Title
Essays in Macroeconomics and Financial Frictions

Author
Ray, Walker

Publication Date
2019

Peer reviewed|Thesis/dissertation

Essays in Macroeconomics and Financial Frictions

By

Walker D Ray

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 Yuriy Gorodnichenko, Chair
Professor Ben Faber
Professor Thibault Fally
Professor Pierre-Olivier Gourinchas
Professor David Romer

Spring 2019

Abstract

Essays in Macroeconomics and Financial Frictions

by

Walker D Ray

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Yuriy Gorodnichenko, Chair

Recent economic experiences have demonstrated the importance of understanding departures from frictionless markets and perfect information-processing for the field of macroeconomics. The global financial crisis has highlighted the importance of financial frictions for macroeconomic policy. With conventional monetary policy unable to stabilize the economy in the wake of the global financial crisis, central banks turned to unconventional tools. Understanding how these tools worked through interactions with financial market disruptions is crucial for designing and implementing policies to deal with the next crisis. The first two chapters show theoretically how unconventional policy worked, and test these predictions in the data.

The rise of political polarization over the past decades raises important questions about how households form macroeconomic beliefs, and how this departs from the typical rationality assumptions embedded in textbook macroeconomic models. The final chapter shows theoretically how imperfect information-processing leads to a divergence of macroeconomic beliefs across households. Empirically, growing disagreement about macroeconomic outcomes is found in survey data; moreover, this disagreement leads to differential consumptions decisions following political shocks.

Chapter 1 embeds a model of the term structure of interest rates featuring market segmentation and limits to arbitrage within a New Keynesian model to study unconventional monetary policy. Because the transmission of monetary policy depends on private agents with limited risk-bearing capacity, financial market disruptions reduce the efficacy of both conventional policy as well as forward guidance. Conversely, financial crises are precisely when large scale asset purchases are most effective. Policymakers can take advantage of the inability of financial markets to fully absorb these purchases, which can push down long-term interest rates and help stabilize output and inflation.

Chapter 2 seeks to understand empirically the effects of large-scale asset purchase programs recently implemented by central banks. In joint work with Yuriy Gorod-

nichenko, we study how markets absorb large demand shocks for risk-free debt. Using high-frequency identification, we exploit the structure of the primary market for U.S. Treasuries to isolate demand shocks. These shocks are sizable, leading to large movements in Treasury yields and impacting corporate borrowing rates. Informed by a preferred habitat model of the term structure, we test for "local" demand effects and find evidence consistent with theoretical predictions. Crucially, this local effect is strongest when financial markets are disrupted. Our estimates are consistent with the view that quantitative easing worked mainly via market segmentation, with a potentially limited role for other channels.

Chapter 3 explores the role of political polarization in shaping the economic expectations and consumption behavior of households. In joint work with Rupal Kamdar, we first develop a rational inattention model in which heterogeneous households must decide how to obtain information. Theoretically, we show there is a "paradox of information" where falling information costs exacerbate disagreement. Next, using survey data, we find evidence that political polarization has increased dispersion in macroeconomic beliefs. Disagreement is particularly acute following a general election when the presidential party switches; moreover, this effect has been increasing since the 1980s. Finally, we also find that polarization feeds into consumption decisions. Using high-frequency spending data at the zip code level in California, we find Republican-leaning regions exhibit substantially larger consumption following the 2016 election.

This dissertation would not have been possible without the incredible support of my family, friends, and mentors during graduate school. I owe a great deal to Yuriy Gorodnichenko, who exemplifies the very best characteristics of what it means to be a graduate advisor. My parents Beth deHamel and Jim Ray have offered constant encouragement of my academic endeavors. Finally, I am incredibly grateful that I was able to share my graduate school experiences with my partner Rupal Kamdar, who has made these past years better than I could have imagined.

Contents

Chapter 0

Introduction

The past several decades have witnessed a secular change in the macroeconomic landscape of the United States and other developed countries. Most notably, the recent global crisis underscores the role of financial frictions in shaping macroeconomic dynamics and propagating shocks. At the same time, the rise of political polarization raises important questions about imperfections in how households process information and form macroeconomic beliefs. Typically, macroeconomic models assume frictionless financial markets and perfect information-processing. These essays develop theoretical tools to study departures from these frictionless baselines, and show empirically that these frictions matter for understanding macroeconomic dynamics and designing policy.

From a policymaking perspective, the financial crisis saw the rollout of so-called "unconventional" monetary policy tools. Conventional monetary policy, conducted through changes in the short-term interest rate, proved unable to stabilize the economy. Central banks began providing forward guidance regarding the path of the policy rate; later, they began a series of large-scale asset purchases (LSAP), the most salient of which was the quantitative easing (QE) programs conducted by the Federal Reserve. Future deployment of these unconventional tools requires policymakers to move beyond the "heat of the moment" policies. Hence, a central question for policymaking and academic research is to develop theoretical models of how these policies work, and to empirically demonstrate which of these theories is the key channel quantitatively. Moreover, understanding how these tools worked through interactions with financial market disruptions is crucial for assessing how shocks propagate through financial markets to the broader economy.

Benchmark models are not amenable to studying these topics. Because of the extreme forward-looking behavior of agents and the lack of any financial frictions in baseline macroeconomic models, forward guidance policies are highly effective to the point of implausibility. Conversely, LSAP policies are completely ineffective in conventional models. These essays develop a macroeconomic model featuring limited arbitrage and market segmentation to show how both conventional and unconventional policies interact with disruptions in financial markets, and empirically find these frictions play a large role in how these policies affect the economy.

Another important departure from baseline macroeconomic models considered in these essays relates to household expectation formation. Expectations play a central role in nearly all macroeconomic models, but the standard approach is to assume full-information rational expectations (FIRE) on the part of private agents. This assumption is theoretically appealing and has proven to be a highly useful analytical tool, but there are important questions about the empirical realism of the assumption. The rise of political polarization over the past decades and the ensuing divergence of macroeconomic expectations across households throws more water on the FIRE assumption. Analyzing how households process economic information is important not only for understanding how these agents form their macroeconomic expectations, but also how this feeds into consumption and savings decisions.

Although these essays consider disparate topics and methodologies, the through line is the importance of the frictions which agents face when making economic decisions. Taken as a whole, these set of essays develop models for rigorously studying these frictions, demonstrate theoretically why these frictions matter, and confirm empirically that these predictions hold in the data.

The purpose of Chapter 1 is to analyze unconventional monetary tools within a tractable, unified theoretical framework. In order to do so, a term structure model featuring market segmentation and limits to arbitrage is embedded within a New Keynesian framework. In contrast to textbook macroeconomic models with frictionless financial markets, the term structure is determined by participants who face limited risk-bearing capacity and are susceptible to demand shocks (as in the *preferred habitat* view of bond markets). Household borrowing depends not only on the policy rate but also on the entire term structure of interest rates, and hence these financial frictions feed into consumption and savings decisions in general equilibrium. The model is used to study conventional and unconventional monetary policy, and in particular how policy actions interact with disruptions in financial markets.

The implications of the model are important for understanding the efficacy of monetary policy. When financial markets are healthy, so that marginal investors in financial markets have high risk-bearing capacity, the "expectation hypothesis" holds: long-term rates are entirely determined by the expected path of short rates. In this case, conventional monetary policy (as well as forward guidance) are effective at stabilizing the economy. Household borrowing responds strongly to shifts in the path of the policy rate, leading to movements in output and inflation, and hence the central bank can dampen macroeconomic fluctuations.

However, the link between expected short rates and the term structure is weakened when financial distress is high. In equilibrium, the model shows that long-term rates under-react to changes in the policy rate. Therefore, during a financial crisis, the monetary authority will find it much more difficult to stabilize the economy using only conventional tools.

The implications are precisely the opposite for LSAPs. Purchases of long-term bonds, as in the various rounds of QE, will have little to no effect on long-term rates when financial markets are healthy. When financial markets are healthy, investors are able to bear a large amount of risk, and so they do not require much excess returns to hold long-term debt securities. Hence, while QE changes the portfolio allocation of the marginal investors in the debt market, this will not have large effects on bond prices and borrowing rates.

Conversely, as financial markets become unable to bear risk, these purchases matter more and more. By changing the riskiness of marginal investors' portfolio allocations, QE leads to changes in equilibrium prices of bonds. The changes in borrowing rates feeds back into the household borrowing decision. In general equilibrium, QE can boost output and stabilize the economy. Note that investors still eliminate risk-free arbitrage opportunities, so that there are no riskless trades left on the table. Any deviations from the expectations hypothesis are due to the risky portfolio allocations chosen by the marginal financial investors. Hence, the channel through which unconventional policies like QE can have aggregate effects is by changing the market prices of risk.

The exact impact of LSAP programs depends on how the purchases are structured. The amounts to be purchased, which maturities are targeted, and the duration of the program all affect the interaction with the sources of risk in the economy and the broader feedback mechanisms in the macroeconomy. As always, one fundamental source of risk is the movement in the short (policy) rate that is set by the central bank. All bonds are exposed to this risk, so as long as arbitrageurs are not perfectly risk-neutral there will be deviations from the expectations hypothesis. All else equal, arbitrageurs will require excess expected returns in order to take non-zero positions in long-term debt. This effect weakens the strength of forward guidance, but opens the door for LSAPs. The central bank is able to change the portfolio allocations of arbitrageurs, which through changes in the price of risk lead to changes in interest rates.

The final section of Chapter 1 estimates the model using U.S. data from before and during the recent financial crisis. Quantitatively, the model predicts that the aggregate output effects of the first round of QE were roughly 40% larger than a 50 basis point expansionary monetary shock during a period of relative financial calm. Further, had the zero lower bound not been binding during the financial crisis, additional rounds of rate cuts would have been 20% less effective than rate cuts during normal times.

Chapter 2 undertakes a rigorous empirical assessment of how the Treasury market responds to large bond purchase shocks such as QE. While QE was clearly successful in reducing short- and long-term interest rates, the mechanisms behind this reaction are still not agreed upon. Chapter 1 developed a rigorous theoretical model where the key channel is market segmentation in bond markets, but other theoretical channels have been proposed. For instance, QE could be effective because it signaled to the markets that the Fed is serious about keeping short-term interest rates low for a long time (a commitment device for forward guidance). Another explanation is that the Fed signaled to that the economy would remain subdued, which pushed interest rates down. Given the paucity of QE events, it has proven remarkably hard to provide clear empirical evidence for each theory, as well as to assess the relative contributions of the proposed channels. Indeed, many channels were likely active during QE rounds and the reactions to QE were observed in a particular state of the economy, which potentially confounds identification and interpretation.

The objective of Chapter 2 is to unbundle QE by focusing on the market segmentation and preferred habitat channel. The empirical approach is to identify shifts in *private* demand for Treasuries that mimic QE, but are independent of other channels. Chapter 1 showed that the key mechanism through which market segmentation and preferred habitat forces operate is not the source of demand shifts *per se*, but rather how marginal investors in the market for Treasury debt absorb these demand shocks. Therefore, the best way to isolate and study the preferred habitat channel of QE is to identify unexpected demand shifts that are unrelated to other possible channels.

This chapter constructs demand shocks with these properties by utilizing the structure and timing of the primary market for Treasury securities. High-frequency (intraday) changes in prices of Treasury futures in small windows around the close of Treasury auctions are used to identify unexpected shocks to demand for Treasuries. The key for identification is that all of the "supply" information (e.g. security characteristics such as the maturity, as well as the amount of newly offered and outstanding securities) is known and priced in by the market. For small enough windows around the close and release of the auction results, any price changes are reactions to information regarding the demand for the Treasury securities from the given auction.

In sharp contrast to QE events, Treasury auctions are frequent. This allows for crisper inference and to study state-dependence in the effect of targeted purchases of assets (e.g. crisis vs. non-crisis states). Because QE events were both infrequent and confounded with a massive financial crisis, having a long time series is instrumental for understanding how QE-like programs can work in normal times. Importantly, because Treasury auctions for specific maturities are spread in time, it is possible to identify changes in demand for government debt of specific maturities. Hence it is possible to trace how a shock in one part of the yield curve propagates to other parts of the yield curve. These natural experiments mimic targeted purchases of the Fed during QE programs. Hence, despite the apparent distance between QE programs and unexpected movements in private demand during regular Treasury auctions, this empirical strategy provides clean identification of demand shifts in order to map out the impact of these shocks.

The results confirm the theoretical predictions of Chapter 1. QE programs can be effective in influencing interest rates for debt at specific maturities when financial markets are disrupted. On the other hand, QE programs are less likely to be effective at this task in normal times when risk-bearing capacity of arbitrageurs is greater.

Quantitatively, the results are consistent with the view that QE worked through a preferred habit channel, with a small *net* effect of other channels.

Chapter 3 changes gears and explores theoretically and empirically how household beliefs and actions interact with political polarization. The chapter develops a model where heterogeneous agents have imperfect information about the state of the economy, and must choose how to acquire costly information in order to inform their decision-making. This leads to a "paradox of information" whereby declining information costs can actually increase ex-post disagreement about the economy; moreover, disagreement persists even with arbitrarily small information costs.

This "paradox of information" can help rationalize the secular increase in political polarization and the simultaneous increase in the ease of acquiring economic information. A naive model would suggest that increased access to information should reduce disagreement about economic fundamentals. The model shows that households are able to learn about the economy more precisely as information costs fall. However, they process information in the manner which is most advantageous for their own idiosyncratic preferences. Since households are not identical, more information can actually exacerbate ex-post disagreement about the economy.

Empirically, political polarization affects both household beliefs and actions. Survey data on U.S. consumer beliefs shows that households generally have persistent and stable economic beliefs and forecasts. However, there are a few striking exceptions following elections where the White House changes party. During these periods, households that were optimistic about economic conditions become more likely to become pessimistic (and vice versa). This effect has been increasing since the 1980s, with the largest impact coming after the 2016 election.

Theoretically, changes in expectations should lead to changes in shifts in consumption and savings patterns. But empirically, the evidence demonstrating this relationship is not as clear; this is particularly true when it comes to politically-motivated changes in economic beliefs. However, the 2016 presidential election in the U.S. shows that these shifts in economic beliefs translate into differential consumption decisions. Disaggregated geographical spending data demonstrates that polarization leads to differential changes in consumer spending: regions with a large Republican voteshare exhibited substantially higher consumption in the wake of the 2016 election.

Chapter 1

Monetary Policy and the Limits to Arbitrage: Insights from a New Keynesian Preferred Habitat Model

1.1 Introduction

Central banks responded aggressively to worsening financial conditions and growing recessionary pressure during the global financial crisis of 2007-8. After steep cuts in policy rates, central banks found themselves constrained by the zero lower bound, and the crisis was followed by a deep recession. Not content to sit on their hands, policymakers undertook various unconventional policy actions such as forward guidance and large scale asset purchases, the most salient of which was the quantitative easing programs carried out by the Federal Reserve.

What was the purpose of these unconventional policies? With policy rates constrained, the immediate goal was to push down long-term interest rates. But more fundamentally, policymakers believed these actions would stimulate the economy by boosting output and stabilizing inflation. As economic conditions have returned to normal, pivotal questions for macroeconomics remain. The emerging view (though not quite a consensus) in the empirical literature surrounding unconventional monetary policies is that large scale asset purchases (LSAPs) were effective at reducing long-term rates. On the other hand, the economy was not as sensitive to forward guidance as implied by some workhorse models. Why was this? And what were the feedback mechanisms of unconventional policy actions to the broader economy?

The purpose of this paper is to study these monetary tools within a tractable, unified framework. To this end, this paper embeds a model of the term structure featuring market segmentation and limits to arbitrage within a New Keynesian framework. There are two key departures from a benchmark model. First, borrowing depends not only on the policy rate but also on the entire term structure of interest rates. Second, the term structure is determined in financial markets whose participants face limited risk-bearing capacity and are susceptible to demand shocks, as in Vayanos and Vila (2009). I use this model as a laboratory to study conventional and unconventional monetary policy. Crucially, the analysis considers policy both during normal times and over increasing degrees of financial crisis, and studies how policy actions interact with disruptions in financial markets.

The empirical literature has highlighted the importance of financial frictions, and in particular market segmentation, for understanding unconventional policy (for example, D'Amico and King (2013), Hamilton and Wu (2012a), Gorodnichenko and Ray (2017)). Assuming that all borrowing takes place frictionlessly at the short (policy) rate is a useful simplification in many settings, but is too strong of an assumption for the purposes of this paper. Adding segmented bond markets to a macroeconomic model allows for more realistic and complicated dynamics in the determination of the term structure of interest rates. This enables the model to accomplish two goals: first, to match the relevant empirical findings regarding the term structure's response to demand shifts; and second, to study how these term structure changes interact with aggregate outcomes in general equilibrium.

The implications of the model are important for understanding the efficacy of monetary policy. When financial markets are healthy, so that marginal investors in financial markets have high risk-bearing capacity, the "expectation hypothesis" holds. That is, long-term rates are (roughly) the average of expected short rates. As a result, both conventional monetary policy and forward guidance are effective at stabilizing the economy. In this situation, household borrowing responds strongly to shifts in the path of the policy rate, leading to movements in output and inflation.

However, the link between expected short rates and the term structure is weakened when financial distress is high. As a result, long-term rates under-react to changes in the policy rate. Therefore, the model predicts that during a financial crisis output and inflation are less responsive to monetary shocks than usual.

Similar logic applies to LSAPs, but the implications are precisely the opposite. Purchases of long-term bonds, as in the various rounds of quantitative easing (QE), will have little to no effect on long-term rates when financial markets are healthy. When the central bank purchases a large amount of debt securities on the secondary market, the purchases change the portfolio allocation of the marginal investors in the debt market. Effectively, QE purchases allow financial investors to offload a source of risk from their portfolios. If financial markets are healthy, these investors are not very concerned with this source of risk to begin with, and so they do not require much excess returns to hold these securities. In this case, policies like QE will have little effect. But as financial markets become unable to bear risk, these purchases may matter a great deal. The shifts in the bond holdings alter the overall riskiness of these portfolios. By changing the riskiness of marginal investors' portfolio allocations, QE leads to changes in equilibrium prices of bonds, which in turn feed back into the household borrowing decision. This general equilibrium channel is akin to the familiar household Euler equation; under the right conditions, QE can boost output and stabilize the economy.

Benchmark models are not amenable to studying unconventional policies. Because of the extreme forward-looking behavior of agents and the lack of any financial frictions in such models, forward guidance policies are highly effective to the point of implausibility. Conversely, LSAP policies are completely ineffective in conventional models. The presence of limited arbitrage in my framework breaks this tight link. But it will still be the case that the term structure is rendered arbitrage-free, so that there are no riskless trades left on the table. Any deviations from the expectations hypothesis are due to the risky portfolio allocations chosen by financial market arbitrageurs (the marginal investors in the model). Hence, the channel through which unconventional policies like QE can have aggregate effects is by changing the market prices of risk.

As always, one fundamental source of risk is the movement in the short (policy) rate that is set by the central bank. All bonds are exposed to this risk, so as long as arbitrageurs are not perfectly risk-neutral there will be deviations from the expectations hypothesis. All else equal, arbitrageurs will require excess expected returns in order to take non-zero positions in long-term debt. This effect weakens the strength of forward guidance, but opens the door for LSAPs. The central bank is able to change the portfolio allocations of arbitrageurs, which through changes in the price of risk lead to changes in interest rates.

The exact impact of LSAP programs depends on how the purchases are structured. The amounts to be purchased, which maturities are targeted, and the duration of the program all affect the interaction with the sources of risk in the economy and the broader feedback mechanisms in the macroeconomy. The model delivers interesting and important interactions in general equilibrium. Because the policy rate responds to shifts in output and inflation, the expected path of short rates is a function of future expected output and inflation. When financial markets exhibit imperfect risk-bearing capacity, there is not a perfect link between longer-term rates and the expected path of short rates. Since these long-term rates affect household borrowing and hence influence output and inflation, the model exhibits rich feedback mechanisms; moreover, the dynamics of the model depend crucially on the health of financial markets. On the other hand, conditional on the term structure dynamics, the aggregate dynamics of the model stay close to benchmark models. The model adds only a handful of additional endogenous parameters which differentiate it from more familiar "three-equation" New Keynesian models. Therefore, the model is amenable to closed-form analysis.

I lay out the main building blocks of my "New Keynesian preferred habitat" framework in Section 1.2. Section 1.3 considers the case where prices are fully rigid. This is of course extreme and rules out important dynamics. However, many of the results can still be obtained and this simplifying assumption allows for a clearer focus on the intuition for the results. The main benefit is that this simplification rules out interesting but tricky determinacy issues (to which I return to later). In the most basic setup, the central bank sets the policy rate according to a Taylor-type rule subject to shocks. This is the only source of uncertainty. I then consider two extensions: first, I study forward guidance by assuming the central bank announces a path of policy rates; second, I study QE by allowing the central bank to directly purchase long-term bonds in the secondary market. The analysis confirms the intuition described above: conventional monetary policy and forward guidance become less effective as financial markets become disrupted (in the sense of both moving long-term rates and of impacting output); while at the same time, LSAP policies become more effective.

Next, Section 1.4 allows for prices to be sticky but not fully fixed. The main results go through here, but only if a determinacy condition is met. This condition is similar to the standard Taylor principle in textbook models, but with a key difference: the determinacy condition depends on the health of financial markets. A novel implication is that as financial markets become more disrupted, the model moves toward the region of indeterminacy. To the extent that model indeterminacy is either a proxy or a cause of excess volatility, this result shows how a purely financial crisis can lead to macroeconomy instability.

The focus of Sections 1.3 and 1.4 is delivering analytical results, but this comes at a cost of empirical realism. Section 1.5 extends the model to allow for many sources of aggregate and financial shocks in order to better match the data. In this section I develop the tools to solve the model numerically and estimate the model using U.S. data from before and during the recent financial crisis. The results confirm the qualitative findings of the more parsimonious models: monetary policy becomes less effective during financial crises; QE becomes more effective. Quantitatively, the model predicts that the aggregate output effects of a policy like the first round of QE were roughly 40% larger than a 50 basis point expansionary monetary shock during a period of relative financial calm. Further, had the zero lower bound not been binding during the financial crisis, additional rounds of rate cuts would have been 20% less effective than rate cuts during normal times.

Section 1.5 also studies more complicated LSAP programs like Operation Twist (where the Federal Reserve bought long-term debt and sold short-term debt). The model predicts these policies may be effective, depending on which maturities are targeted for purchase and the overall structure of risk in financial markets. In particular, when financial markets are relatively healthy, Operation Twist will have the net effect of pushing down interest rates across the entire term structure. However, the estimated model shows that when financial frictions are very high, Operation Twist-style policies will push down long-term rates but push up short-term rates. To the extent output is most sensitive to intermediate yields, this can have the net effect of raising effective borrowing rates of households, leading to contractionary outcomes.

I also show how determinacy can be restored if the central bank follows an endogenous rule for QE purchases. The main result is that, as financial frictions increase, the standard Taylor rule is less effective at stabilizing the economy. Formally, a Taylor rule eventually becomes unable to guarantee determinacy (for any parameterization).

But the same forces that make standard policy ineffective also make a QE rule more effective, hence this carries some of the weight of the determinacy issues and restores stability. Finally, Section 1.5 also conducts optimal policy experiments and finds that the endogenous policy response to inflation should become more aggressive in financial crises. Section 3.5 discusses additional extensions of the model and concludes.

