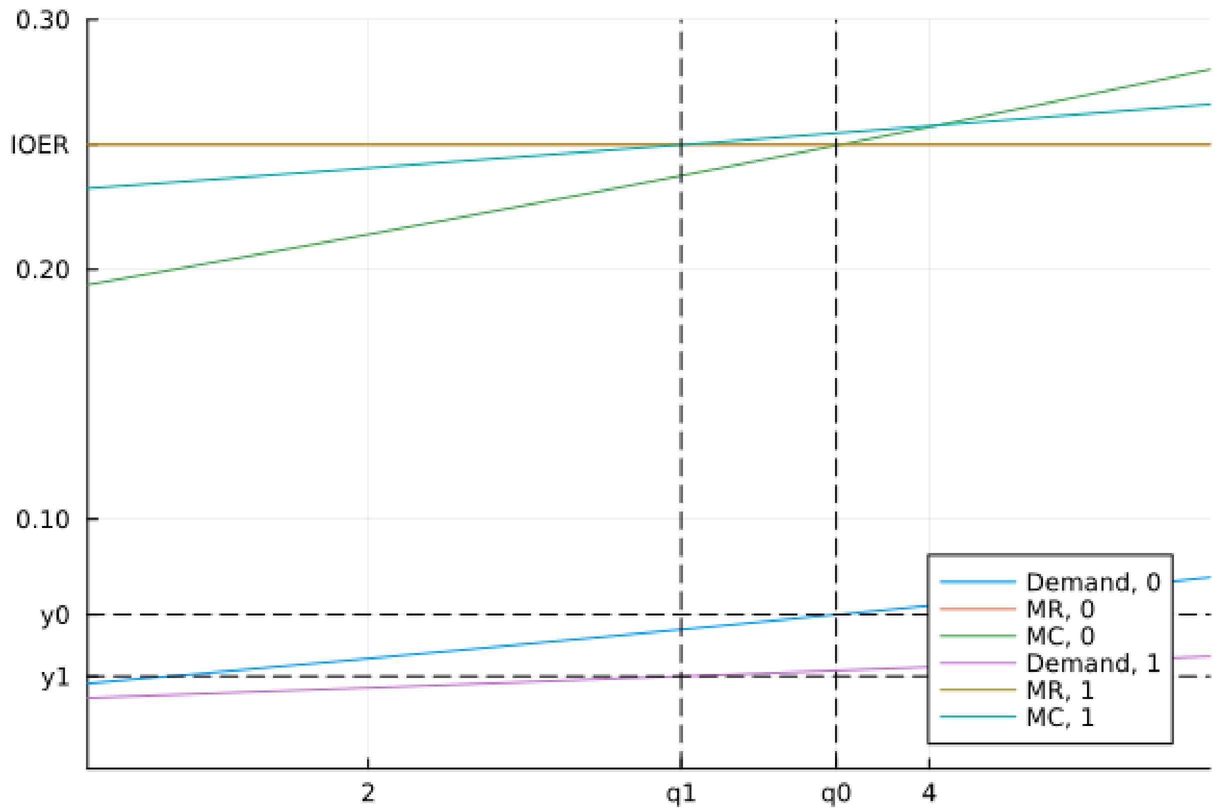


Figure 3.7: Impact of window dressing on commercial paper issuance under the full asset demand system.

To build intuition for this new setup, Figure 3.7 below reproduces some key components from Figure 3.2 in the new framework. The Figure shows the key curves governing a given FCP issuer's choice of issuing quantity: the marginal revenue $MR$, which is fixed at IOER; the marginal cost curve $MC$, which accounts for both the increased interest cost and increased balance sheet costs of increasing issuance; and the demand curve, which calculates the market-clearing yield at a given quantity of issuance. As before, each issuer chooses a quantity of commercial paper such that $MC = MR$, and pays a yield that is consistent with their demand curve.

In Figure 3.7, time "0" refers to a trading date just *before* the end of a month, before many firms in the market receive a window dressing supply shock. While each of these curves is close to linear, an important difference between this framework and that of Figure 3.2 is that these curves are not entirely linear. Rather, they are dictated by the market clearing conditions suggested by the logit-style demand system.

In the Figure, time "1" is meant to simulate a window dressing event, in which this firm is one of the window dressers. At time "1", the balance sheet costs for the depicted firm rise, consistent with this firm receiving a window dressing incentive. However, as in Figure 3.2, this shock being delivered to *many* firms in the market causes the demand curve for this firm's overnight FCP to shift *down* substantially – in a way that is now fully consistent with the logit-style demand system. In the numerical exercise shown in this figure, the shock to this firm's balance sheet costs outweigh the equilibrium-driven, beneficial shift of the firm's demand curve, such that the firm still chooses to issue less overnight FCP (but at a lower yield).

Under our identifying assumption, and with correct guesses of $\phi_{MMF}$, $\phi_{NonMMF}$, and $\beta_S$, the object $\varepsilon^S_{d,t}$ constructed via equation 3.20 is a valid instrument for estimating $\phi_{MMF}$ and $\phi_{NonMMF}$ via equation 3.18. That is, with estimates of $\varepsilon^S_{d,t}$, one can use the *cross-section* of portfolio shares in different overnight CP, in period $t$, across MMFs and non-MMFs, to estimate the yield sensitivity of different investors using balance sheet costs as an instrument for yields.

More formally, our moment conditions become

$$\begin{aligned} E_{MMF}(e^{\perp}_{f,d,t}\left(\Theta)\varepsilon^S_{d,t}(\Theta)\right) &= 0 \\ E_{NonMMF}(e_{f,d,t}\left(\Theta)\varepsilon^S_{d,t}(\Theta)\right) &= 0, \\ \text{where } e_{f,d,t} = \delta_{f,t} - 1 - \tilde{\phi}_f(y_{d,t} - y_{ONRRP,t}) \end{aligned}$$

where $E_{MMF}$ refers to a sample average, taken across dealers and across money market funds, and where $E_{NonMMF}$ is the same, taken across non-MMF investors, and $e^{\perp}_{f,d,t}$ refers to $e_{f,d,t}$, orthgonalized with respect to any observable asset characteristics $x_n$ that the econometrician views as relevant for demand (that is, relevant for a fund determining the asset's contribution to portfolio-wide risk). $\Theta = \{\tilde{\phi}_{MMF}, \tilde{\phi}_{NonMMF}, \beta_S\}$ represents the three parameters being estimated in this framework, and the notation $\varepsilon_{f,d,t}(\Theta)$ and $\varepsilon^S_{d,t}(\Theta)$ is meant to remind that these objects rely in part on a given guess of the parameters in $\Theta$.

Taken together, these two moment conditions, along with the window dressing moment condition described in equation 3.19, describe a just-identified system with three moment conditions and three parameters, $\Theta$. Intuitively, solving these moment conditions for a parameter estimate $\hat{\Theta}$ means finding parameter draws that *jointly* account for why window dressers see larger decreases in yields at the end of each month; for why money market mutual funds choose to hold smaller portfolio shares in FCP from issuers with larger balance sheet costs; and why *non-MMF* investors choose to do the same. Given that the model implies the same asset demand system outlined above should govern *both* the time series variation of yields at the end of each month *and* the cross-sectional variation of portfolio shares across assets in a given date, this empirical strategy is able to use both sources of variation to learn about the model parameters.

### 3.4.6 Data and Feasibility

This more developed empirical strategy – aside from the obvious goal of relying on plausible identifying assumptions – was designed with specific attention paid to current universe of available data on US money markets. This is so that this strategy might be used, modulo some adaptations for minor data practicalities, in future work estimating this system.

The data that we use in our simplified empirical work in section 3 includes detailed (e.g. CUSIP-level) public holdings data from every US money market mutual fund, as reported to the SEC on the last trading date of each month. Importantly, this data does *not* include money fund holding data as of the penultimate trading date in each month. While useful for the results presented in section 3, this dataset is not necessary for the fuller approach outlined in this section.

Instead, this new empirical approach first relies on the Federal Reserve's dataset of daily, non-anonymized trading volumes by counterparty at the Federal Reserve's Overnight Reverse Repo facility. This data occupies an important role in this empirical strategy because, in the logit-style demand setup, an investor's position in the outside asset (ONRRP) acts as a sufficient statistic for the total desirability of the inside (non-ONRRP) options in their investment universe. This data is important for evaluating the moment condition in Equation 3.19.

The more developed empirical strategy of this section further relies on having detailed, daily-level data available on the financial commercial paper market in the United States. This data is necessary to identify the likeliest window dressing CP issuers (via searching for obvious patterns in end-of-month issuance by certain firms); to measure CUSIP-level yields in commercial paper; and to measure total outstanding quantity in commercial paper at the CUSIP level. This data is important in evaluating all three of the proposed moment conditions.

Thankfully, this data exists and is commercially available from the Depository Trust and Clearing Corporation. [8] This private dataset includes trade-level information on volumes and yields, for nearly all CUSIPs sold in the primary commercial paper market in the United States.

## 3.5 Conclusion and Future Work

In this paper, we have developed a framework to use currently-available, existing data to estimate a full asset demand system for US money markets. This framework takes advantage of both the advantages and disadvantages of the existing data, where certain narrow objects (like ONRRP volume) are available are very high frequencies and other very-detailed objects (like detailed MMF holdings) are only available at month end. The estimation strategy proposed in this paper relies on the existence of issuer-level, exogenous heterogeneity in balance sheet "costs of leverage", and on the plausible assumption that changes in these costs drive end-of-month changes in the quantities and yields of the likeliest window dressers in the FCP market.

Naturally, this paper opens the door for future work that applies this framework to estimate this asset demand system, using the more comprehensive data from the clearing house DTCC. While this will undoubtedly entail its own, unexpected implementation challenges, our analysis in this paper's final section suggests that the key dimension of heterogeneity required by this approach does exist in the data – that larger average end-of-month decreases in outstanding quantity along geographic lines comes with larger average decreases in yield, consistent with geographic lines capturing a *supply* shock dimension. This is consistent with commentary in the empirical literature and the business press, which often describe certain geographic jurisdictions in which window dressing incentives are particularly salient.

[8] See a description of the dataset at https://www.dtcc.com/settlement-and-asset-services/settlement/dtcc-commercial-paper-and-institutional-certificates-of-deposit-data.

# Bibliography


Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson (2017): "Measuring systemic risk," *The Review of Financial Studies*, 30, 2–47.

Adrian, T. and M. K. Brunnermeier (2016): "CoVaR," *The American Economic Review*, 106, 1705–1741.

Adrian, T., E. Etula, and T. Muir (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance*, 69, 2557–2596.

Ahnert, T. and M. Macchiavelli (2021): "Safe Assets , Intermediation and Fragility : Theory and Evidence ," *Federal Reserve Bank of New York Staff Report*, 1026.

Allen, F. and D. Gale (2000): "Financial contagion," *Journal of political economy*, 108, 1–33.

Anbil, S. and Z. Senyuz (2018): "The Regulatory and Monetary Policy Nexus in the Repo Market," *Federal Reserve Board Finance and Economics Discussion Series*, 027.

Angeletos, G.-M., F. Collard, and H. Dellas (2023): "Public Debt as Private Liquidity: Optimal Policy," *Journal of Political Economy*, 131.

Bacchetta, P., S. Tièche, and E. van Wincoop (2023): "International Portfolio Choice with Frictions: Evidence from Mutual Funds," *The Review of Financial Studies*, 36, 4233–4270.

Bayer, C., B. Born, and R. Luetticke (2023): "The liquidity channel of fiscal policy," *Journal of Monetary Economics*, 134, 86–117.

Blanchard, O. (2019): "Public debt and low interest rates," *American Economic Review*, 109, 1197–1229.

Boss, M., H. Elsinger, M. Summer, S. Thurner, et al. (2004): "An empirical analysis of the network structure of the Austrian interbank market," *Oesterreichesche Nationalbanks Financial stability Report*, 7, 77–87.

Brownlees, C. and R. F. Engle (2016): "SRISK: A conditional capital shortfall measure of systemic risk," *The Review of Financial Studies*, 30, 48–79.

Brunnermeier, M., E. Farhi, R. S. J. Koijen, A. Krishnamurthy, S. C. Ludvigson, H. Lustig, S. Nagel, and M. Piazzesi (2021): "Review Article: Perspectives on the Future of Asset Pricing," *The Review of Financial Studies*, 34, 2126–2160.

Carlson, M., B. Duygan-bump, F. Natalucci, and R. William (2014): "The Demand for Short-Term, Safe Assets and Financial Stability: Some Evidence and Implications for Central Bank Policies," *Federal Reserve Board*, 102.

Cashin, D. (2023): "Treasury Safety, Liquidity, and Money Premium Dynamics: Evidence from Debt Limit Impasses," *Journal of Money, Credit and Banking*, 55, 1475–1506.

Craig, B. and G. von Peter (2014): "Interbank tiering and money center banks," *Journal of Financial Intermediation*, 23, 322–347.

D'Amico, S. and T. B. King (2013): "Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply," *Journal of Financial Economics*, 108, 425–448.

D'Avernas, A. and Q. Vandeweyer (2023): "Treasury Bill Shortages and the Pricing of Short-Term Assets *," *Journal of Finance Forthcoming*.

Diamond, W. (2020): "Safety Transformation and the Structure of the Financial System," *Journal of Finance*, 75, 2973–3012.

Drechsler, I., A. Savov, and P. Schnabl (2018): "A Model of Monetary Policy and Risk Premia," *Journal of Finance*, 73, 317–373.

Drehmann, M. and N. Tarashev (2013): "Measuring the systemic importance of interconnected banks," *Journal of Financial Intermediation*, 22, 586–607.

Du, W., A. Tepper, and A. Verdelhan (2018): "Deviations from Covered Interest Rate Parity," *Journal of Finance*, 73.

Duffie, D. (2010): "Presidential address: Asset price dynamics with slow-moving capital," *Journal of Finance*, 65, 1237–1267.

Eisenberg, L. and T. H. Noe (2001): "Systemic risk in financial systems," *Management Science*, 47, 236–249.

Elsinger, H., A. Lehar, and M. Summer (2006a): "Risk assessment for banking systems," *Management science*, 52, 1301–1314.

——— (2006b): "Using market information for banking system risk assessment," *International Journal of Central Banking*, 2, 137–165.

Feldhütter, P. and D. Lando (2008): "Decomposing swap spreads," *Journal of Financial Economics*, 88, 375–405.

Fleming, M. J. and A. Sarkar (2014): "The Failure Resolution of Lehman Brothers," *Federal Reserve Bank of New York Economic Policy Review*, 20, 175.

Freixas, X., B. M. Parigi, and J.-C. Rochet (2000): "Systemic risk, interbank relations, and liquidity provision by the central bank," *Journal of money, credit and banking*, 611–638.

Gabaix, X. and R. S. J. Koijen (2021): "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis," *NBER WORKING PAPER SERIES*, 28967.

Garbade, K. D. (2007): "The Emergence of Regular and Predictable as a Treasury Debt Management Strategy," *FRBNY Economic Policy Review*, 53–71.

Gârleanu, N. and L. H. Pedersen (2016): "Dynamic portfolio choice with frictions," *Journal of Economic Theory*, 165, 487–516.

Glasserman, P., A. Sirohi, and A. Zhang (2017): "The Effect of 'Regular and Predictable' Issuance on Treasury Bill Financing," *Economic Policy Review*, 43–56.

Glasserman, P. and H. P. Young (2015): "How likely is contagion in financial networks?" *Journal of Banking & Finance*, 50, 383–399.

——— (2016): "Contagion in Financial Networks," *Journal of Economic Literature*, 54, 779–831.

Gorodnichenko, Y. and W. Ray (2018): "The Effects of Quantitative Easing: Taking A Cue From Treasury Auctions," *NBER Working Paper Series*, 24122.

Gorton, G. (2017): "The history and economics of safe assets," *Annual Review of Economics*, 9, 547–586.

Greenwood, R., S. Hanson, J. Rudolph, and L. Summers (2015a): "The Optimal Maturity of Government Debt," in *The $13 Trillion Question: Managing the U.S. Government's Debt*, ed. by D. Wessell, Brookings Institution Press, chap. 1, 1–42.

Greenwood, R., S. G. Hanson, and G. Y. Liao (2018): "Asset price dynamics in partially segmented markets," *Review of Financial Studies*, 31.

Greenwood, R., S. G. Hanson, and J. C. Stein (2015b): "A Comparative-Advantage Approach to Government Debt Maturity," *Journal of Finance*, 70, 1683–1722.

Gropp, R., M. lo Duca, and J. Vesala (2009): "Cross-border bank contagion in Europe," *International Journal of Central Banking*, 97–139.

Hawkesby, C., I. Marsh, and I. Stevens (2007): "Comovements in the equity prices of large complex financial institutions," *Journal of Financial Intermediation*, 2, 391–411.

He, Z., B. Kelly, and A. Manela (2017): "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 126, 1–35.

He, Z. and A. Krishnamurthy (2013): "Intermediary asset pricing," *American Economic Review*.

He, Z., S. Nagel, and Z. Song (2022): "Treasury inconvenience yields during the COVID-19 crisis," *Journal of Financial Economics*, 143, 57–79.

Infante, S. (2020): "Private money creation with safe assets and term premia," *Journal of Financial Economics*, 136, 828–856.

Kacperczyk, M., C. Pérignon, and G. Vuillemey (2021): "The Private Production of Safe Assets," *Journal of Finance*, 76, 495–535.

Kekre, R. and M. Lenel (2023): "The Flight to Safety and International Risk Sharing," *NBER Working Paper*, 29238.

Klingler, S. and S. Sundaresan (2023): "Diminishing treasury convenience premiums: Effects of dealers' excess demand and balance sheet constraints," *Journal of Monetary Economics*, 135, 55–69.

Koijen, R. S. and M. Yogo (2019): "A demand system approach to asset pricing," *Journal of Political Economy*, 127, 1475–1515.

Koijen, R. S. J. and M. Yogo (2020): "Exchange Rates and Asset Prices in a Global Demand System," *NBER Working Paper Series*, 1–53.

Krishnamurthy, A. and W. Li (2023): "The Demand for Money, Near-Money, and Treasury Bonds," *The Review of Financial Studies*, 36, 2091–2130.

Krishnamurthy, A. and A. Vissing-Jorgensen (2011): "The Effects of quantitative easing on interest Rates: Channels and implications for policy," *Brookings Papers on Economic Activity*, 287, 215–287.

——— (2012): "The aggregate demand for Treasury debt," *Journal of Political Economy*, 120, 233–267.

——— (2015): "The impact of Treasury supply on financial sector lending and stability," *Journal of Financial Economics*, 118, 571–600.

Longstaff, F. A. (2004): "The Flight to Liquidity Premium in U . S . Treasury Bond Prices," *The Journal of Business*, 77, 511–526.

Lou, D., H. Yan, and J. Zhang (2013): "Anticipated and repeated shocks in liquid markets," *Review of Financial Studies*, 26, 1890–1912.

Mehrotra, N. R. and D. Sergeyev (2021): "Debt sustainability in a low interest rate world," *Journal of Monetary Economics*, 124, S1–S18.

Mian, A. R., L. Straub, and A. Sufi (2022): "A GOLDILOCKS THEORY OF FISCAL DEFICITS We are grateful to A Goldilocks Theory of Fiscal Deficits," *NBER WORKING PAPER SERIES*.

Munyan, B. (2015): "Regulatory Arbitrage in Repo Markets," *OFR Working Paper Series*, 15.

Nagel, S. (2016): "The liquidity premium of near-money assets," *Quarterly Journal of Economics*, 131, 1927–1971.

Nazeran, P. and D. Dwyer (2015): "Credit Risk Modeling of Public Firms: EDF9," .

Olea, J. L. M. and M. Plagborg-Møller (2021): "Local projection inference is simpler and more robust than you think," *Econometrica*, 89, 1789–1823.

Phillot, M. (2023): "U.S. Treasury Auctions : A High Frequency Identification of Supply Shocks," *American Economic Journal: Macroconomics Forthcoming.*

Reis, R. (2021): "The constraint on public debt when r < g but g < m," *BIS Working Paper Series.*

Santoro, P. J. (2012): "The Evolution of Treasury Cash Management During the Financial Crisis," *Federal Reserve Bank of New York Current Issues in Economics and Finance*, 1917.

Sunderam, A. (2015): "Money creation and the shadow banking system," *Review of Financial Studies*, 28, 939–977.

Upper, C. (2011): "Simulation methods to assess the danger of contagion in interbank markets," *Journal of Financial Stability*, 7, 111–125.

Upper, C. and A. Worms (2004): "Estimating bilateral exposures in the German interbank market: Is there a danger of contagion?" *European Economic Review*, 48, 827–849.

van Lelyved, I. and F. Liedorp (2006): "Interbank Contagion in the Dutch Banking Secgtor: A Sensitivity Analysis," *International Journal of Central Banking*, 2, 99–133.

Vayanos, D. and J.-L. Vila (2021): "A PreferredHabitat Model of the Term Structure of Interest Rates," *Econometrica*, 89, 77–112.

Vissing-Jorgensen, A. (2023): "Balance Sheet Policy Above the Effective Lower Bound," *ECB Forum on Central Banking 2023 Working Paper.*

# Appendix A

# Appendix to Chapter 1

## A.1 Seasonality Instruments

The most common practice for estimating $\beta$ in this literature is using a single equation method, either via OLS or via 2SLS with seasonality instruments. First, I show that single equation methods likely mask substantial *horizon* heterogeneity in the convenience yield response to supply changes. Even simple OLS results suggest that the short-run response to a supply change is substantially larger than the medium-run response. Next, I show that the seasonality instruments of Greenwood et al. (2015b), which provided sensible results in a pre-crisis sample, do not have the theoretically-suggested effect on $\beta$ estimates in a post-crisis sample, compared to the simple OLS estimation.

### A.1.1 Misspecification in Single Equation Estimates

A single equation, seasonal instrumental variables approach over the pre-crisis sample shows large effects of T-bill supply movements on convenience yields, but little effect in the post-crisis era. In both the pre-crisis and post-crisis sample, there is evidence that a single equation estimated in 4-week differences masks information about the horizons over which these effects matter.