This paper makes a number of contributions to the literature. One theoretical contribution is to extend the logic of Vayanos and Vila (2009) to a general equilibrium macroeconomic setting. This setup is amenable to analyzing state-dependent responses to policy changes while maintaining a relatively tractable framework. The paper adds to the large literature exploring the importance of macroeconomic factors in explaining the term structure, such as Ang and Piazzesi (2003). Other papers that tie reduced-form term structure modeling to New Keynesian macroeconomic dynamics include Hördahl et al. (2006) and Rudebusch and Wu (2008). My model contributes to this literature and can be viewed as a microfoundation for an affine term structure in macroeconomic factors.

From a partial equilibrium perspective, preferred habitat models can rationalize the interest rate response of QE but are silent on the aggregate effects of QE on inflation and output. Moreover and crucially, even the term structure response is conditional on the path of the short rate and independent of all other possible macroeconomic determinants of the term structure. My model is able to study both the direct and indirect effects of QE on the term structure and the transmission channels to the aggregate economy. Beyond that, the model makes additional important predictions: the frictions that imply QE is effective also imply that monetary policy conducted through changes in the short rate is less effective, and leads to increased aggregate instability. The model also highlights the centrality of these financial frictions as opposed to the zero lower bound constraint on conventional policy. The zero lower bound is neither necessary nor sufficient: QE is effective even away from the ZLB, but only when financial markets are imperfect; conversely, even at the ZLB, if financial markets are healthy then QE will have no effect.

From a theoretical perspective, this paper also introduces a relatively tractable model in which aggregate demand explicitly depends on long-term rates, and demonstrates how to solve the model with or without the expectations hypothesis. The model is also one in which aggregate dynamics are approximated linearly while still demonstrating sensitivity to risk.

I focus on limits to arbitrage and preferred habitat as important mechanisms for understanding conventional and unconventional monetary policy, based not only on empirical work looking at QE, but also studies of the determinants of the yield curve more generally (e.g. D'Amico and King (2013), Hamilton and Wu (2012a), Gorodnichenko and Ray (2017), Greenwood and Vayanos (2014), Beraja et al. (2015)). In this way, my paper adds to the literature studying how market segmentation interacts with unconventional monetary policy (e.g. Alvarez et al. (2002), Gertler and Karadi (2013), Chen et al. (2012), Carlstrom et al. (2017)). This overcomes the irrelevance results derived in Wallace (1981) and contrasts with an alternative view that treats QE as a signalling tool of the central bank (e.g. Christensen and Rudebusch (2012), Bauer and Rudebusch (2014), or Bhattarai et al. (2015)).

In addition, I move beyond studying QE and explore the implications of bond market frictions for monetary policy more broadly defined. Hence, the paper also falls into a broad class of macroeconomic models focusing on financial frictions (for a recent paper see Adrian and Duarte (2018); see Brunnermeier et al. (2012) for a survey). The approach in this paper differs from recent work focusing on borrowing constraints on the part of households (e.g. Kaplan et al. (2018)). The model also has a similar flavor as recent work which focuses on breaking the tight link between the path of future expected policy and current economic responses (e.g. McKay et al. (2016), Farhi and Werning (2017), Gabaix (2016), Angeletos and Lian (2018)). These papers make agents less responsive to expected future shocks. My approach has similar implications, but the degree of under-reaction is governed by the risk-bearing capacity of financial markets. In other words, the key frictions I focus on are those which mitigate the transmission of monetary policy through financial markets to the broader macroeconomy.

1.2 A New Keynesian Preferred Habitat Framework

I work with continuous time New Keynesian models that are largely characterized by a Phillips curve (relating current inflation π_t to current and future output gaps x_t), an IS curve (relating output growth to the real borrowing rate), and a monetary policy rule that governs how the nominal policy rate r_t reacts to macroeconomic variables. The setup is similar to Werning (2011).

The key difference between my model and a textbook New Keynesian model is that output depends on some "effective" nominal borrowing rate that depends not only on short rates, but also on longer rates. Assume that there is a continuum of zero-coupon nominal bonds with maturities $\tau \in (0, T]$, with time t price $P_{t,\tau}$ and yield given by

$$R_{t,\tau} = -\frac{\log P_{t,\tau}}{\tau}.$$

I assume that the effective nominal rate is

$$\tilde{r}_t \equiv \int_0^T \eta(\tau) R_{t,\tau}\, \mathrm{d}\tau \tag{1.1}$$

where $\eta(\tau)$ is a positive but otherwise arbitrary weighting function. This is a flexible way to allow for borrowing to depend explicitly on long-term interest rates. Assuming that all borrowing takes place at the short rate is a useful simplification but is too strong an assumption for the purposes of studying unconventional policies like forward guidance and QE. My specification aims to capture aspects of investment and savings decisions that are typically abstracted away in simple models. Although I do not explicitly model durable consumption or housing, in reality these are drivers of household borrowing that depend a great deal on long-term interest rates. Capital investment is also outside the model, but firm investment is similarly sensitive to long-term interest rates.[1]

Formally, these weights may arise for lifecycle borrowing reasons, or due to household's limited access to debt markets. Appendix A.3 presents microfoundations based on the latter setup. I relegate the derivations to the appendix due to the similarity with the derivation of a benchmark New Keynesian model and start the exposition with the familiar linearized aggregate equations governing the dynamics of the macroeconomy. The IS curve is modified such that the output gap evolves according to

$$\mathrm{d}x_t = \varsigma^{-1} \left(\tilde{r}_t - \pi_t - \bar{r}\right) \mathrm{d}t \tag{1.2}$$

where ς^{-1} is the intertemporal elasticity of substitution and $\bar{r}$ is the "natural" real borrowing rate (assumed constant in this model). As in a benchmark model, changes in the output gap are increasing in the nominal borrowing rate and decreasing in inflation. The only difference is that now the growth rate of output depends explicitly on the entire term structure of interest rates.

The Phillips curve, relating current inflation to current and future output, is unchanged relative to a standard New Keynesian model. The dynamics of inflation are governed by

$$\mathrm{d}\pi_t = \left(\rho\pi_t - \delta x_t\right) \mathrm{d}t \tag{1.3}$$

where ρ is the discount rate and δ governs the degree of price stickiness. $\delta \to \infty$ implies fully flexible prices, while $\delta \to 0$ implies fully rigid prices.

Finally, the central bank controls the instantaneous short rate r_t (the policy rate). At first, I will assume that this takes the form of a Taylor-type rule where the policy rate reacts to current levels of output and inflation subject to shocks. When studying forward guidance, I will also consider models in which the central bank announces a path of the policy rate.

For modeling the macroeconomic data generating process, I attempt to stay as close as possible to a benchmark "three equation" New Keynesian model for two reasons. First, even in partial equilibrium solving the term structure model becomes quite complicated. Embedding this setup in a dynamic general equilibrium model adds to the complexity. Keeping the model tractable and deriving analytical results is only possible if the underlying dynamics of the macroeconomy are kept simple. Second, because the mechanisms of simple New Keynesian models are still present in

[1]See Kaplan and Violante (2014) for a discussion of household portfolio allocations across short- and long-term securities in the United States. Note that I focus purely on maturity, rather than differences between liquid and illiquid savings vehicles.

more sophisticated models, it's reasonable to expect that the findings of my model will hold in more detailed extensions.

Before diving into the details, some intuitive results are immediate. The central bank sets only the short rate, while the rest of the term structure of interest rates is an equilibrium object. Since the effective borrowing rate depends on the entire yield curve, specifying only a rule for the policy rate does not necessarily close the model. However, if the expectations hypothesis holds, then long-term yields are fully determined by the expected path of short rates. In this case, without fully solving the model it's possible to see how unconventional policies will work. Forward guidance is powerful: by announcing a path of short rates, long-term rates will react immediately, and therefore the effective borrowing rate will also move sharply. This implies that consumption (and output) will also respond sharply. On the other hand, QE is ineffective: purchases of long-term bonds have no effect on the path of short rates, and hence do not change long-term yields.

But if the expectations hypothesis does not hold, a monetary policy rule no longer closes the model and it becomes necessary to specify how the entire term structure of interest rates are determined. For this purpose, I embed a "preferred habitat" model of the term structure, as in Vayanos and Vila (2009). Interest rates are determined by the interaction of two types of investors: clientele investors, who have idiosyncratic demand for bonds of specific maturities, and arbitrageurs with limited risk-bearing capacity, who integrate bond markets. The preferred habitat view of the term structure has long been of relevance to practitioners, but less so in academic models. Partly this is due to the fact that naïve forms of preferred habitat models conflict with no-arbitrage conditions: if the term structure were determined only by clientele investors with extreme preferences for bonds of specific maturities, bonds that are close to one another in maturity space could have large price differences. By allowing for arbitrageurs to integrate bonds of different maturities, the model avoids the unrealistic outcome of extreme segmentation and ensures that no-arbitrage conditions are satisfied. However, when arbitrageurs do not have perfect risk-bearing capacity, deviations from the expectations hypothesis arise.

I follow Vayanos and Vila (2009) in setting up the preferences for arbitrageurs and idiosyncratic preferred habitat investors. Arbitrageurs choose how much of each bond to hold (denoted by $b_{t,\tau}$) in order to maximize an instantaneous mean-variance trade-off of the change in wealth, subject to their budget constraint:

$$\max_{b_{t,\tau}} \; \mathrm{E}_t \, \mathrm{d}W_t - \frac{a}{2} \operatorname{Var}_t \mathrm{d}W_t$$
$$\text{s.t.} \;\; \mathrm{d}W_t = \left(W_t - \int_0^T b_{t,\tau} \, \mathrm{d}\tau \right) r_t \, \mathrm{d}t + \int_0^T b_{t,\tau} \frac{\mathrm{d}P_{t,\tau}}{P_{t,\tau}} \, \mathrm{d}\tau \, .$$

By holding $b_{t,\tau}$ of a τ bond, arbitrageurs receives the instantaneous return $\frac{\mathrm{d}P_{t,\tau}}{P_{t,\tau}}$. The remainder of their wealth not invested in long-term bonds is held at the risk-free rate

r_t (the short rate). The risk-aversion parameter $a \geq 0$ is fixed, but should be thought of as a proxy for limited risk-bearing capacity of financial markets.

The other side of the bond market is the demand from idiosyncratic clientele preferred habitat investors, which is given by

$$\tilde{b}_{t,\tau} = \alpha(\tau)\tau(R_{t,\tau} - \beta_{t,\tau}). \tag{1.4}$$

The function $\alpha(\tau) > 0$ is the semi-elasticity of preferred habitat demand (note $\tau R_{t,\tau} = -\log P_{t,\tau}$), and hence governs how sensitive these investors are to returns for τ bonds. This function is otherwise unrestricted; but the sign restriction implies that demand from preferred habitat investors is downward-sloping (increasing in yields). $\beta_{t,\tau}$ is a demand shifter, and can be thought of as a target yield. When rates are above the target, a τ investor increases their demand for τ bonds; and vice versa for when rates fall below the target. Note that, holding $R_{t,\tau}$ fixed, an increase in $\beta_{t,\tau}$ implies that τ investors *reduce* their holdings of τ bonds. I will consider different forms of this shifter (both deterministic and stochastic) in later sections. In equilibrium, bond prices must adjust so that arbitrageurs absorb the demand from preferred habitat investors ($b_{t,\tau} = -\tilde{b}_{t,\tau}$) while satisfying their mean-variance portfolio allocation problem.

This bond market setup is stylized: arbitrageurs are infinitesimally lived, and a τ-bond preferred habitat investor cares only about a specific slice of maturity space. Nevertheless, the model captures important facets of segmented markets, and how limited arbitrage smooths out idiosyncratic demand shocks. The preferred habitat setup is a natural way to study the affects of LSAP programs. Moreover, private investors such as pension funds and insurance companies often have demand for long-term bonds that arise from the need to match their long-term liabilities; these important sources of demand are not captured by intertemporal consumption substitution decisions that drive the term structure in more standard models. But these investors are not the only participants; otherwise debt markets would exhibit extreme segmentation. Arbitrageurs integrate debt markets and eliminate risk-free arbitrage opportunities, but are risk-averse and face limits to their trading activities. How conventional and unconventional policy affects the entire term structure will depend heavily on the limits to arbitrageurs' risk-bearing capacity.

In general, the term structure will be determined by complicated interactions between these two classes of investors and the general equilibrium dynamics of the macroeconomy. However, two limiting cases can be analyzed immediately. First, if arbitrageurs are risk-neutral ($a = 0$, so they only care about expected returns), then equilibrium can only be achieved if $\mathrm{E}_t\left[\frac{\mathrm{d}P_{t,\tau}}{P_{t,\tau}}\right] = r_t$. And if expected instantaneous returns of all bonds are equalized at the short rate, then risk-neutral arbitrageurs are indifferent between any bond allocation. In this case, they will absorb any demand shifts from preferred habitat investors without any equilibrium price changes. In other words, idiosyncratic demand shifts will not affect the term structure of interest rates.

In the other extreme, if arbitrageurs abandon the bond market (allocating the entirety of their wealth to the risk-free short rate), then equilibrium is only satisfied ($0 = b_{t,\tau} = -\tilde{b}_{t,\tau}$) if yields satisfy $R_{t,\tau} = \beta_{t,\tau}$. This would imply that bonds of very close maturity could have very different yields (and would potentially evolve unrelated to the short rate). Again, note that in this extreme case an increase in the demand shifter $\beta_{t,\tau}$ would push up the τ-bond yield. But this extreme segmentation does not occur in equilibrium because arbitrageurs do take non-zero positions in long-term bonds. The impact of changes in preferred habitat demand (if any) will depend on how arbitrageurs adjust their portfolio allocations. In turn, this will depend on the equilibrium dynamics of the short rate and other macroeconomics variables.

Intuitively, what does general equilibrium look like in this model? From the perspective of households, the key factor is how sensitive their effective borrowing rate is to the short rate. The model reduces to a benchmark New Keynesian model when these rates move one-for-one, but in general $\tilde{r}_t \neq r_t$. Suppose that the effective rate is highly responsive to the policy rate. Then household borrowing is also highly sensitive to the policy rate, and therefore the growth rate of consumption will also react strongly to the policy rate. On the other hand, when the effective rate is insensitive to the policy rate, the pass-through of changes in the policy rate to households is weakened. Through the borrowing decisions of the household, the growth rate of consumption is less responsive to the policy rate. That is,

$$\frac{\partial \, \mathrm{d}x_t}{\partial r_t} \propto \frac{\partial \tilde{r}_t}{\partial r_t}$$

and moreover, the sensitivity of the change in the output gap to the policy rate will determine the equilibrium reversion rate of monetary shocks. The higher the sensitivity, the quicker output gaps revert to steady state. Inflation will follow a similar path since, due to standard Phillips curve dynamics, inflation is the present discounted value of future output gaps. Thus, through the central bank's endogenous reactions to either output or inflation, the policy rate also reverts back to steady state quickly.[2]

However, the sensitivity of the effective borrowing rate to the policy rate is an *equilibrium* object, which also depends on financial markets. Bond prices will adjust in order to achieve equilibrium in bond markets, such that arbitrageurs' portfolio allocation satisfies their mean-variance tradeoff while also clearing the market given the demand from preferred habitat investors. In this model, arbitrage is imperfect and the term structure will not be characterized by the expectations hypothesis except under special circumstances. Therefore, it is the risk-adjusted dynamics of the

[2]The sensitivity of the change in the output gap to the policy rate plays a similar role as the intertemporal elasticity of substitution (ς^{-1}). In this case, the responsiveness of household borrowing to the policy rate is governed by preferences, namely the willingness to tolerate large changes in consumption across short periods of time. But the outcome is the same: when changes in the output gap are sensitive to the policy rate, the equilibrium rate at which shocks dissipate is high.

macroeconomy which determine bond prices in financial markets, rather than the actual dynamics of the short rate only. For example, in response to a monetary shock, the response of the yield of a τ bond will be roughly equal to the risk-adjusted average of the short rate over the life of the bond. Thus, if the short rate has a very high risk-adjusted mean reversion rate, long-term bond yields will not respond much to shocks to the short rate. This force implies that increases in the risk-adjusted reversion rate of monetary shocks lead to decreases in the sensitivity of all bond yields to the policy rate. In particular, the effective borrowing rate also becomes less responsive.

General equilibrium is obtained when these two forces balance. Thus, characterizing equilibrium involves two key steps: first, understanding the differences between the actual and risk-adjusted dynamics of the economy; and second, linking household savings and consumption choices with the bond prices determined in imperfect financial markets.

1.3 A Rigid Price Model

This section simplifies the model by assuming prices are fully fixed. This is an obviously extreme assumption, but many of the key results can still be obtained in this case. The upside is that the solution is considerably simpler, and I can sidestep determinacy issues that arise when prices can adjust. I return to these questions in Section 1.4.

As discussed in Section 1.2, the difficulty in solving the model relative to standard New Keynesian models is the mismatch between the effective borrowing rate and the policy rate. The approach I take to solving the model is as follows. First, start with the conjecture that bond prices are affine functions of the macroeconomic state variables. This allows the macroeconomic dynamics to be transformed into a system of linear differential equations. Next, conditional on the affine coefficients, solving for the rational expectations equilibrium is straightforward. Then, I turn to solving for these affine coefficients by solving the arbitrageur's portfolio problem. Finally, putting both sides of the model together, I characterize the unique general equilibrium solution.

After solving the model, I explore the implications for monetary policy. First I focus on conventional policy, where the central bank sets the policy rate through a Taylor rule. Then I study forward guidance by assuming the central bank instead announces a long-lived interest rate peg. Finally, I study LSAPs by allowing the central bank to purchase longer-term bonds on the secondary market. Throughout, I focus on how these policies affect the macroeconomic dynamics of the model, and in particular how these effects depend on the risk-bearing capacity of investors in financial markets.

1.3.1 Macroeconomic Dynamics

Prices are fully fixed when the parameter $\delta \to 0$. In this case, eq. (1.2) is simply

$$\mathrm{d}x_t = \varsigma^{-1} \left(\tilde{r}_t - \bar{r}\right) \mathrm{d}t \,. \tag{1.5}$$

Again, consumption (and output) growth is increasing in the borrowing rate, but now borrowing depends on some weighted average of the term structure of interest rates, given by eq. (1.1).

I assume the central bank follows a Taylor rule with persistence:

$$\mathrm{d}r_t = -\kappa_r(r_t - \phi_x x_t - r^*)\, \mathrm{d}t + \sigma_r \, \mathrm{d}B_{r,t} \,, \tag{1.6}$$

where $B_{r,t}$ is a standard Brownian motion and σ_r governs the size of the shocks. For now, I assume that changes in the short rate are the only source of uncertainty. ϕ_x govern the feedback rule for changes in output to changes in the policy rate. κ_r is a mean-reversion parameter; if $\kappa_r \to \infty$, eq. (1.6) simplifies to a (non-stochastic) Taylor rule with no gradual adjustments in the policy rate.[3] r^* is the central bank's target policy rate, which it sets in order to deliver a steady state with zero output gap. In a benchmark model where the borrowing rate is the same as the policy rate, this is accomplished by setting $r^* = \bar{r}$, but in this setup the optimal target is more complicated. I return to this in later sections.

Unlike a standard New Keynesian model, the interest rate rule does not close the model; it is necessary to specify how the entire term structure of interest rates is determined in equilibrium. I conjecture that the model features an affine term structure (which I will confirm in the next section):

$$-\log P_{t,\tau} = A_r(\tau) r_t + C(\tau).$$

Bond prices are sensitive to changes in the short rate; the sensitivity of a τ-maturity bond is governed by the coefficient function $A_r(\tau)$. Note this also implies the effective borrowing rate can be written as

$$\begin{aligned} \tilde{r}_t &= \left[\int_0^T \frac{\eta(\tau)}{\tau} A_r(\tau)\, \mathrm{d}\tau\right] r_t + \left[\int_0^T \frac{\eta(\tau)}{\tau} C(\tau)\, \mathrm{d}\tau\right] \\ &\equiv \hat{A}_r r_t + \hat{C} \end{aligned}$$

and hence the IS curve given by eq. (1.5) becomes

$$\mathrm{d}x_t = \varsigma^{-1} \left(\hat{A}_r r_t + \hat{C} - \bar{r}\right) \mathrm{d}t \,. \tag{1.7}$$

[3] The inertia term is present because empirically, central banks rarely change the policy rate by large jumps. Rather than assume that the policy rule is subject to long-lasting shocks, I instead assume that the policymaker prefers to smooth out changes over time (see Coibion and Gorodnichenko (2012) for an overview).

In terms of the macroeconomic dynamics, the coefficients $\hat{A}_r$ and $\hat{C}$ are the only difference between my model and a standard New Keynesian model; setting $\hat{A}_r = 1$ and $\hat{C} = 0$ recovers the standard dynamics. From the perspective of understanding aggregate dynamics, it is seen from eq. (1.7) that $\hat{A}_r$ is the key determinant of the responsiveness of consumption growth to the policy rate. Additionally, the coefficient functions $A_r(\tau)$ and $C(\tau)$ are equilibrium objects which will depend on the interplay of arbitrageurs and preferred habitat investors in imperfect financial markets. In particular, $A_r(\tau)$ governs the sensitivity of the price of a τ-bond to short-rate. Since the short rate is the only risk factor in the model, $\hat{A}_r$ is the weighted average of risk sensitivity of the entire term structure. In general equilibrium, $\hat{A}_r$ will have to satisfy both of these roles.

The affine functional form implies that the macroeconomic dynamics are governed by a linear stochastic differential equation. I solve for the rational expectations equilibrium following Buiter (1984), the continuous time analogue of Blanchard and Kahn (1980). In general, let $\mathbf{Y}_t = [\mathbf{y}_t \ \mathbf{x}_t]^T$ where $\mathbf{x}_t$ are the "jump" variables and $\mathbf{y}_t$ are the state variables. Writing the model in general matrix form (in terms of deviations from steady state) gives

$$\mathrm{d}\mathbf{Y}_t = -\Upsilon \left(\mathbf{Y}_t - \mathbf{Y}^{SS}\right) \mathrm{d}t + \mathbf{S}\, \mathrm{d}\mathbf{B}_t\,. \tag{1.8}$$

Under certain determinacy conditions, the rational expectations dynamics are given by

$$\mathrm{d}\mathbf{y}_t = -\Gamma \left(\mathbf{y}_t - \mathbf{y}^{SS}\right) \mathrm{d}t + \mathbf{S}\, \mathrm{d}\mathbf{B}_t \tag{1.9}$$

$$\mathbf{x}_t - \mathbf{x}^{SS} = \Omega \left(\mathbf{y}_t - \mathbf{y}^{SS}\right) \tag{1.10}$$

where Γ and Ω are given by eqs. (C10) and (C11).[4]

In the current rigid price model, the output gap x_t is the only jump variable, while the interest rate rule implies that r_t is the only state variable. The dynamics matrix Υ is given by

$$\Upsilon = \begin{bmatrix} \kappa_r & -\kappa_r \phi_x \\ -\varsigma^{-1}\hat{A}_r & 0 \end{bmatrix}.$$

Then the rational expectations equilibrium is determinate if and only if Υ has one stable eigenvalue λ_1; the general equilibrium dynamics simplify to

$$\begin{aligned} \mathrm{d}r_t &= -\lambda_1 (r_t - r^{SS})\, \mathrm{d}t + \sigma_r\, \mathrm{d}B_{r,t} \\ x_t - x^{SS} &= \omega_x (r_t - r^{SS}). \end{aligned}$$

The equilibrium mean-reversion properties of the short rate are governed by λ_1, while the equilibrium response of the output gap to changes in the short rate are characterized by ω_x. The following Lemma characterizes the conditions under which the

[4]With an abuse of notation, I use the same symbols for the shocks $\mathbf{S}$ and $\mathbf{B}_t$ in the general dynamics eq. (1.8) and the state space representation eq. (1.9).

model is determinate under rational expectations. Additionally, it characterizes the relationship between $\hat{A}_r$, λ_1, and ω_x. All proofs are in Appendix C.1.