The most common specification in the literature for using high frequency data to estimate $\beta$ involves an estimating, via OLS or instrumental variables, a four-week differenced variant of the long-run convenient asset demand curve, Equation (1.1) above. That is, most researchers begin with the equation[1]

$$\frac{i_t - i_t^B}{1+i_t} - \frac{i_{t-4} - i_{t-4}^B}{1+i_{t-4}} = \alpha + \beta\left(B_t - B_{t-4}\right) + \xi_t - \xi_{t-4} \tag{A.1}$$

Of course, this specification will be inconsistent with my structural framework above in the event that $B_t - B_{t-4} \neq B_t^* - B_{t-4}^*$ (as is the case in my model of sluggish adjustment). From Equation (**??**) above we can see that under my model, even in the case where $B_t$ and

[1] #Note to self: Check that this is true for the non-GHS, recent papers, or if the specification is slightly different

$\xi$ are uncorrelated, the estimate $\hat{\beta}$ in this equation will be a consistent estimator only of a complicated function of $\beta$, $\phi$, $\mu$, and the population autocorrelation of $B_t$. This specification also implies a permanence of the effect captured in $\hat{\beta}$. This equation suggests that the convenience yield response will only tend to reverse if convenient asset supplies themselves reverse.

OLS estimates in Table 1 below offer initial suggestive evidence, from two non-overlapping subsamples, that estimates of $\hat{\beta}$ from these single equation specifications are driven primarily by a very strong effect at the shortest horizons. I show this by first estimating Equation by OLS, in pre-crisis and post-crisis samples in columns 1-2. I then estimate two variants of Equation (A.1) by OLS: One in columns 3-4 mimicking the specification with one-week differences, and another in columns 5-6 conducting a horse race between $B_t - B_{t-1}$ and $B_{t-1} - B_{t-4}$.

| | Δ4.OIS-Bill, 4w | Δ4.OIS-Bill, 4w | Δ4.OIS-Bill, 4w | Δ4.OIS-Bill, 4w |
|---|---|---|---|---|
| Δ4.All Bills /GDP | $-2.18^{***}$ | | $-12.14^{***}$ | |
| | 0.68 | | 4.59 | |
| Δ.All Bills /GDP | | $-13.62^{***}$ | | $-51.94^{***}$ |
| | | 1.69 | | 8.76 |
| *L*.Δ3.All Bills /GDP | | 1.14 | | −0.39 |
| | | 0.90 | | 4.89 |
| $R^2$ | 0.02 | 0.09 | 0.03 | 0.11 |
| N | 575 | 575 | 311 | 311 |
| Sample | 2009 – 2019 | 2009 – 2019 | 2002 – 2008 | 2002 – 2008 |
| Estimation | *OLS* | OLS | OLS | OLS |

Table A.1: Single Equation OLS Estimates

The results suggest that forcing $\Delta B_t$ and $B_{t-1} - B_{t-4}$ to share the same parameter estimate in this setting is a misspecification. Instead, it appears that the entirety of the single equation, four-week difference coefficient estimate is driven by a very strong, very short-lived effect. This finding is shared in both the pre-crisis and post-crisis samples.

### A.1.2 Seasonal Instrumental Variables

Instrumenting for seasonal supply in regressions of this form will produce inconsistent estimates if there is a seasonal component to convenient asset demand that is correlated with the seasonality in T-bill supply. Single equation IV results estimated over the post-crisis period are consistent with a positive correlation between these seasonalities, which would tend to attenuate estimates of $\beta$.

A standard estimating approach in the literature is to instrument for T-bill supply using seasonality, in the single equation approach. This practice comes from a simultaneity concern that I will refer to in this paper as *opportunistic issuance*.[2] If the US Treasury tends to

[2]An equivalent way of phrasing this worry is that the US Treasury's T-bill *supply* curve may be upward-

respond to positive convenience demand shocks (which raise the convenience yield) by issuing more T-bills, then that will induce a positive correlation between $B_t$ and the unobserved structural residual in this case, $\xi_t - \xi_{t-4}$. Estimating $\hat{\beta}$ using the seasonal variation in $B_t - B_{t-4}$ alone is meant to combat this worry, because the seasonality in T-bill supply comes from a well-understood source – the timing of concentrated fiscal cash receipts around predictable calendar dates. The Treasury's typical practice of decreasing T-bill supplies after April 15th, for instance, is almost certainly a result of the influx of federal cash on tax day and not a response to demand conditions. In practice, this seasonal IV approach is typically implemented by using 52 week of year dummy variables as instruments when estimating Equation 3.

While a seasonal IV approach addresses worries about opportunistic issuance, it does not address other concerns of omitted variable bias that threaten to violate the exclusion restriction. For seasonality to serve as a valid instrument in this setting, it must be true that seasonal variation in T-bill supply is uncorrelated with changes in demand conditions. Violations of this exclusion restriction need not be causal. To threaten the exclusion restriction, it need only be true that there exists some seasonality in convenient asset demand that is correlated with the seasonality in T-bill supplies.

There is ample reason to believe that there exists a seasonal component in T-bill demand, especially in the post-crisis sample of this paper. Firms preparing for cash outflows related to payroll or year-end bonuses may desire a T-bill's perfect nominal safety or convenience more in the run-up to those payments. Regulated financial institutions that value convenient assets holdings for regulatory reasons may desire those holdings more strongly before a month-end of quarter-end regulatory filing. Or, financial institutions attempting to window dress their leverage ratios at quarter-end or month-end may limit their issuance of private sector substitutes for T-bills at those times.[3]

In the post-crisis sample of this paper, the way in which the seasonal IV affects estimates of $\hat{\beta}$ is no longer consistent with the assumption that justified their use, that seasonal variation is less positively correlated with demand than nonseasonal variation. In column 1 of Table 2, we estimate $\beta$ by regressing one-week differences in the convenience yield on one-week changes in T-bill supplies, on a post-crisis weekly sample from 2009-2019. The results are negative and significant, and indicate that a one-week increase in T-bill supplies equal to 1% of nominal GDP decreases the T-bill convenience yield by 2.84 basis points. In column 2, we estimate the same equation by 2SLS, using 52 week-of-year dummy variables as instruments for $\Delta B_t$. Unlike in Greenwood et al. (2015b), which is estimated on a pre-crisis sample, this procedure *attenuates* the observed effect and causes it to lose statistical significance. In column 3, I take the opposite approach to column 2 by controlling for 52 week-of-year dummy variables. Column 2 relies on sesaonal variation for identification, while column 3 removes seasonal variation in T-bill supplies in identification. In the opportunistic issuance story that justifies the seasonality instrument, seasonal variation in supplies is the variation that should create a larger magnitude effect on convenience yields. The results of Table 2 instead suggest the opposite - that sesaonal variation in T-bill supplies has a weaker connection to convenience yields. These results are consistent with the interpretation that,

sloping in the convenience yield.

[3] Window dressing of this sort has been convincingly documented in Du et al. (2018).

in a post-crisis sample, seasonality in T-bill supplies is positively correlated with seasonality in unobserved convenience demand shifters.[4]

These results suggest that a new approach is needed for estimating $\beta$ using post-crisis data.

| | (1) | (2) | (3) |
|---|---|---|---|
| | $\Delta$.OIS-Bill, 4w | $\Delta$.OIS-Bill, 4w | $\Delta$.OIS-Bill, 4w |
| $\Delta$.All Bills /GDP | $-2.84^{**}$ | $-0.45$ | $-4.54^{***}$ |
| | 1.14 | 1.86 | 1.55 |
| $R^2$ | 0.01 | 0.00 | 0.18 |
| N | 575 | 575 | 575 |
| Sample | 2009 − 2019 | 2009 − 2019 | 2009 − 2019 |
| Estimation | OLS | Seasonal IV | Seasonal Controls |

Table A.2: Seasonal Instrumental Variables

### A.1.3 Structural Estimates of $\xi$ and Seasonality

Another way to understand the relationship between my results and those suggested by a seasonality-based instrument is to study the properties of $\hat{\xi}_t = i_t - i_t^B - \hat{\beta}\left(\frac{1}{\hat{\mu}} - 1\right)\Delta \frac{B}{GDP}_t - \hat{\beta}\frac{B}{GDP}_t$ – that is, the implied structural demand residual implied by my setup and my estimates of the flow and stock effect of T-bill supplies. Understanding any apparently *seasonality* in this estimated structural residual shows what my parameter estimates suggest for the implied bias from using seasonal variation in T-bill supplies as an instrument, in this setting.

Figure A.1 below plots coefficients from the following two regression equations, estimated via OlS

$$\hat{\xi}_t - \hat{\xi}_{t-4} = \phi_0 + \sum_{w=1}^{52} \phi_w \mathbf{1}\left(\text{week}(t) = w\right) + e_t$$
$$\frac{B}{GDP}_t - \frac{B}{GDP}_{t-4} = \psi_0 + \sum_{w=1}^{52} \psi_w \mathbf{1}\left(\text{week}(t) = w\right) + w_t$$

with $\phi_w$ and $\psi_w$ plotted for each week $w$. This shows one transparent, if simple, estimate of the *seasonality* in four-week differences in these two object (recall from above that four-week difference specifications are common in this literature). Figure A.1 shows that my parameter estimates suggest a *positive* correlation between the seasonality in T-bill supply changes and the seasonality in unobserved demand changes. [5]

[4] Of course, it is also consistent with a story that non-seasonal variation in T-bill supply is *negatively* correlated with unobserved demand. Differentiating between these two stories is challenging, of course.

[5] Indeed, the correlation between these two seasonal components is 35.37%. The $R^2$ measure from an OLS regression of $\phi_w$ on $\psi_w$ is 12.5%.

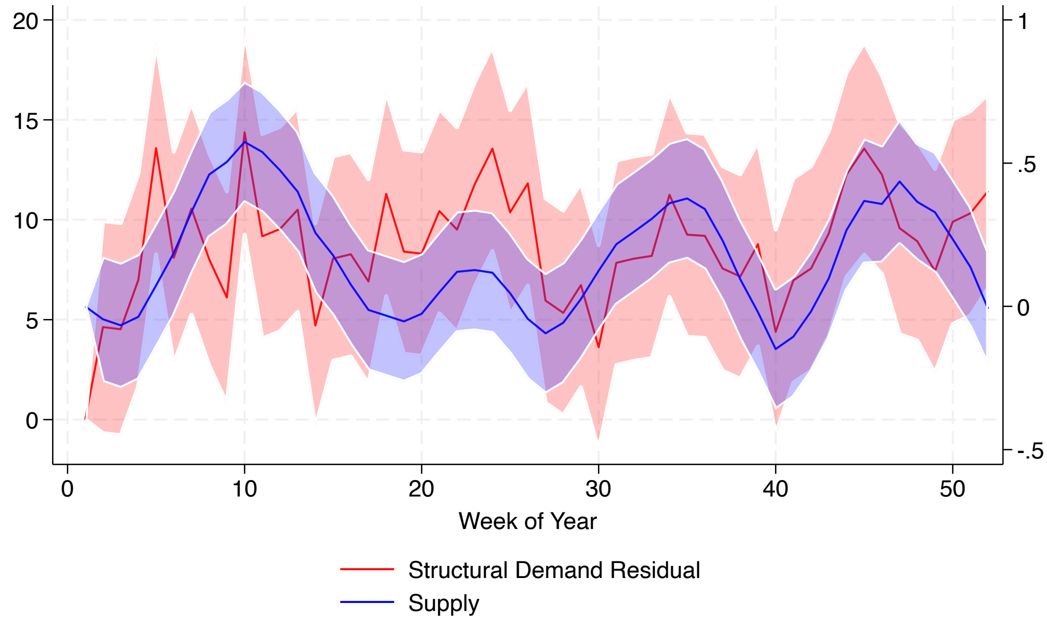


Figure A.1: Seasonal Components of Changes in T-bill Supplies and Structural Residual

## A.2 GMM Estimates with 2-Parameter Investor Inertia

In this appendix, I show GMM results with a more complicated assumed structure for the investor inertia leading to different stock and flow effects of T-bill supply changes. In this appendix, we will assume that investor T-bill holdings follow

$$B_t - B_{t-1} = \mu \left( \frac{B_t^*}{GDP_t} - B_{t-1}^* \right) + \phi \left( B_{t-1}^* - B_{t-1} \right)$$

This parameterization has one additional parameter $\phi$ than the baseline representation in the main text. This additional parameter allows the convenience yield response to a permanent T-bill supply change to converge to new long-run level only gradually. Figure A.2 shows the GMM "fit" for the LP-IV empirical moments, with this three parameter setup (including $\beta$, $\mu$ and $phi$), using the power-preserving instrument setup described in Section 1.7.5. Unsurprisingly, given that another parameter is available to improve it, the fit is slightly better than the fit of the baseline representation, shown in Figure 1.9 from the main text.

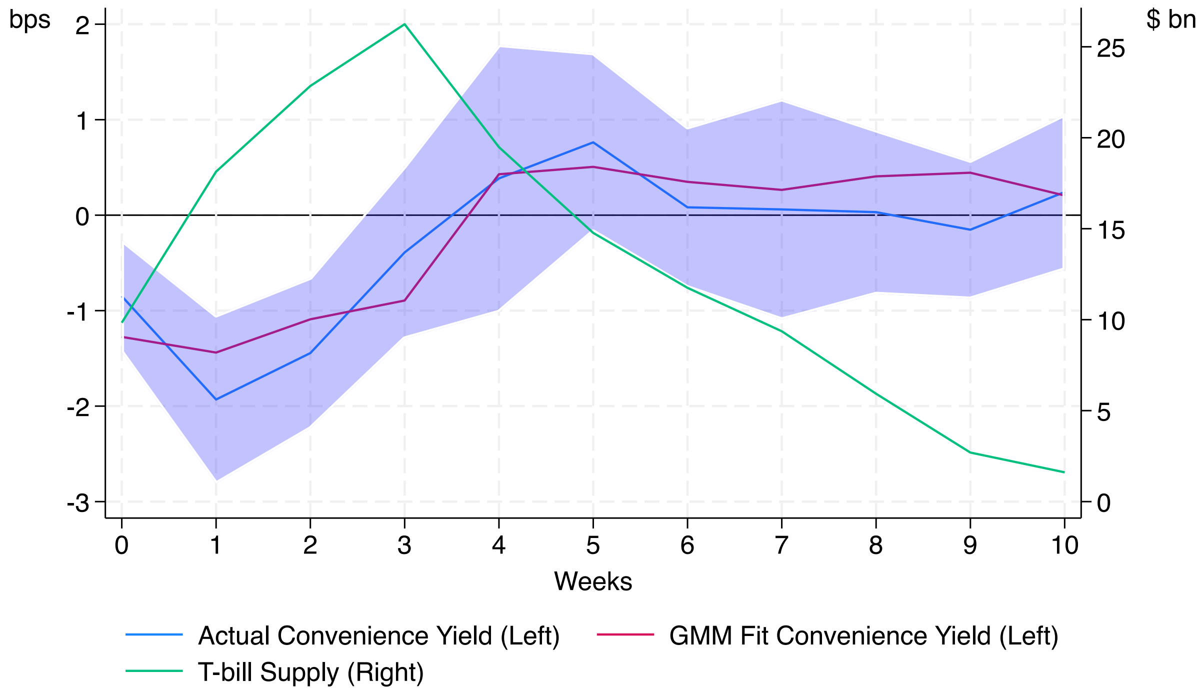


Figure A.2: GMM Fit with Two-Parameter Investor Inertia Model

*Note:*

However, the point estimates of the estimate parameters in this case do not suggest fundamentally different results than the baseline. The estimates underlying the fit in Figure A.2 have $\hat{\beta} = -1.26\frac{\text{bp}}{\$100bn}$, where $\beta$ still dictates the long-run convenience yield effect as before. I estimate $\hat{\mu} = 0.097$, which suggests that the same-week impact of a \$100bn increase in T-bill supply equals 14.2 basis points – only moderately larger than the effect estimated in the baseline. The estimate for $\phi$ is 0.076. [6]

However, the standard errors on these estimates are notably larger than in the baseline setup. Standard errors on the $\beta hat$ estimate are over twice as large in this 3-parameter setup, compared to the baseline. When not using the power-preserving setup of Section 1.7.5, these multiples are even larger, with the 3-parameter setup having $\hat{\beta}$ standard errors over three times as large as the baseline case.

In sum, the point estimates of this more complicated parameterization suggest that estimates from this setup are not fundamentally different from those in the baseline. However, the inclusion of this extra complication into the estimation decreases precision considerably.

[6]Note that this parameterization collapses to the baseline parameterization when $\mu = \phi$. These parameter estimates suggest that the estimated levels of these two parameters are indeed similar.

## A.3 Additional Information on Treasury Issuance Policy

In this Appendix, I present a separate slide from the same Treasury presentation as Figure 1.3. This slide, shown in Figure A.3, shows some additional information about the Treasury's "regular and predictable" issuance philosophy". This discussion is useful for those who are puzzled by a perceived inconsistency between the Treasury's refusal to respond to demand shoccks and their mandate to borrow at the least cost to taxpayers.

This slide describes how the Treasury believes that avoiding opportunistic issuance *does* lead to lower funding costs *in the long run.* The slide clarifies that the Treasury's stategy does allow for "slowly adjusting to shifts in expected costs" over time. However, the Treasury appears to believe that doing so at higher frequencies might compromise liquidity in a way that raises the Treasury's long-run funding costs.

As the main text discusses, my identification strategy does *not* require that these views of the Treasury's be *accurate* (although they may well be). It only requires that the Treasury believe them to be so, such that the Treasury is incentivized to adhere to their policy of avoiding opportunistic issuance.

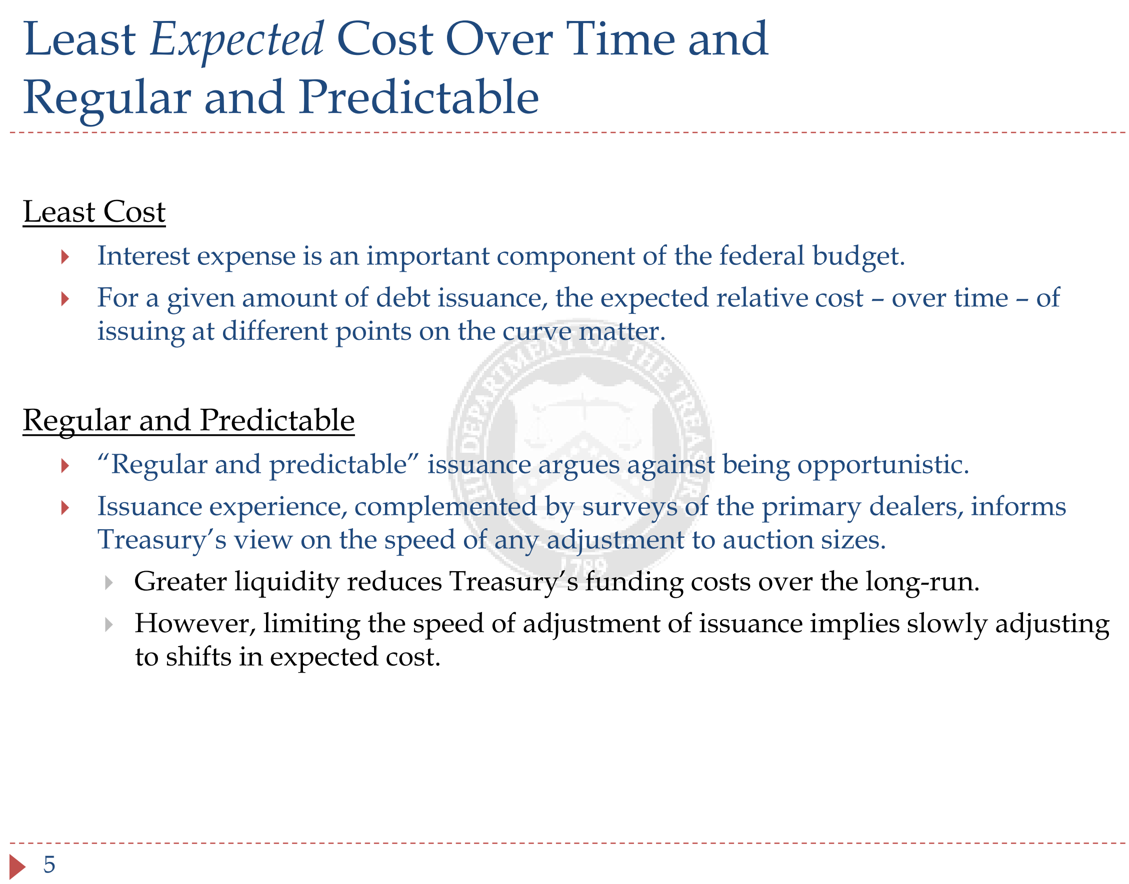


Figure A.3: Additional Informative Slide from Treasury Presentation

*Note:*

## A.4 GMM Estimates with Pre-Specified Weight Matrix

Here, I present the GMM fit and parameter estimates for a specification with a prespecified GMM weight matrix. In the baseline case, the theoretical "optimal" GMM weight matrix is used, given Newey-West estimates of the variance of the fitted moment conditions. Here, the prespecified weight matrix is the identity matrix.

When the identity matrix is used to weight the moments, the GMM fitting procedure equally prioritizes matching each input moment. As such, the visual fit of this procedure looks modestly better (in the sense that fewer of the LP-IV moments are fitted poorly) than in the baseline.