Lemma 1.1 (Characterizing $\hat{A}_r$, rigid prices)**.** *Consider the rigid price model.*

1. *Υ has exactly one eigenvalue with positive real part if and only if $\hat{A}_r > 0$. Further, this stable root is real: $\lambda_1 > 0$.*
2. *$\hat{A}_r = h(\lambda_1)$ where $h : \mathbb{R}_+ \to \mathbb{R}$:*

$$h(\lambda) \equiv \frac{\lambda(\lambda - \kappa_r)}{\varsigma^{-1}\kappa_r\phi_x}. \tag{1.11}$$

3. *The output gap dynamics are given by*

$$\omega_x = -\frac{\varsigma^{-1}\hat{A}_r}{\lambda_1} = \frac{\kappa_r - \lambda_1}{\kappa_r\phi_x}.$$

The first result in Lemma 1.1 says that the model is determinate when consumption growth moves in the same direction as changes in the policy rate. This is a natural conjecture, which I will confirm holds in general equilibrium. Since prices are fully fixed, changes in the policy rate coincide with changes in the real short rate. Therefore, through standard intertemporal decisions of the household, one would expect that a higher real borrowing rate would lead to higher savings and thus increasing consumption into the future. However, this result is not immediate in my model because of the disconnect between the policy rate and the effective borrowing rate (even when nominal and real rates coincide). To show that this result holds requires explicitly determining the term structure of interest rates.

The second and third results characterize $\hat{A}_r$ and ω_x in terms of the macroeconomic parameters and the equilibrium eigenvalues of the model. Note that when the model is determinate, $\omega_x < 0$; this says that a positive (contractionary) shock to the policy rate leads to an immediate decline in the output gap.

As discussed previously, $\hat{A}_r$ is pulling double duty: it governs both the sensitivity of the term structure to risk, as well as the sensitivity of changes in the output gap to the policy rate. Lemma 1.1 characterizes $\hat{A}_r$ in terms of the latter interpretation. Thus, the function $h(\cdot)$ should be thought of as a mapping between $\frac{\partial \mathrm{d}x_t}{\partial r_t}$ (how sensitive the output gap is to the policy rate) and λ_1 (the equilibrium mean reversion rate of monetary shocks). This result, together with the results in the next section, are used to solve for $\hat{A}_r$ in general equilibrium.

Focusing on the determinate case, a high degree of sensitivity of consumption growth to the policy rate is associated with high mean reversion of the policy rate. Intuitively, in this case borrowing decisions are highly responsive to the policy rate. An increase in the level of the policy rate implies that households increase savings significantly; this leads to an immediate drop in the level of consumption (hence output) and then a rapid increase in consumption growth leading back to steady state. Through the endogenous response of the monetary rule, this implies that the policy rate also quickly returns to steady state. That is, a high $\hat{A}_r$ is associated with high λ_1 and large (negative) ω_x.

Conversely, as the sensitivity of output growth falls ($\hat{A}_r$ decreases towards zero), so does the sensitivity of the level of the output gap (ω_x approaches zero). Additionally, the equilibrium reversion rate of the short rate λ_1 approaches κ_r, the inertial term of the Taylor rule. Note that $h(\lambda_1)$ is negative when $0 < \lambda_1 < \kappa_r$. That is, model indeterminacy implies that in equilibrium monetary shocks mean revert *slower* than the inertial term of the Taylor rule κ_r.

The general equilibrium value of the sensitivity of output growth and the mean reversion rate of monetary shocks are of course endogenous. However, studying how the dynamics of the model change as these parameters vary exogenously is illuminating. Because the model only features two macroeconomic variables, phase diagrams can be used to study the aggregate dynamics.

Figure 1.1 plots phase diagrams of the rigid price models for different values $\hat{A}_r$, the sensitivity of changes in the output gap to the policy rate. This leads to variation in the equilibrium mean reversion rate of monetary shocks λ_1. First, I consider a benchmark New Keynesian model where the Euler equation depends only on the short rate. Equivalently, in my model this corresponds to $\hat{A}_r = 1$; with the given parameterization I consider this implies $\lambda_1 \approx 0.9$. This model is shown in Panel A. Note that the stable arm (the solid blue line) slopes downwards; after a contractionary monetary shock (increase in the short rate), the output gap jumps down and then the economy moves along this arm back to steady state. The dashed and dot-dashed lines show example (unstable) trajectories for the initial conditions $r = -1, x = 0$ and $r = 0, x = 1$.

What happens if $\hat{A}_r$ is higher than this benchmark case? This would mean that in equilibrium, monetary shocks revert faster towards steady state. In Panel B, I set $\hat{A}_r = 2.5$, which implies from Lemma 1.1 that $\lambda = 1.1$. Now the stable arm (the solid blue line) is much steeper (more negatively sloped) than the benchmark case, meaning that monetary shocks move output more than the benchmark prediction. The opposite is true when equilibrium monetary shocks mean revert slower: for $\hat{A}_r = 0.2$ so that $\lambda_1 = 0.8$ in Panel C, the stable arm is flattened.

Finally, at this point I cannot rule out the theoretical possibility that the model is in the region of indeterminacy. In Panel D I set $\hat{A}_r = -0.3$; that is, output growth is decreasing as the policy rate increases. This implies that monetary shocks revert towards steady state slowly ($\lambda_1 = 0.6$). Now the "stable" arm is upward-sloping (which would imply a contractionary monetary shock increases output), but note that this is no longer the unique stable path. Any initial level of the short rate and the output gap will return to steady state. The trajectories from initial conditions $r = -1, x = 0$ and $r = 0, x = 1$ shown by the dotted and dash-dotted lines are seen to return to steady state, whereas in the previous examples these were unstable paths. This is

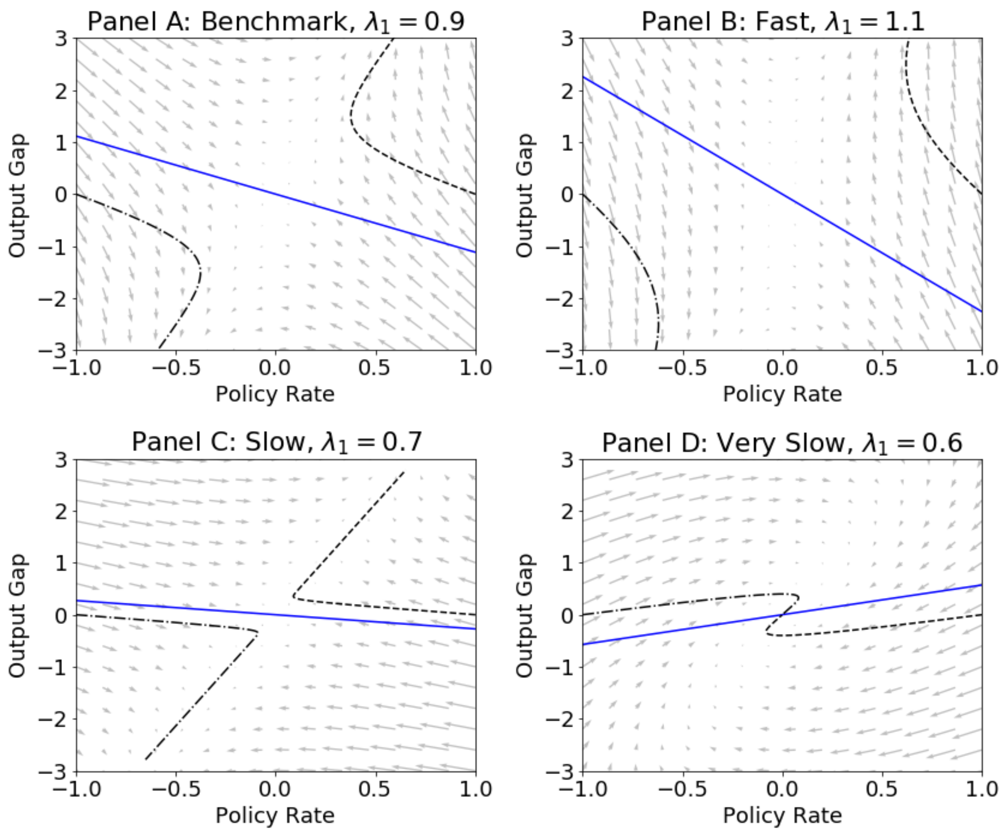

Figure 1.1: Phase Diagrams, Varying Output Sensitivity to Policy Rate

Notes: phase diagrams of the rigid price model (in terms of deviations from steady state). The solid blue line is the stable arm, while dashed black lines correspond to the trajectory of output and the policy rate given two different initial values. Each panel corresponds to different values of $\hat{A}_r$ and λ_1. Higher sensitivity of output to the policy rate ($\uparrow \hat{A}_r$) implies a faster mean equilibrium reversion of monetary shocks ($\uparrow \lambda_1$, the stable eigenvalue). The other parameters are set to $\kappa_r = 0.7$, $\phi_x = 0.25$, and $\varsigma^{-1} = 1$.

because, when $\hat{A}_r$ is negative, household borrowing moves in the opposite direction as the policy rate, leading to model indeterminacy.

Which of these cases will occur in general equilibrium? When I bring in the term structure side of the model, it will turn out that the relevant parameter space is one in which monetary shocks mean revert more slowly than a benchmark model (where the effective rate and policy rate coincide), but will still guarantee determinacy. In order to derive these results, I now turn to explicitly modeling the term structure.

1.3.2 Term Structure Determination

To study limited arbitrage, I embed a "preferred habitat" model of the term structure based on Vayanos and Vila (2009). Interest rates are determined by the interaction of two types of investors: clientele investors with idiosyncratic demand for bonds, and risk-averse arbitrageurs.

Given the affine functional form conjecture, I derive the optimality conditions of the arbitrageurs. Ito's lemma allows for the calculation of the instantaneous return of a τ bond, and hence the mean and variance of the change of arbitrageur's wealth. Arbitrageurs' expectations are rational and hence in equilibrium they expect the state variables to evolve according to eq. (1.9). For the case of a scalar state, the dynamics are given by

$$\mathrm{d}r_t = -\lambda(r_t - r^{SS})\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t}\,. \tag{1.12}$$

In equilibrium it will be the case that $\lambda = \lambda_1$, the (only) positive eigenvalue of the matrix Υ described in the previous section. However, since arbitrageurs take as given the dynamics of the short rate, it is useful to study how the term structure is determined for arbitrary dynamics of the short rate, governed by any $\lambda > 0$.

Lemma 1.2 (Arbitrageur optimality conditions, scalar state)**.** *Suppose the short rate is characterized by eq.* (1.12) *for some $\lambda > 0$. Then arbitrageurs choose a portfolio allocation such that*

$$\mu_{t,\tau} - r_t = A_r(\tau)\zeta_t \tag{1.13}$$

$$\zeta_t \equiv a\sigma_r^2 \int_0^T b_{t,\tau} A_r(\tau)\,\mathrm{d}\tau \tag{1.14}$$

where $\mu_{t,\tau}$ is the expected instantaneous return of a τ-maturity bond, given by eq. (A1)*.*

In this simple rigid price model, conditional on the dynamics of the policy rate the optimality conditions are the same as in Vayanos and Vila (2009). What this says is that a τ bond's expected excess return, $\mu_{t,\tau} - r_t$, is proportional to its sensitivity to the short rate, as measured by $A_r(\tau)$. This measure of proportionality is the same across all bonds, and follows solely from the absence of (risk-free) arbitrage.

No-arbitrage implies that the amount of excess return per unit of risk is the same for all bonds; this is ζ_t, the market price for risk. The expected excess return compensates arbitrageurs for taking on additional risk, which in this case comes from the short rate and so is fully characterized by the coefficient function $A_r(\tau)$. This compensation is higher when risk aversion is high (a), volatility is high (σ_r^2); or their portfolio is already sensitive to risk (the integral term in eq. (1.14)). The optimality conditions immediately imply that when arbitrageurs are risk neutral ($a = 0$), expected excess returns of all bonds are zero, and arbitrageurs are indifferent between holding any amount of bonds.

Arbitrageurs must hold the opposite of preferred habitat investors, whose demand is given by eq. (1.4). In this section, I assume that the demand shifter $\beta_{t,\tau}$ is independent of time and deterministic: $\beta_{t,\tau} \equiv \bar{\beta}(\tau)$. I will relax this assumption in later sections.

In equilibrium, prices must adjust such that arbitrageurs absorb the demand from preferred habitat investors and satisfy their mean-variance tradeoff. When arbitrageurs are risk-neutral and only care about expected returns, they will want to buy (sell) any bond with expected return greater than (less than) the risk-free rate. In this case, in equilibrium all bonds will have expected excess returns of zero and arbitrageurs will accomodate any shifts in demand from preferred habitat investors, recovering the results of a standard model of the term structure. But when arbitrageurs are risk averse, arbitrageurs will require non-zero excess expected returns to accomodate preferred habitat investors' demand. In this case, prices will depend on the arbitrageurs' portfolio allocations.

Thus, when $a > 0$, their portofolio is endogenous. The demand for bonds from preferred habitat investors along with the condition $b_{t,\tau} = -\tilde{b}_{t,\tau}$ leads to equilibrium conditions that characterize the coefficient function $A_r(\tau)$ and $C(\tau)$ in terms of the dynamics of the short rate.

Lemma 1.3 (Affine coefficients, scalar state term structure)**.** *Suppose the short rate is characterized by eq. (1.12) for some $\lambda > 0$. Then $A_r(\tau)$ is given by*

$$A_r(\tau) = \tau f\left(\nu(\lambda)\tau\right) \tag{1.15}$$

where $f(x) = \frac{1-e^{-x}}{x}$ and

$$\nu(\lambda) = \lambda + a\sigma_r^2 \int_0^T \alpha(\tau)\tau^2 f\left(\nu(\lambda)\tau\right)^2 \mathrm{d}\tau\,. \tag{1.16}$$

$C(\tau)$ is given by eq. (A2). Then $\hat{A}_r = g(\lambda)$ where $g : \mathbb{R}_+ \to \mathbb{R}$

$$g(\lambda) = \int_0^T \eta(\tau) f\left(\nu(\lambda)\tau\right) \mathrm{d}\tau\,. \tag{1.17}$$

The exponential function $f(x) = \frac{1-e^{-x}}{x}$ is a common functional form which occurs as solutions to differential equations such as these. The results from Lemma 1.3 help to characterize the responsiveness of the term structure to the policy rate. It follows that $A_r(\tau) > 0$, and moreover the affine functional form of bond prices implies the following:

$$\begin{aligned}\frac{\partial \log P_{t,\tau}}{\partial r_t} &= -A_r(\tau) = \frac{e^{-\nu\tau} - 1}{\nu} \\ \frac{\partial R_{t,\tau}}{\partial r_t} &= \frac{1}{\tau} A_r(\tau) = \frac{1 - e^{-\nu\tau}}{\nu\tau} \\ \mu_{t,\tau} - r_t = A_r(\tau)\zeta_t &= \frac{1 - e^{-\nu\tau}}{\nu}\zeta_t.\end{aligned}$$

Therefore, taking derivatives with respect to maturity τ gives:

$$\left|\frac{\partial^2 \log P_{t,\tau}}{\partial r_t \partial \tau}\right| > 0, \quad \frac{\partial^2 R_{t,\tau}}{\partial r_t \partial \tau} < 0, \quad \left|\frac{\partial \mu_{t,\tau} - r_t}{\partial \tau}\right| > 0.$$

That is, as maturity increases, bond (log) prices become more sensitive while bond yields become less sensitive to the short rate. And as maturity increases, excess expected returns grow in magnitude (the sign depends on ζ_t, the market price of short-rate risk).

Recall that from the perspective of aggregate dynamics, $\hat{A}_r$ is the key determinant of the sensitivity of output growth to the policy rate (determined by the function $h(\cdot)$ derived in Lemma 1.1). Additionally, from the perspective of term structure determination, $\hat{A}_r$ also governs the weighted average sensitivity of bond yields to short-rate risk. Lemma 1.3 shows that this is determined by the function $g(\cdot)$.

The key to characterizing the behavior of $g(\cdot)$ is the parameter ν. Intuitively, what is ν, and how does it compare to λ? λ governs the actual dynamics of the short rate, while ν governs the dynamics of the short rate under the risk-neutral measure. When risk aversion is non-zero, these parameters do not coincide. Eq. (1.16), which determines the parameter ν, is a fixed point problem and so will not have a simple solution except in the case when $a = 0$. But the proof of Lemma 1.3 shows that $\nu \geq \lambda$ and is increasing in λ, with equality if and only if $a = 0$. Since ν is increasing in λ, it follows that $g'(\lambda) < 0$. That is, bond yields become less responsive to policy rate movements when monetary shocks revert faster. Thus, the effective borrowing rate also becomes less sensitive.

The fact that $\nu \geq \lambda$ says that the risk-adjusted average of the short rate over the life of a long-term bond is lower than the expected average short rate over the same period. Thus

$$\frac{1}{\tau} E_t\left[\int_0^\tau \frac{\partial r_{t+u}}{\partial r_t}\,\mathrm{d}u\right] = f(\lambda\tau) \geq f(\nu\tau) = \frac{\partial R_{t,\tau}}{\partial r_t}.$$

The expectations hypothesis implies that these two responses should be identical, which occurs only when arbitrageurs are perfectly risk-neutral. Thus, this result says that long-term yields under-react to changes in the short rate.

Putting everything together, Lemma 1.3 shows how the interaction of arbitrageurs and preferred habitat investors characterizes the term structure. The arbitrageur's portfolio problem leads to a disconnect between the actual and risk-adjusted dynamics of the model. This leads to the under-reaction of longer-term rates (and therefore the effective borrowing rate) to changes in the policy rate. Further, if monetary shocks mean revert more quickly, then long-term rates are less responsive to the policy rate. In turn, this implies that the response of the effective borrowing rate to the policy rate is muted.

However, these are partial equilibrium results, which characterizes the affine coefficients in terms of the parameters from the term structure side of the model, taking as given the dynamics of the policy rate. The previous section documented a general equilibrium effect: increases in the sensitivity of household borrowing to the policy rate leads to longer-lasting monetary shocks. Formally, the results in these sections imply that increases in λ_1 lead to an increase in $h(\lambda_1)$ but a decrease in $g(\lambda_1)$. In general equilibrium, it must be that $\hat{A}_r = h(\lambda_1) = g(\lambda_1)$, so these forces must balance. The next section characterizes this interaction and solves for the general equilibrium solution.

1.3.3 General Equilibrium Solution

The results of Lemma 1.1 and 1.3 lead to a solution for $\hat{A}_r$, which characterizes general equilibrium. In this model, household borrowing decisions depend on the weighted average of longer-term yields. Hence consumption growth also depends on the weighted average of longer-term yields. Therefore, the sensitivity of output growth to the policy rate must coincide with the sensitivity of the effective borrowing rate to the policy rate. When output growth reacts strongly to the policy rate, monetary shocks mean revert quickly. On the other hand, due to the interaction of arbitrageurs and preferred habitat investors in imperfect financial markets, short-lived monetary shocks lead to a low degree of sensitivity of longer-term yields to the policy rate.

Recall $h(\cdot)$ determines the sensitivity of output growth to the policy rate, while $g(\cdot)$ determines the sensitivity of longer-term yields to the policy rate. Both are a function of the equilibrium rate of mean reversion of monetary shocks. In equilibrium, $\hat{A}_r$ will have to satisfy both of these jobs. Prop. 1.1 characterizes λ_1 and $\hat{A}_r$.

Proposition 1.1 (General equilibrium, rigid prices)**.** *Consider the rigid price model. There exists a unique positive eigenvalue of* Υ $\lambda_1 > 0$ *for which* $g(\lambda_1) = h(\lambda_1)$*, which fully characterizes the model equilibrium. Further, this implies* $0 < \hat{A}_r < 1$.

Figure 1.2 illustrates the equilibrium obtained in Prop. 1.1. In a benchmark New Keynesian model, implicitly it is always the case that $\hat{A}_r = 1$. This is because household borrowing takes place entirely at the short rate, so through the lens of my model the effective borrowing rate and the policy rate coincide. Once household borrowing depends on longer-term rates, this one-to-one correspondence breaks down.

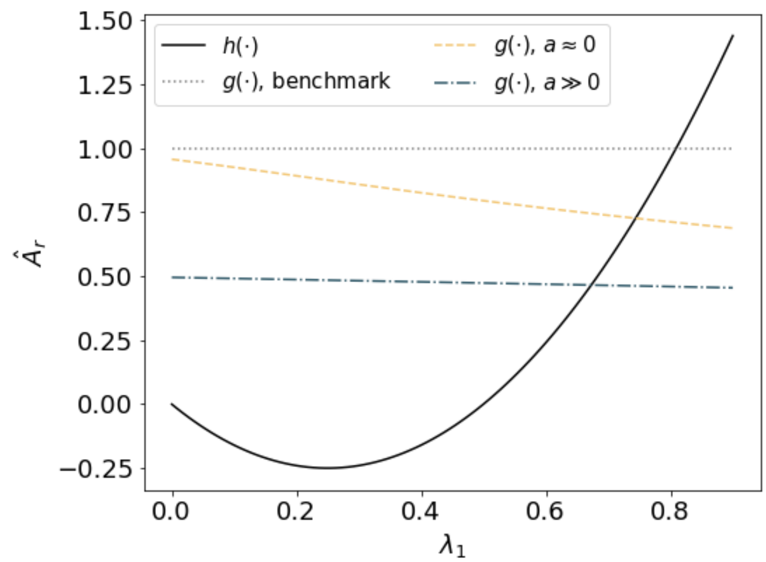

Figure 1.2: Intersection of $g(\lambda)$ and $h(\lambda)$, rigid prices

Notes: intersection of the functions $g(\lambda)$ and $h(\lambda)$, which determine λ_1 and $\hat{A}_r$ in equilibrium. The black line is $h(\cdot)$; the dashed light orange and dash-dotted dark teal lines are $g(\cdot)$ for low and high levels of risk aversion, respectively. The dotted grey line is the equivalent function $g(\cdot)$ for a benchmark New Keynesian model (fixed at unity). The parameters are set to $\kappa_r = 0.4$, $\phi_x = 0.25$, $\varsigma^{-1} = 1$, $T = 10$, $\alpha(\tau) = e^{-0.1\tau}$, $\sigma_r = 0.1$, $\eta(\tau)$ is the pdf of a (truncated) Gamma distribution, and $a \in \{0, 5\}$.

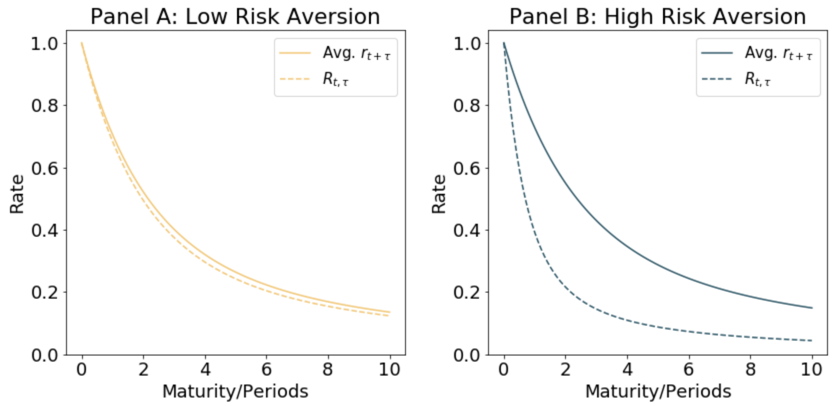

Figure 1.3: Rate Responses to Monetary Shock

Notes: responses of average short rates (solid lines) and spot rates (dotted lines) in response to a unit policy rate shock. The first panel shows the response when risk aversion is low ($a = 1$), while the second panel plots the responses when risk aversion is high ($a = 150$).