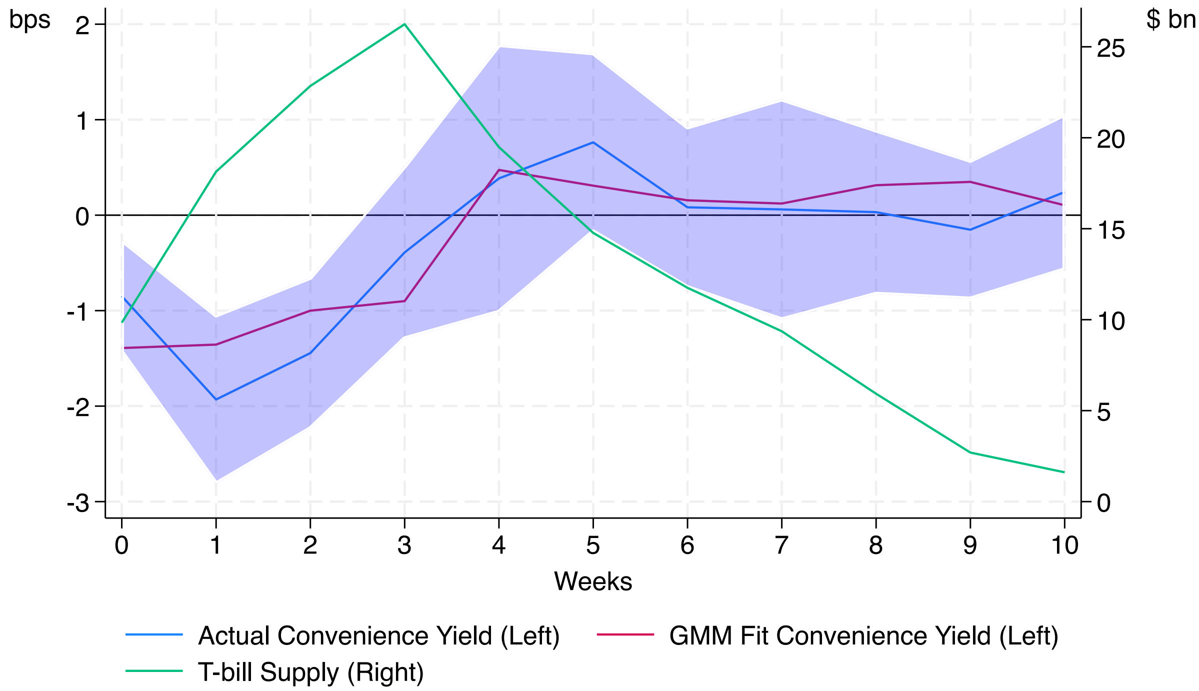


Figure A.4: GMM Fit with Pre-specified GMM Weight Matrix, Identity

*Note:*

Figure A.5 shows the parameter estimates underlying the fit of Figure A.4. These values are quite similar to the baseline numbers presented in Figure 1.10. However, standard errors are somewhat larger – consistent with the fact that the "optimal" GMM is meant to weight moments in the way that least to the smallest asymptotic variance.

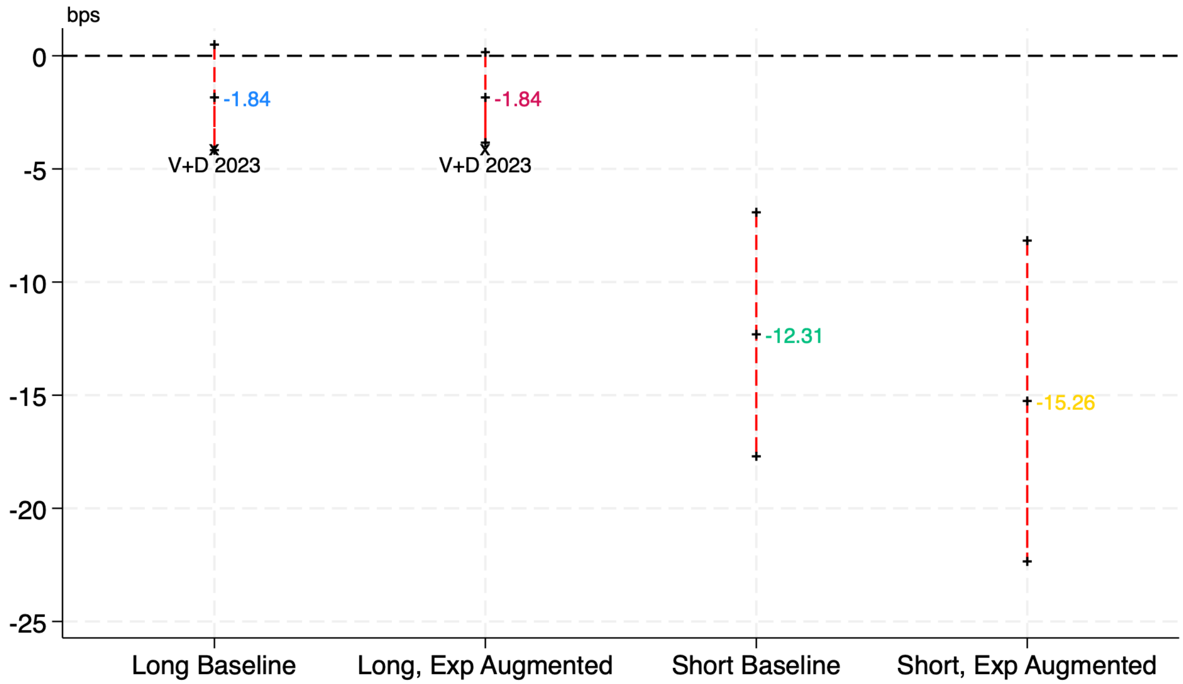


Figure A.5: GMM Stock and Flow Estimates with Pre-specified GMM Weight Matrix, Identity

## A.5 Cyclicality of Fiscal Surprises

A channel through which our exclusion restriction could be violated is if changes in convenient asset demand are correlated with revisions to market expectations of future receipts or expenditures. That is, an alternate set of conditions that will tend to violate the exclusion restriction is

$$\mathbb{C}(\mathbb{E}_{t,Priv}(e_{t+k}) - \mathbb{E}_{t-\delta,Priv}(e_{t+k}), \varepsilon^{D}_{t+h}) \neq 0$$
$$\mathbb{C}(\mathbb{E}_{t,Priv}(e_{t+k}) - \mathbb{E}_{t-\delta,Priv}(e_{t+k}), B_t - \mathbb{E}_{t-\delta,Priv}B_t) \neq 0$$

and analagously for receipts $r_{t+h}$. The second of these is almost surely true: A surprise in T-bill issuance will almost certainly cause private actors to update their expectations of future receipts or expenditures.

My first evidence in favor of the exclusion restriction is that federal government receipt and expenditure surprises at high frequencies have very little power in predicting other macroeconomic variables that could plausibly drive the cyclicality in convenience demand. I demonstrate this using a sensible proxy for well-informed agents' expectations of government receipts, government expenditures, and other macroeconomic variables. As my proxy, I use pre-FOMC projections by Federal Reserve Board economists from the Federal Reserve Tealbooks. Using this data, I wish to show whether quarterly surprises (relative to the

Tealbook projections) in government expenditures or receipts tend to be correlated with quarterly surprises in other macroeconomic variables of interest.

To show this, I estimate a simple set of OLS regressions meant to transparently show whether this correlation exists, and at what frequencies of surprise it tends to be most severe. Specifically, I estimate OLS regressions of the form

$$Z_t - \mathbb{E}_{t-h,FOMC}(Z_t) = \alpha + \gamma(G_t - \mathbb{E}_{t-h,FOMC}(G_t)) + e_t$$

where $G \in \{\text{Receipts}, \text{Expenditures}\}$, and $Z$ is some other macroeconomic object of interest as indicated. The object $\mathbb{E}_{t-h,FOMC}$ represents the Tealbook projections for some variable's realization in quarter $t$, as of the first available projection in quarter $t-h$. That is, $\mathbb{E}_{t,FOMC}(X_t)$ indicates a start-of-quarter projection of a quarterly realization that will not be known until the end of the quarter and so, in general, $X_t \neq \mathbb{E}_{t,FOMC}(X_t)$.

Consistent with the sensible intuition that fiscal flows have a cyclical component, there is evidence that receipt and expenditures surprises are tied to business cycle innovations at horizons around 1-year. Evidence to this effect is show in Table A.3, which shows results from estimation equation 4 vis OLS, with $Z \in \{\text{Nominal GDP Growth}, \text{Real GDP Growth}\}$ and $h \in \{0, 1, 2, 3, 4\}$. The most striking components of this table are the $R^2$ measures. At the 1-year horizon, government receipt and expenditure surprises have a moderately predictive quality for real and nominal GDP growth surprises. The peak is for tax receipt surprises, which help explain 10% and 5% of nominal and real GDP growth surprises, respectively. Although note that, even at this horizon, the $R^2$ for tax receipts is notably less than that for any of the other righthand side macroeconomic variable shown.

This evidence disappears at shorter horizons, so that same-quarter receipt and expenditure projection errors appear wholly disconnected from other macroeconomic surprises. The bottom panel of Table A.3 shows this most clearly, where the $R^2$ measure for same-quarter receipt and expenditure surprises predicting same-quarter GDP growth surprises has fallen nearly to zero. The progression from the top $h = 4$ panel to the bottom $h = 0$ panel shows the steady decline, both in the $R^2$ measure and the statistical significance of the regression coefficients, as we study surprises at higher and higher frequencies.

The most relevant panel in Table A.3 for considering the plausibility of this paper's exclusion restriction is the bottom panel, studying same-quarter surprises. This is for two reasons. First, my proxy for T-bill issuance surprises will be based on a private actor's projections as of just hours before the associated announcement. As such, any information known by private actors at any longer frequencies should already be incorporated into the private expectations. Second, I will show below that, while high frequency T-bill surprises should indeed shift private agents' expectations of future receipts and expenditures, they only do so are relatively short horizons of the following several weeks. The conclusion to draw from these facts is that the receipt and expenditure information content of T-bill issuance surprises in this paper share the most in common with the highest-frequency surprises shown in Table A.3.[7]

[7]# Note to preliminary reader: I am open to suggestions about how to alter this table to better fit onto a single page. Shrinking fonts beyond their current level seems like a bad idea. But I do like display all of the $h \in \{0, 1, 2, 3, 4\}$ results, so readers can see the progression of coefficients and $R^2$.

$h = 4$

| | NGDP | NGDP | NGDP | NGDP | NGDP | RGDP | RGDP | RGDP | RGDP | RGDP |
|---|---|---|---|---|---|---|---|---|---|---|
| Housing Starts | $5.78^{***}$ | | | | | $5.64^{***}$ | | | | |
| | 1.2 | | | | | 1.1 | | | | |
| Unemployment | | $-1.31^{***}$ | | | | | $-0.90^{**}$ | | | |
| | | 0.4 | | | | | 0.4 | | | |
| Receipts | | | $2.43^{***}$ | | | | | $1.56^{**}$ | | |
| | | | 0.6 | | | | | 0.6 | | |
| Outlays | | | | $-0.17^{***}$ | | | | | $-0.07^{***}$ | |
| | | | | 0.0 | | | | | 0.0 | |
| Outlays - OMF | | | | | $-0.18^{***}$ | | | | | $-0.07^{***}$ |
| | | | | | 0.0 | | | | | 0.0 |
| N | 144 | 144 | 140 | 140 | 140 | 144 | 144 | 140 | 140 | 140 |
| $R^2$ | 0.20 | 0.14 | 0.10 | 0.08 | 0.08 | 0.22 | 0.08 | 0.05 | 0.01 | 0.01 |

$h = 3$

| | NGDP | NGDP | NGDP | NGDP | NGDP | RGDP | RGDP | RGDP | RGDP | RGDP |
|---|---|---|---|---|---|---|---|---|---|---|
| Housing Starts | $7.02^{***}$ | | | | | $6.39^{***}$ | | | | |
| | 1.4 | | | | | 1.3 | | | | |
| Unemployment | | $-1.83^{***}$ | | | | | $-1.31^{***}$ | | | |
| | | 0.4 | | | | | 0.4 | | | |
| Receipts | | | $2.28^{***}$ | | | | | $1.38^{**}$ | | |
| | | | 0.7 | | | | | 0.7 | | |
| Outlays | | | | $-2.21^{**}$ | | | | | $-1.86^{*}$ | |
| | | | | 1.1 | | | | | 1.0 | |
| Outlays - OMF | | | | | $-1.92^{**}$ | | | | | $-1.58^{**}$ |
| | | | | | 0.9 | | | | | 0.8 |
| N | 157 | 157 | 151 | 151 | 151 | 157 | 157 | 151 | 151 | 151 |
| $R^2$ | 0.18 | 0.16 | 0.06 | 0.05 | 0.06 | 0.20 | 0.11 | 0.03 | 0.04 | 0.05 |

$h = 2$

| | NGDP | NGDP | NGDP | NGDP | NGDP | RGDP | RGDP | RGDP | RGDP | RGDP |
|---|---|---|---|---|---|---|---|---|---|---|
| Housing Starts | $8.80^{***}$ | | | | | $7.99^{***}$ | | | | |
| | 1.5 | | | | | 1.4 | | | | |
| Unemployment | | $-2.61^{***}$ | | | | | $-1.98^{***}$ | | | |
| | | 0.5 | | | | | 0.4 | | | |
| Receipts | | | $2.01^{**}$ | | | | | 1.12 | | |
| | | | 0.8 | | | | | 0.8 | | |
| Outlays | | | | $-1.29$ | | | | | $-1.22$ | |
| | | | | 1.0 | | | | | 0.9 | |
| Outlays - OMF | | | | | $-1.48^{*}$ | | | | | $-1.32^{**}$ |
| | | | | | 0.8 | | | | | 0.7 |
| N | 159 | 159 | 152 | 152 | 152 | 159 | 159 | 152 | 152 | 152 |
| $R^2$ | 0.23 | 0.22 | 0.04 | 0.02 | 0.04 | 0.24 | 0.16 | 0.01 | 0.02 | 0.04 |

$h = 1$

| | NGDP | NGDP | NGDP | NGDP | NGDP | RGDP | RGDP | RGDP | RGDP | RGDP |
|---|---|---|---|---|---|---|---|---|---|---|
| Housing Starts | $9.17^{***}$ | | | | | $8.79^{***}$ | | | | |
| | 1.3 | | | | | 1.3 | | | | |
| Unemployment | | $-3.41^{***}$ | | | | | $-2.67^{***}$ | | | |
| | | 0.5 | | | | | 0.5 | | | |
| Receipts | | | 1.09 | | | | | 0.49 | | |
| | | | 1.0 | | | | | 0.9 | | |
| Outlays | | | | $-0.60$ | | | | | $-0.79$ | |
| | | | | 1.2 | | | | | 0.9 | |
| Outlays - OMF | | | | | $-0.81$ | | | | | $-0.80$ |
| | | | | | 0.8 | | | | | 0.7 |
| N | 159 | 159 | 152 | 152 | 152 | 159 | 159 | 152 | 152 | 152 |
| $R^2$ | 0.24 | 0.23 | 0.01 | 0.00 | 0.02 | 0.25 | 0.16 | 0.00 | 0.01 | 0.02 |

$h = 0$

| | NGDP | NGDP | NGDP | NGDP | NGDP | RGDP | RGDP | RGDP | RGDP | RGDP |
|---|---|---|---|---|---|---|---|---|---|---|
| Housing Starts | $9.05^{***}$ | | | | | $8.47^{***}$ | | | | |
| | 2.4 | | | | | 2.1 | | | | |
| Unemployment | | $-4.82^{***}$ | | | | | $-3.88^{***}$ | | | |
| | | 1.0 | | | | | 0.9 | | | |
| Receipts | | | 0.42 | | | | | $-0.29$ | | |
| | | | 1.4 | | | | | 1.3 | | |
| Outlays | | | | 0.95 | | | | | 1.07 | |
| | | | | 0.8 | | | | | 0.7 | |
| Outlays - OMF | | | | | 0.46 | | | | | 0.49 |
| | | | | | 1.0 | | | | | 1.0 |
| N | 159 | 159 | 152 | 152 | 152 | 159 | 159 | 152 | 152 | 152 |
| $R^2$ | 0.14 | 0.22 | 0.00 | 0.01 | 0.00 | 0.15 | 0.17 | 0.00 | 0.01 | 0.00 |

Table A.3: Greenbook Fiscal Surprise Regressions

There is some evidence that same-quarter tax receipt surprises today can help predict GDP growth surprises in the following quarter, but there is no evidence for the same from expenditures. To reach this conclusion, we estimate a similar set of equations via OLS, with instead

$$Z_{t+1} - \mathbb{E}_{t,FOMC}(Z_{t+1}) = \alpha_2 + \gamma_2(G_t - \mathbb{E}_{t,FOMC}(G_t)) + w_t$$

This specification instead asks whether a same-quarter receipt or expenditure surprises this quarter should shift the Tealbook forecasters' projections for GDP growth in the following quarter. While $R^2$ values are quite modest, and below those of more-closely watched macroeconomic indicators such as housing starts and industrial production, there is some evidence that tax receipt surprises are informative for the following quarter's nominal GDP projections (although not real GDP growth projections). However, we see no such result for government outlays.

| $h = 1QForward$ | NGDP | NGDP | NGDP | NGDP | NGDP | NGDP | RGDP | RGDP | RGDP | RGDP | RGDP | RGDP |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Receipts | $3.74^{***}$ | | $3.54^{**}$ | | $3.52^{**}$ | | 1.99 | | 1.87 | | 1.80 | |
| | 1.4 | | 1.5 | | 1.4 | | 1.4 | | 1.4 | | 1.4 | |
| Outlays | | 1.19 | 0.70 | | | | | 0.70 | 0.44 | | | |
| | | 1.0 | 1.0 | | | | | 1.1 | 1.1 | | | |
| Outlays - OMF | | | | $1.43^{*}$ | $1.24^{*}$ | | | | | $1.23^{*}$ | 1.13 | |
| | | | | 0.7 | 0.7 | | | | | 0.7 | 0.7 | |
| $fund_{n}eed_{t}emp$ | | | | | | 0.24 | | | | | | 0.52 |
| | | | | | | 0.8 | | | | | | 0.7 |
| N | 151 | 151 | 151 | 151 | 151 | 151 | 151 | 151 | 151 | 151 | 151 | 151 |
| $R^2$ | 0.04 | 0.01 | 0.04 | 0.02 | 0.05 | 0.00 | 0.01 | 0.00 | 0.01 | 0.02 | 0.03 | 0.00 |

Table A.4: Fiscal Surprises, Forward Looking

Taking these analyses as evidence that any links between convenient asset demand and the fiscal information content of T-bill surprises is likely to be low, I proceed next to introducing my measure of T-bill surprises and presenting my principal results. Further exploring the likely quantitative implications of the final result of this section, that tax receipts but not expenditures may be informative of next quarter's future GDP growth, is the subject of Section A.6 of this paper. In very brief summary, restricting my shock measure to only those T-bill surprises likely to be least informative about future tax receipts has nearly no effect on any of the quantitative results of this paper.

## A.6 Treasury Information Effects

The quarterly analysis of fiscal surprises in Section **??** left open the possibility of a moderate capacity for same-quarter surprises in government receipts to predict next quarter's macroeconomic fundamentals. I show that any added macroeconomic informativeness of Wrightson T-bill issuance surprises is unlikely to drive any sizable share of my convenience yield impulse response results. I do this by separately estimating impulse responses for T-bill surprises that are *more* and *less* informative about future government tax receipts, and show that the two impulse responses are quantitatively very similar.

### A.6.1 Intuition

The empirical test to follow will assess whether Wrightson issuance surprises that occur in weeks where T-bill issuance is especially *informative* about future tax receipts have meaningfully different impulse response estimates than surprises in tax uninformative weeks. I identify tax informative weeks as weeks in the data where relatively more variation in near-future treasury net cash flows comes from tax receipts, for seasonal reasons. This test determines whether the information content about tax receipts has a separate effect on T-bill convenience yields that differs from information content about expenditures. This is informative about the plausible *total* magnitude of tax receipt information effects because the low-frequency analysis of Section **??** suggests that the macroeconomic informativeness of high frequency government expenditure surprises is close to zero.

The quarterly analysis of section **??** suggests that the ideal experiment for studying T-bill supply shocks with zero demand-shifting information content would study T-bill issuance surprises that are unambiguously interpreted by market participants as arising from government *expenditure* surprises. Whatever modest correlation that exists between same-quarter fiscal surprises and future macroeconomic outcomes appears to be restricted to government receipts and does not extend to expenditures. Thus, a T-bill issuance surprise that market participants understand is driven by expenditures alone is most-likely to be free of any information content that would affect how investors value asset convenience.

That ideal experiment is not empirically feasible, given my data. The US Treasury does not give an accounting of the precise reasons why they issue their chosen quantities of T-bills. While Wrightson's Treasury commentaries sometimes will offer an interpretation of what might have driven a given issuance surprise, those discussions are not sufficiently common to construct a large sample of surprises widely interpreted as arising from expenditures alone. While Wrightson does now publicize their Treasury cash flow projections, they only began doing so in 2017, leaving limited overlap with my 2009-2019 sample.

However, the well-known seasonal pattern in federal government receipts (and, indeed, for many expenditures) makes it feasible to isolate a subsample of T-bill issuance surprises that are substantially *more* informative about the Treasury's private information regarding future tax receipts. Three qualities of the Treasury's T-bill issuance program allow this. First, strong seasonality in tax receipts centered around federal tax deadlines mean that a greater share of the *uncertainty* in Treasury net cash flows comes from tax receipts in the weeks surrounding these deadlines. Second, the Treasury's desire to smooth T-bill issuance sizes over time means that they will tend to respond *today* to private information about near-future cash flows. Third, because greater (or smaller) issuance of 4-Week T-bills directly affects the Treasury's cash balance only over the next four weeks, the Treasury's 4-Week T-bill issuance decision today will tend to be most informative about the Treasury's expectations of cash flows over the next four weeks.

Motivated by these qualities, I define a "tax informative" T-bill issuance surprise as a Wrightson issuance surprise occurring in a week where more of the near-term variation in Treasury cash flows likely comes from government tax receipts. I define "near-term" variation in Treasury cash flows as variation over the following four weeks. I measure which weeks have greater near-term receipt variation by studying the sample variance in cash flows for a given week, but across many years, in the sample.

Figure A.6 illustrates the procedure. In the left panel of Figure A.6 is an example weekly time series of Treasury net tax and net non-tax cash flows, for the year 2016. That is, it plots $(\text{Net Tax Flow})_t$ and $(\text{Net Nontax Flow})_t$ for every week $t$ in the year. Even in a single year of data, the spikes associated with tax deadlines are readily apparent. These spikes typically fall on the 15th of the final month in each quarter.