In particular, since long-term rates respond less than one-to-one with short rates, then $\hat{A}_r < 1$. Figure 1.2 also shows that in equilibrium, monetary shocks last longer than a benchmark model (lower λ_1). Since monetary shocks move long-term rates (and hence the effective borrowing rate) less than one-for-one, output also moves less. This means that the endogenous response of the central bank is muted relative to the benchmark, hence the shock mean reverts slower in equilibrium.

How does the sensitivity of long-term rates depend on the health of financial markets? Figure 1.2 shows that $\hat{A}_r$ (the weighted sensitivity of long-term rates to changes in the short rate) declines as risk aversion increases. In fact, the entire curve $g(\lambda)$ shifts down. Recall from the discussion of Lemma 1.3, there is a partial equilibrium effect that leads to an under-reaction of long-term rates to changes in the policy rate. Due to financial market imperfections, long-term rates under-react to monetary shocks relative to the predictions of the expectations hypothesis. Moreover, this under-reaction becomes more severe as the risk-bearing capacity of arbitrageurs declines. What are the general equilibrium implications of this under-reaction?

Figure 1.3 plots of the responses the short rate and long-term rates to a monetary shock, and explores the effects of increasing risk aversion graphically. Panel A corresponds to a low level of risk aversion, while Panel B sets a higher value of risk aversion. In both panels, the solid line is the average change in the average short rate over the course of τ periods, while the dotted line is the immediate change in the yield of a τ bond. Under the expectations hypothesis, the two responses would be identical. Hence as expected, in Panel A when risk aversion is close to zero, the responses are very similar. But in Panel B where risk aversion is high, the responses differ by quite a bit. The immediate response of long-term rates lies well below the expected path of average short rates (which is the response of long-term rates that would occur under the expectations hypothesis).

Interestingly, the response of the term structure under the expectations hypothesis differs between the two experiments. This is due to the fact that the expectations hypothesis implies that the response of long-term rates is solely determined by λ_1, the mean reversion of the short rate. But this is an equilibrium object, which depends on the risk aversion of arbitrageurs. In the parameterizations considered in these two experiments, in equilibrium $\lambda_1 = 0.74$ when risk aversion is low vs. $\lambda_1 = 0.67$ when risk aversion is high. This can be seen comparing the solid lines in each panel: the monetary shock lasts longer when risk aversion is high in the second panel compared to the first.

1.3.4 Conventional Policy

What are the implications for monetary policy, and how do they differ from benchmark models? Note that (from Lemma 1.3), the parameters governing the macroeconomic dynamics (for example, κ_r and ϕ_x) only enter through the eigenvalue λ_1. Fixing λ, the term structure side of the model is independent of these parameters. Similarly, from Lemma 1.1, the term structure parameters (for example, a and σ_r^2) only enter through the coeffcient function $A_r(\tau)$. Fixing $A_r(\tau)$, the aggregate dynamics of the model are independent of these parameters. Hence, for comparative statics, it is possible to make substantial progress despite the complexity of the model.

Corollary 1.1.1 (Comparative statics, rigid prices)**.** *Consider the rigid price model. In general equilibrium:*

1. $\frac{\partial \lambda_1}{\partial a} < 0$, $\frac{\partial \hat{A}_r}{\partial a} < 0$, $\frac{\partial \omega_x}{\partial a} > 0$. *Moreover,* $\hat{A}_r \to 0$, $\omega_x \to 0$, *and* $\lambda_1 \to \kappa_r$ *as* $a \to \infty$. *Further, if* $a \neq 0$, *the same results hold for* σ_r.

2. $\frac{\partial \lambda_1}{\partial \kappa_r} > 0$, $\frac{\partial \hat{A}_r}{\partial \kappa_r} < 0$, $\frac{\partial \omega_x}{\partial \kappa_r} > 0$.

3. $\frac{\partial \lambda_1}{\partial \phi_x} > 0$, $\frac{\partial \hat{A}_r}{\partial \phi_x} < 0$, $\frac{\partial \omega_x}{\partial \phi_x} > 0$.

4. *Consider two different weighting functions* $\eta^s(\tau)$ *and* $\eta^\ell(\tau)$, *such that for some* T^*, $\eta^s(\tau) \geq \eta^\ell(\tau) \iff \tau \leq T^*$. *Then* $\lambda_1^s > \lambda_1^\ell$, $\hat{A}_r^s > \hat{A}_r^\ell$, $\omega_x^s < \omega_x^\ell$ *where superscripts denote the equilibrium outcomes under the corresponding weighting functions.*

Note that when household borrowing depends on long-term interest rates (so $\eta(\tau) > 0$ for some $\tau > 0$) it will always be the case that $\hat{A}_r < 1$. This is because the effective borrowing rate does not respond one-to-one with the policy rate. Partly this is somewhat mechanical, and will be true even when arbitrageurs are risk neutral. In this case, the response of long-term rates to a change in the policy rate will be equal to the average change of expected future policy rates, and these shocks to short rates are not permanent but instead mean-revert.

The first result in Cor. 1.1.1 says something more interesting: as the risk aversion of arbitrageurs increases, household borrowing becomes less responsive to changes in the policy rate. That is to say, $\hat{A}_r$ is decreasing in the risk aversion of arbitrageurs. This occurs because long-term rates under-react to shifts in the expected path of short rates. In other words, now the response of long-term rates to a change in the policy rate will be smaller than the average change of expected future policy rates. This implies that current borrowing (and hence output) responds less than it would when financial markets exhibit perfect risk-bearing capacity.

The fact that current macroeconomic outcomes are less responsive to future expected policy changes bears a similarity to recent work that derives a "discounted Euler equation," as in Gabaix (2016) or Farhi and Werning (2017). But the mechanism here is not that, for behavioral reasons, household forecasts deviate from the actual paths of future outcomes. Rather, expected changes in policy are transmitted through imperfect financial markets, which endogenously decreases the responsiveness of current macroeconomic variables to future expected policy changes.

The model makes clear that monetary policy is effective only to the extent that policy changes are transmitted through financial markets. In the model, the health of financial markets is proxied by the risk aversion of arbitrageurs, which is a fixed parameter. But more generally, risk aversion may increase during periods of financial panics (e.g. as in Kyle and Xiong (2001a)). This implies that monetary policy becomes less effective during financial crises.

The next result relates to the persistence of the central bank's policy rule (governed by the mean reversion in the Taylor rule, κ_r). This governs the level of inertia in the central bank's policy rate (a higher value implies the rate returns to the target rate faster). Recall that λ_1 determines the equilibrium mean reversion behavior of the policy rate. So unsurprisingly, if the central bank reduces the inertia in its policy rule (increases κ_r) then the policy rate in equilibrium mean reverts faster (higher λ_1); because policy rate gaps persist for less time, the effective borrowing rate responds less to these monetary shocks (lower $\hat{A}_r$).

The intuition regarding the central bank's sensitivity to output (ϕ_x) is somewhat similar to the inertia parameter. The central bank responds more forcefully to output gaps, so in equilibrium the policy rate deviations subside faster (higher λ_1). Because they are shorter lived, output responds less as well.

How does the model depend on the weighting function, $\eta(\tau)$, that determines the effective borrowing rate? The final result in Cor. 1.1.1 answers this question. The two weighting functions correspond to two models where the effective borrowing rate is more geared towards short-term rates (η^s) or long-term rates (η^ℓ). The results says that as the effective borrowing rate becomes more dependent on long-term rates, the model is less sensitive to the policy rate. Further, in equilibrium this implies that monetary shocks persist longer.

1.3.5 Optimal Long-Run Monetary Target

Before turning to unconventional policy, I also solve for the "optimal" target in the central bank's policy rule r^*, which guarantees a steady state with zero output gap. In a benchmark model the central bank simply sets $r^* = \bar{r}$, the natural short rate. But when the effective borrowing rate depends on long-term rates, this is no longer the case. The central bank still should set their target to the natural short rate; but the natural short rate differs from the household's natural effective rate. Moreover, as with the transmission of monetary shocks, the optimal target also depends on the risk-bearing capacity of financial markets.

Corollary 1.1.2 (Optimal long-run target, rigid prices)**.** *Consider the rigid price model. Then the optimal target short rate that delivers $x^{SS} = 0$ is the "natural" short rate, given by*

$$r^* = \frac{\bar{r} - \hat{C}}{\hat{A}_r}. \tag{1.18}$$

Further, when $a > 0$, a higher level of habitat demand $\bar{\beta}(\tau)$ leads to decreases in the optimal target.

Note that the natural effective borrowing rate (which is the steady state value of the effective borrowing rate) is $\tilde{r}^{SS} = \bar{r}$ and is determined by factors outside of the control of the central bank (assumed to be fixed in this model). The effective borrowing rate is an affine function of the short rate (both in transition and in steady state), and hence the natural short rate is the steady state value of the short rate that delivers the natural effective rate:

$$\tilde{r}^{SS} = \bar{r} = \hat{A}_r r^{SS} + \hat{C}.$$

When arbitrageurs are not perfectly risk-neutral, the constant term is affected by shifts in habitat demand. Cor. 1.1.2 shows that as overall demand increases, the optimal central bank target falls whenever arbitrageurs are not perfectly risk-neutral.

1.3.6 Forward Guidance

Given the key under-reaction result when financial markets exhibit limited risk-bearing capacity, it is natural to expect that more explicit forward guidance policies will also prove less effective. This section studies this policy and confirms this intuition.

Instead of following a Taylor rule for setting the short rate, suppose instead that the central bank announces a target peg for interest rates $r^\diamond$, which will last for a set period of time $t^\diamond$ before returning to a standard Taylor rule. That is, the short rate evolves according to

$$\mathrm{d}r_t = \begin{cases} -\kappa_r^\diamond(r_t - r^\diamond)\,\mathrm{d}t + \sigma_r^\diamond\,\mathrm{d}B_{r,t} & \text{if } 0 < t < t^\diamond \\ -\kappa_r(r_t - \phi_x x_t - r^*)\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t} & \text{if } t \geq t^\diamond \end{cases}$$

and initially, the short rate at $t = 0$ is at the peg: $r_0 = r^\diamond$. Note that this setup implies that the peg is the target, but the policy rate may deviate from this target. The shocks to the short rate are again the only source of uncertainty in order to keep the solution tractable.

The output gap still evolves according to eq. (1.5), but the dynamics of the effective borrowing rate $\tilde{r}_t$ differ from the previous section. Conjecturing once again that bond prices are affine implies

$$-\log P_{t,\tau} = \begin{cases} A_r^\diamond(\tau)r_t + C^\diamond(\tau) & \text{if } 0 < t < t^\diamond \\ A_r(\tau)r_t + C(\tau) & \text{if } t \geq t^\diamond \end{cases} \implies \tilde{r}_t \;=\begin{cases} \hat{A}_r^\diamond r_t + \hat{C}^\diamond & \text{if } 0 < t < t^\diamond \\ \hat{A}_r r_t + \hat{C} & \text{if } t \geq t^\diamond \end{cases}.$$

Now there are two sets of affine coefficient functions to solve for, corresponding to the two different monetary regimes. But note that after the peg ends ($t \geq t^\diamond$), the model reduces to the one considered in Section 1.3. Further, during the peg ($0 < t < t^\diamond$), the results derived previously can be utilized to solve for the term structure coefficients. Lemmas 1.2 and 1.3 still apply, with $\lambda = \kappa_r^\diamond$.

With this characterization of the term structure during the two monetary regimes, I now turn to solving for the rational expectations equilibrium dynamics, and in particular the initial level of the output gap x_0.

Proposition 1.2 (Forward guidance, rigid prices)**.** *Consider the forward guidance rigid price model. In general equilibrium:*

1. $\frac{\partial x_0}{\partial r^\diamond} \leq 0$, *is increasing in* a, *and approaches* 0 *as* $a \to \infty$.

2. $\frac{\partial^2 x_0}{\partial r^\diamond \partial t^\diamond} \leq 0$, *is increasing in* a, *and approaches* 0 *as* $a \to \infty$.

As in a benchmark model of forward guidance, my model predicts that if the central bank sets a very low interest rate peg, then the output gap falls ($\frac{\partial x_0}{\partial r^\diamond} \leq 0$), and that this effect grows as the duration of the peg lengthens ($\frac{\partial^2 x_0}{\partial r^\diamond \partial t^\diamond} \leq 0$). The more interesting results from Prop. 1.2 is how these forces interact with the level of risk aversion of arbitrageurs. The first result says that the current output gap becomes less sensitive to the size of the forward guidance shock as the risk-bearing capacity of arbitrageurs falls. Moreover, output eventually becomes completely insensitive as arbitrageurs become infinitely risk averse. The second part of Prop. 1.2 is an analogous finding for the interaction of the duration of an interest rate peg and risk aversion. The effectiveness of lengthening the peg also diminishes as arbitrageur risk aversion increases, eventually becoming completely ineffective.

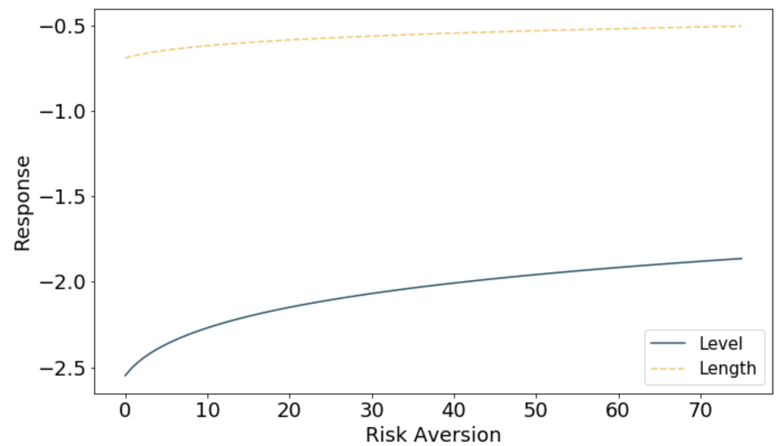

Figure 1.4: Output Responses to Forward Guidance

Notes: plots of $\frac{\partial x_0}{\partial r^\diamond}$ ("level"; the interaction of the level of the peg and output) and $\frac{\partial^2 x_0}{\partial r^\diamond \partial t^\diamond}$ ("length"; the interaction of the length of the peg and output). These objects are plotted for various levels of risk aversion (x-axis).

Figure 1.4 shows this graphically. The dark "level" line corresponds to $\frac{\partial x_0}{\partial r^\diamond}$, while the lighter "length" line corresponds to $\frac{\partial^2 x_0}{\partial r^\diamond \partial t^\diamond}$. As risk aversion increases, both of this effects are mitigated.

Intuitively, when financial markets are disrupted, the sensitivity of output to forward guidance falls. Note that in this model, households are still very forward-looking: households are still very responsive to far-off changes in their borrowing rates, and so conditional on the expected path of the effective borrowing rate the model delivers similar predictions as a benchmark New Keynesian model. The mitigating effect comes from the mismatch between the policy rate and the effective borrowing rate, which is bridged by imperfect financial markets.

1.3.7 Quantitative Easing

While forward guidance is less effective when financial markets are disrupted, this imperfection opens the door to LSAP policies. Given that the expectations hypothesis does not hold, purchases by the central bank may have price effects; moreover, it's natural to think that QE-type policies would push down long-term rates. This section studies this policy and confirms this result.

I now suppose that in addition to setting the short rate, the central bank also directly purchases longer term bonds through open market operations. In the model, these purchases take place through the demand shifter $\beta_{t,\tau}$ in preferred habitat investor demand given by eq. (1.4). I assume that

$$\begin{aligned} \beta_{t,\tau} &= \bar{\beta}(\tau) + \theta(\tau)\beta_t \\ \mathrm{d}\beta_t &= -\kappa_\beta \beta_t \, \mathrm{d}t \,. \end{aligned} \tag{1.19}$$

Note that this formulation treats LSAP programs as a zero-probability shock β_t to preferred habitat investor demand, which returns to zero at a rate according to κ_β. The function $\theta(\tau)$ governs where in maturity space the purchases are targeted. To capture the essence of a QE shock, I assume that $\theta(\tau) \geq 0$ for all maturities (and strictly positive for some maturities). This means that the central bank is only seeking to purchase positive amounts of long-term bonds. This rules out LSAP programs like Operation Twist; I return to this type of policy in an extension of the model in Section 1.5.[5]

Now the affine functional form of bond prices implies that

$$\begin{aligned} -\log P_{t,\tau} &= A_r(\tau) r_t + A_\beta(\tau)\beta_t + C(\tau) \\ \implies \tilde{r}_t &= \hat{A}_r r_t + \hat{A}_\beta \beta_t + \hat{C}, \end{aligned}$$

which introduces a new coefficient function $A_\beta(\tau)$.

This change complicates the model, because there are now two state variables. However, monetary policy remains the only source of uncertainty, so solving the model is similar to the previous case. First, I solve the macroeconomic dynamics, taking as given the affine coefficients. Write the model in matrix form according to eq. (1.8), where

$$\Upsilon = \begin{bmatrix} \kappa_r & 0 & -\kappa_r \phi_x \\ 0 & \kappa_\beta & 0 \\ -\varsigma^{-1}\hat{A}_r & -\varsigma^{-1}\hat{A}_\beta & 0 \end{bmatrix}.$$

Lemma 1.4 (Characterizing $\hat{A}_r$ and $\hat{A}_\beta$, rigid prices)**.** *Consider the rigid price QE model.*

1. *Υ has exactly two eigenvalues with positive real part if and only if $\hat{A}_r > 0$. Further, these stable roots are real. One of these eigenvalues is κ_β, the other is $\lambda_1 > 0$.*

2. *$\hat{A}_r$ is given by eq. (1.11).*

[5] Were the QE programs demand or a supply shocks? By treating the central bank as another "preferred habitat" investor, the model implicitly assumes that the QE purchases are (positive) demand shocks. But if the central bank is thought of as a conglomerate with the fiscal authority, then QE purchases are perhaps more naturally thought of as (negative) supply shocks: by buying bonds, the central bank acts on behalf of the fiscal authority and removes these bonds from the secondary market. Whichever the preferred labeling, in either case the outcome is the same. As the model makes clear, the effects are all about how the purchases impact the marginal investors, the arbitrageurs. Whether QE is thought of as a demand or supply shock, the result is that QE leads to changes in the portfolio allocation of arbitrageurs.

3. *The rational expectations equilibrium dynamic matrices (from eqs. (1.9) and (1.10)) are given by*

$$\Gamma = \begin{bmatrix} \lambda_1 & \frac{\hat{A}_\beta \varsigma^{-1} \kappa_r \phi_x}{\lambda_1 + \kappa_\beta - \kappa_r} \\ 0 & \kappa_\beta \end{bmatrix} \tag{1.20}$$

$$\Omega = \begin{bmatrix} \frac{\kappa_r - \lambda_1}{\kappa_r \phi_x} & \frac{\varsigma^{-1} \hat{A}_\beta}{\kappa_r - \lambda_1 - \kappa_\beta} \end{bmatrix}. \tag{1.21}$$

Next, given how the state evolves in general equilibrium, I solve the arbitrageurs' optimality conditions and equilibrium allocations, which solves for the affine term structure coefficient functions.

Lemma 1.5 (Affine coefficients, demand factor term structure)**.** *Suppose the short rate and demand factor are characterized by*

$$\begin{aligned} \mathrm{d}r_t &= -\left(\gamma_1(r_t - r^{SS}) + \gamma_{12}\beta_t\right)\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t} \\ \mathrm{d}\beta_t &= -\gamma_2\beta_t\,\mathrm{d}t\,. \end{aligned}$$

Then $\hat{A}_r$ is given by eq. (1.17) (and ν given by eq. (1.16)). $\hat{A}_\beta$ is given by

$$\hat{A}_\beta = \frac{\nu_{12}}{\nu - \gamma_2} \int_0^T \eta(\tau)(f(\nu\tau) - f(\gamma_2\tau))\,\mathrm{d}\tau \tag{1.22}$$

where the coefficient ν_{12} is given by eq. (A3).

Note that, conditional on the equilibrium value of λ_1, $\hat{A}_r$ is the same as the baseline rigid price model. This also implies that $\hat{A}_\beta$, the sensitivity of the effective borrowing rate to QE shocks, does not affect the determinacy of the model. In fact, if $\beta_t = 0$, then the dynamics of the model are identical to the baseline rigid price model. This is unsurprising, as in this case QE is a zero-probability shock and hence the model evolves as if QE will never occur.

Proposition 1.3 (QE, rigid prices)**.** *Consider the QE rigid price model. Suppose that $\theta(\tau) \geq 0$. In general equilibrium:*

1. *$\hat{A}_\beta \geq 0$ and $\frac{\partial x_t}{\partial \beta_t} \leq 0$, with equality if and only if $a = 0$.*

2. *$\hat{A}_\beta \to 0$ as $\kappa_\beta \to \infty$.*

The coefficient $\hat{A}_\beta$ determines the effects of QE in general equilibrium. The key is under what conditions this coefficient is non-zero, and what sign. In words, Prop. 1.3 says: if arbitrageurs are perfectly risk-neutral (financial markets exhibit perfect arbitrage), QE has no effect. This is the standard result: the expectations hypothesis holds, and since QE purchases do not change the expected path of short rates, there is no change in long-term rates. More explicitly, with perfect risk-neutrality, arbitrageurs (the marginal investors in bond markets) only care about expected instantaneous returns. In equilibrium it therefore must be the case that these are equalized for all bonds (and equal to the policy rate). When expected excess returns are always zero, arbitrageurs are happy to accomodate any shifts in demand, hence in equilibrium yields are unchanged. Since long-term rates are unaffected by QE purchases, there is no direct effect on household consumption or savings decisions, therefore no indirect effect on the expected path of the policy rate.

But whenever arbitrageurs are risk-averse, LSAPs push down interest rates and boost output. It is still the case that QE purchases do not have a (direct) effect on the expected path of the policy rate. But when arbitrageurs care about risk, expected excess returns of long-term bonds are not necessarily zero. Arbitrageurs demand compensation for taking on risk, and in equilibrium the market price of risk is not zero. The price of risk depends on the portfolio allocations of arbitrageurs, so the more concentrated the arbitrageurs' portfolio is in risky long-term bonds, the higher this compensation is required to be. By purchasing long-term debt, QE effectively reduces the amount of risk arbitrageurs are required to hold, which puts downward pressure on returns of all bonds.

This partial equilibrium effect is mitigated by a general equilibrium effect: when long-term rates fall, so does the effective borrowing rate of households. Through the standard Euler equation dynamics, this implies that consumption (hence output) will rise. This indirect effect puts countervailing upward pressure on long-term rates, as the expected path of short rates is higher than before. But since arbitrageurs are risk-averse, this upward pressure is weakened relative to the predictions of the expectations hypothesis. Prop. 1.3 shows that this indirect effect does not outweigh the direct effect, and it will still be the case that in general equilibrium QE purchases will push down effective borrowing rates, leading to an increase in output.

Together, these results show that the effectiveness of the two major unconventional monetary policy tools are mirror images of one another. In either case, passthrough to households only occurs to the extent that arbitrageurs respond to the policy changes. For forward guidance, healthy financial markets are key as arbitrageurs only care about the future path of the policy rate. But as financial markets become disrupted, arbitrageurs become more concerned with risk and less responsive to future changes in the short rate. Then this is precisely the time when LSAPs are most effective: by removing risk from the portfolio of arbitrageurs, these purchases push down long-term interest rates and boost output.

The second result in Prop. 1.3 shows that the effect of QE depends critically on the mean reversion properties of purchases. Even when financial markets are highly disrupted, the aggregate effects will be minimal if the purchases are undone very quickly. While in the case of QE these purchases were not directly unwound quickly, Greenwood et al. (2016) provide some evidence that Treasury actions had an equivalent effect. While QE was removing long-term debt from private portfolios, the Treasury was extending the average maturity of newly issued debt, perhaps partially offsetting the impact of QE.

1.4 Allowing for Sticky Prices

This section extends the analysis from Section 1.3 to allow for inflation. I confirm the results in the case of fully rigid prices go through when prices are sticky but not fully fixed. The main difference between the two models is the conditions for determinacy. In the rigid price model, determinacy in general equilibrium was guaranteed. Once prices are not fixed, determinacy is only guaranteed for some parameterizations of the model.

This result is also present in a benchmark New Keynesian model, and is often stated as the following: the central bank must move the nominal rate more than one-for-one with inflation, in order to move real rates. In benchmark models this is achieved by a simple inequality condition on the Taylor rule coefficients (frequently $\phi_\pi > 1$). But because aggregate dynamics depend on the household's real effective borrowing rate, the determinacy condition involves the entire term structure of interest rates. And since monetary policy is transmitted to the term structure through imperfect financial markets, this condition will implicitly depend on the risk-bearing capacity of arbitrageurs. As I will show, increasing limits to arbitrage moves the model towards the region of indeterminacy.

1.4.1 Macroeconomic Dynamics

Now I assume that prices are not fully rigid, so $\delta > 0$ and inflation and the output gap evolve according to eqs. (1.2) and (1.3). The central bank follows a Taylor rule with persistence:

$$\mathrm{d}r_t = -\kappa_r(r_t - \phi_\pi \pi_t - \phi_x x_t - r^*)\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t} \tag{1.23}$$

where $B_{r,t}$ is a standard Brownian motion. ϕ_π and ϕ_x govern the feedback rule for changes in inflation and output to changes in the policy rate, and κ_r is the mean-reversion parameter. r^* is the central bank's target policy rate, which is set to deliver zero inflation and output gap in steady state.

Again start with the conjecture that the model features an affine term structure in the state variables. Both inflation π_t and the output gap x_t are jump variables, while as before the interest rate rule implies that r_t is the only state variable, so the term structure is characterized by two coefficient functions $A_r(\tau)$ and $C(\tau)$. Writing the model in matrix form according to eq. (1.8) gives

$$\Upsilon = \begin{bmatrix} \kappa_r & -\kappa_r\phi_\pi & -\kappa_r\phi_x \\ 0 & -\rho & \delta \\ -\varsigma^{-1}\hat{A}_r & \varsigma^{-1} & 0 \end{bmatrix}.$$

The rational expectations equilibrium is determinate if and only if Υ has one stable eigenvalue λ_1; under rational expectations the dynamics of the short rate and inflation and the output gap are given by:

$$\begin{aligned} \mathrm{d}r_t &= -\lambda_1(r_t - r^{SS})\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t} \\ \pi_t &= \omega_\pi(r_t - r^{SS}), \quad x_t = \omega_x(r_t - r^{SS}). \end{aligned}$$

The following Lemma characterizes the equilibrium object $\hat{A}_r$ in terms of this eigenvalue.

Lemma 1.6 (Characterizing $\hat{A}_r$, sticky prices)**.** *Consider the sticky price model.*

1. Υ *has exactly one eigenvalue with positive real part if and only if*

$$\hat{A}_r > \frac{\delta}{\delta\phi_\pi + \rho\phi_x}. \tag{1.24}$$

2. $\hat{A}_r$ *is given by the function* $h : \mathbb{R} \to \mathbb{R}$*:*

$$h(\lambda_1) = \frac{(\lambda_1 - \kappa_r)(\lambda_1^2 + \lambda_1\rho - \varsigma^{-1}\delta)}{\varsigma^{-1}\kappa_r\left(\delta\phi_\pi + \rho\phi_x + \lambda_1\phi_x\right)}. \tag{1.25}$$

3. *The inflation and output gap dynamics are given by*

$$\omega_\pi = \frac{\delta(\kappa_r - \lambda_1)}{\kappa_r\left(\delta\phi_\pi + \rho\phi_x + \lambda_1\phi_x\right)}, \quad \omega_x = \frac{(\lambda_1 + \rho)(\kappa_r - \lambda_1)}{\kappa_r\left(\delta\phi_\pi + \rho\phi_x + \lambda_1\phi_x\right)}.$$

The macroeconomic dynamics continue to nest the benchmark New Keynesian model, where the affine coefficients are simply $\hat{A}_r = 1$ and $\hat{C} = 0$. When this is the case, if the central bank only cares about inflation (so $\phi_x = 0$), the determinacy condition eq. (1.24) simplifies to the standard condition that $\phi_\pi > 1$. But in this model, $\hat{A}_r$ is a general equilibrium object which will depend on the risk aversion of arbitrageurs. Thus, whether the model satisfies determinacy will also depend on the level of risk aversion.

Note that as $\delta \to 0$, all the above results simplify to what was found in the case of fully rigid prices studied in the previous section.

1.4.2 General Equilibrium Solution

The policy rate is the only state variable; hence, conditional on the equilibrium dynamics of the policy rate, the arbitrageur optimality conditions and the characterization of the affine coefficients is the same as when prices are fully fixed. Thus, the results of Lemma 1.6 and 1.3 together solve for $\hat{A}_r$ and λ_1 in general equilibrium. This also allows for a characterization of when the model has a unique equilibrium.

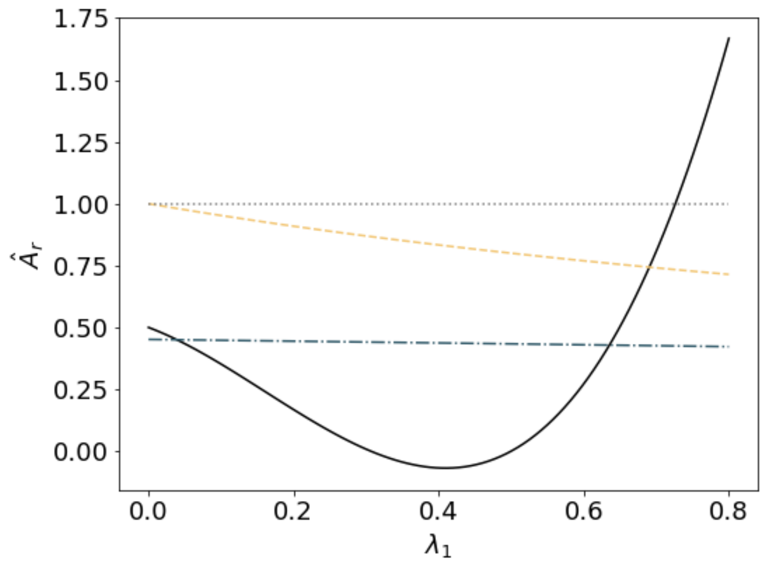

Figure 1.5: Intersection of $g(\lambda)$ and $h(\lambda)$, sticky prices
Notes: intersection of the functions $g(\lambda)$ and $h(\lambda)$, which determine λ_1 and $\hat{A}_r$ in equilibrium. The black line is $h(\cdot)$; the dotted light orange and dark teal lines are $g(\cdot)$ for low and high levels of risk aversion, respectively. The dotted grey line is the equivalent function $g(\cdot)$ for a benchmark New Keynesian model.

Proposition 1.4 (General equilibrium, sticky prices)**.** *Consider the sticky price model.*

1. *There exists some positive eigenvalue of* Υ $\lambda_1 > 0$ *for which* $g(\lambda_1) = h(\lambda_1)$*.*

2. *If the model is determinate (the inequality in eq. (1.24) is satisfied), then* $\hat{A}_r < 1$ *and is unique, which fully characterizes the model equilibrium.*

Note that if eq. (1.24) is violated, there exists another eigenvalue $\lambda_2 > 0$ such that $\hat{A}_r = h(\lambda_2)$. This implies the model is indeterminate. $\hat{A}_r$ may still be unique, but it may also be the case that there is some $0 < \lambda' < \lambda_1$ such that $g(\lambda') = h(\lambda')$. Hence, rather than attempting to define some selection criteria to re-establish determinacy, I choose to only focus on models that satisfy the conditions for determinacy.

Figure 1.5 illustrates the equilibrium obtained in Prop. 1.4. The dotted light orange and dark teal lines are $g(\cdot)$ for low and high levels of risk aversion, respectively. Unlike the rigid price model, this plot illustrates that for some parameterizations (in this case, when risk aversion is very high), equilibrium is no longer unique. Finally, I also plots the equivalent function $g(\cdot)$ for a benchmark New Keynesian model (the dotted grey line). Note that the determinacy condition of the benchmark model depends only on the properties of the function $h(\cdot)$ (which is determined by the macroeconomic parameters); financial frictions are not present.

1.4.3 Conventional Policy

This section studies how conventional monetary policy works in the sticky price model.

Corollary 1.4.1 (Comparative statics, sticky prices)**.** *Consider the sticky price model. In general equilibrium:*

1. $\frac{\partial \lambda_1}{\partial a} < 0,\ \frac{\partial \hat{A}_r}{\partial a} < 0,\ \frac{\partial \omega_\pi}{\partial a} > 0,\ \frac{\partial \omega_x}{\partial a} > 0.$
2. $\frac{\partial \lambda_1}{\partial \kappa_r} > 0,\ \frac{\partial \hat{A}_r}{\partial \kappa_r} < 0,\ \frac{\partial \omega_\pi}{\partial \kappa_r} > 0,\ \frac{\partial \omega_x}{\partial \kappa_r} > 0.$
3. $\frac{\partial \lambda_1}{\partial \phi_\pi} > 0,\ \frac{\partial \hat{A}_r}{\partial \phi_\pi} < 0,\ \frac{\partial \omega_\pi}{\partial \phi_\pi} > 0,\ \frac{\partial \omega_x}{\partial \phi_\pi} > 0.$
4. $\frac{\partial \lambda_1}{\partial \phi_x} > 0,\ \frac{\partial \hat{A}_r}{\partial \phi_x} < 0,\ \frac{\partial \omega_\pi}{\partial \phi_x} > 0,\ \frac{\partial \omega_x}{\partial \phi_x} > 0.$
5. *Consider two different weighting functions $\eta^s(\tau)$ and $\eta^\ell(\tau)$, such that for some T^*, $\eta^s(\tau) \geq \eta^\ell(\tau) \iff \tau \leq T^*$. Then $\lambda_1^s > \lambda_1^\ell$, $\hat{A}_r^s > \hat{A}_r^\ell$, $\omega_\pi^s < \omega_\pi^\ell$, $\omega_x^s < \omega_x^\ell$ where superscripts denote the equilibrium outcomes under the corresponding weighting functions.*

The results in Cor. 1.4.1 confirm the results of the rigid price model carry over to the case of sticky prices. The first result shows that the responsiveness of the (nominal) effective borrowing rate to changes in the short rate declines as risk aversion increases. In general equilibrium, when the model is determinate, this will imply that the real effective borrowing rate also becomes less responsive, hence inflation and output also respond less to changes in the short rate as risk aversion increases. This also implies that monetary shocks become less persistent. The final result, which shows what happens when household borrowing becomes more concentrated towards long-term rates, has similar results. Borrowing becomes less responsive to the short rate, and thus inflation and output also respond less to changes in the short rate. In equilibrium, this implies that monetary shocks are not as persistent.

The other set of results relate to the central bank's policy rule. Increases in the mean-reversion parameter (κ_r), the sensitivity to inflation (ϕ_π), and the sensitivity to output (ϕ_x) all cause shocks to the policy rate in equilibrium to return to steady state faster (higher λ_1) and therefore the (nominal) effective borrowing rate responds less to these monetary shocks (lower $\hat{A}_r$). In general equilibrium when the model is determinate, this also leads to smaller changes in inflation and output.

1.4.4 Determinacy

In the rigid price model, determinacy was guaranteed. This is no longer the case once prices are not fixed. Moreover, the determinacy condition is more complicated than a benchmark New Keynesian model.

Corollary 1.4.2 (Determinacy, sticky prices). *Consider the determinacy condition of the sticky price model. If $\delta > 0$, there exists some upper bound $\overline{a}$ such that, whenever $a > \overline{a}$, the model is indeterminate. Similarly, there are upper bounds for κ_r and σ_r (if $a \neq 0$) above which the model is indeterminate. For ϕ_π and ϕ_x there exists upper bounds above which the model is determinate.*

From a macroeconomic perspective, the reason for indeterminacy in this model is analogous to a standard model. Suppose inflation increases. The central bank responds by increasing the policy rate, but if the real borrowing rate faced by households does not rise then the policy response is unable to stabilize the macroeconomy. This logic holds in my model, but the relationship between the policy rate and the household effective borrowing rate is complicated by the fact that policy changes are passed through by risk-averse arbitrageurs.

The first result in Cor. 1.4.2 says that if financial markets become severely disrupted, the model moves into the region of indeterminacy. This is for the same reason that conventional policy becomes less effective: the passthrough of policy rate changes to the (nominal) effective borrowing rate becomes dampened, and this dampening can become severe enough that the real effective borrowing rate no longer moves to stabilize the economy.

On the other hand, the central bank can guarantee determinacy given any level of (finite) risk aversion by responding more aggressively to inflation (or output). Also, if the policy rule does not exhibit enough inertia (that is, the mean reversion rate κ_r is very high), this can also induce model indeterminacy. This is perhaps one theoretical reason why central banks seem to pursue inertial policy rules. A high degree of monetary policy inertia implies that changes in the policy rate are endogenously passed to long-term rates. When household borrowing depends on a mix of short- and long-term rates, this leads to larger responses of output and inflation.

What does it mean in reality for a macroeconomic model to exhibit indeterminacy? As discussed in e.g. Clarida et al. (2000), one practical way of thinking about indeterminacy is that it induces excess volatility. When policy is such that it gives rise to an indeterminate model, this opens up the possibility of self-fulfilling expectations. This may increase the amount of volatility in the model.

In the model, the parameters governing both the central bank's policy rule as well as the risk aversion of arbitrageurs are fixed across time. But stepping outside the model, what do the results in Cor. 1.4.2 say? One interpretation is that if financial crises are thought of as large disruptions in financial markets, with big increases in risk aversion of investors (or the risk-bearing capacity of investors), then financial crises can lead to macroeconomic instability, even if the cause of the crises was unrelated to the macroeconomy.

Central bankers can induce stability again, but they can do so by becoming more aggressive. This runs counter to the idea that central banks should become more passive in crises, and make up for it after the crisis passes.

1.4.5 Forward Guidance

This section extends the previous analysis of forward guidance to the case of sticky prices. As before, the central bank announces a target peg for interest rates $r^\diamond$, which will last for a set period of time $t^\diamond$ before returning to a standard Taylor rule. That is, the short rate is evolves according to

$$\mathrm{d}r_t = \begin{cases} -\kappa_r^\diamond (r_t - r^\diamond)\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t} & \text{if } 0 < t < t^\diamond \\ -\kappa_r (r_t - \phi_\pi \pi_t - \phi_x x_t - r^*)\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t} & \text{if } t \geq t^\diamond \end{cases}$$

and initially, the short rate at $t = 0$ is at the peg: $r_0 = r^\diamond$. Inflation and the output gap still evolve according to eqs. (1.2) and (1.3).

Since the policy rate is the only state variable, the affine functional form is the same as in the rigid price model. Using the same approach as before, I now turn to solving for the rational expectations equilibrium dynamics and the initial level of inflation π_0 and the output gap x_0.

Proposition 1.5 (Forward guidance, sticky prices)**.** *Consider the forward guidance sticky price model. In general equilibrium:*

1. $\frac{\partial \pi_0}{\partial r^\diamond} \leq 0$ *and* $\frac{\partial x_0}{\partial r^\diamond} \leq 0$*. Both are increasing in* a*, and approach* 0 *as* $a \to \infty$*.*
2. $\frac{\partial^2 \pi_0}{\partial r^\diamond \partial t^\diamond} \leq 0$ *and* $\frac{\partial^2 x_0}{\partial r^\diamond \partial t^\diamond} \leq 0$*. Both are increasing in* a*, and approach* 0 *as* $a \to \infty$*.*

Note that inflation and the output gap fall if the central bank increases the level of the peg, and this effect grows with the length of the peg. So the first result of Prop. 1.5 says that current inflation and output become less sensitive to the size of the forward guidance shock as the risk-bearing capacity of arbitrageurs falls, and eventually become completely insensitive; while the second result says the same regarding length of the peg.

Figure 1.6 shows the results regarding inflation and output graphically. The dark "level" line in the first panel corresponds to $\frac{\partial \pi_0}{\partial r^\diamond}$, while the lighter "length" line corresponds to $\frac{\partial^2 \pi_0}{\partial r^\diamond \partial t^\diamond}$. The second panel plots the same objects for the output gap x_0. As risk aversion increases, both of this effects are mitigated.

1.4.6 Quantitative Easing

I now modify the sticky price model to allow for QE shocks. Suppose that the central bank also purchases long-term bonds, so that the demand shifter in eq. (1.4) evolves according to eq. (1.19).

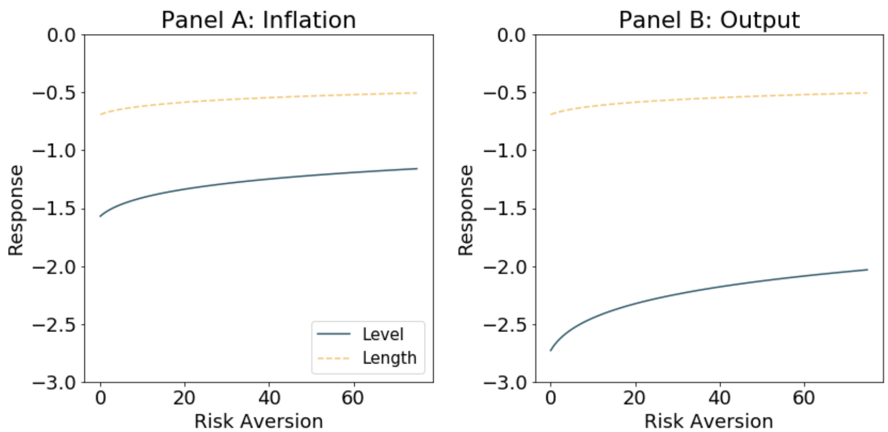

Figure 1.6: Inflation and Output Responses to Forward Guidance

Notes: panel A plots $\frac{\partial \pi_0}{\partial r^\diamond}$ ("level"; the interaction of the level of the peg and inflation) and $\frac{\partial^2 \pi_0}{\partial r^\diamond \partial t^\diamond}$ ("length"; the interaction of the length of the peg and inflation). Panel B plots the corresponding level and length interaction terms for output x_0. These objects are plotted for various levels of risk aversion (x-axis).

First, I solve the macroeconomic dynamics, taking as given the affine coefficients. Write the model in matrix form according to eq. (1.8), where

$$\Upsilon = \begin{bmatrix} \kappa_r & 0 & -\kappa_r \phi_\pi & -\kappa_r \phi_x \\ 0 & \kappa_\beta & 0 & 0 \\ 0 & 0 & -\rho & \delta \\ -\varsigma^{-1} \hat{A}_r & -\varsigma^{-1} \hat{A}_\beta & \varsigma^{-1} & 0 \end{bmatrix}.$$

Lemma 1.7 (Characterizing $\hat{A}_r$ and $\hat{A}_\beta$, sticky prices)**.** *Consider the sticky price QE model.*

1. *Υ has exactly two eigenvalues with positive real part if and only if the condition eq. (1.24) is satisfied. Further, these stable roots are real. One of these eigenvalues is κ_β, the other is $\lambda_1 > 0$.*

2. *$\hat{A}_r$ is given by eq. (1.25).*

3. *The rational expectations equilibrium dynamic matrices are given by eqs. (A5) and (A6).*

Given the equilibrium value of λ_1, $\hat{A}_r$ is the same as the sticky price model. The determinacy condition is also equivalent. As in the rigid price model, this is because I treat QE as a zero-probability event. Moreover, conditional on how the state evolves in general equilibrium, the arbitrageurs' optimality conditions and portfolio problem is the same as considered in Lemma 1.5. Therefore, putting everything together shows how QE shocks impact the economy in general equilibrium.

Proposition 1.6 (QE, sticky prices)**.** *Consider the QE sticky price model. In general equilibrium, the model is determinate then*

$$\hat{A}_\beta \geq 0 \implies \frac{\partial \pi_t}{\partial \beta_t} \leq 0, \; \frac{\partial x_t}{\partial \beta_t} \leq 0$$

with equality if and only if $a = 0$.

Prop. 1.6 confirms that in the sticky price model, when the determinacy condition is satisfied QE works in the same way as in the rigid price model. Expansionary QE shocks move both inflation and output in the same direction, but only when arbitrageurs are not risk-neutral.

1.5 General Numerical Model

The analysis thus far has focused on delivering analytical results, but this comes at the cost of realism. This section generalizes the model in order to move closer to the data and allows me to take a first step towards quantifying the effects of unconventional policies in and out of financial crises.

This section first extends the model to allow for a richer set of shocks and develops the tools to solve the model numerically. Since this requires taking a stance on parameter values, I next turn to estimating the model. I then use the estimated model to quantify the effects of unconventional policies. The extended model allows for the study of not only standard QE policies, but also more complicated LSAP programs such as Operation Twist, where the Federal Reserve bought long-term bonds while selling shorter term securities. I also study a counterfactual LSAP policy whereby the central bank conducts QE endogenously according to a Taylor-type of rule. Finally, I use the model to study optimal policy as a function of financial frictions.

1.5.1 Macroeconomic Dynamics and Term Structure Determination

I assume prices are sticky but not fully rigid. Besides monetary policy shocks, I add demand shocks and cost-push shocks. I also include shocks that shift demand for bonds coming from the idiosyncratic preferred habitat investors. This model allows me to explore the robustness of the results in the previous section, as well as explore the implications for unconventional monetary policy. Adding multiple demand factors not only makes the model more realistic, but also allows me to explore more complicated LSAP programs such as Operation Twist, where the Fed simultaneously purchased and sold Treasuries of long- and short-term maturities, respectively.

The aggregate dynamics of the extended model are as follows:

$$\mathrm{d}\pi_t = (\rho\pi_t - \delta x_t - z_{\pi,t})\,\mathrm{d}t \tag{1.26}$$
$$\mathrm{d}x_t = \varsigma^{-1}(\tilde{r}_t - \pi_t - \bar{r} - z_{x,t})\,\mathrm{d}t \tag{1.27}$$
$$\mathrm{d}r_t = -\kappa_r(r_t - \phi_\pi\pi_t - \phi_x x_t - r^*)\,\mathrm{d}t + \sigma_r\,\mathrm{d}B_{r,t}\,. \tag{1.28}$$

The new aggregate variables are a cost-push shock to the Phillips curve ($z_{\pi,t}$) and a demand shock to the Euler equation $z_{x,t}$.These shocks follow simple Ornstein-Uhlenbeck processes:

$$dz_{i,t} = -\kappa_{z_i} z_{i,t}\,\mathrm{d}t + \sigma_{z_i}\,\mathrm{d}B_{z_i,t}\,.$$

Additionally, I assume there are factors $\beta_{k,t}$ that affect the demand of preferred habitat investors. These shocks also follow simple Ornstein-Uhlenbeck processes:

$$d\beta_{k,t} = -\kappa_{\beta_k}\beta_{k,t}\,\mathrm{d}t + \sigma_{\beta_k}\,\mathrm{d}B_{\beta_k,t}\,.$$

This is the only change to the term structure side of the model: the demand shifter $\beta_{t,\tau}$ from idiosyncratic preferred habitat investors is given by

$$\beta_{t,\tau} = \bar{\beta}(\tau) + \sum_k \beta_{k,t}\theta_k(\tau)$$

where $\theta_k(\tau)$ governs how shifts in the demand factor $\beta_{k,t}$ affects the level of demand for τ bonds. Unlike the QE models explored in previous sections, the habitat demand factors are now stochastic.