The right panel of Figure A.6 shows the criteria by which I select "tax informative" weeks in the data. For each of the 52 weeks in a year, I calculate
$\mathbb{V}(\sum_{\ell=0}^{4}(\text{Net Tax Flow})_{w+\ell})$ and $\mathbb{V}(\sum_{\ell=0}^{4}(\text{Net Non-Tax Flow})_{w+\ell})$ in each week $w$, where $\mathbb{V}$ signifies a sample variance of each week $w$ in my sample (i.e. variance within week $w$'s, but across *years* in the data). These sample variances of forward-looking cumulative cash flows are my proxy for near-future cash flow *uncertainty* coming from the two fundamental components of cash flows: receipts and expenditures. The object plotted at each week $w$ in the right panel of Figure A.6 is $\frac{\mathbb{V}(\sum_{\ell=0}^{4}(\text{Net Tax Flow})_{w+\ell})}{\mathbb{V}(\sum_{\ell=0}^{4}(\text{Net Non-Tax Flow})_{w+\ell})}$, the ratio of tax and non-tax variance.

A week is selected as a "tax informative" week if this ratio is greater than the median ratio across all 52 weeks. In both the left and right panels, red dots signify the weeks of the year that are selected as "tax informative" by this metric. The left panel lends the intuition – tax informative weeks are typically the four weeks preceding (and sometimes a single week after) each of the five major spikes in tax receipts. The fact that the selection criteria, which is based off an empirical moment *across* years in the data, selects weeks in this particular year that sensibly correspond to these spikes is an indicator that this tax deadline seasonality is roughly calendar-constant across years in the sample.

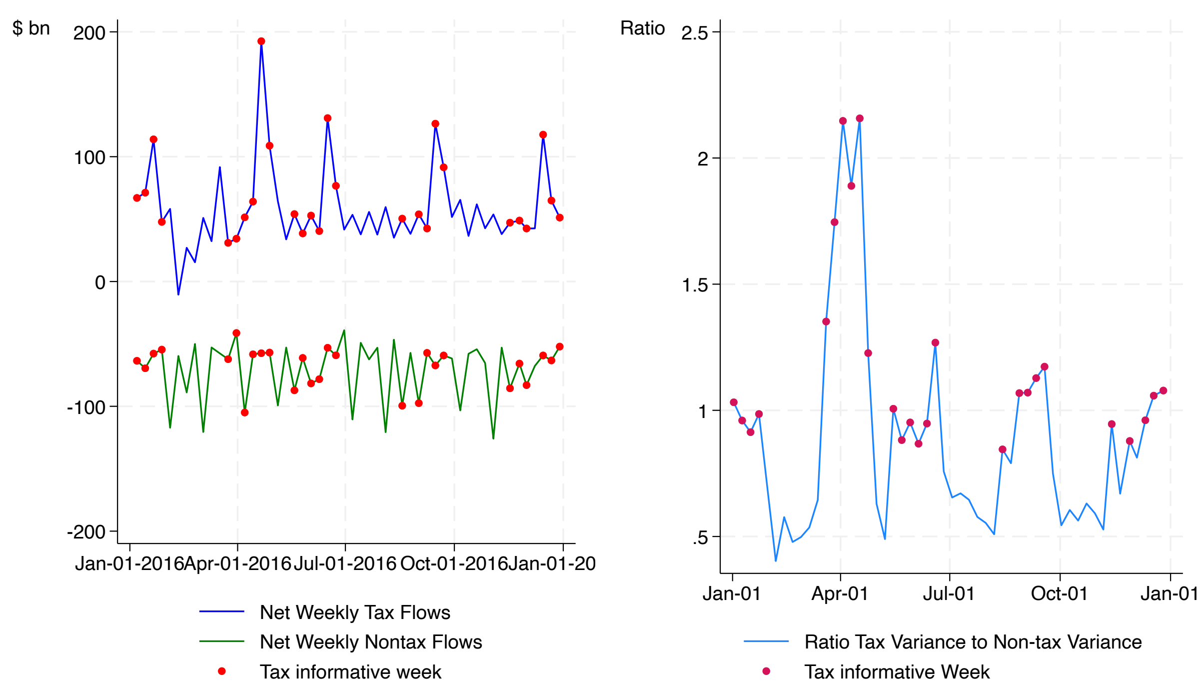


Figure A.6: Tax Informative Weeks

### A.6.2 An Illustrative Model

I construct a simple but realistic Treasury T-bill issuance model, in which the Treasury partially reveals its private information about future cash flows via its public T-bill issuance decisions. I use to model to show that a rational private agent should indeed view T-bill issuance decisions in the weeks before a seasonal tax deadline as more informative about the Treasury's private receipt information.[8]

In the model, the Treasury balances two objectives: to keep its cash balances close to some target level, and to avoid changing T-bill issuance sizes too quickly. The first is consistent with the Treasury's tendency to discuss "target" cash balances for particular dates. The second is consistent with the Treasury's stated objective to keep issuance "regular and predictable". We represent these with the objective function

$$\min_{B_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(c_t^2 + \gamma_1 (B_t - B_{t-1})^2\right) \tag{A.2}$$

$$c_t = c_{t-1} + B_t - B_{t-4} + r_t - e_t - r_{t-H} + e_{t-H} \tag{A.3}$$

Here, $c_t$ represents the deviation of the Treasury's cash balance from some target level. $B_t$ is T-bill issuance this week. In the model, every T-bill issued by the Treasury has a 4-week maturity. As such, $B_{t-4}$ represents the T-bills *maturing* this week, which must be redeemed for investors and thus contribute negatively to this week's cash flows. $r_t$ are today's cash receipts, which contribute positively to the cash balance. $e_t$ are today's expenditure, which contribute negatively.

The objects $r_{t-H}$ and $e_{t-H}$ represent a tendency by the Treasury to eventually finance government surpluses using issuance of unmodeled, longer-term coupon *bond* issuance. In the model, this means that a \$1 net cash flow today will affect T-bill issuances in the near-term, but not the long-term. While a simplifying abstraction, these objects ensure that this model embodies the Treasury's short-term *cash management* problem, which is likely most important at the high frequencies of this paper. The longer-term Treasury *debt management* problem remains unmodeled.

I model the law of motion of the fiscal objects $r_t$ and $e_t$ as

$$r_t = \phi_{w(t)} y_t \tag{A.4}$$

$$y_t = \rho_y y_{t-1} + \varepsilon_t^y \tag{A.5}$$

$$e_t = \varepsilon_t^e \tag{A.6}$$

$$\varepsilon_t^y \sim N(0, \sigma_y^2) \tag{A.7}$$

$$\varepsilon_t^e \sim N(0, \sigma_e^2) \tag{A.8}$$

$$0 = \mathbb{C}(\varepsilon_t^y, \varepsilon_e^e) \tag{A.9}$$

[8]#Note to preliminary reader: This paragraph used to have another sentence in it, to reflect a result that I couldn't get ready in time for this draft. It was meant to read: "Via a simulated time series, I show that estimated convenience yield impulse responses from surprises in high and low tax informative weeks should exhibit different impulse responses, if tax receipt information does indeed move convenience yields in an economically meaningful way."

Receipts in the model are a simple linear function of the persistent state variable $y_t$, which is observable to the Treasury but will be unobservable to private market participants. I will consider the possibility that the state variable $y_t$ is some macroeconomic indicator that is relevant for investors' demand for safe, convenient assets. However, the loading $\phi_{w(t)}$ which maps the macroeconomic state $y_t$ into receipts $r_t$ is *seasonal*, and will depend on the week $t$'s placement in the quarter (i.e. $w(t) \in \{1, 2, \cdots, 13\}$). In my illustrative numerical example, I will have $\phi_i = 0$ for $i \in \{1, 2, \cdots, 12\}$ and $\phi_{13} >> 0$. That is, all of the quarter's tax receipts arrive in the cash account in the last week of each quarter.

Private agents in the model have an informational disadvantage compared to the Treasury. For private agents, the observable state variables in the model include $c_{t-1}$; the history of past receipts $r_{t-\ell}\ \forall \ell \geq 1$; the history of past expenditures $e_{t-\ell}\ \forall \ell \geq 1$; current T-bill issuance $B_t$; and all past T-bill issuances $B_{t-\ell}\ \forall \ell \geq 1$. However, the private agents *cannot* observe the state variable $y_t$, current receipts $r_t$ or current expenditures $e_t$.

The private agents in the model know the parameters of the Treasury's objective function, meaning that they know the mapping between state variables (both observable and unobservable) and the Treasury's choice variable $B_t$. Because the Treasury's problem has the well-known, linear quadratic setup, the Treasury's T-bill issuance policy function is a linear function of the state variables observable to the Treasury. As such, the information structure of the model admits a linear state space representation with Gaussian shocks.

Like other linear state space models with normally distributed shocks, private agents' *updates* to their expected values of each state variable can be solved for iteratively via the Kalman filter. A departure of this setup from the simplest, canonical Kalman filter application is that uncertainty in the model is *seasonal*, driven by the seasonal loading of receipts on the unknown (to private agents) state variable $y_t$. In this setting, the Kalman gain that maps surprises in observables into updated beliefs about next-period's state variables are season (i.e. week-of-quarter) specific.

For any choice of the four model parameters $\Gamma_0$, $\rho_y$, $\sigma_e^2$, and $\sigma_y^2$ it is straightforward to solve numerically for the Kalman gain matrix for each week of the quarter. At an illustrative choice of model parameters listed in the Appendix, I solve for this Kalman gain for each of the 13 weeks in a quarter. The top panel of Figure A.7 shows $\frac{d\mathbb{E}_t(y_t)}{d(B_t - \mathbb{E}_{t-1}B_t)}$ – that is, how the rational private agent of the model will update their beliefs about the persistent, unobserved, underlying macroeconomic state variable $y_t$ in response to a surprise in T-bill issuance.

The pattern of macroeconomic informativeness of T-bill issuance surprises from the model follows essentially the same pattern as my empirical implementation of the "tax informativeness" classifications. That is – the private rational agent will tend to view T-bill issuance surprises in the last four weeks before each major tax receipt inflow date (week 13 of the quarter, in the model) as particularly informative about future receipts, and thus macroeconomic fundamentals.

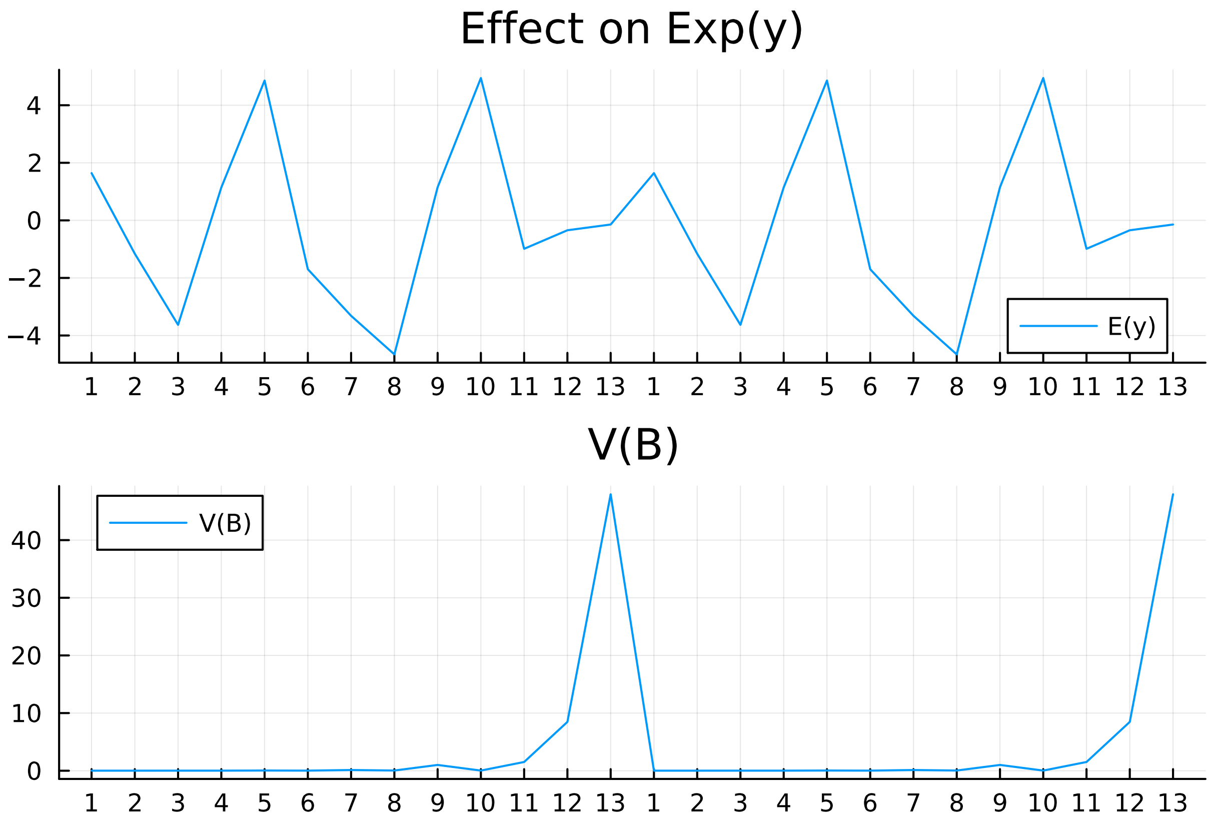


Figure A.7: Seasonal Kalman Gain and T-bill Issuance Uncertainty from Model

We can also use the illustrative model to understand the implications of a hypothetical correlation between expected macroeconomic fundamentals and convenient asset demand. Recalling our notation from Section 2.2 that $\xi_t$ represents factors *shifting* the convenient asset demand curve, we say that

$$\xi_t = \xi_t^e + \omega \mathbb{E}_t(y_t) \tag{A.10}$$
$$\mathbb{C}(\xi_t^e, \mathbb{E}_t(y_t)) = 0 \tag{A.11}$$

That is, we consider the possibility that $\mathbb{C}(B_t - \mathbb{E}_{t-1}B_t, \xi_t) \neq 0$, via the mechanism that T-bill issuance surprises shift private agent understandings of macroeconomic fundamentals in a way that is relevant for demand.

### A.6.3 Empirical Test Results

The reduced form impulse response function estimates show that T-bill issuance surprises in "tax informative" periods have substantially more predictive power for future tax receipts than surprises outside the "tax informative" weeks, consistent with my story. The estimated impulse response function for T-bill convenience yields to shocks in the two periods look exceedingly similar. This is inconsistent with large tax receipt information effects driving a quantitatively meaningful share of my results.

To find suggestive evidence for this test's underlying hypothesis that T-bill issuance surprises in "tax informative" weeks have more predictive value for future Treasury net tax flows, I use a local projection approach similar to that from my earlier results. Specifically, at each horizon $h$ I estimate

$$\begin{aligned}\sum_{\ell=0}^{h} r_{t+\ell} = \alpha_h &+ \beta_h^I \mathbf{1}\left(\text{week}(t) \in \text{Informative}\right) \times \varepsilon_t^s \\ &+ \beta_h^{NI} \mathbf{1}\left(\text{week}(t) \in \text{Not Informative}\right) \times \varepsilon_t^s \\ &+ \phi_h' X_{t-\delta} + e_{t+h} \end{aligned} \tag{A.12}$$

Where $r_t$ show net Treasury tax *flows* in week $t$. This means that $\sum_{\ell=0}^{h} r_{t+\ell}$ measures *cumulative* Treasury tax flows between week $t$ and week $t+h$. In the results to come, I plot both $\{\beta_h^I\}_{h=0}^H$ and $\{\beta_h^{NI}\}_{h=0}^H$. These estimates show how T-bill issuance surprises shift expectations of cumulative tax flows differently, depending on the seasonal informativeness of week $t$ in which the surprise was realized.

T-bill issuance surprises delivered in one of my "tax informative" weeks do indeed predict future tax receipts much more than T-bill surprises in "tax uninformative" weeks. These results are shown in Figure A.8. The estimates suggest that a \$10 billion T-bill issuance surprise in an informative week should shift expectations of cumulative tax receipts over the next three weeks by -\$2.5 billion. In contrast, a a \$10 billion T-bill issuance surprise in an uninformative week should shift expectations of cumulative tax receipts over the next three weeks by only -\$0.5 billion.[9]

[9] Admittedly, confidence intervals are sufficiently wide that I cannot reject the null hypothesis that the 3-week impulse response is equal across the two subsamples.

Figure A.8: Receipt Informativeness of Issuance Surprises, Informative-vs-Uninformative Weeks

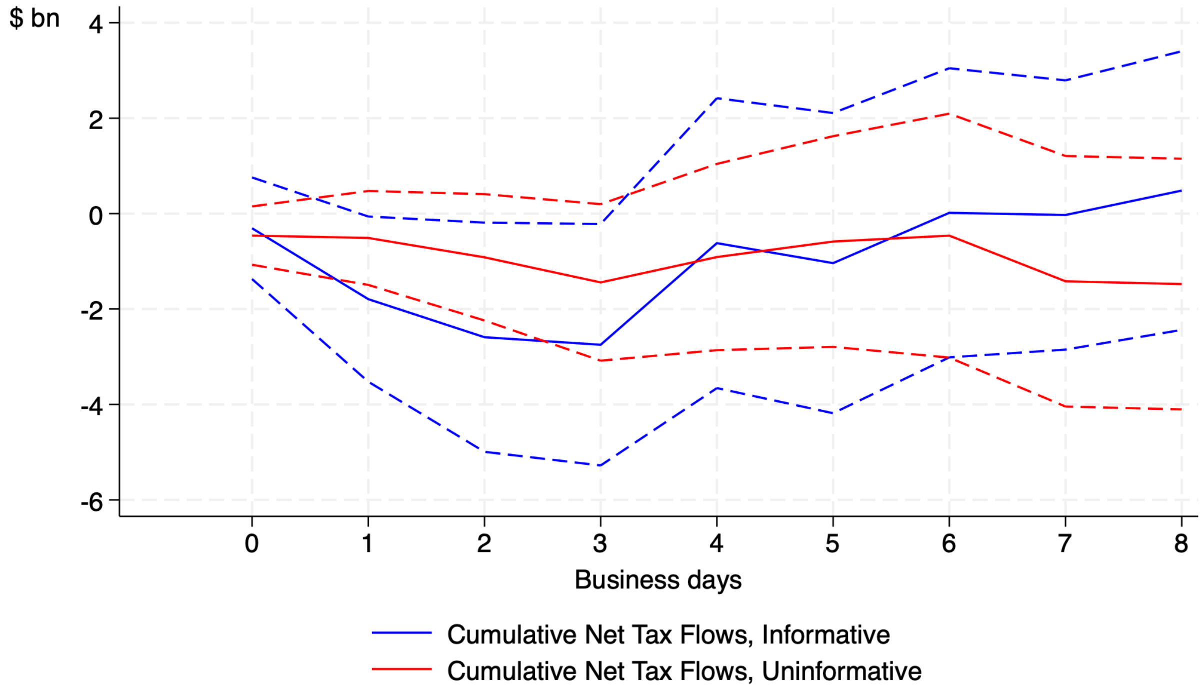


To estimate different impulse responses for T-bill convenience yields in these two tax informativeness periods, I again estimate impulse response functions that allow for separate parameter estimates for shocks realized in the two periods. To do this, I estimate an alternate version of equation A.12, with the lefthand side variable replaced with future realizations of the T-bill convenience yield.

Considering that these two response are estimated with nonoverlapping subsamples of shocks, the impulse response of T-bill convenienece yields across informativeness periods looks remarkably similar. These results are shown in Figure A.9, where the T-bill convenience yield response to issuance surprises in a "tax informative" week are shown in red and the response to issuance surprises in a "tax uninformative" week are shown in blue. With the possible exception of an approximately weeklong period in the 4-5 week horizon, these responses look quite similar.

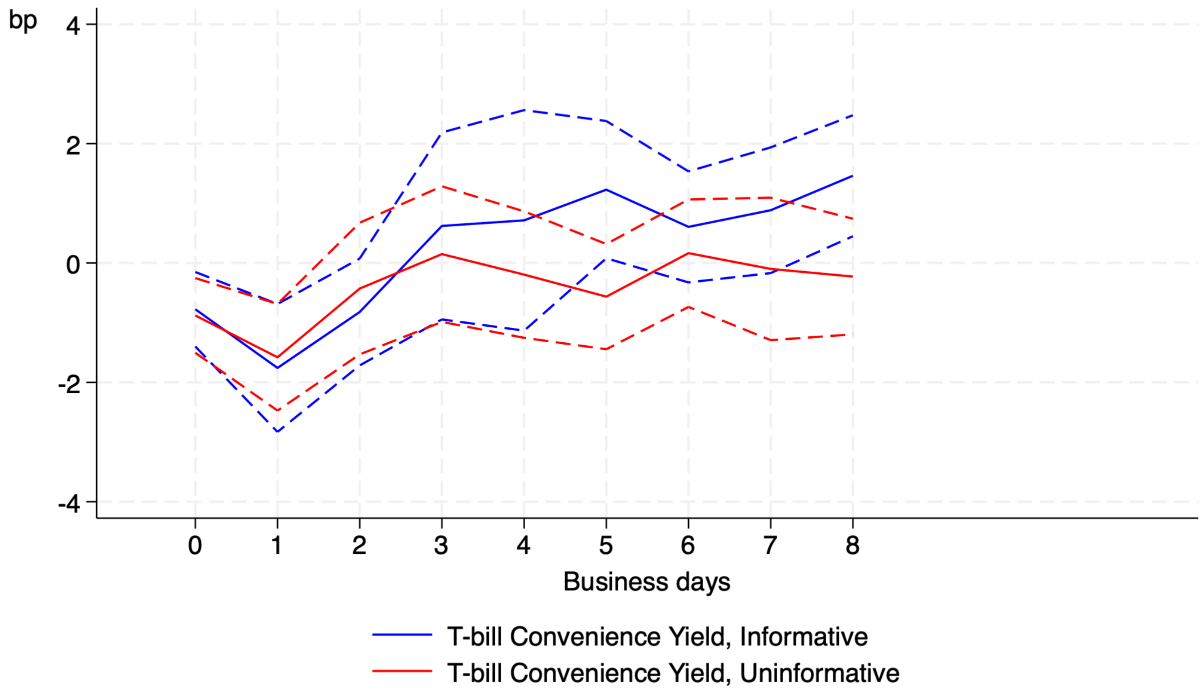


Figure A.9: T-bill Convenence Yield Response By Subsample, Tax Informative-vs-Uninformative

Most importantly, these responses look *incredibly* similar at the earliest horizons, in the first few days after an issuance surprise is realized. This horizon is especially important, because the endogeneity mechanism that I am considering relies on the *information effect* of the surprise itself. If there is indeed any such meaningful information effect, we would expect to see sizable differences in the convenience yield effect near the time when that information is revealed to market participants. We surely do not see any substantial difference in convenience yield responses at these shorter horizons.