As before, start with the conjecture that bonds are affine in the state variables. Now the state consists of the short rate r_t, the demand shock and cost push shocks $z_{x,t}$ and $z_{\pi,t}$, and the preferred habitat demand factors $\beta_{k,t}$. If $\mathbf{Y}_t$ is the vector of all variables and the state is denoted by a vector $\mathbf{y}_t$, then effective rate is therefore also affine in the state variables:

$$\tilde{r}_t = \mathbf{y}_t^T\hat{\mathbf{A}} + \hat{C}, \tag{1.29}$$

and the preferred habitat demand shifter is

$$\beta_{t,\tau} = \bar{\beta}(\tau) + \mathbf{y}_t^T\Theta(\tau)$$

where $\Theta(\tau)$ is a vector of that collects the $\theta_k(\tau)$ functions corresponding to each $\beta_{k,t}$ demand shocks (and is zero for the other state variables).

Using matrix notation, the dynamics matrix from eq. (1.8) is a function of the affine coefficients $\hat{\mathbf{A}}$ from eq. (1.29). Therefore, so is the state dynamics matrix $\Gamma(\hat{\mathbf{A}})$, given by eq. (C10). Prop. 1.7 characterizes the affine coefficients and general equilibrium solution in this setup.

Proposition 1.7 (General equilibrium characterization)**.** *Suppose the aggregate economy evolves according to eq.* (1.8)*, where the dynamics matrix is a function of affine coefficients* $\hat{\mathbf{A}}$ *in eq.* (1.29)*. Define the matrix*

$$\mathbf{M} = \Gamma(\hat{\mathbf{A}})^T - a\left[\int_0^T \alpha(\tau)\left(\tau\Theta(\tau) - \mathbf{A}(\tau)\right)\mathbf{A}(\tau)^T\,\mathrm{d}\tau\right]\boldsymbol{\Sigma} \tag{1.30}$$

as a function of $\hat{\mathbf{A}}$*. Letting* $\mathbf{e}_1$ *be the first standard basis coordinate vector, if* $\mathbf{M}$ *is diagonalizable and invertible then* $\mathbf{A}(\tau)$ *solves*

$$\mathbf{A}(\tau) = \mathbf{G}\mathbf{D}^{-1}\left[\mathbf{I} - \exp(-\mathbf{D}\tau)\right]\mathbf{1} \tag{1.31}$$

$$\implies \hat{\mathbf{A}} = \mathbf{G}\mathbf{D}^{-1}\int_0^T \frac{\eta(\tau)}{\tau}\left[\mathbf{I} - \exp(-\mathbf{D}\tau)\right]\mathrm{d}\tau\,\mathbf{1} \tag{1.32}$$

where $\mathbf{D}$ *is the diagonal matrix of the eigenvalues of* $\mathbf{M}$*, and* $\mathbf{G}$ *is the matrix of corresponding eigenvectors, normalized such that* $\mathbf{G}\mathbf{1} = \mathbf{e}_1$*, where* $\mathbf{1}$ *is a vector of ones and* $\mathbf{e}_1$ *is the first standard basis coordinates. Then the general equilibrium solution is a fixed point of the matrix function defined in eq.* (1.30)*.*

Note that the matrix $\mathbf{M}$ is nothing more than the multidimensional generalization of the scalar ν from Lemma 1.3. Thus Prop. 1.7 is similar to the fixed point problem which defines the parameter ν, but is now complicated by the fact that the problem is no longer scalar-valued. Except in special cases such as when $a = 0$, the problem no longer lends itself to tractable solutions but instead must be solved numerically. An algorithm for solving the model is described in Appendix A.2.

1.5.2 Calibration

In order to implement the numerical solution method I need to parameterize the model. I calibrate the model by separating the parameters into two groups: the macroeconomic dynamics parameters and the term structure preferred habitat parameters. The macroeconomic parameters consist of: the weighting function in the effective borrowing rate ($\eta(\tau)$); the preference parameters (ρ and ς^{-1}); nominal rigidity (δ); the Taylor rule coefficients (ϕ_π and ϕ_x); the mean reversion of monetary shocks, cost-push shocks, and demand shocks (κ_r, κ_{z_π}, and κ_{z_x}); and the volatility of these shocks (σ_r^2, $\sigma_{z_\pi}^2$, and $\sigma_{z_x}^2$). The term structure parameters consist of: arbitrageur risk aversion (a); preferred habitat demand elasticities ($\alpha(\tau)$); the number and location of demand factors ($\theta_k(\tau)$); and the mean reversion and volatility of these demand factors (κ_{β_k} and $\sigma_{\beta_k}^2$).[6]

[6]The remaining parameters $\bar{\beta}(\tau)$, $\bar{r}$, and r^* affect the steady state but play no role governing the dynamics of the model; I focus on equilibrium dynamics linearized around a zero steady state. Formally, I assume that given any parameterization, the central bank sets the target rate r^* to deliver a zero steady state as discussed in Cor. 1.1.2.

I estimate the first group of parameters using data from the U.S. from 1985-2007. Since 1985-2007 was largely a period of financial calm in the U.S., I estimate these parameters by assuming $a \approx 0$. Given estimates of the macroeconomic parameters, I estimate the term structure parameters in a second step by focusing on the term structure response to QE during 2009.

Effective Borrowing Rate Weights

A key input is the weighting function in the effective borrowing rate $\eta(\tau)$. This function governs the household allocation of borrowing across the term structure, and therefore the sensitivity of the effective borrowing rate is to short- and long-term rates.

I set $\eta(\tau)$ to match the average maturity structure of outstanding U.S. Treasury Bills, Notes, and Bonds from 1985-2007; data is from the monthly CRSP U.S. Treasury database. While it would be possible to match this distribution non-parametrically, as explained in Appendix A.2 the numerical algorithm to solve the model requires being able to solve closed-form solutions to many integral expressions involving $\eta(\tau)$. To that end, I assume that $\eta(\tau)$ is equal to the probability density function of a (truncated) Gamma distribution with shape parameter 2 and rate parameter η_1 (or scale parameter $1/\eta_1$), so $\eta(\tau) \propto \tau \exp(-\eta_1 \tau)$. I estimate the rate parameter η_1 in order to minimize the distance between the parameterized $\eta(\tau)$ and the distribution of outstanding U.S. Treasuries.

Macroeconomic Parameters

In order to estimate the parameters governing the macroeconomic dynamics, I take a moments-matching approach. I target 9 moments in the data: the variance of the short rate, inflation, and the output gap; the respective covariances; and the respective one-year autocovariances. I use data from the U.S. from 1985-2007. For the short rate I use the 3-month Treasury Bill rate; for inflation I use PCE; and for the output gap I use the CBO's nominal potential GDP. The data series are from FRED.

I choose the period 1985-2007 because this was largely a period of financial calm. Hence, I solve the model and compute the model analogues of the variance-covariances assuming that risk aversion $a = 0$. This also allows me to defer estimating many of the parameters from the finance side of the model. I additionally set $\varsigma^{-1} = 1$ and the discount factor $\rho = 0.04$. This leaves 9 parameters to be estimated: the inertia terms κ_r, κ_{z_π}, and κ_{z_x}; the Taylor rule coefficients ϕ_π and ϕ_x; the nominal rigidity term δ; and the shock variances σ_r^2, $\sigma_{z_\pi}^2$, and $\sigma_{z_x}^2$.

Since I have 9 parameters to estimate, the model is able to match the target 9 moments perfectly. Although each moment is sensitive to each parameter, intuitively it is useful to discuss which moments respond most strongly to which parameters. The auto-covariances are largely dependent on the inertia terms, while the overall volatility of the economy is a function of the volatility of the fundamental shocks. Finally, the covariances between the short rate, inflation, and the output gap depend on the degree of nominal rigidity and how the central bank changes the policy rate in response to deviations in inflation and output.

Term Structure Parameters

The other set of parameters to estimate come from the preferred habitat side of the model. I estimate these parameters in order to match the change in the yield curve following the FOMC announcement regarding QE1 on March 18, 2009. Since the model makes predictions about zero-coupon yields, the response I target is the daily change in zero-coupon yields as taken from Gurkaynak et al. (2007).

Before estimating the model, I make some simplifying assumptions. First, I assume that there are two demand factors $\beta_{s,t}$ and $\beta_{\ell,t}$ that are otherwise identical (same inertia parameter κ_β and variance σ_β^2) but concentrated at short and long maturities (different functions $\theta_s(\tau)$ and $\theta_\ell(\tau)$). I assume that these functions are entirely concentrated at maturities of length 2 and 10 years, respectively.[7] Second, I assume that the preferred habitat demand elasticities $\alpha(\tau)$ are constant across maturities. Since the risk aversion coefficient a always enters multiplicatively with $\alpha(\tau)$, I normalize $\alpha(\tau) = 1$. Hence, the estimated coefficient a should be interpreted as a mix of both the arbitrageur's preferences for risk as well as the preferred habitat investor's sensitivity to price movements. The remaining parameters to be estimated are the inertia parameter κ_β, variance σ_β^2, and the risk aversion term a. In the model, I assume that QE1 was a ten standard deviation shock; this coincides with the findings in Gorodnichenko and Ray (2017) that QE purchases were roughly ten times larger than typical private demand shocks for long-term Treasuries.

Calibration Results

Table 1.1 summarizes the results of the calibration. In this section I briefly discuss the results in more detail.

During 1985-2007, on average roughly 40% of Treasury debt was less than one year, 35% was between 1 and 5 years, and the remaining 25% was between 5 and 30 years. Targeting these moments, I estimate that the rate parameter in the (truncated) Gamma distribution is $\eta_1 \approx 1.7$. As Figure 1.7 shows, the model analogue matches the short end of the distribution but somewhat understates the fraction of long-term debt. Although the concentration is highest for shorter maturities, the weighting implies that over 60% of the distribution of borrowing is weighted towards maturities over 1 year. This will imply substantial deviations from benchmark models where borrowing is entirely concentrated at the short rate.

[7] Formally, I assume that the functions $\theta_s(\tau) = \delta(\tau - 2)$ and $\theta_\ell(\tau) = \delta(\tau - 10)$, where $\delta(\cdot)$ is the Dirac delta function. The Dirac delta function can be interpreted as the limit of a mean=zero normal distribution as the variance approaches 0.

Table 1.1: Calibration Results

Parameter	Value	Description	Target
Effective Borrowing Rate			
η_1	1.7069	Weight Scaling Factor	Treasury Maturity Distribution
Macroeconomic Dynamics			
ρ	0.0400	Discount Factor	Long-Run Interest Rate
ς^{-1}	1.0000	Intertemporal Elasticity	Balanced Growth
κ_r	0.9473	Monetary Policy Inertia	$\text{Cov}[r_t, r_{t-1}] = 3.5013$
κ_{z_π}	0.5863	Cost-Push Shock Inertia	$\text{Cov}[\pi_t, \pi_{t-1}] = 0.9141$
κ_{z_x}	0.2554	Demand Shock Inertia	$\text{Cov}[x_t, x_{t-1}] = 2.2908$
ϕ_π	2.0420	Inflation Taylor Coeff.	$\text{Cov}[r_t, \pi_t] = 1.0006$
ϕ_x	0.9709	Output Taylor Coeff.	$\text{Cov}[r_t, x_t] = 0.7722$
δ	0.0459	Nominal Rigidity	$\text{Cov}[\pi_t, x_t] = -0.3015$
σ_r	0.0116	Monetary Shock Vol.	$\text{Var}[r_t] = 2.7066$
σ_{z_π}	0.0068	Cost-Push Shock Vol.	$\text{Var}[\pi_t] = 0.5097$
σ_{z_x}	0.0126	Demand Shock Vol.	$\text{Var}[x_t] = 1.5192$
Term Structure			
$\theta_s(\tau)$	$\delta(\tau - 2)$	Short Factor Location	LSAP Targets
$\theta_\ell(\tau)$	$\delta(\tau - 10)$	Long Factor Location	LSAP Targets
$\alpha(\tau)$	1.0000	Habitat Elasticity	Normalized
κ_β	0.1710	Habitat Factor Inertia	QE1 Yield Curve Response
σ_{z_β}	0.0142	Habitat Factor Vol.	QE1 Yield Curve Response
a	1559.7	Risk Aversion	QE1 Yield Curve Response

Notes: results of the calibration exercise. The effective borrowing weight term η_1 is the rate factor in a (truncated) Gamma distribution: $\eta(\tau) \propto \tau \exp(-\eta_1 \tau)$. For the macroeconomic dynamics coefficients, each parameter is listed alongside the covariance target which is most sensitive to changes in the given parameter; however, the parameters are jointly estimated in order to match all the target moments. Variances and covariances are expressed in percentage points. The short and long demand factor location functions are Dirac delta functions $\delta(\cdot)$.

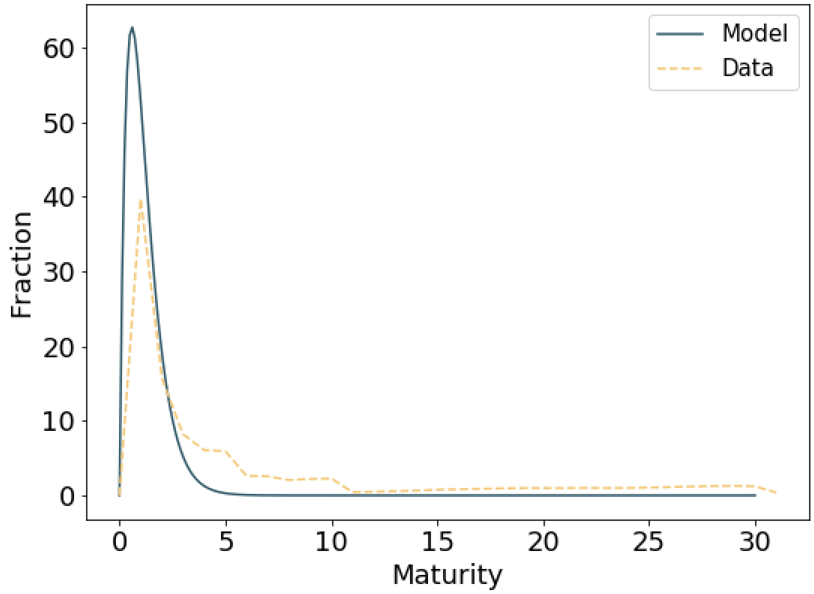

Figure 1.7: Estimated Borrowing Weights $\eta(\tau)$
Notes: the estimated effective borrowing weights $\eta(\tau)$ and the distribution of outstanding Treasury debt in the data.

Figure 1.8 compares the yield curve response to QE1 in the data vs. the model.

1.5.3 Responses to Conventional and Unconventional Monetary Shocks

Using the estimated parameters, I now explore the implications for monetary policy in general equilibrium. I study the model for varying degrees of risk aversion; from very low $a \approx 0$ to even higher than the the level estimated during the QE1 period, in order to understand how policy interacts with different degrees of financial crisis.

Expansionary Monetary Shock

I first study the macroeconomic response to a standard monetary shock. The shock is a 50 basis point fall in the policy rate (an expansionary shock).

The first panel of Figure 1.9 plots the immediate response of the yield curve (in terms of deviations from steady state). Lighter lines plot the response for low levels of risk aversion; the darker lines correspond to high levels of risk aversion. When risk aversion is very low, the entire yield curve shifts down significantly. But as risk aversion increases, the darker lines begin to move closer to zero. While short-term rates still respond strongly, long-term rates become nearly unresponsive. The expected path of the policy rate is not that different, but arbitrageurs do not equalize

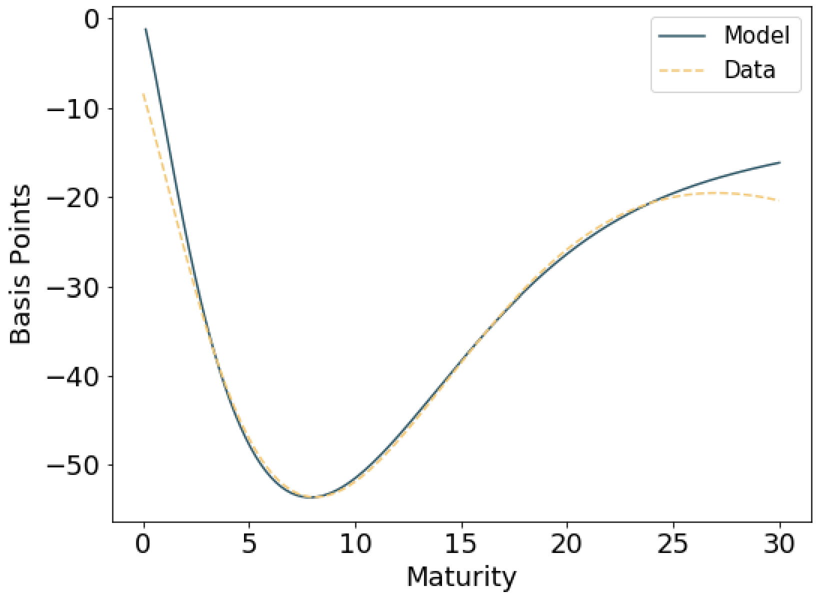

Figure 1.8: Yield Curve Response to QE1, Model vs. Data
Notes: the estimated yield curve response to a QE shock as compared to the actual response.

the expected returns of all bonds and hence long-term rates under-react to the change in the policy rate.

The bottom panels of Figure 1.9 plot the immediate responses of output and inflation, with the level of risk aversion on the x-axis. These results show that as risk aversion increases, the immediate macroeconomic effects of a monetary shock fall. This follows from the results regarding the yield curve. Since the entire yield curve becomes less responsive to a given monetary shock, the household effective borrowing rate similarly responds less. Hence the consumption response (and output gap response) is smaller. Smaller output gaps into the future imply that inflation also responds less.

The exercise in Figure 1.9 also allows for a quantification of the relative effectiveness of monetary policy in and out of financial crises. Comparing the output response to a monetary shock during normal times ($a \approx 0$) to periods of high financial distress (the dotted black line, corresponding to the estimated value of risk aversion a), output responds by about 20% less to a monetary shock in a financial crisis than out. As risk aversion increases, the effectiveness falls even further. This suggests that had the zero lower bound on the policy rate had not been binding (or alternatively, had the Fed pursued a policy of negative rates), monetary policy still would have struggled to boost output during the recent financial crisis.

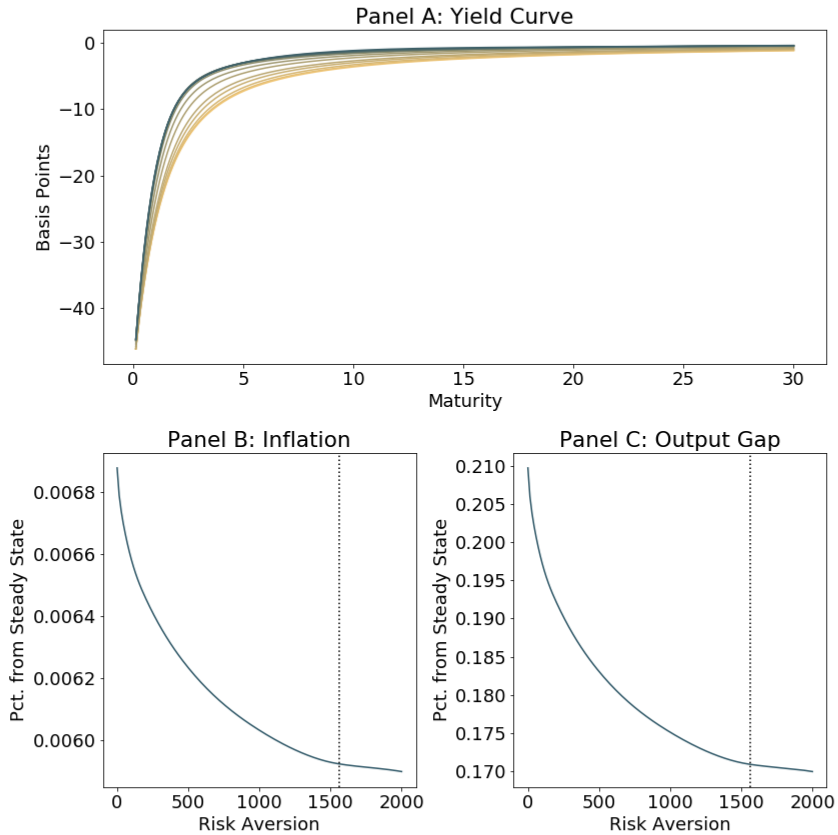

Figure 1.9: Monetary Policy Shock

Notes: Panel A is the contemporaneous response of the yield curve to a 50 basis point monetary policy shock. Responses are plotted as deviations from steady state, in terms of basis points. The x-axis is maturity. Lighter lines correspond to models where risk aversion is low; darker lines to models with high risk aversion. Panels B and C are the contemporaneous response of inflation and the output gap to the same shock. Responses are plotted as deviations from steady state, in terms of percentage points. The x-axis is level of risk aversion. The dotted black line corresponds to the estimated level of risk aversion.

QE Shocks

I now study the response to two types of QE shocks. In the model, these correspond to shocks to the demand factors which I label as $\beta_{s,t}$ (short-term shock, purchases concentrated at 2-year maturities) and $\beta_{\ell,t}$ (long-term shock, purchases concentrated at 10-year maturities).

Figure 1.10 plots the response to a shock to the short demand factor $\beta_{s,t}$. The top panel plots how the yield curve changes immediately, where darker lines correspond to increasing levels of risk aversion. The bottom panels plot immediate responses of output and inflation (with the level of risk aversion on the x-axis).

When risk aversion is very low, there is little to no response to a QE shock. This is because arbitrageurs are able to fully absorb the purchases without requiring any changes in the returns of their portfolio. This implies that there is very little reaction in the yield curve, hence there is little feedback to household borrowing. But as risk aversion increases, the effects increase. Again this is due to the behavior of arbitrageurs. QE purchases offload some risk from the portfolio of arbitrageurs. Hence, they require smaller excess returns to hold other bonds, which pushes down interest rates. This leads to a decline in the effective borrowing rate of households, which boosts consumption (and hence output) on impact.

Figure 1.11 repeats the above exercise for shocks to the long demand factor $\beta_{\ell,t}$. Once again, there is nearly no response when arbitrageurs are close to risk neutral; as risk aversion increases, so too does the magnitude of the response of the yield curve and macroeconomic variables.

The intuition is similar for both "short" and "long" QE shocks, but comparing the differential impacts also reveals interesting results. Studying the yield curve responses in Figures 1.10 and 1.11 reveals that there are four "regimes" where the yield curve responses are qualitatively different. The first regime corresponds to very low levels of risk aversion, where QE shocks have essentially no effect. In the second regime, when financial markets are somewhat disrupted, both QE shocks put downward pressure on interest rates but the response of the yield curve to both short-term and long-term purchases is very similar in shape. The response is hump-shaped, with the peak response occurring at shorter to intermediate maturities, and pushes both short- and long-term spot rates in the same direction. Only the magnitude differs, with the long-term shock leading to somewhat larger responses.