The estimates of $\beta$, $\phi$, and $\mu^{-1}$ from Section 1.7.4 also effectively summarize the shape of the convenience yield impulse response for both of these subperiods. In Figure A.9, the blue dots show the fitted convenience yield impulse response, using the estimates of $\beta$, $\phi$, and $\mu^{-1}$ from Section 1.7.4 (i.e. the pooled sample estimates) and the estimates (not shown) of future T-bill supply to a T-bill issuance shock in a "tax informative" period. The red dots do the same, for the tax uninformative period. These dots are especially useful for understanding some differences in the impulse responses at the longer horizons of 30-45 business days, where the "informative" response lies consistently above the "uninformative" response. The dots show us that this pattern can be explained via differences in the impulse response of T-bill *supply* (quantity) at those horizons, given my estimates of $\beta$ and the sluggishness parameters $\phi$ and $\mu^{-1}$. The fact that the blue dots are above the red dots at those horizons reflects the fact that T-bill supplies are typically still meaningfully *falling* at those horizons after an "informative" shock, but not after an "uninformative" shock.[10]

[10]In other words – the different convenience yield responses at those horizons can be explained by observed

## A.7 LP-IV Results: Future Cash Flows

On average, T-bill supply surprises today predict offsetting, non-debt net cash flows into the US Treasury's cash account in the coming weeks. This is consistent with my overall story that these surprises are driven by differences in Wrightson's and the Treasury's near-future cash flow projections.

To show that T-bill supply surprises today appear to come with 1-to-1 offsetting net Treasury cash flows in the near future, I now estimate equations (1.9) and (1.10), setting $Y_{t+h} = \sum_{\ell=0}^{h} \left(\text{Non-debt Receipts}_{t+\ell} - \text{Non-debt Outflows}_{t+\ell}\right)$. "Non-debt receipts" and "Non-debt Outflows" refer to the observable inflows and outflows into the US Treasury's General Account at the Federal Reserve. Data on inflows and outflows is publicly available, published daily at a one-day lag.

If it is indeed the case that T-bill supply surprises today are generally driven by the near-term cash needs of the federal government, then we would expect our impulse response estimates $\{\beta_{t+h}\}_{h=0}^{H}$ in this specification to be *negative*. When the Treasury surprises market participants by issuing \$10 billion more in T-bills today, that means it is raising approximately \$10 billion more today via debt issuance. If the Treasury does this to fully offset cash flows in the coming weeks, then we might expected \$10 billion more to leave the US Treasury's general account than was originally expected (i.e. negative net flows) for non-debt related reasons in the coming weeks.

Future predicted Treasury cash flows almost perfectly offset the predicted changes in near-term T-bill supply from a surprise T-bill issuance today. These results are shown in Figure A.10 below. In the left panel, I plot future cumulative non-debt cash flows, along with the T-bill supply impulse response function in gray. The units on these two results are the same. The sign of near-term cumulative cash flows is negative, as expected. To facilitate comparing the magnitude of this response to the magnitude of the T-bill response, the right panel *negates* the cumulative cash flow response from the left panel. The fact that the T-bill supply response and the cumulative cash flow response lie nearly atop one another in the right panel suggests a simple, and empirically justified story of the drivers of T-bill issuance shocks: the Treasury surprises the market with increased T-bill issuance to offset future net cash flow surprises essentially one-for-one.

*supply* variation. It does not require a correlation with unobserved *demand* at those horizons.

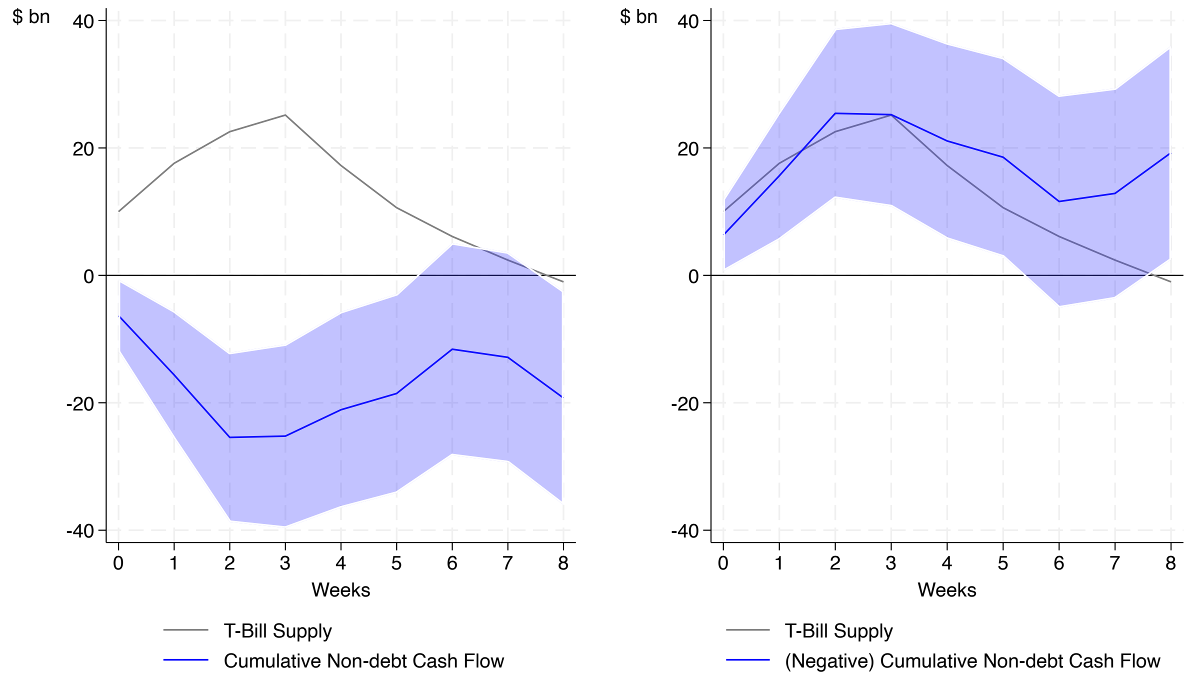


Figure A.10: Impulse Response of Future Treasury Net Cash Flows

## A.8 Understanding Temporary Effects

A crucial ingredient of most limits to arbitrage models that allow for temporary effects after a supply shock is that the universe of investors who are able to absorb an increase in supply widens over time after a shock.[11] I show that holdings data from the money market mutual fund sector alone shows this quality. Treasury-only money market mutual funds appear to absorb a larger share of *surprise* 4-Week T-bill issuance in the week after a Wrightson issuance surprise, compared with surprises more than one week old.

One model with limits to arbitrage that can produce transitory price impacts after a supply shock that are smaller than the persistent component is the slow moving capital model of Duffie (2010). In that model, a sector of fast-moving investors rebalances their portfolio each period, while additional cohorts of slower-moving investors rebalance only periodically. There, an increase in the supply of some security is absorbed disproportionately by the fast investor sector in the initial periods after a shock, until successive cohorts of slow investors arrive to trade, responding to then-higher expected returns on the security since their last rebalance.

In this sector, I show suggestive empirical evidence that the distribution of increased holdings after a T-bill issuance surprise soon after the surprise differs from the distribution more than one week after the shock. Thinking of slow-moving capital such as Duffie (2010),

[11]#Note to preliminary reader: This subsection is in the roughest state of all subsections of the paper, because it was written most recently. Apologies for things like sloppy notation and clunky equations.

this suggests that the investor sector with disproportionately *large* holdings after a T-bill supply shock serve as the fast investors in this setting.

To conduct this analysis, I use CUSIP-level end of month holdings data on money market mutual funds, reported to the SEC via the N-MFP reporting form, and made publicly available since 2012. Using end of month holdings data with my high frequency, weekly T-bill issuance surprise measurement has the disadvantage of fewer observations and some awkwardness in reconciling data sources of different frequency. However, the granular, CUSIP-level data in the N-MFP makes it possible to track individual money market funds' holdings of the *exact* T-bill that experiences surprise issuance. This is critical for my application, because a large portion of money market funds' T-bill holdings are in longer-term bills (such as the 6-month of 1-year maturity) that only very rarely experience issuance surprises.[12].

My empirical strategy in this section is to compare the money fund holdings effects of 4-Week T-bill issuance surprises from the very-recent past to those from the more-distant past. Because my money fund holdings data is at the end-of-month frequency, this will be an inherently monthly analysis. To manage this, I construct two different, monthly-frequency extensions of the weekly Wrightson 4-Week T-bill surprises. The first, Last Week 4W Surprise$_m$, is the end-of-month realization of the weekly T-bill surprise instrument used above. This is my proxy for more-recent T-bill supply shocks, relative to the end-of-month timing for holdings reports in the N-MFP form. The second, Monthly 4W T-Bill Surprise$_m$ is a measure of the *total* surprise in 4-Week T-bill issuance this month, relative to Wrightson's start-of-month projections. This construction leverages the fact that each week of Wrightson projections offers projections over the next 7-8 weeks of issuance, permitting construction of a sensible "monthly" issuance projection as of the information set of the beginning of the month. To formalize, I construct

$$
\begin{aligned}
&\text{Monthly 4W T-bill Surprise}_m = \\
&\text{First Week of Month, 4W Bill Issuance}_m \\
&- \mathbb{E}_{\text{Wrightson, m-1}}\left(\text{First Week of Month, 4W Bill Issuance}_m\right) \\
&+ \dots \\
&+ \text{Last Week of Month, 4W Bill Issuance}_m \\
&- \mathbb{E}_{\text{Wrightson, m-1}}\left(\text{Last Week of Month, 4W Bill Issuance}_m\right)
\end{aligned}
$$

I then estimate two different OLS regression results, using these measures. I estimate

$$
\begin{aligned}
&\text{Original-Maturity 4-Week T-bill Holdings}_{s,m} = \\
&\alpha_{s,1} + \rho_{s,1}\text{Original-Maturity 4-Week T-bill Holdings}_{s,m-1} + \gamma_{s,1}' X_{s,m-1} \\
&+ \beta_{s,1}\text{Last Week 4w Surprise}_m + b_{s,1}\text{Expected Last Week 4w Issuance}_m + e_{s,m}
\end{aligned}
$$

and

[12]See the forecasting exercises of Section 1.6.3

$$
\begin{aligned}
&\text{Original-Maturity 4-Week T-bill Holdings}_{s,m} = \\
&\alpha_{s,2} + \rho_{s,2}\text{Original-Maturity 4-Week T-bill Holdings}_{s,m-1} + \gamma'_{s,2}X_{s,m-1} \\
&+ \beta_{s,2}\text{Monthly 4W T-bill Surprise}_m \\
&+ b_{s,2}\text{Start of month expected 4w T-bill issuance}_m + e_{s,m}
\end{aligned}
$$

for money market fund subsector $s$ and month $m$. I aggregate the fund-level N-MFP holdings into sectors $s$ of Prime, Treasury-only, and General Government funds. General government funds are able to hold short-term securities issued by the US Treasury or by US government agencies such as the FHLB, Fannie Mae, and Freddie Mac. They are also able to hold repurchase agreements that are collateralized by US Treasury or agency securities. I classify "Treasury only" funds as money market funds that hold, on average, more than 80% of their holdings in US Treasuries each month.[13] Prime money market funds can hold any security that a General Government fund can hold, as well as a host of private sector, highly-rated short term debt, such as commercial paper and certificates of deposit.[14]

Table A.5 shows results for $\beta_{s,1}$ and $\beta_{s,2}$. The results suggest that the distribution of increased holdings after a very-recent T-bill issuance surprise is very different from the lower-frequency, monthly surprise measure. When looking only at 4-Week T-bill issuance surprises of the past week, a \$10 billion higher surprise T-bill issuance predicts a \$6.9 billion increase in 4-Week T-bill holdings of Treasury-only money market funds. This suggests that Treasury-only funds provide almost 70% of the total market elasticity, in the first week after a supply surprise. When looking at the monthly surprise measure, the numbers look very different. Treasury-only funds' holdings can account for 13% of the surprise increased 4-Week T-bill supply, at the monthly frequency. General government-only funds account for 23% of the surprise increase. General government-only funds tend to hold similar magnitudes of T-bills as Treasury-only funds, but hold far greater total assets overall. Prime funds account for a smaller, but nontrivial, 7% of increase holdings.

[13]Treasury only money market funds exist to cater to investors with mandates to only invest in US Treasury securities. Often, those investor mandates will have some restriction against investing in repurchase agreements collateralized in US Treasury securities.

[14]The lefthand side variable of Original-Maturity 4-Week T-bill Holdings$_{s,m}$ adds up the sector's holdings in all T-bills that were issued as 4-Week T-bills in that month. That is, it is the sum of the sector's holdings in the then-current 4-Week bill issued that week; the then-current 3-Week bill that was issued as a 4-week bill in the prior week; the then-current 2-week bill that was issued as a 4-Week bill two weeks prior; and the then-current 1-week bill that was issued as a 4-Week bill three weeks prior.

| | Treas | Treas | Gov | Gov | Prime | Prime |
|---|---|---|---|---|---|---|
| Last Week 4w Surpise | 0.69** | | 0.03 | | −0.17 | |
| | 0.32 | | 0.42 | | 0.17 | |
| Monthly 4w Surprise | | 0.13** | | 0.23*** | | 0.07** |
| | | 0.06 | | 0.05 | | 0.03 |
| $R^2$ | 0.43 | 0.44 | 0.77 | 0.80 | 0.65 | 0.66 |
| N | 108 | 109 | 108 | 109 | 108 | 109 |

Table A.5: Money Market Fund Holdings Results

Importantly, Table A.5 suggests that it is indeed the case that the universe of investors absorbing a supply shock varies in the *horizon* since the shock. The results even suggest an identity for the fast-moving investors described in Duffie (2010), as Treasury-only money market funds.

## A.9 Robustness of Core Results

In this Appendix, I present LP-IV estimates for T-bill supply and convenience yields to a Wrightson surprise, under several different specification changes from the baseline. In general, the results presented in the main text are quite robust. The one notable exception is that the impulse response of future T-bill supply to a Wrightson surprise is sensitive to including Wrightson's pre-surprise *projections* of future supplies as controls in the regression.

In Figure A.11 and Figure A.12, four lagged realizations of Wrightson surprises are included as regression controls. This makes very little difference in any of the results

Figure A.11: LP-IV Impulse Response Estimates, Controlling for Lagged Surprises: Convenience Yield

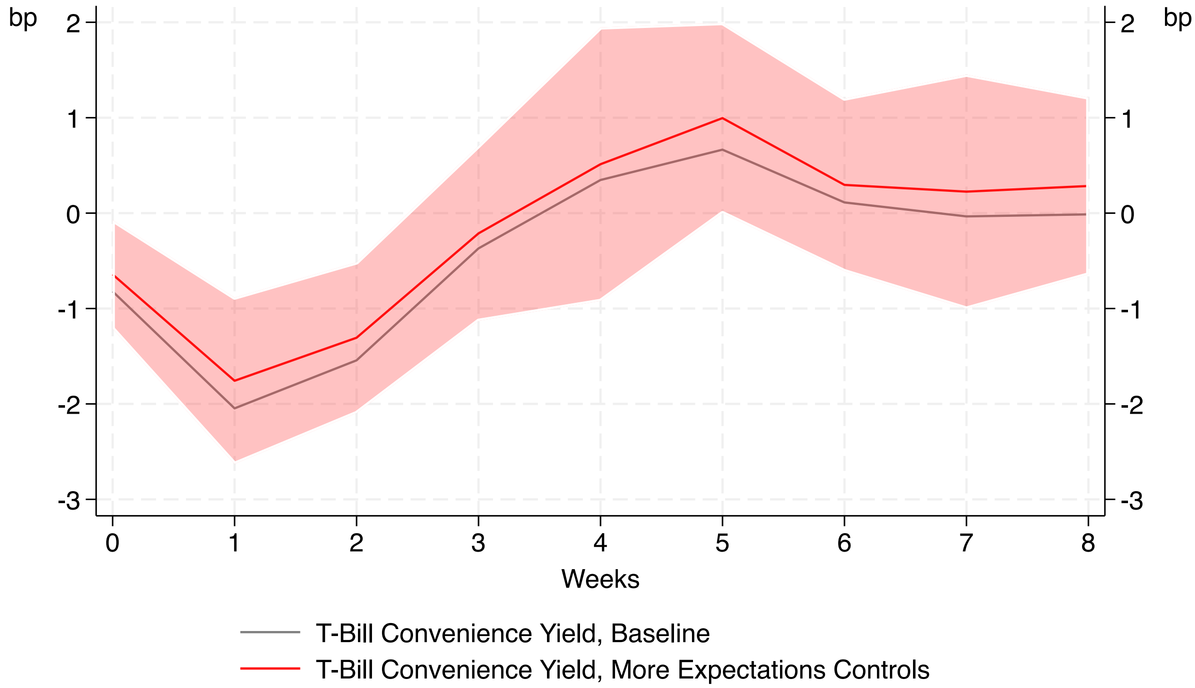

Figure A.12: LP-IV Impulse Response Estimates, Controlling for Lagged Surprises: T-bill Supply

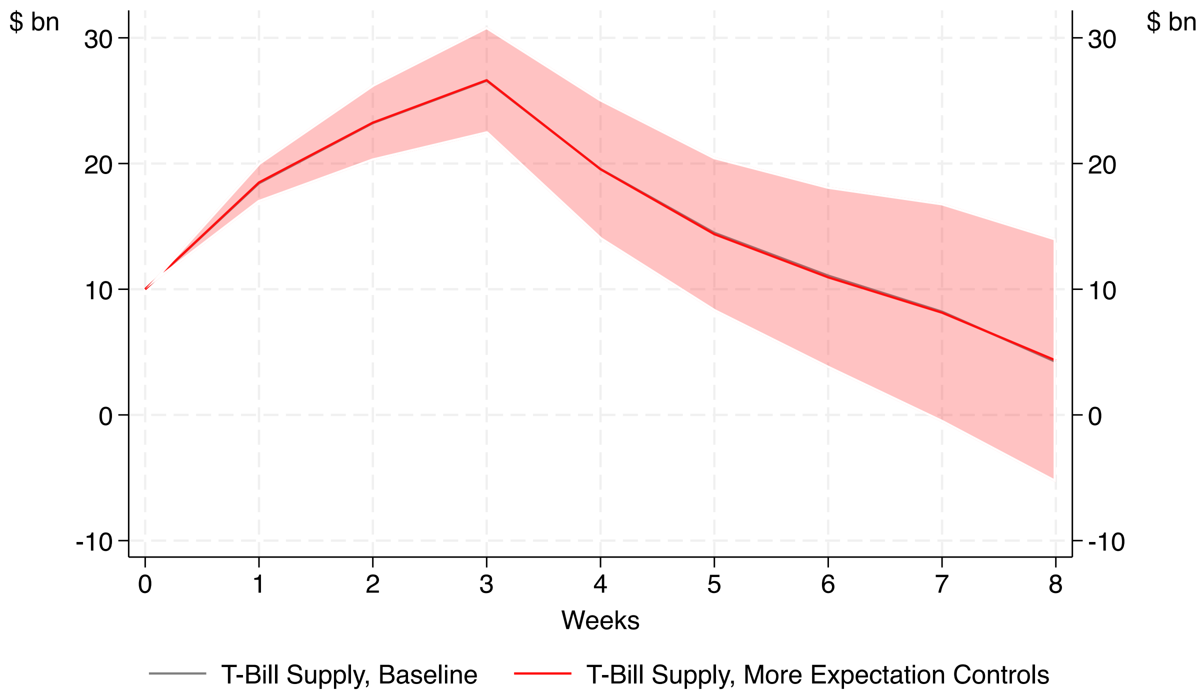


In Figure **??** and Figure A.13, Wrighton's pre-surprise projections of future T-bill supply are *removed* as regression controls. Pre-surprise projections are included as controls in the baseline. While the impact of this removal on the convenience yield response is modest, its effect on the response of future T-bill supply is more substantial. Without pre-surprise projections as controls, Wrightson surprises appear to predict more persistent changes in T-bill supply.

It is worth noting that a more persistent estimate of T-bill supply and a largely unchanged convenience yield response will tend to *strengthen* the qualitative results of this paper that flow effects in this market are substantial but stock effects are modest. In other words, if Wrightson surprises predict a permanent change in T-bill supply, it is all the more puzzling (without flow effects) that the point estimate for convenience yield responses returns to zero so quickly.

Figure A.13: LP-IV Impulse Response Estimates, No Forward Projection Controls: T-bill Supply

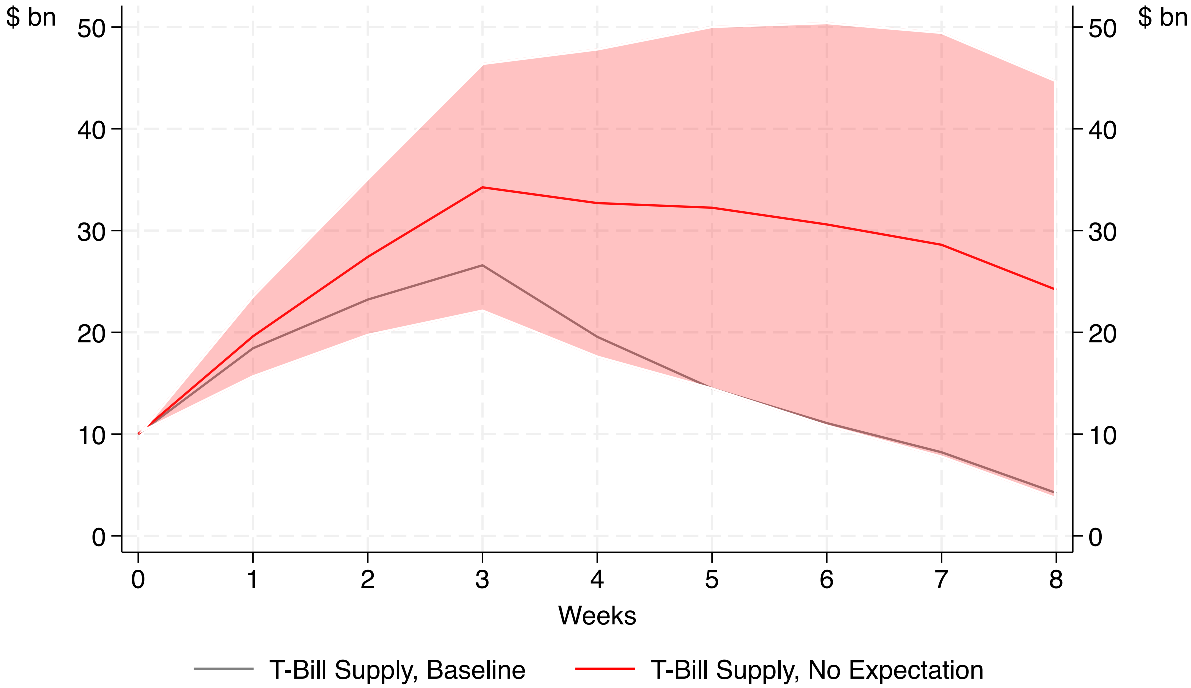

Figure A.14: LP-IV Impulse Response Estimates, No Forward Projection Controls: Convenience Yield

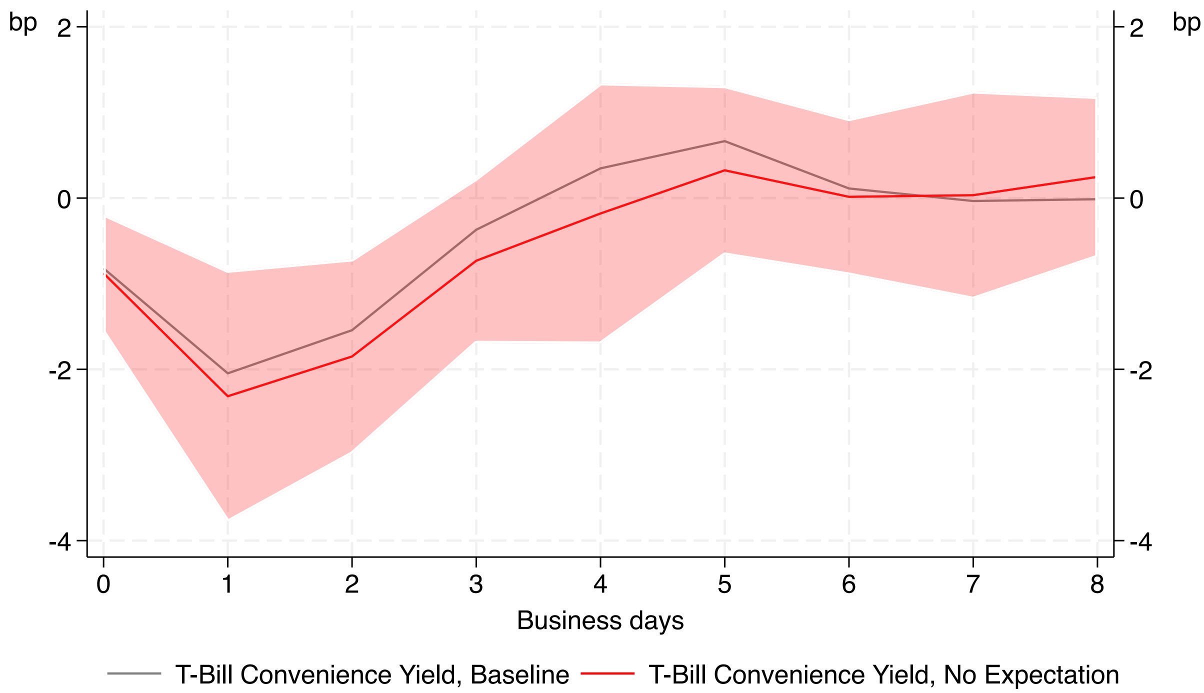


## A.10 Alternate Depictions of Convenience Yield Response

Here I present estimates of some alternate ways to represent the T-bill convenience yield response to a Wrightson surprise. In the baseline, I present the response of then-current on-the-run 4-Week T-bill convenience yields to a Wrightson surpsrise. This means that each point measures the convenience yield of a T-bill at a fixed maturity (four weeks). However, each point measures the convenience yield of a *different* T-bill (i.e. CUSIP), because a different bill becomes on-the-run in each week.

In Figure A.15, I instead track constant CUSIPs over time in the impulse response estimates. In each panel "Then-Current 4W" is the baselin result, described above. "Week 1 OTR" tracks the convenience yield on the T-bill that becomes the on-the-run 4-Week bill in week 1 of the impulse response. The week 0 response of the "Week 1 OTR" estimate measure the convenience yield of a 5-Week maturity bill. The week 1 response of the "Week 1 OTR" estimate measures the convenience yield of a 4-week maturity bill (as such, this response is very similar to the baseline result). The week 2 response for "Week 1 OTR" measures 3-week maturity bill, and so on. Each "Week X OTR" response is constructed analagously. Note that each "Week X OTR" bill *matures* at a different point in the impulse response horizon – thus disappearing from the estimates.

The results look fairly similar across all of these specifications. In one sense, this is reassuring, and shows that several different notions of constructing the short-term T-bill convenience yield produce very similar results. In another sense, however, this similarity itself produces a result of potentially independent interest : supply shocks in short-term convenience assets have short-run effects on convenience yields that are *not* limited to the convenience yield of the most-liquid (i.e. possibly most convenient), on-the-run bill.

Figure A.15: LP-IV Impulse Response Estimates, Tracking Individual CUSIPs

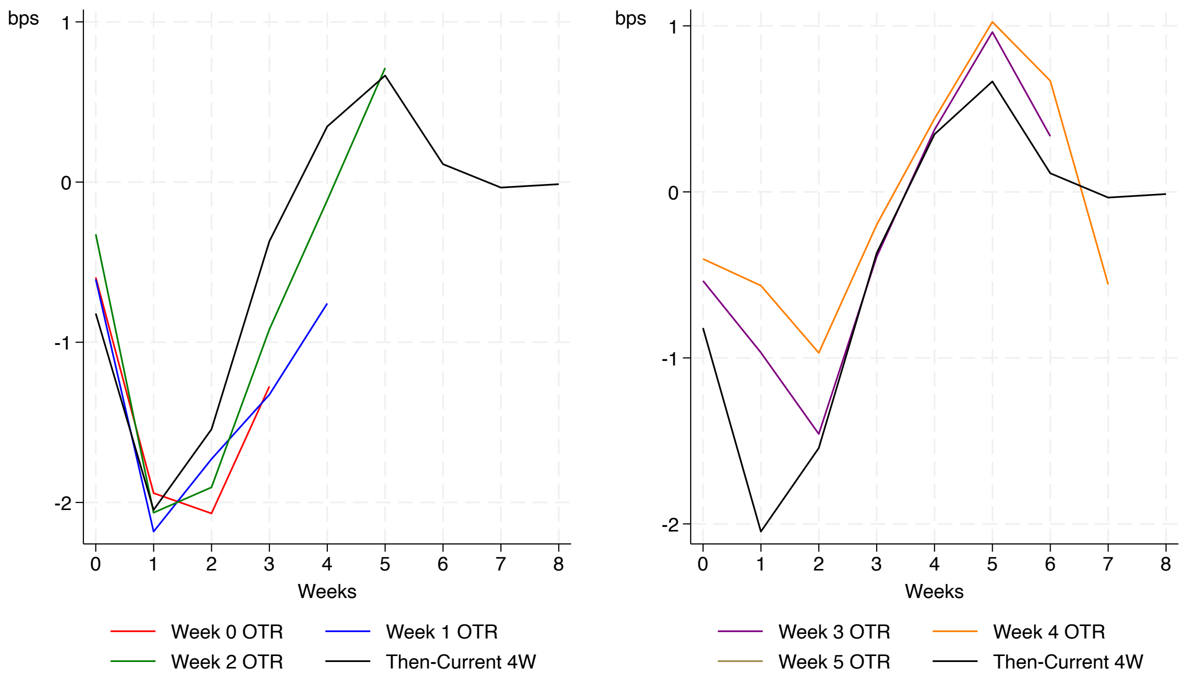


Figure A.16 differs from the baseline in another way. Here, I again fix the *maturity* of the bills being studied at each point in the impulse response horizon. However, I show results when I differ the maturity of that fixed point. These results indicate that the convenience yield response shown in the baseline appears strongest at maturities of six weeks or less. Convenience yield responses at the 8-week and 7-week maturities are more modest. And results at the 13-week horizon are more modest still.

Figure A.16: LP-IV Impulse Response Estimates, Across the Term Structure

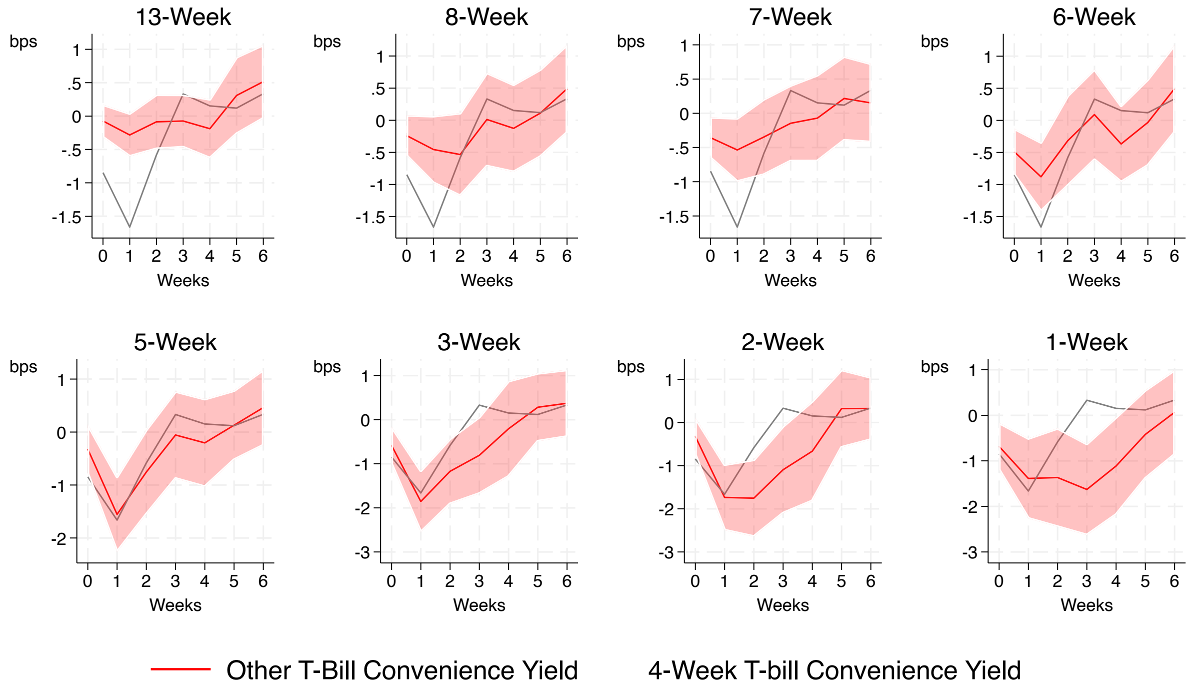

# Appendix B

# Appendix to Chapter 2

## B.1 Additional Robustness Exercises

What follows is a continuation of the robustness exercises in Section 2.5, showing how our central Network Vulnerability Index (NVI) measure changes with different empirical decisions.

### Financial Subsectors Included

Figure B.1 shows how the NVI changes with the addition of each new financial subsector using the approximation procedure outlined in Section 2.3.4. The mostbasic version of our measure, which only uses data from the FR-Y9C form (that is, only large Bank Holding Companies) serves as a base sample, to which other subsectors are individually added through the equation 2.4 approximation. As FR-Y9C data is used for us to find the maximum liability connectivity $\beta^+$ for the sample, those firms must be included in each of the configurations. Our 'other' category is the only sector whose addition into the NVI alters the measure in any substantial way.

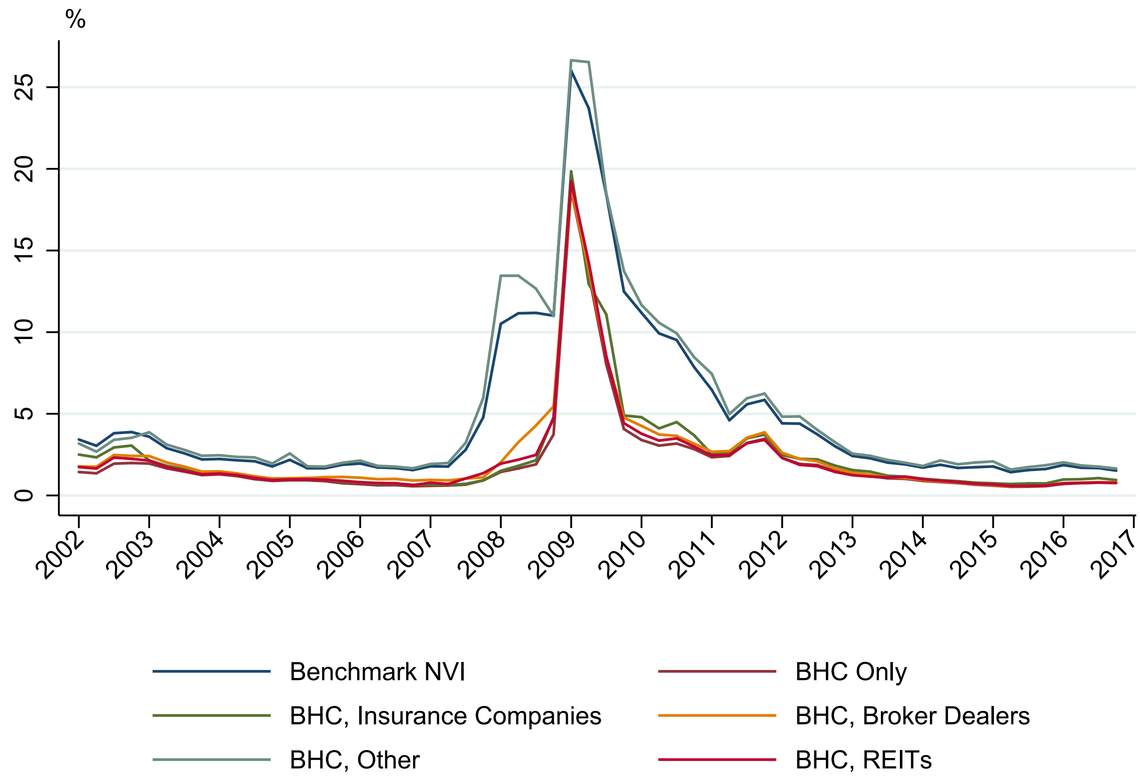


Figure B.1: **Network Vulnerability Index in Different Subsamples.** This figure shows the NVI where only FR-Y9C firms are included in the network, and then a series of other configurations with other financial subsectors added to the network. The addition of the 'other' category, with its moderately large quantity of assets and very high probabilities of default, has the most impact on the NVI.

## Moody's EDF Version

As Nazeran and Dwyer (2015) describe in some detail, there were several notable changes made to the Moody's EDF methodology between versions 8 and 9 of the data. A non-exhaustive list include changes to: the maxmium allowable EDF for financial firms, the assumed informational value of financial firms' balance sheets, and a large increase in financial firm defaults with which to inform estimation of the final empirical fitting of the model. While we are convinced that these changes improve the EDF measure's applicability for our purposes, they do mean that the EDFs can look notably different depending on which version is being used.

Figure B.2 shows the benchmark NVI (using EDF version 9), as well as an NVI series computed identically save for a switch from EDF version 9 to EDF version 8. When compared to an NVI calculated with version 8, the benchmark NVI increases earlier and more dramatically in the crisis, peaks somewhat higher in 2009, but then is lower from the end of the financial crisis until 2014. Figure B.3 shows a subsector breakdown of how average default probabilities change between the new and old data versions. BHCs, in particular, have very different EDF magnitudes in the peak of the crisis, and most other subsectors show some differences in the timing and duration of crisis EDFs.

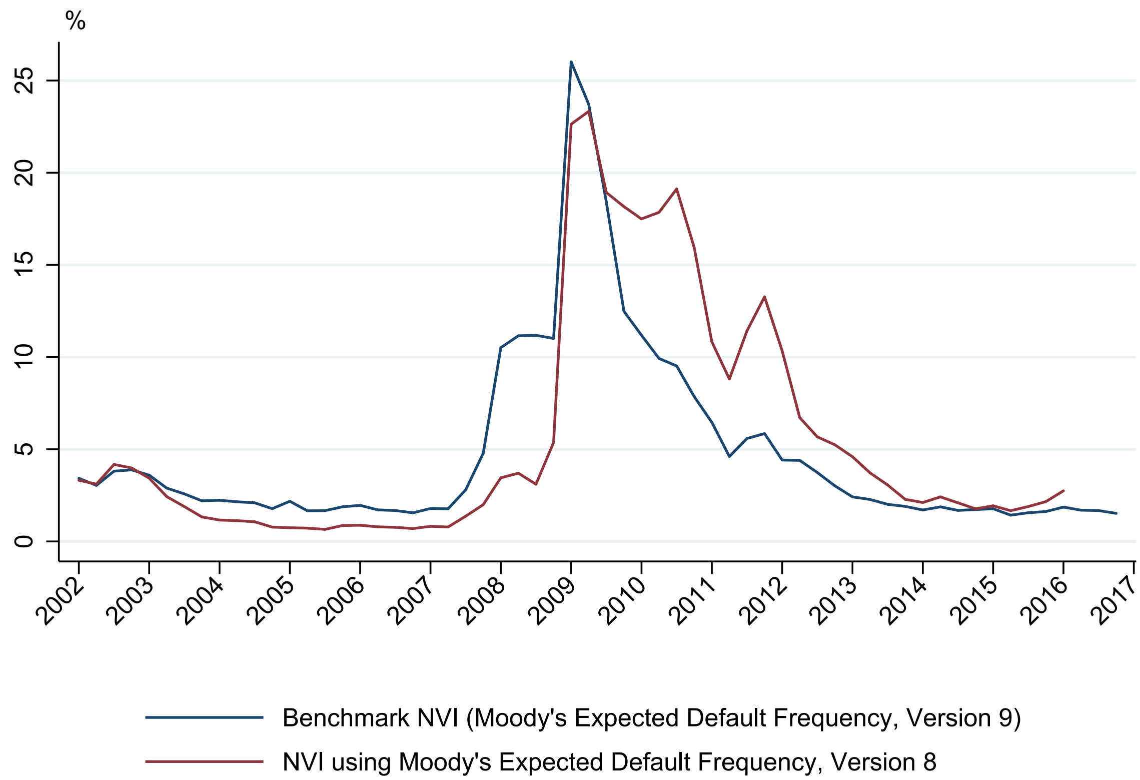


Figure B.2: **Network Vulnerability Index with Different Versions of Moody's Expected Default Frequency.** The changes that Moody's Analytics implemented to their Expected Default Frequency (EDF) series between versions 8 and 9 have a material effect on our spillover measure. Particularly, under version 9 the measure rises earlier leading to the 2008 financial crisis, reaches higher magnitudes, then drops more rapidly after the most-severe parts of the crisis have passed.

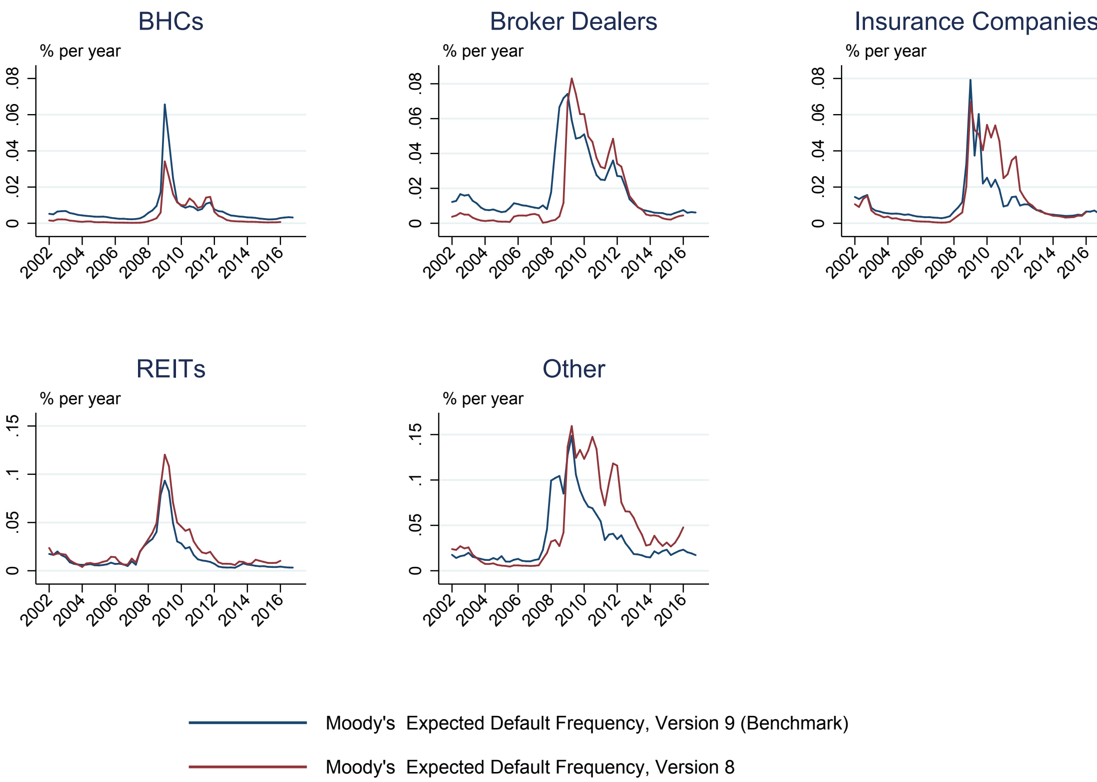


Figure B.3: **Sector-Wide Asset-Weighted Default Probabilities with Different Versions of Moody's Expected Default Frequency.** Different versions of Moody's Analytics' Expected Default Frequency series suggest somewhat different default probability dynamics around the Financial Crisis. For certain types of firms, the new version gives much higher probabilities in the peak of the crisis. For other firm types, general magnitudes remain similar, but the timing and duration of high EDF spells change.