However, when risk aversion becomes sufficiently high, the effects become more localized. This is the third regime, when financial markets start to become severely disrupted; this regime corresponds to the level of risk aversion estimated from the QE1 data. In this case, both long and short QE shocks push down all interest rates, but now short-term QE purchases have larger effects on short-term yields than long-term QE purchases. Finally, the fourth regime corresponds to extreme levels of financial crisis (higher than what was estimated). The localization of demand shocks becomes so extreme that long-term demand shocks actually lead to increases in short-term rates, while short-term demand shocks also put upward pressure on long-term rates.

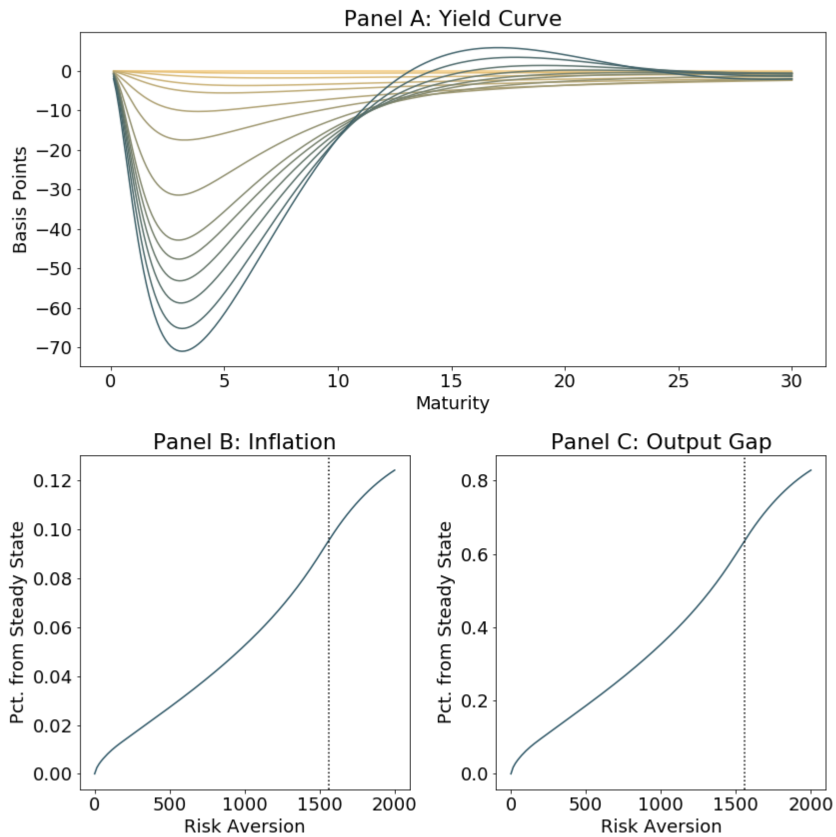

Figure 1.10: QE Shock (short-term purchases)

Notes: Panel A is the contemporaneous response of the yield curve to a QE shock, where purchases are concentrated at shorter term bonds. Responses are plotted as deviations from steady state, in terms of basis points. The x-axis is maturity. Lighter lines correspond to models where risk aversion is low; darker lines to models with high risk aversion. Panels B and C are the contemporaneous response of inflation and the output gap to the same shock. Responses are plotted as deviations from steady state, in terms of percentage points. The x-axis is level of risk aversion. The dotted black line corresponds to the estimated level of risk aversion.

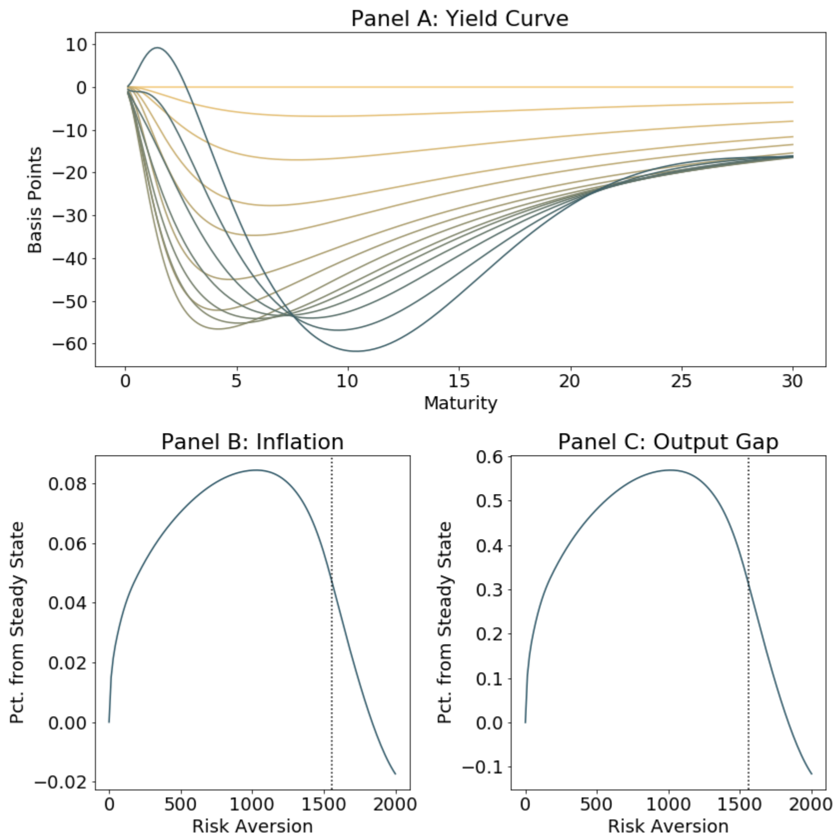

Figure 1.11: QE Shock (long-term purchases)

Notes: Panel A is the contemporaneous response of the yield curve to a QE shock, where purchases are concentrated at longer term bonds. Responses are plotted as deviations from steady state, in terms of basis points. The x-axis is maturity. Lighter lines correspond to models where risk aversion is low; darker lines to models with high risk aversion. Panels B and C are the contemporaneous response of inflation and the output gap to the same shock. Responses are plotted as deviations from steady state, in terms of percentage points. The x-axis is level of risk aversion. The dotted black line corresponds to the estimated level of risk aversion.

The increasing localization of demand shocks is a key partial equilibrium result from Vayanos and Vila (2009); this exercise shows that the finding holds in general equilibrium. What causes this localization? Intuitively, consider the limiting case when arbitrageurs only want to minimize the variance of the change in their wealth. Hence they allocate their entire portfolio to the short (risk-free) rate, taking no positions in any other longer-term bonds. In response to a long-term demand shock from preferred habitat investors, arbitrageurs would be unwilling to make any changes their portfolio allocation to accomodate the shift. The only way for the net supply condition to be satisfied is if the prices of the bonds that are affected by the demand shock respond. Moreover, these are the only bonds that see price changes. In other words, the demand shock has only local effects.

Of course, even when arbitrageurs are very risk-averse, the bond market will not exhibit this type of extreme segmentation. But the qualitative behavior of the term structure response is different for low and high levels of risk aversion. When risk aversion is low, demand shocks have a (small) effect, but the location of the shock does not matter. A the key source of risk that arbitrageurs are concerned with is short-rate risk. Every bond has some sensitivity to this source of risk, hence shifts in demand from preferred habitat investors change the portfolio allocations of arbitrageurs. Thus this changes the market price of short-rate risk, and as this impacts all bonds this has a global effect regardless of where the demand shock originates. For example, in response to positive demand shifts, arbitrageurs sell bonds, reducing their exposure to short-rate risk. Hence they require lower expected returns to hold bonds, pushing down rates. The bonds that respond the most are those most sensitive to short-rate risk, which does not depend on the location of the shock, but only on the stochastic (mean-reversion) properties of the shocks.

But when risk aversion is very high, the location of the shock matters. When risk aversion is very high and there are multiple sources of risk (short-rate risk, multiple demand factors, and other structural shocks), arbitrageurs try to limit their exposure to these sources of risk. This means that demand shocks must have mostly local effects, since otherwise arbitrageurs would be exposing themselves to other sources of risk. In other words, if a demand shock for long-term debt had big effects on short-term debt, this would be because arbitrageurs were making large changes to their holdings of short-term debt. But this changes their exposure to other sources of risk. Hence they choose not to integrate the markets across maturities, and the shocks stay more localized.

Operation Twist Shock

Next, I look at a policy mimicking "Operation Twist," where the Federal Reserve purchased long-term securities and sold short-term securities. To model this, I analyze the effect of a simultaneous increase in long-term purchases ($\beta_{\ell,t}$) and decrease in short-term purchases ($\beta_{s,t}$). Figure 1.12 plots the yield curve and real responses.

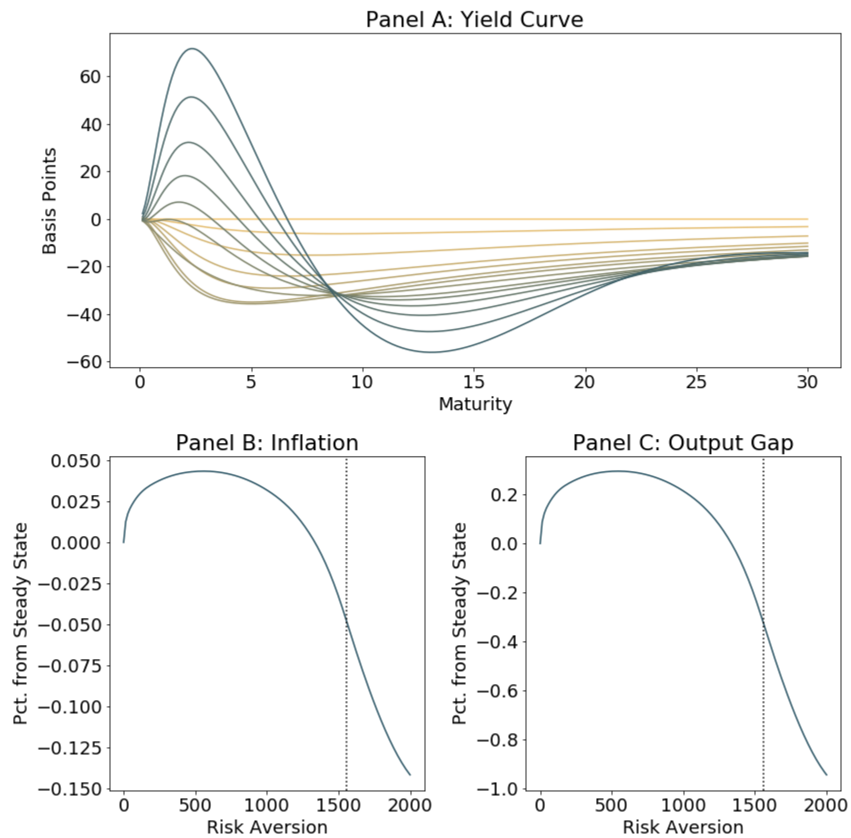

Figure 1.12: QE Shock (long-term purchases)

Notes: Panel A is the contemporaneous response of the yield curve to an Operation Twist shock, where longer term bonds are purchased and shorter term bonds are sold. Responses are plotted as deviations from steady state, in terms of basis points. The x-axis is maturity. Lighter lines correspond to models where risk aversion is low; darker lines to models with high risk aversion. Panels B and C are the contemporaneous response of inflation and the output gap to the same shock. Responses are plotted as deviations from steady state, in terms of percentage points. The x-axis is level of risk aversion. The dotted black line corresponds to the estimated level of risk aversion.

The previous QE exercises showed that when risk aversion is low, the yield curve reacts similarly to the short and long demand factors; the only difference is the magnitude of the response to the long demand factor is larger. When the two shocks are combined but with opposite signs, this implies that the long demand factor dominates, leading to a decline in the term structure. The macroeconomic response is thus a similar but muted version of the QE responses observed previously.

Recall that in the calibration, the effective borrowing rate is weighted mostly towards maturities between 1 and 5 years; there is little weight on interest rates 10 years and above. However, when financial markets are relatively healthy ("regime 2" discussed above), it is still the case that QE purchases concentrated at long-term rates is more effective at push up output than short-term purchases.

But as can be seen in Figure 1.12, when risk aversion is high enough, the short and long demand factors have differential impacts on the yield curve. The combination of the two shocks leads to declining long-term rates, but at some point leads to increasing short-term rates. For intermediate values of risk aversion, real effective borrowing rates still fall, leading to increases in output and inflation. However, as risk aversion continues to increase, the magnitude of the macroeconomic effects declines. For high values of risk aversion, the sign eventually flips: after the Operation Twist shock, both inflation and the output gap actually decline.

The exercise implies that if Operation Twist had taken place during March 2009 when risk aversion was very high (the dotted black line in the bottom panels of Figure 1.12), the aggregate effect would have actually been contractionary. Moreover, the actual yield curve responses around September 21, 2011 (when Operation Twist was announced by the FOMC) saw long-term rates fall while short-term rates increased slightly. While financial market disruptions in 2011 had subsided relative to the peak of the crisis, this suggests that financial frictions were still high. Hence, the aggregate effects of Operation Twist were likely muted at best.

One key result in all of these exercises is that the transmission of these policies depends crucially on the health of financial markets. As with conventional monetary policy explored in previous sections, the effectiveness of QE policies varies with the level of risk aversion of arbitrageurs. But the interaction is the opposite of conventional policy. QE has no effect when arbitrageurs are risk neutral; as risk aversion increases, QE becomes more and more effective at moving the yield curve and hence boosting output.

Policies like Operation Twist may have more ambiguous effects. When financial markets are highly disrupted, bond markets become much more segmented. Targeted buying and selling of securities throughout the term structure will have highly localized effects. Depending on which maturities are the most important for household borrowing, this could end up having the opposite of the intended effect, leading to falling consumption.

1.5.4 Can LSAPs Be Stabilizing? Endogenous QE Rules and Determinacy

The results above show that, under the right financial conditions, QE can move output and inflation. A separate question is: can the central bank use QE to achieve determinacy and stabilize the economy?

Suppose that the central bank conducts QE operations in a similar manner to how it sets the policy rate. I study this by modifying the idiosyncratic demand factor from eq. (1.19) so that it endogenously reacts to inflation:

$$\mathrm{d}\beta_t = -\kappa_\beta \left(\beta_t - \phi_\pi^\beta \pi_t\right) \mathrm{d}t\,. \tag{1.33}$$

To simplify, I assume this is the only demand factor. Further, I assume that the central bank does not react to the output gap ($\phi_x = 0$). This will imply that the model moves towards the region of indeterminacy for intermediate to high levels financial frictions. However, by choosing $\phi_\pi^\beta > 0$, it is possible to re-establish determinacy.

Figure 1.13 shows how the determinacy condition varies with different levels of risk aversion (a) and QE responsiveness to inflation (ϕ_π^β). When risk aversion is very low, the model is determinate for all values of ϕ_π^β. Moreover, changes in the responsiveness to either inflation or output have no effect on the determinacy condition (the value of the unstable eigenvalues). As risk aversion increases, the baseline model with no endogenous QE responses quickly moves towards the region of indeterminacy. But allowing for a stabilizing QE policy that reacts to inflation ($\phi_\pi^\beta > 0$) moves the model back towards the region of determinacy. For strong enough responses of QE to inflation, the model regains determinacy.

Intuitively, an endogenous QE rule works by picking up the slack left by conventional policy. When financial markets are healthy, a standard Taylor rule is stabilizing. As financial frictions increase, the pass-through of conventional policy deteriorates. If this is the only policy rule, then eventually the model becomes unstable. However, as seen above, QE becomes more effective as financial frictions increase. Hence, just as conventional policy is failing, the endogenous QE rule becomes more and more effective.

1.5.5 Optimal Policy

A full treatment of a planner's optimal policy problem is beyond the scope of this paper. Allowing for the central bank to choose a path of the policy rate that differs from the Taylor-type rules considered in this paper will in general lead to bond prices that are not affine functions of the state variables. Instead, I consider a simpler optimal policy problem. This section analyzes how the central bank would choose to set the response to inflation ϕ_π and the inertia term κ_r optimally, as a function of financial market health (and further simplifies the problem by assuming no response

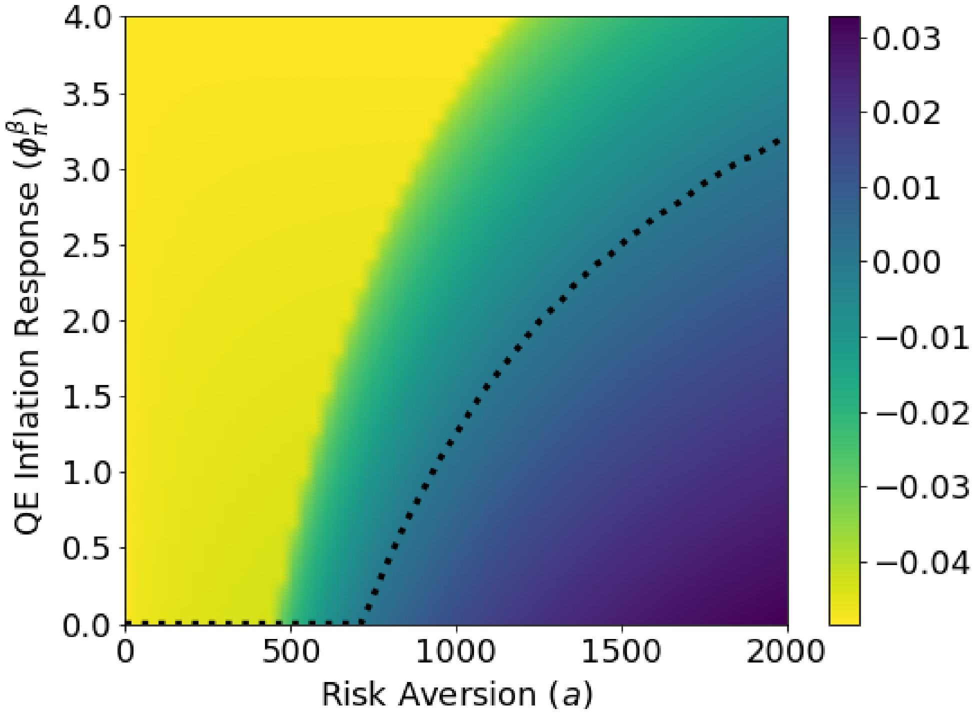

Figure 1.13: Regions of Determinacy

Notes: the heatmaps show the region of determinacy (lighter regions, upper left) and indeterminacy (darker regions, lower right). The x-axis is the level of risk aversion a, and the y-axis is the level of the parameter ϕ_π^β, which governs how strongly QE reacts to inflation. The values correspond to the smallest unstable (real) eigenvalue; the model is determinate when this is negative. The dotted black line delineates the region of indeterminacy from determinacy.

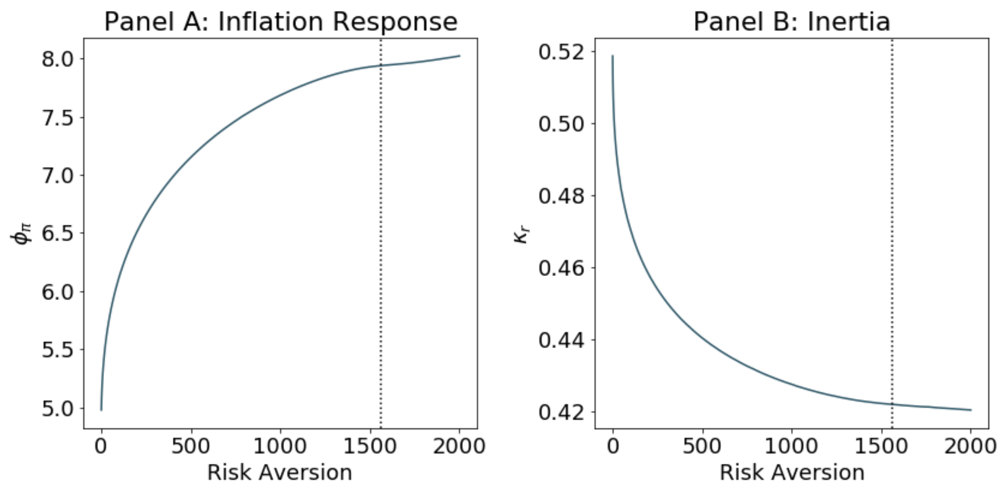

Figure 1.14: Optimal Response to Shocks

Notes: optimal policy coefficients as risk aversion increases. The planner weights are set to $w_\pi = 1, w_x = 0.02$. Panel A plots the optimal Taylor rule coefficient on inflation (ϕ_π), while Panel B plots the optimal inertia term (κ_r). The response Taylor rule coefficient on output (ϕ_x) is set to 0. The dotted black line corresponds to the estimated value of risk aversion.

to output: $\phi_x = 0$). I assume a quadratic loss function for the planner, given by

$$\min_{\phi_\pi, \kappa_r} E_0 \int_0^\infty e^{-\rho t} \left(w_\pi \pi_t^2 + w_x x_t^2 \right) \mathrm{d}t$$

where $w_\pi \geq 0$ and $w_x \geq 0$ are the weights that the planner assigns to inflation and output, respectively. Appendix A.3.4 derives the general expression for this present discounted value of future second moments of the jump variables in this class of models.

Figure 1.14 plots the optimal coefficients ϕ_π (Panel A) and κ_r (Panel B) as a function of risk aversion. I suppose that the planner cares more about inflation relative to output by setting the planner weights as $w_\pi = 1, w_x = 0.02$.[8]

The optimal response to inflation is increasing (higher ϕ_π) as arbitrageur risk aversion increases (higher a). Further, optimal inertia is also increasing (lower κ_r). Since monetary policy relies on the pass-through provided by financial markets, as the risk-bearing capacity of arbitrageurs becomes disrupted, monetary policy efficacy is weakened. Moreover, as financial market disruptions increase, the economy becomes less stable and the output and inflation responses to these shocks increase. In order to effectively stabilize the economy (minimize the volatility of inflation and output), the optimal response becomes more agressive and longer-lasting.

[8]Note that in a full optimal policy experiment, the quadratic loss function is an approximation to the welfare loss function, and the planner weights would be derived from this welfare function.

1.6 Concluding Remarks

This paper studies conventional and unconventional monetary policy through the lens of a general equilibrium "preferred habitat" New Keynesian model. When output depends on both short- and long-term borrowing rates, the transmission of monetary policy to the macroeconomy depends on imperfect financial markets. As a result, standard monetary policy becomes less effective when financial markets are disrupted. For the same reason, the efficacy of forward guidance is weakened. However, financial crises open the door to other unconventional policies such as quantitive easing, which can push down long-term rates and stabilize output and inflation.

The framework considered in this paper suggests promising avenues for future work. The results relating financial health and the determinacy of the model are important for understanding how policymakers can achieve macroeconomic stability. The model suggests revisiting the empirical work studying the stability properties of central bank policy rules, as the requirements for stability are state-dependent. Furthermore, the findings provide a possible justification for macro-prudential policies. The model assumes that risk aversion is fixed, but in reality the risk-bearing capacity of financial markets is endogenous. When this capacity is too low, the result is macroeconomic instability. If investors in financial markets do not internalize this, then there is an externality that can be addressed by policy. Extending the model to allow for endogenous risk-bearing capacity is useful for understanding how such macro-prudential policies should be carried out.

The model can also be extended to allow for a more realistic treatment of household borrowing. For example, the model does not allow for default risk, but in reality the key borrowing rates for households will not be default-free. Moreover, a major part of some QE programs was targeting mortgage-backed securities. The localization results suggest that during financial disruptions, LSAP programs will be most effective when they target borrowing markets in which households are most active. But the sensitivity of real activity to different borrowing rates, and the stability of this relationship over time, is an open empirical question.