## Balanced vs Unbalanced FR-Y9C Panel

Throughout 2002-2016, a number of firms enter and exit our FR-Y9C sample. Notable changes include the additions of Morgan Stanley and Goldman Sachs during the financial crisis, the departure of Metlife from the sample in 2012, and the temporary inclusion of American International Group in 2013 and 2014. We check whether the path or magnitudes of our NVI change when we restrict ourselves to a balanced panel of firms for the portions of our measure relying on the FR-Y9C data.

As Figure B.4 illustrates, different balanced panel treatments have noticeable but relatively modest effects on the measure. In one Figure B.4 configuration, we entirely drop any FR-Y9C balance sheet information unless the firm has filed the FR-Y9C for the entirety of our sample period. In the second alternate setup, we use all available FR-Y9C balance sheet information (as in the benchmark case), but we restrict our selection of maximum connectivity $\beta^+$ to those BHCs whose data is available for the entire time series. As Goldman Sachs or Morgan Stanley occupy the position of most-connected firm in the benchmark case after their inclusion in the FR-Y9C sample in 2008, this second change does have a noticeable effect[1].

Figure B.5 shows how subsector average default probabilities (which, aside from the maximum connectivity changes described above, are the primary way that a balanced panel can change our measure) change when a balance panel restriction is imposed. Most subsectors show very similar average default probabilities under the benchmark setup and a balanced panel treatment. The major exceptions are securities brokers and dealers, whose average default probability decreases much quicker after the height of the 2008 financial crisis. This shows the the effect outlined above - in a balanced panel setup, the default probabilities and assets of Morgan Stanley and Goldman Sachs are permitted to factor into the subsector average, which has a stabilizing effect on its magnitude.

[1]Although this configuration is identical to the benchmark setup before 2008, as JP Morgan (which has FR-Y9C data in each quarter) is selected as the most-connected firm in those quarters.

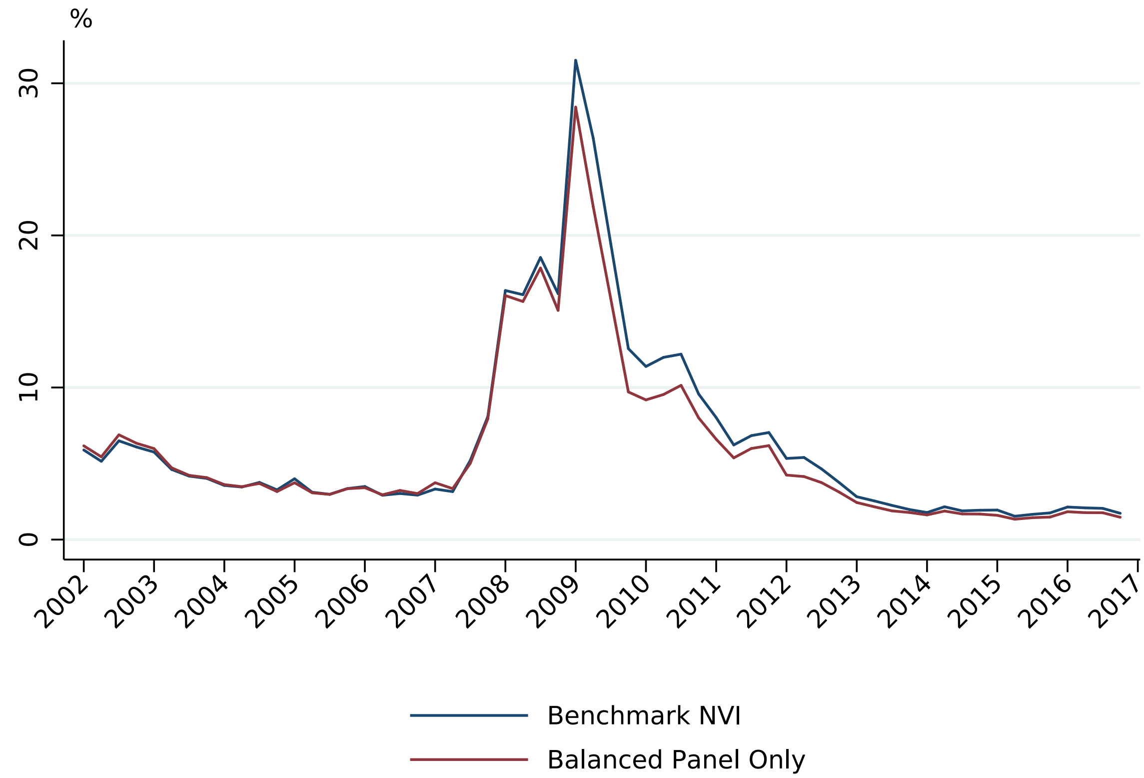


Figure B.4: **Network Vulnerability Index with Balanced FR-Y9C Panels.** The red line of the figure shows our network spillover measure when we restrict our FR-Y9C sample to only those firms where data is available for our entire sample period, 2002-Q1 to 2016-Q4. The green line shows the spillover measure when we restrict the firms eligible to have their connectivity chosen as the maximum connectivity for the measure in equation 2.3 to the same balanced panel. Both treatments have noticeable, but relatively modest effects.

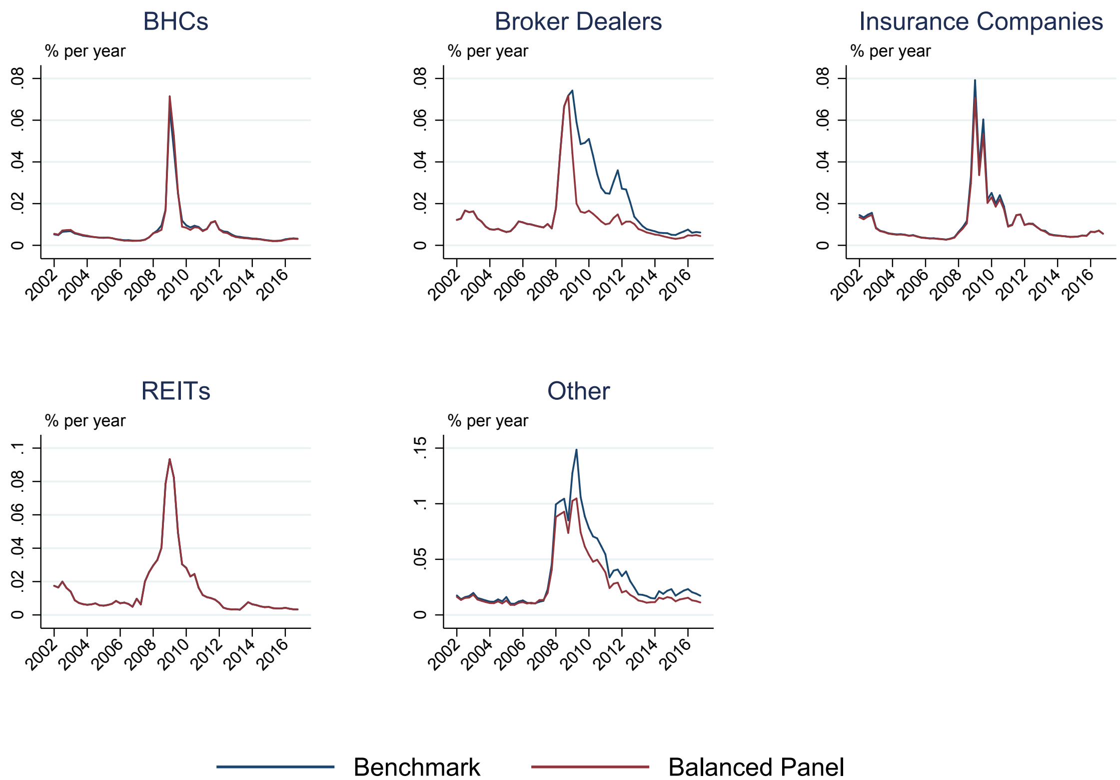


Figure B.5: **Sector-Wide Asset-Weighted Default Probabilities with Balanced FR-Y9C Panels.** The subsector with the largest default probability change under a balanced panel treatment are security brokers and dealers. This happens because a balanced panel allows several large broker dealers who became BHCs during the Financial Crisis to remain in the subsector sample after 2008.

## High, Fixed Default Probability

Next we show the behavior of our NVI when we assume crisis-like default conditions in every quarter of the sample - fixing default probabilities for all firms (in the FR-Y9C and in approximated subsector nodes) at 6% - which is close to the maximum average default probability in the BHC subsample. Figure B.6 shows that the NVI remains relatively high throughout the entirety of the sample when this restriciton is imposed. This shows that, while some variation in the NVI comes from connectivity dynamics reflected through $\beta^+$, that majority of variation over time comes from the credit risk of firms.

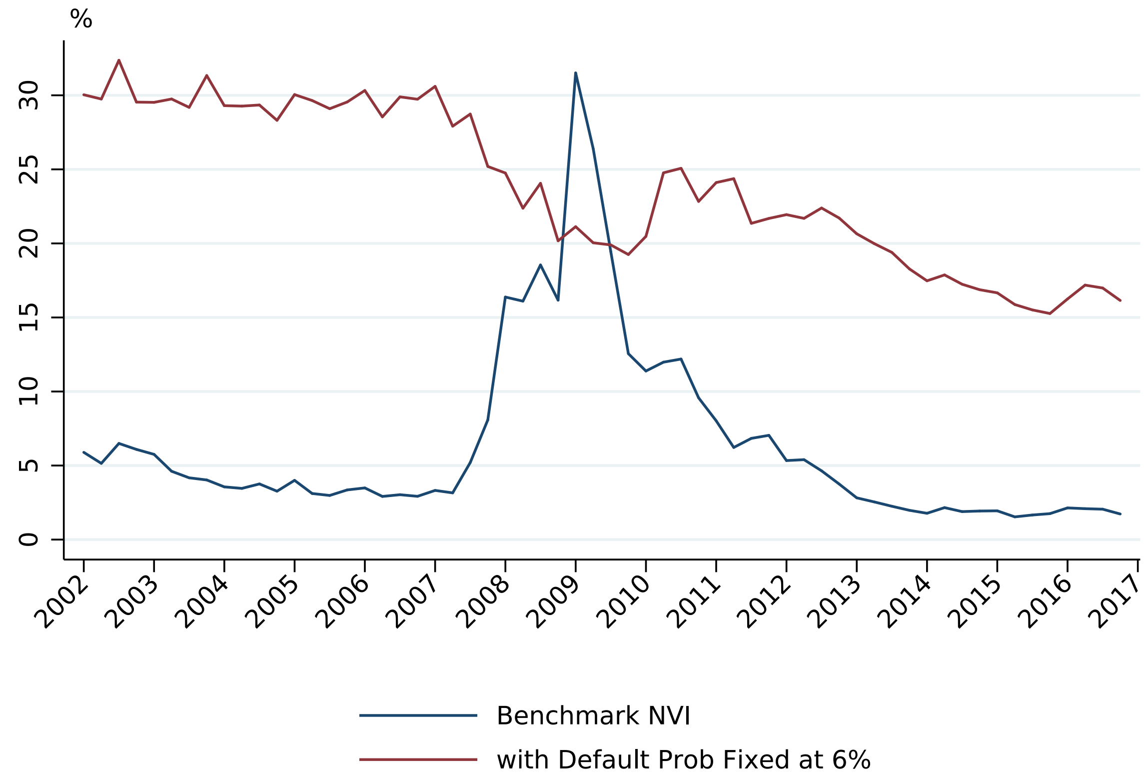


Figure B.6: **Network Vulnerability Index Under a Fixed and High Default Probability of 6%.** The red line above shows the NVI when we assume that all firms in the network have a constant default probability of 6%. This shows that most time variation in the measure is driven by firm credit risk dynamics. As the NVI becomes linear in default probabilities when they are uniform across firms, the second line can be scaled to represent the NVI at any possible fixed default probability.

## B.2 Subsector Firm Sample

As Section 2.3.4 describes, we use our Expected Default Frequency database from Moody's Analytics to compute an asset-weighted average probability of default for those firms where firm-level balance sheet data is less readily-accessible. Tables B.6,B.4,B.2,and B.3 show the asset weightings assigned to firms in each included subsector for 2016-Q4, that last period in our sample. For display purposes, any firms assigned less than a 1% weighting are not included in the table, although each table's final line shows what portion of total weights are assigned to all such firms. Note that these samples exclude any firms whose data is already included in our FR-Y9C sample (note the absence of Goldman Sachs or Morgan Stanley in the sample for security brokers and dealers, for instance).

| Security Broker and Dealer | Asset Weighting |
|---|---|
| Ameriprise Financial Inc | 0.34 |
| Intercontinental Exchange | 0.18 |
| Cme Group Inc | 0.15 |
| Interactive Brokers Group | 0.13 |
| Td Ameritrade Holding Corp | 0.07 |
| Nasdaq Omx Group Inc | 0.03 |
| Kcg Holdings Inc | 0.02 |
| Intl Fcstone Inc | 0.01 |
| Bgc Partners Inc | 0.01 |
| Lpl Financial Holdings Inc | 0.01 |
| Number of Firms in Sample | 39 |
| Weighting from Rest of Sample | 0.05 |

Table B.1: **Asset Weighting in Average Default Probabilities for Securities Brokers and Dealers, 2016-Q4.** Note that any firms with less than a 1% weighting are not displayed.

| Insurance Company | Asset Weighting |
|---|---|
| Metlife Inc | 0.20 |
| Prudential Financial Inc | 0.18 |
| American International Group | 0.11 |
| Lincoln National Corp | 0.06 |
| Hartford Financial Services | 0.05 |
| Voya Financial Inc | 0.05 |
| Aflac Inc | 0.03 |
| Unitedhealth Group Inc | 0.03 |
| Genworth Financial Inc | 0.02 |
| Allstate Corp | 0.02 |
| Travelers Cos Inc | 0.02 |
| Loews Corp | 0.02 |
| Aetna Inc | 0.02 |
| Anthem Inc | 0.01 |
| Unum Group | 0.01 |
| Cigna Corp | 0.01 |
| Cna Financial Corp | 0.01 |
| American Eqty Invt Life Hldg | 0.01 |
| Reinsurance Group America Inc | 0.01 |
| American Financial Group Inc | 0.01 |
| Number of Firms in Sample | 93 |
| Weighting from Rest of Sample | 0.11 |

Table B.2: **Asset Weighting in Average Default Probabilities for Insurance Companies, 2016-Q4.** Note that any firms with less than a 1% weighting are not displayed.

| Real Estate Investment Trust | Asset Weighting |
|---|---|
| Annaly Capital Management | 0.06 |
| Starwood Property Trust Inc | 0.06 |
| American Capital Agency Corp | 0.04 |
| Simon Property Group Inc | 0.02 |
| Prologis Inc | 0.02 |
| American Tower Corp | 0.02 |
| Health Care Reit Inc | 0.02 |
| Ventas Inc | 0.02 |
| Crown Castle Intl Corp | 0.02 |
| Two Harbors Investment Corp | 0.02 |
| Equity Residential | 0.02 |
| Hcp Inc | 0.02 |
| Vornado Realty Trust | 0.02 |
| Boston Properties Inc | 0.01 |
| New Residential Inv Cp | 0.01 |
| Avalonbay Communities Inc | 0.01 |
| Chimera Investment Corp | 0.01 |
| Invesco Mortgage Capital Inc | 0.01 |
| American Rlty Cap Ppty Inc | 0.01 |
| Sl Green Realty Corp | 0.01 |
| Capstead Mortgage Corp | 0.01 |
| Number of Firms in Sample | 230 |
| Weighting from Rest of Sample | 0.54 |

Table B.3: **Asset Weighting in Average Default Probabilities for Real Estate Investment Trusts, 2016-Q4.** Note that any firms with less than a 1% weighting are not displayed.

| Other Financial Firm | Asset Weighting |
|---|---|
| Principal Financial Group Inc | 0.22 |
| Navient Corp | 0.12 |
| Blackrock Inc | 0.07 |
| Visa Inc | 0.06 |
| Oaktree Capital Group Llc | 0.05 |
| Santander Consumer Usa Hldgs | 0.04 |
| Kkr Co Lp | 0.04 |
| Nelnet Inc | 0.03 |
| Invesco Ltd | 0.02 |
| Blackstone Group Lp | 0.02 |
| Marsh Mclennan Cos | 0.02 |
| Springleaf Holdings Inc | 0.02 |
| Slm Corp | 0.02 |
| Walter Investment Mgmt Corp | 0.02 |
| Mastercard Inc | 0.02 |
| Franklin Resources Inc | 0.02 |
| Federal Agriculture Mtg Cp | 0.02 |
| Nationstar Mortgage Holdings | 0.02 |
| Fidelity Natl Finl Fnf Group | 0.01 |
| Arthur J Gallagher Co | 0.01 |
| Number of Firms in Sample | 128 |
| Weighting from Rest of Sample | 0.18 |

Table B.4: **Asset Weighting in Average Default Probabilities for Other Financial Firms, 2016-Q4.** Note that any firms with less than a 1% weighting are not displayed.

| Top 10 Dealers by Assets | Asset Weighting |
|---|---|
| Goldman Sachs Group Inc | 0.41 |
| Morgan Stanley | 0.37 |
| Ameriprise Financial Inc | 0.07 |
| Intercontinental Exchange | 0.04 |
| Cme Group Inc | 0.03 |
| Interactive Brokers Group | 0.03 |
| E Trade Financial Corp | 0.02 |
| Raymond James Financial Corp | 0.01 |
| Td Ameritrade Holding Corp | 0.01 |
| Stifel Financial Corp | 0.01 |
| Number of Firms in Sample | 10 |

Table B.5: **Asset Weighting in Average Default Probabilities for Top 10 Securities Brokers and Dealers, 2016-Q4.**

| Top 1125 Dealers by Assets | Asset Weighting |
|---|---|
| Nasdaq Omx Group Inc | 0.27 |
| Kcg Holdings Inc | 0.12 |
| Intl Fcstone Inc | 0.11 |
| Bgc Partners Inc | 0.09 |
| Lpl Financial Holdings Inc | 0.08 |
| Virtu Financial Inc | 0.06 |
| Oppenheimer Holdings Inc | 0.05 |
| Piper Jaffray Cos Inc | 0.04 |
| Gain Capital Holdings Inc | 0.03 |
| Evercore Partners Inc | 0.03 |
| Waddellreed Finl Inc Cl A | 0.03 |
| Bats Global Markets Redh | 0.02 |
| Fxcm Inc | 0.02 |
| Jmp Group Llc | 0.02 |
| Bankrate Inc | 0.02 |
| Number of Firms in Sample | 15 |

Table B.6: **Asset Weighting in Average Default Probabilities for Top 11-25 Securities Brokers and Dealers, 2016-Q4.**

## B.3 Balance Sheet Asset and Liability Classifications

Table B.7: Asset Variables

| FR Y-9C Variable | 2002-Q2 - 2005-Q4 | 2006-Q1 - 2007-Q4 | 2008-Q1 - 2009-Q1 | 2009-Q2 - 2010-Q4 | 2011-Q1 - Present | %In | %Out |
|---|---|---|---|---|---|---|---|
| Interest-Bearing Balances, Foreign Offices | BHCK0397 | BHCK0397 | BHCK0397 | BHCK0397 | BHCK0397 | 100 | 0 |
| HTM: Other Residential MBS | (BHCK1709+ BHCK1733)* $F_1$ | (BHCK1709+ BHCK1733)* $F_1$ | (BHCK1709+ BHCK1733)* $F_1$ | BHCKG320 | BHCKG320 | 100 | 0 |
| HTM: ABS | BHCKB838+ BHCKB842+ BHCKB846+ BHCKB850+ BHCKB854+ BHCKB858 | BHCKC026 | BHCKC026 | BHCKC026+ BHCKG336+ BHCKG340+ BHCKG344 | BHCKC026+ BHCKG336+ BHCKG340+ BHCKG344 | 100 | 0 |
| For-Sale: Other Residential MBS | (BHCK1713+ BHCK1736) * $F_2$ | (BHCK1713+ BHCK1736)* $F_2$ | (BHCK1713+ BHCK1736)* $F_2$ | BHCKG323 | BHCKG323 | 100 | 0 |
| Available-For-Sale: ABS | BHCKB841+ BHCKB845+ BHCKB849+ BHCKB853+ BHCKB857+ BHCKB861 | BHCKC027 | BHCKC027 | BHCKC027+ BHCKG339+ BHCKG343+ BHCKG347 | BHCKC027+ BHCKG339+ BHCKG343+ BHCKG347 | 100 | 0 |
| Fed Funds | BHDMB987 | BHDMB987 | BHDMB987 | BHDMB987 | BHDMB987 | 100 | 0 |
| Repo Purchases | BHCKB989 | BHCKB989 | BHCKB989 | BHCKB989 | BHCKB989 | 100 | 0 |
| Trading: Structured Financial Products | BHCK3537* $F_3$ | BHCK3537* $F_3$ | BHCM3537* $F_3$ | BHCKG383+ BHCKG384+ BHCKG385 | BHCKG383+ BHCKG384+ BHCKG385 | 100 | 0 |
| Subsidiary Investments | BHCK2130 | BHCK2130 | BHCK2130 | BHCK2130 | BHCK2130 | 100 | 0 |
| Domestic Interest-Bearing Deposits | BHCK0395 | BHCK0395 | BHCK0395 | BHCK0395 | BHCK0395 | 50 | 50 |
| Available-For-Sale: Mutual Fund Investments | BHCKA511 | BHCKA511 | BHCKA511 | BHCKA511 | BHCKA511 | 50 | 50 |
| Goodwill | BHCK3163 | BHCK3163 | BHCK3163 | BHCK3163 | BHCK3163 | 50 | 50 |
| Available-For-Sale: Other Residential MBS, Govt Guaranteed | (BHCK1717+ BHCK1732)* $F_4$ | (BHCK1717+ BHCK1732)* $F_4$ | (BHCK1717+ BHCK1732)* $F_4$ | BHCKG319 | BHCKG319 | 50 | 50 |
| HTM: Other Residential MBS, Govt Guaranteed | (BHCK1714+ BHCK1718)* $F_5$ | (BHCK1714+ BHCK1718)* $F_5$ | (BHCK1714+ BHCK1718)* $F_5$ | BHCKG316 | BHCKG316 | 50 | 50 |
| Interest Strips, Mortgages | BHCKA519 | BHCKA519 | BHCKA519 | BHCKA519 | BHCKA519 | 50 | 50 |

Table B.7: Asset Variables

| FR Y-9C Variable | 2002-Q2 - 2005-Q4 | 2006-Q1 - 2007-Q4 | 2008-Q1 - 2009-Q1 | 2009-Q2 - 2010-Q4 | 2011-Q1 - Present | %In | %Out |
|---|---|---|---|---|---|---|---|
| Life Insurance, General Account | - | - | - | - | BHCKK201 | 50 | 50 |
| Life Insurance, Separate Account | - | - | - | - | BHCKK202 | 50 | 50 |
| Life Insurance, Hybrid Account | - | - | - | - | BHCKK270 | 50 | 50 |
| Trading: Other Loans | - | - | BHCKF618 | BHCKF618 | BHCKF618 | 50 | 50 |
| HTM: Non-Agency Pass-Through MBS | (BHCK1709+ BHCK1733)* $F_6$ | (BHCK1709+ BHCK1733)* $F_6$ | (BHCK1709+ BHCK1733)* $F_6$ | (BHCKG308+ BHCKG320+ BHCKG324+ BHCKG328)* $F_6$ | BHCKG308+ BHCKK146+ BHCKK154 | 50 | 50 |
| Available-For-Sale: Non-Agency Pass-Through MBS | (BHCK1713+ BHCK1736)* $F_7$ | (BHCK1713+ BHCK1736)* $F_7$ | (BHCK1713+ BHCK1736)* $F_7$ | (BHCKG311+ BHCKG323+ BHCKG327+ BHCKG331)* $F_7$ | BHCKG311+ BHCKK149+ BHCKK157 | 50 | 50 |
| Trading: Other MBS | BHCK3536 | BHCK3536 | BHCM3536 | (BHCKG381+ BHCKG382) | BHCKG381+ BHCKK198 | 50 | 50 |
| Trading: Other | BHCK3541 | BHCK3541 | BHCM3541 | BHCM3541 | BHCM3541 | 50 | 50 |
| Trading: Derivatives | (BHCK3543+ BHFN3543) | (BHCK3543+ BHFN3543) | BHCM3541 | BHCM3543 | BHCM3543 | 50 | 50 |
| Mortgage Servicing Assets | BHCK6438 | BHCK6438 | BHCK6438 | BHCK6438 | BHCK6438 | 50 | 50 |
| Credit Card Relationships | BHCKB026 | BHCKB026 | BHCKB026 | BHCKB026 | BHCKB026 | 50 | 50 |
| Other Intangible | BHCK5507 | BHCK5507 | BHCK5507 | BHCK5507 | BHCK5507 | 50 | 50 |
| Accrued Interest | BHCKB556 | BHCKB556 | BHCKB556 | BHCKB556 | BHCKB556 | 50 | 50 |
| Other Interest-only Strips | BHCKA520 | BHCKA520 | BHCKA520 | BHCKA520 | BHCKA520 | 50 | 50 |
| Other | BHCK2168 | BHCK2168 | BHCK2168 | BHCK2168 | BHCK2168 | 50 | 50 |
| Noninterest-Bearing Deposits | BHCK0081 | BHCK0081 | BHCK0081 | BHCK0081 | BHCK0081 | 0 | 100 |
| HTM: Treasuries | BHCK0211 | BHCK0211 | BHCK0211 | BHCK0211 | BHCK0211 | 0 | 100 |
| HTM: Agency Debt | BHCK1289 | BHCK1289 | BHCK1289 | BHCK1289 | BHCK1289 | 0 | 100 |
| HTM: GSE Debt | BHCK1294 | BHCK1294 | BHCK1294 | BHCK1294 | BHCK1294 | 0 | 100 |
| HTM: State Debt | BHCK8496 | BHCK8496 | BHCK8496 | BHCK8496 | BHCK8496 | 0 | 100 |
| HTM: GSE MBS | BHCKK1698+ BHCK1703 | BHCKK1698+ BHCK1703 | BHCKK1698+ BHCK1703 | BHCKG300+ BHCKG304 | BHCKG300+ BHCKG304+ BHCKK142 | 0 | 100 |
| HTM: Agency MBS | BHCK1714+ BHCK1718- BHCKG316 | BHCK1714+ BHCK1718- BHCKG316 | BHCK1714+ BHCK1718- BHCKG316 | BHCKG312 | BHCKG312+ BHCKK150 | 0 | 100 |

Table B.7: Asset Variables

| FR Y-9C Variable | 2002-Q2 - 2005-Q4 | 2006-Q1 - 2007-Q4 | 2008-Q1 - 2009-Q1 | 2009-Q2 - 2010-Q4 | 2011-Q1 - Present | %In | %Out |
|---|---|---|---|---|---|---|---|
| HTM: Other Domestic Debt Securities | BHCK1737 | BHCK1737 | BHCK1737 | BHCK1737 | BHCK1737 | 0 | 100 |
| HTM: Foreign Debt Securities | BHCK1742 | BHCK1742 | BHCK1742 | BHCK1742 | BHCK1742 | 0 | 100 |
| Available-For-Sale: Treasuries | BHCK1287 | BHCK1287 | BHCK1287 | BHCK1287 | BHCK1287 | 0 | 100 |
| Available-For-Sale: Agency Debt | BHCK1293 | BHCK1293 | BHCK1293 | BHCK1293 | BHCK1293 | 0 | 100 |
| Available-For-Sale: GSE Debt | BHCK1298 | BHCK1298 | BHCK1298 | BHCK1298 | BHCK1298 | 0 | 100 |
| Available-For-Sale: State Debt | BHCK8499 | BHCK8499 | BHCK8499 | BHCK8499 | BHCK8499 | 0 | 100 |
| Available-For-Sale: GSE MBS | BHCK1702+ BHCK1707 | BHCK1702+ BHCK1707 | BHCK1702+ BHCK1707 | BHCKG303+ BHCKG307 | BHCKG303+ BHCKG307+ BHCKK145 | 0 | 100 |
| Available-For-Sale: Agency MBS | BHCK1717+ BHCK1732- BHCK319 | BHCK1717+ BHCK1732- BHCK319 | BHCK1717+ BHCK1732- BHCK319 | BHCKG315 | BHCKG315+ BHCKK153 | 0 | 100 |
| Available-For-Sale: Other Domestic Debt Securities | BHCK1741 | BHCK1741 | BHCK1741 | BHCK1741 | BHCK1741 | 0 | 100 |
| Available-For-Sale: Foreign Debt Securities | BHCK1746 | BHCK1746 | BHCK1746 | BHCK1746 | BHCK1746 | 0 | 100 |
| Loans and Leases Held for Sale | BHCK5369 | BHCK5369 | BHCK5369 | BHCK5369 | BHCK5369 | 0 | 100 |
| Loans and Leases, Net Unearned Income and Loss Allowance | BHCKB529 | BHCKB529 | BHCKB529 | BHCKB529 | BHCKB529 | 0 | 100 |
| Trading: Treasuries | BHCK3531 | BHCK3531 | BHCM3531 | BHCM3531 | BHCM3531 | 0 | 100 |
| Trading: Agency Debt | BHCK3532 | BHCK3532 | BHCM3532 | BHCM3532 | BHCM3532 | 0 | 100 |
| Trading: State Debt | BHCK3533 | BHCK3533 | BHCM3533 | BHCM3533 | BHCM3533 | 0 | 100 |
| Trading: GSE MBS | BHCK3534+ BHCK3535 | BHCK3534+ BHCK3535 | BHCK3534+ BHCK3535 | BHCKG379+ BHCKG380 | BHCKG379+ BHCKG380+ BHCKK197 | 0 | 100 |
| Trading: Other Debt Securities | BHCK3537- BHCKG383- BHCKG384- BHCKG385 | BHCK3537- BHCKG383- BHCKG384- BHCKG385 | BHCM3537- BHCKG383- BHCKG384- BHCKG385 | BHCKG386 | BHCKG386 | 0 | 100 |
| Trading: Real Estate Loans | - | - | BHCKF610 | BHCKF610 | BHCKF610 | 0 | 100 |
| Trading: Commercial Loans | - | - | BHCKF614 | BHCKF614 | BHCKF614 | 0 | 100 |
| Trading: Credit Cards | - | - | BHCKF615 | BHCKF615 | BHCKF615 | 0 | 100 |

Table B.7: Asset Variables

| FR Y-9C Variable | 2002-Q2 - 2005-Q4 | 2006-Q1 - 2007-Q4 | 2008-Q1 - 2009-Q1 | 2009-Q2 - 2010-Q4 | 2011-Q1 - Present | %In | %Out |
|---|---|---|---|---|---|---|---|
| Trading: Revolving Credit | - | - | BHCKF616 | BHCKF616 | BHCKF616 | 0 | 100 |
| Trading: Consumer Loans | - | - | BHCKF617 | BHCKF617 | BHCKK199+ BHCKK210 | 0 | 100 |
| Premises | BHCK2145 | BHCK2145 | BHCK2145 | BHCK2145 | BHCK2145 | 0 | 100 |
| Other Real Estate Owned | BHCK2150 | BHCK2150 | BHCK2150 | BHCK2150 | BHCK2150 | 0 | 100 |
| Real Estate Venture Investment | BHCK3656 | BHCK3656 | BHCK3656 | BHCK3656 | BHCK3656 | 0 | 100 |
| Deferred Tax Assets | BHCK2148 | BHCK2148 | BHCK2148 | BHCK2148 | BHCK2148 | 0 | 100 |
| Equities of Indeterminable Value | BHCK1752 | BHCK1752 | BHCK1752 | BHCK1752 | BHCK1752 | 0 | 100 |

Source: FR-Y9C, FFIEC031. Sample 2002-Q2 - 2016-Q4. More information on each variable is available from the Board of Governors of the Federal Reserve System at https://www.federalreserve.gov/apps/mdrm/. Cells filled with '-' denote fields that are unavailable (and with no discernible substitute) in particular timespans.

Table B.8: Liabilities Variables

| FR Y-9C Variable | 2002-Q2 - 2005-Q4 | 2006-Q1 - 2007-Q4 | 2008-Q1 - 2008-Q4 | 2009-Q1 - Present | %In | %Out |
|---|---|---|---|---|---|---|
| Noninterest-bearing deposits (uninsured) | BHCB2210 | BHCB2210 | BHCB2210 | BHCB2210 | 100 | 0 |
| Noninterest-bearing deposits (insured) | BHCB2210 | BHCB2210 | BHCB2210 | BHCB2210 | 0 | 100 |
| Noninterest-bearing deposits (Non-Commercial Banks) | BHOD3189 | BHOD3189 | BHOD3189 | BHOD3189 | 50 | 50 |
| Interest-bearing deposits (uninsured) | BHCB3187 | BHCB3187 | BHCB3187 | BHCB3187 | 100 | 0 |
| Interest-bearing deposits (insured) | BHCB3187 | BHCB3187 | BHCB3187 | BHCB3187 | 0 | 100 |
| Interest-bearing deposits (Non-Commercial Banks) | BHOD3187 | BHOD3187 | BHOD3187 | BHOD3187 | 50 | 50 |
| Savings deposits (uninsured) | BHCB2389 | BHCB2389 | BHCB2389 | BHCB2389 | 100 | 0 |
| Savings deposits (insured) | BHCB2389 | BHCB2389 | BHCB2389 | BHCB2389 | 0 | 100 |
| Savings deposits (Non-Commercial Banks) | BHOD2389 | BHOD2389 | BHOD2389 | BHOD2389 | 50 | 50 |
| Time deposits <$100,000 (uninsured) | BHCB6648 | BHCB6648 | BHCB6648 | BHCB6648 | 100 | 0 |
| Time deposits <$100,000 (insured) | BHCB6648 | BHCB6648 | BHCB6648 | BHCB6648 | 0 | 100 |
| Time deposits <$100,000 (Non-Commercial Banks) | BHOD6648 | BHOD6648 | BHOD6648 | BHOD6648 | 50 | 50 |
| Time deposits > $100,000 (uninsured) | BHCB2604 | BHCB2604 | BHCB2604 | BHCB2604 | 100 | 0 |
| Time deposits > $100,000 (insured) | BHCB2604 | BHCB2604 | BHCB2604 | BHCB2604 | 0 | 100 |
| Time deposits > $100,000 (Non-Commercial Banks) | BHOD2604 | BHOD2604 | BHOD2604 | BHOD2604 | 50 | 50 |
| Subordinated debt | BHCK4062 | BHCK4062 | BHCK4062 | BHCK4062 | 0 | 100 |
| Net deferred taxes | BHCK3049 | BHCK3049 | BHCK3049 | BHCK3049 | 0 | 100 |
| Other short term debt | BHCK2332 | BHCK2332 | BHCK2332 | BHCK2332 | 50 | 50 |
| Other longer term debt | BHCK2333 | BHCK2333 | BHCK2333 | BHCK2333 | 50 | 50 |

Table B.8: Liabilities Variables

| FR Y-9C Variable | 2002-Q2 - 2005-Q4 | 2006-Q1 - 2007-Q4 | 2008-Q1 - 2008-Q4 | 2009-Q1 - Present | %In | %Out |
|---|---|---|---|---|---|---|
| Noninterest-bearing deposits, foreign | BHFN6631 | BHFN6631 | BHFN6631 | BHFN6631 | 50 | 50 |
| Interest-bearing deposits, foreign | BHFN6636 | BHFN6636 | BHFN6636 | BHFN6636 | 50 | 50 |
| Commercial paper | BHCK2309 | BHCK2309 | BHCK2309 | BHCK2309 | 50 | 50 |
| Subordinated notes to trusts | - | BHCKC699 | BHCKC699 | BHCKC699 | 50 | 50 |
| Off-balance sheet credit losses | BHCKB557 | BHCKB557 | BHCKB557 | BHCKB557 | 50 | 50 |
| Other | BHCKB984 | BHCKB984 | BHCKB984 | BHCKB984 | 50 | 50 |
| Repos | BHCKB995 | BHCKB995 | BHCKB995 | BHCKB995 | 100 | 0 |
| Liabilities for short positions | BHCK3546 | BHCK3546 | BHCK3546 | BHCKG209+ BHCKG210+ BHCKG211 | 100 | 0 |
| Derivatives with Negative Fair Value | BHCK3547 | BHCK3547 | BHCK3547 | BHCK3547 | 100 | 0 |
| Other trading liabilities | - | - | BHCKF624 | BHCKF624 | 100 | 0 |
| Revaulation losses | BHCK3547 | BHCK3547 | BHCK3547 | BHCK3547 | 100 | 0 |
| Fed funds | BHDMB993 | BHDMB993 | BHDMB993 | BHDMB993 | 100 | 0 |

Source: FR-Y9C, FFIEC031. Sample 2002-Q2 - 2016-Q4. More information on each variable is available from the Board of Governors of the Federal Reserve System at https://www.federalreserve.gov/apps/mdrm/. Cells filled with '-' denote fields that are unavailable (and with no discernible substitute) in particular timespans.

Table B.9: Proportional Adjustments

| Code | Fraction |
|---|---|
| $F_1$ | $\frac{BHCKG320}{BHCKG320+BHCKG308+BHCKG324+BHCKG328}$ |
| $F_2$ | $\frac{BHCKG323}{BHCKG311+BHCKG323+BHCKGBHCKG327+BHCKG331}$ |
| $F_3$ | $\frac{BHCKG383+BHCKG384+BHCKG385}{BHCKG383+BHCKG384+BHCKG385+BHCKG386}$ |
| $F_4$ | $\frac{BHCKG319}{BHCKG319+BHCKG315}$ |
| $F_5$ | $\frac{BHCKG316}{BHCKG316+BHCKG312}$ |
| $F_6$ | $\frac{BHCKG308+BHCKK146+BHCKK154}{BHCKG308+BHCKG320+BHCKK146+BHCKK154}$ |
| $F_7$ | $\frac{BHCKG311+BHCKK149+BHCKK157}{BHCKG311+BHCKK149+BHCKK157+BHCKG323}$ |

The final number used in Table B.7 is constructed as the average of these values over every quarter where all relevant series are available.

| Balance Sheet Classification | %In | %Out | % Of Sector Assets or Liabilities |
|---|---|---|---|
| Insurance Companies (Assets) | | | |
| Corporate and Foreign Bonds | 50 | 50 | 35.2 |
| Mutual Fund Shares | 50 | 50 | 18.0 |
| Corporate Equities | 50 | 50 | 9.3 |
| Municipal Securities | 0 | 100 | 6.0 |
| Mortgages | 0 | 100 | 5.5 |
| Agency and GSE-backed securities | 0 | 100 | 5.4 |
| Treasury Securities | 0 | 100 | 3.8 |
| US direct investment | 0 | 100 | 2.0 |
| Checkable Deposits and Currency | 100 | 0 | 1.1 |
| Market Paper | 50 | 50 | 0.7 |
| Money Market Mutual Fund Shares | 50 | 50 | 0.7 |
| Deferred and Unpaid Life ins Premiums | 0 | 100 | 0.4 |
| Security Repurchase Agreements | 100 | 0 | 0.0 |
| Equity in FHLB | 0 | 100 | 0.0 |
| Other Loans and Advances | 50 | 50 | 1.9 |
| Other/Unallocated Claims | 50 | 50 | 10.0 |
| Real Estate Investment Trusts (Assets) | | | |
| Nonfinancial Assets | 0 | 100 | 57.5 |
| Loans | 0 | 100 | 14.3 |
| Agency and GSE-backed securities | 0 | 100 | 12.6 |
| Checkable Deposits and Currency | 100 | 0 | 2.4 |
| Corporate and Foreign Bonds | 50 | 50 | 2.5 |
| Other | 50 | 50 | 10.7 |
| Other: Credit Unions (Assets) | | | |
| Loans | 0 | 100 | 70.6 |
| Agency and GSE-Backed | 0 | 100 | 13.7 |
| Reserves | 0 | 100 | 5.3 |
| Treasury Securities | 0 | 100 | 1.1 |
| Corporate and Foreign Bonds | 50 | 50 | 0.9 |
| Municipal Bonds | 0 | 100 | 0.4 |
| Mutual Fund Shares | 0 | 100 | 0.2 |
| Fed Funds and Repos | 100 | 0 | 0.0 |
| Open Market Paper | 50 | 50 | 0.0 |
| Other | 50 | 50 | 7.8 |
| Other: Finance Companies (Assets) | | | |
| Consumer Credit | 0 | 100 | 35.9 |
| Other Loans | 50 | 50 | 25.4 |
| US Direct Investment Abroad | 0 | 100 | 16.5 |
| Mortgages | 0 | 100 | 9.1 |
| Corporate and Foreign Bonds | 50 | 50 | 4.8 |
| Time and Savings Deposits | 100 | 0 | 4.1 |
| Checkable Deposits and Currency | 100 | 0 | 1.4 |

| | | | |
|---|---|---|---|
| Other | 50 | 50 | 25.4 |
| Other: Funding Corporations (Assets) | | | |
| Investment in Brokers and Dealers | 100 | 0 | 41.1 |
| Money Market Fund Shares | 0 | 100 | 32.8 |
| Open Market Paper | 50 | 50 | 15.1 |
| Investment in Foreign Banks | 0 | 100 | 5.9 |
| Corporate and Foreign Bonds | 50 | 50 | 4.1 |
| Loans | 50 | 50 | 1.0 |
| Corporate Equities | 50 | 50 | 0.0 |
| Security Repurchase Agreements | 100 | 0 | 0.0 |
| Other: ABS Issuers (Assets) | | | |
| Mortgages | 0 | 100 | 72.8 |
| Other Loans | 50 | 50 | 13.0 |
| Consumer Credit | 0 | 100 | 4.1 |
| Trade Credit | 0 | 100 | 2.8 |
| Treasury Securities | 0 | 100 | 1.6 |
| Agency and GSE-Backed Securities | 0 | 100 | 0.0 |
| Other | 50 | 50 | 13.0 |
| Top 25 Dealers, FOCUS (Liabilities) | | | |
| Repurchase Agreements | 100 | 0 | 33.3 |
| Payables to Customers | 0 | 100 | 26.3 |
| Payables to BDs, Clearing | 100 | 0 | 12.2 |
| Securities Sold Short | 100 | 0 | 8.7 |
| Obligation to Return Securities | 100 | 0 | 4.7 |
| Notes and Mortgages | 0 | 100 | 4.6 |
| Subordinated Liabilities | 0 | 100 | 3.8 |
| Accounts Payable and Accrued Liabilities | 0 | 100 | 3.4 |
| Payables to Non-Customers | 0 | 100 | 1.5 |
| Bank Loans Payable | 100 | 0 | 1.4 |
| Special Liabilities | 0 | 100 | 0.1 |

Table B.10: **Classification of Financial Accounts of the United States Asset Classes Into 'Inside' or 'Outside' the Financial System** Using variables from the Financial Accounts of the United States, we categorize assets of each financial subsector as either 'inside' or 'outside' the financial system.