This paper shows that LSAPs should be a tool in central bankers' arsenal to stabilize the economy during financial crises. However, the predictions of the model are not only dependent on the health of financial markets but also sensitive to interactions with the location of the purchases in maturity space, how long the purchases last, how sensitive the real economy is to long-term rates, and the structure of the other fundamental shocks in the economy. Policymakers are likely to face a great deal of uncertainty about these factors, which raises the possibility that a better approach may be to target specific long-term rates and allow for flexibility in the open-market operations used to achieve these targets. More generally, the model should serve as the basis for more quantitative analysis in order to design optimal policies.

Chapter 2

Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions

2.1 Introduction

Demand for safe assets, and U.S. Treasuries in particular, plays a central role in the macro-financial landscape.[1] To offset the negative effects of the recent financial crisis, central banks have implemented various large scale asset purchases, representing a sharp increase in demand for these assets. The most salient of these is the quantitative easing (QE) programs carried out by the Federal Reserve, which involved two trillion dollars of Treasury security purchases. Apart from the massive scale of these purchases, the Federal Reserve disproportionately bought long-term government debt, thus departing from the practice of having the distribution of its portfolio close to the distribution of outstanding debt (Figure 2.1).

While evaluating the program, Ben Bernanke, the chair of the Fed at the time, observed, "The problem with QE is it works in practice but it doesn't work in theory." Indeed, QE was successful in reducing short- and long-term interest rates, but the mechanism behind this reaction is still not well understood. For example, standard macro-financial models imply that the demand for assets such as Treasuries is determined solely by economic agents' intertemporal consumption decisions, which does not capture the sources of demand shifts initiated by the Fed. Although the workhorse macroeconomic models cannot readily explain the workings of the QE, several explanations have been put forth. For instance, QE could be effective because it signaled to the markets that the Fed is serious about keeping short-term interest rates low for a long time (forward guidance). Or, perhaps the Fed exploited frictions (limited arbitrage and market segmentation) in the financial markets by purchasing securities in a particular segment. Finally, by buying assets on a massive scale, the Fed could signal a poor state of the economy which pushed interest rates down ("Delphic" effect;

[1]This chapter is based on my joint paper with Yuriy Gorodnichenko.

see Campbell et al. (2012) for more details).

Future deployment of unconventional tools such as QE requires policymakers to move beyond the "heat of the moment" policies, and hence a central question for policymaking and academic research is which of these theories is the key channel. But given the paucity of QE events, it has proven remarkably hard to provide clear empirical evidence for each theory, as well as to assess the relative contributions of the proposed channels. Indeed, many channels were likely active during QE rounds and the reactions to QE were observed in a particular state of the economy, which potentially confounds identification and interpretation.

The objective of this paper is to unbundle QE by focusing on one channel: market segmentation and *preferred habitat*, which posits that certain investors have preferences for specific maturities. Our approach is to identify shifts in *private* demand for Treasuries that mimic QE, but are independent of the other channels discussed above. The key mechanism through which market segmentation and preferred habitat forces operate is not the source of demand shifts, but rather how marginal investors in the market for Treasury debt absorb these demand shocks. Therefore, the best way to isolate and study the preferred habitat channel of QE is to identify unexpected demand shifts that are unrelated to other channels of QE.

In order to construct demand shocks with these properties, we utilize the structure and timing of the primary market for Treasury securities. Similar to the empirical monetary policy literature (e.g., Bernanke and Kuttner (2005), Gurkaynak et al. (2007), Gorodnichenko and Weber (2016)), we look at high-frequency (intraday) changes in prices of Treasury futures in small windows around the close of Treasury auctions to identify unexpected shocks to demand for Treasuries. The key for identification is that all of the "supply" information (e.g. security characteristics such as the maturity, as well as the amount of newly offered and outstanding securities) is known and priced in by the market. For small enough windows around the close and release of the auction results, any price changes are reactions to information regarding the demand for the Treasury securities from the given auction. We interpret these price changes as demand shocks. Utilizing high-frequency changes in asset prices along with the timing of Treasury auctions in this manner allows us to rule out confounding factors and identify unexpected shifts in demand in a model-free way.

Treasury auctions have a number of properties that can help us understand the workings of QE. First, although the auctions are not as large as the QE rounds, the Treasury sells about $150 billion in notes and bonds per month in recent years. Because the primary market for Treasuries is a convenient venue for investors who wish to purchase large amounts of government securities, the release of Treasury auction results can reveal potentially large shifts in demand for Treasuries. The surprise movements in the yields are reasonably large: a typical (one standard deviation) shock is equivalent to a yield change of roughly 2 basis points, which is much larger than similar changes on non-auction dates. For comparison, Chodorow-Reich (2014) estimates that the first round of the QE program in the U.S. cut Treasury rates (five-year maturity) by 9 basis points following the announcement from Chairman Bernanke on December 1, 2008.

Second, we document that demand shocks are driven by institutional investors such as foreign monetary authorities, investment funds, insurance companies and the like. Moreover, these shocks are not driven by changes in expectations about inflation, output, or other broad market conditions. Therefore, variation during Treasury auctions can help us to isolate the effect of idiosyncratic purchases in specific asset segments on the level and shape of the yield curve, which is difficult to achieve by examining only QE events.

Finally, in sharp contrast to QE events, Treasury auctions are frequent and information is available back to 1979. This gives us an opportunity for crisper inference and to study state-dependence in the effect of targeted purchases of assets (e.g., crisis vs. non-crisis states). Because QE events were both infrequent and confounded with a massive financial crisis, having a long time series is instrumental for understanding how QE-like programs can work in normal times.

Importantly, because Treasury auctions for specific maturities are spread in time, we can identify changes in demand for government debt of specific maturities. As a result, we can trace how a shock in one part of the yield curve propagates to other parts of the yield curve. In this sense, we have natural experiments which can mimic targeted purchases of the Fed during QE programs. Hence, despite the apparent distance between QE programs and unexpected movements in private demand during regular Treasury auctions, this empirical strategy provides clean identification of demand shifts and allows us to map out the impact of these shocks.

Although we do not have a structural interpretation of unexpected changes in demand, we can still use shocks in demand for specific maturities of government debt to investigate how these shocks spread to other maturities. Specifically, we examine reactions across maturities through the lens of a formal preferred habitat theory of the term structure. Building on Vayanos and Vila (2009), we present a series of numerical simulations to provide qualitative predictions about how the location of the shock in maturity space affects the relative change in the term structure and how the reaction depends on the risk-bearing capacity of marginal investors. Informed by theory, we test these predictions using daily changes in spot rates for government debt in response to our measures of surprise movements in private demand at particular maturities. We find evidence consistent with our theoretical predictions.

Our results suggest that QE programs can be effective in influencing interest rates for debt at specific maturities when financial markets are disrupted. On the other hand, QE programs are less likely to be effective at this task in normal times when risk-bearing capacity of arbitrageurs is greater. In this case, demand shocks at a specific maturity likely move the whole yield curve rather than a specific segment, and the response may peak at a maturity other than the targeted maturity. Furthermore, if the Fed attempts to use purchases of debt with specific maturities to shift down the whole yield curve during a crisis, this exercise may be ineffective and the Fed should intervene at multiple maturities.

Furthermore, our results provide a quantitative sense of *how much* QE programs could influence interest rates through the preferred habitat channel. Specifically, using our regression estimates, we show that the amount of government debt purchases during the QE1 program should generate declines in yields similar to what was observed in the data. In other words, given the reaction of yields to surprise movements in *private* demand during Treasury auctions, we can account for most of the market reaction to QE1 announcements. This result is consistent with the view that QE worked mainly via market segmentation and preferred habit, and that the *net* effect of other channels was small.

Our study contributes to several strands of previous research. First, we provide new evidence to the literature examining theoretically (e.g., Vayanos and Vila (2009)) and empirically (e.g., Greenwood and Vayanos (2014), Krishnamurthy and Vissing-Jorgensen (2012), Hamilton and Wu (2012b)) determinants of demand for government debt. In particular, we add to the literature departing from the "expectations hypothesis" (e.g., Kuttner (2006)) of the term structure of interest rates, and provide evidence for alternative explanations such as limited arbitrage and market segmentation. Our findings are complementary to Lou et al. (2013) and Fleming and Liu (2016) who also utilize Treasury auctions to explore how *supply* shocks interact with these forces. Second, we contribute to the rapidly growing literature studying the effects of QE programs in the U.S. and other countries (see Martin and Milas (2012) for a survey) and in particular the literature studying how market segmentation interacts with QE programs (e.g. **?**). While most of these studies focus on market movements around QE announcements (e.g., Krishnamurthy and Vissing-Jorgensen (2012), Chodorow-Reich (2014)), we instead focus on market movements around Treasury auctions that can also give us an opportunity to investigate market reactions to unexpected changes in demand for government debt not only in crisis but also in normal times. Third, our paper is methodologically related to earlier studies (e.g., Kuttner (2001), Bernanke and Kuttner (2005)) utilizing high-frequency data to construct surprise movements in policy. Although we do not measure unexpected movements in policy, we construct shocks in private demand that inform us about how markets can react to changes in policy.

2.2 Data and Institutional Details

In this section we describe the primary sources of our data and present basic statistics. First, we describe the U.S. Treasury auctions for U.S. government notes and bonds (coupon-bearing nominal securities). Second, we describe the details of futures contracts for these Treasury securities.

2.2.1 Primary Market for Treasury Securities

The Treasury sells newly issued securities to the public on a regular basis through auctions. Currently, 2-, 3-, 5- and 7-year notes are auctioned monthly. 10-year notes and 30-year bonds are auctioned in February, May, August and November with reopenings in the other 8 months. The frequency of auctions evolved over time. For example, 30-year bonds were not issued between 1999 and 2006 and were issued only twice a year between 1993 and 1999.

There are two types of bids: noncompetitive and competitive. Noncompetitive bidders agree to accept the terms settled at the auction, and are typically limited to $5 million per bidder. Competitive bidders submit the amount they would like to purchase and the price (the interest rate) at which they would like to make the purchase. For each competitive bidder, the submitted amount cannot be greater than 35% of the amount offered at the auction.

Auction participants include primary dealers, other non-primary brokers and dealers, investment funds (for example, pension, hedge, mutual), insurance companies, depository institutions, foreign and international entities (governmental and private), the Federal Reserve (System Open Market Account), and individuals. These participants are classified into three groups. The first group is Primary Dealers (brokers and banks) that trade on their accounts with the Federal Reserve Bank of New York. This group typically buys the largest share of auctioned debt and is required to participate in every Treasury auction. The second group is Direct Bidders: non-primary dealers submitting bids for their own proprietary accounts. The third group is Indirect Bidders who submit competitive bids via a direct submitter, including Foreign and International Monetary Authorities placing bids through the Federal Reserve Bank of New York.[2]

Additionally, the Treasury divides investors into the following classes: Investment Funds (mutual funds, money market funds, hedge funds, money managers, and investment advisors); Pension and Retirement Funds and Insurance Companies (pension and retirement funds, state and local pension funds, life insurance companies, casualty and liability insurance companies, and other insurance companies); Depository Institutions (banks, savings and loan associations, credit unions, and commercial bank investment accounts); Individuals (individuals, partnerships, personal trusts, estates, non-profit and tax-exempt organizations, and foundations); Dealers and Brokers (primary dealers, other commercial bank dealer departments, and other non-bank dealers and brokers); Foreign and International (private foreign entities, non-private foreign

[2] Additionally the Federal Reserve System purchases securities for its System Open Market Account (SOMA). Starting in 1997, the SOMA amount was changed from being listed within the announced offering amount to being additions to the announced offering amount. That is, if the Treasury auctions $15 billion in bonds and the Federal Reserve would like to purchase $1 billion in the auction, the Treasury issues $16 billion in bonds. This change was made so that the Treasury would be able to provide better information to the market about the amount of securities actually available for sale to the public.

entities placing tenders external of the Federal Reserve Bank of New York (FRBNY), and official foreign entities placing tenders through FRBNY); Federal Reserve System (the Federal Reserve Banks System Open Market Account (SOMA)); Other (represents the residual from categories not specified in investor class descriptions above). Fleming (2007) describes in greater detail the breakdown by types and class of bidders.

As detailed in Figure 2.2, there are four stages of a Treasury auction:[3]

1. *Announcement*: A few days before an auction, the Treasury releases all the pertinent information regarding the upcoming auction. An announcement includes security information (maturity, CUSIP identifier, schedule of coupon payments, etc.) as well as the amount offered, the bidding closing times, which class of bidders can participate, and other information describing the rules of the auction.

   Figure 2.3 presents a typical announcement. At this auction, the Treasury offers \$16 billion in 30-year bonds. This is a new auction (that is, the Treasury does not reopen a previous auction) with the maximum award (that is, maximum allocation to a bidder) of \$5.6 billion.

2. *Bidding*: After the announcement, individuals and institutions may submit bids up until the closing times of the auction. The announcement in Figure 2.3 stipulated that non-competitive bids should be submitted by 12:00 p.m., while the deadline for competitive bids is 1:00 p.m.

3. *Results*: Most Treasury note and bond auctions close at 1:00 p.m. Competitive bids are accepted in ascending order (in terms of yields) after the auction closes until the quantity meets the amount offered minus the amount of non-competitive bids. All bidders receive the same yield as the highest accepted bid.[4] Once the auction closes and the winning bids are determined, the information regarding the results is released immediately. Besides the winning yield, the Treasury announces various aggregate statistics regarding the bidding. Beginning in the early 2000s, auction results are released within minutes of the close of the auction (see Garbade and Ingber (2005)).

   Figure 2.4 presents a typical announcement about auction results, which corresponds to the auction announcement presented in Figure 2.3. The demand (tendered) for the security was \$33.3 billion, most of the bids came from primary dealers (\$23.7 billion), \$489.9 million was bought by the Federal Reserve (SOMA), and a relatively low amount was bought via non-competitive bids (\$14.8 million). The "bid-to-cover", the ratio of all bids received to all bids accepted, was

[3]See Driessen (2016) for details on the design of Treasury auctions. Garbade (2007) provides historical details regarding the manner in which the Treasury has conducted auctions.

[4]Between 1970 and 1992, Treasury did not charge a uniform price. Instead, allocation of bonds was made at the individual yields stipulated by the bidders.

$33.3/$16.0=2.08. The interest rate, corresponding to the winning yield, was set at 3.75 percent per year.

4. *Issuance*: A few days after the close of an auction, the Treasury delivers the securities and charges the winning bidders for payment of the security. At this point the winning bidders can hold the security to maturity and receive coupon payments, or sell the security on the secondary market.

Data from the announcements and results of every auction since late 1979 are available from TreasuryDirect.gov. Data regarding amounts accepted and tendered by bidder type (Primary Dealer, Direct, and Indirect) are available starting in 2003. Additionally, the Treasury provides information regarding allotment by investor class (Investment Funds, Individuals, etc) starting in 2000.

2.2.2 Treasury Futures

We use Treasury futures prices in order to construct market-based measures of demand surprises occurring during Treasury auctions. Treasury futures are standardized contracts that obligate the seller to deliver a valid Treasury security to the buyer at a later date. Futures contracts for 30-year Treasury bonds were introduced in 1977, followed later by 10-year, 5-year, and 2-year Treasury note futures. Treasury futures currently trade on the Chicago Mercantile Exchange (CME), and intraday tick-level data are available starting in 1995. The market for Treasury futures is deep: the average daily volume of trade in 2012 was more than 2 million contracts with more than $100 billion of notional value.

The futures contracts close in March, June, September, and December. We focus on the "closest" contract, i.e. the contract that closes within 1-3 months as these are by far the most traded. For example, in February we use the March expiry, while in March we use the June expiry. Although contracts that close in a given month can still be traded, the volume of trades is substantially lower.

Note that futures are not tied to any specific bond issue (CUSIP). Each futures contract allows for a range of deliverable Treasury securities. 2-year futures contracts allow for delivery of Treasury notes with remaining maturity between 1-year 9-months to 2 years; 5-year futures allow for remaining maturity between 4-year 2-months to 5-year 3-months; 10-year futures allow for remaining maturity between 6-year 6-months to 10-years; and 30-year futures allow for delivery of Treasury bonds with remaining maturity of at least 15 years.[5] In principle any permissible Treasury security can

[5]The 30-year futures contract is also known as the "classic T-Bond" future. This contract originally allowed for delivery of bonds with remaining maturity between 15 years and 30 years. In 2009, the CME Group introduced "Ultra T-Bond Futures" which uses Treasury bonds with remaining maturity of at least 25 years but no more than 30 years, and changed the range of deliverable maturities to the classic contract to bonds with remaining maturity between 15 years and 25 years. While the "Ultra" futures contract provides a better match for long maturities, the time series for

be delivered into a futures contract, but as explained in Lauszweski et al. (2014) in practice a so-called "cheapest to deliver" (CTD) security emerges for a given futures contract. Although which Treasury security is used for payment can vary over time, this variation happens at relatively low frequencies (weekly or monthly) and therefore our analysis at high frequencies should not be materially affected by this peculiarity of Treasury futures contracts.

Because futures cannot be matched to a specific CUSIP, we link a given auction to Treasury futures using the maturity offered in the auction. For example, if the Treasury auctions 7-year notes, we use the 10-year futures contract which allows delivery of securities maturing in at least 6.5 years years and no more than 10 years. While this linking introduces a mismatch in terms of maturities, the difference between the maturity of matched futures contracts and the maturity of the auctioned government debt is relatively small.

We use Treasury futures prices for a number of reasons. First and foremost, Treasury futures provide a natural market-based measure of unexpected shifts in Treasury prices. Further, Treasury futures trade on a standardized exchange rather than over the counter. Another option would be to use "when-issued" prices of the Treasury securities being auctioned. Although this option has the benefit of matching perfectly the security being auctioned, the downside is that the when-issued market systematically trades at lower yields than the winning yield at the auction (see e.g. Fleming and Liu (2016)).[6]

2.2.3 Summary Statistics

In our analysis we focus on Treasury note and bond auctions. We exclude inflation protected securities (TIPS), floating rate notes (FRNs), cash management bills (CMBs), and callable bonds (the last of which was issued in 1984), because these securities have different structural arrangements than simple coupon-bearing nominal securities. We also exclude Treasury bills (zero-coupon securities with maturity one year or less) because the QE programs mainly bought long-term nominal U.S. government debt.[7] Further, Treasury futures contracts exist for 2-year, 5-year, 10-year, and 30-year nominal Treasury notes and bonds, but not for shorter term bills.

Figure 2.5 plots the number of note and bond auctions per year in our sample, broken up by term length. The number of auctions is relatively stable throughout the 1980s to mid 1990s. In the face of declining government debt, the number of auctions temporarily fell in the late 1990s and early 2000s, which also coincides with

the contract is relatively short and the volume of trades is small relative to the other longer-running futures. For these reasons we use the "Classic T-Bond Futures" to ensure consistency over time.

[6]In 2015, the U.S. Department of Justice launched a probe to investigate whether various financial companies (most of them are primary dealers of U.S. Treasury securities) participated in a conspiracy to manipulate the "when issued" market for Treasuries.

[7]For example, between November 3, 2010 and June 29, 2011 (QE2), the Fed bought \$750 billion in Treasury securities, of which TIPS purchases were only approximately \$26 billion.

the termination of new issuances of 30-year bonds. After the Great Recession, the number of auctions increased significantly.

Table 2.1 presents summary statistics for note and bond auctions between 1979 and 2015, and the subsample 1995-2015 for which we have intraday Treasury futures prices. Since 1995, a typical offering is about \$20 billion which generates more than \$50 billion in demand so that the bid-to-cover ratio is approximately 2.6. The largest source of demand for Treasuries is primary dealers (their bid-to-cover ratio is $\approx$2) but other types of bidders also account a large fraction of auction offerings. Primary dealers purchase approximately 60 percent of auctioned Treasuries with the rest split equally between investment funds and foreign buyers.

There is considerable variation in the offered amounts (standard deviation is $\approx$\$9 billion) as well as the level and composition of demand (standard deviation for the bid-to-cover ratio is $\approx$0.5 and the standard deviation of bid-to-cover ratio for primary dealers is 0.35). In our sample of Treasury note and bond auctions, the median maturity is 5 years. The winning yield ("high yield") is on average close to 3.2 percent per year with standard deviation of 1.9 percentage points. The distribution of submitted bids tends to be fairly compressed: the high-median yield spread is approximately 3 basis points with standard deviation of 2 basis points. However, on some occasions the spread can be as high as 10 basis points.[8]

2.3 Quantifying Demand Shocks

In this section, we describe how we measure the surprise movements in prices of Treasure futures around Treasury auctions and document properties of these surprises. Our key assumption is that within small enough windows around the close and release of Treasury auction results, shifts in the prices of Treasury futures reflect unexpected changes in market beliefs about the demand for Treasuries with a specific maturity. Indeed, since the Treasury announces an offered amount well before an auction happens thus fixing supply, between the announcement and close of the auction futures prices should move only in response to changes in demand conditions. By focusing our analysis on a narrow window around the close time of an auction, we likely isolate variation only due to unexpected shifts in demand for this specific auction. As a result, we can identify a demand shock for a specific maturity and then use this shock to trace the reaction of Treasury futures prices for the given maturity and for other maturities as well as reactions for other parts of the financial market.

[8] Between 1999-2015 when the data is available, the Fed purchased Treasuries through SOMA in approximately two thirds of auctions; when doing so they purchased an average of \$2.3 billion (standard deviation of \$2 billion).

2.3.1 Shock Construction

Let $P^{(m)}_{t,pre}$, $P^{(m)}_{t,post}$ be the futures prices before and after the close of the auction on date t with maturity $m = 2, 5, 10, 30$. We measure the surprise movements in the futures prices as:

$$D^{(m)}_t = \log P^{(m)}_{t,post} - \log P^{(m)}_{t,pre}. \tag{2.1}$$

These surprises are computed for all maturities at date t irrespective of what maturity is being auctioned on date t. In other words, we compute $D^{(2Y)}_t$ (surprise movement in the 2-year Treasury futures) not only for auctions that offer 2-year government notes but also for auctions that offer Treasuries with other maturities.

For all auctions, $P_{t,pre}$ is the last price observed 30 minutes before the close of the auction. For auctions taking place between 1995 and 1999, $P_{t,post}$ is the first price observed 1.5 hours after the close of the auction; after 2000, we use the first price observed 30 minutes after the close of the auction. The Treasury began releasing results much faster in the early 2000s, but in the 1990s auction results frequently took over an hour after the close of the auction to be released. Unlike the close of the auction, the time at which the results are released is not reported by the Treasury. However, wire reports from Bloomberg allow for an upper bound on the release time. Note that we use small symmetric windows around the events to eliminate predictable movements in prices identified in Lou et al. (2013) and Fleming and Liu (2016). Indeed, Fleming and Liu (2016) show that these predictable movements extend to the hours before and after the auction, but near the close of the auction and release of the results the price movements are reactions to the surprises regarding the demand observed at the auction. Hence, the use of small intraday windows is key to identifying unanticipated demand shocks.

Figure 2.6 plots the time series of our constructed shock measures, with summary statistics presented in Table 2.2. Panel A of Table 2.2 reports summary statistics for $D^{(m)}_t$ shocks during auction dates (our main sample). The mean values of the shocks are close to zero suggesting that surprises are not systematic and do not contain predictable movements. The standard deviation of $D^{(m)}_t$ increases in maturity m. To verify that these shocks are not spurious we also report (Panel B of Table 2.2) movements in futures prices on non-auction days (for days without auctions, the same "pre" and "post" windows are used as auctions in the same year). In all cases, the variance of the shocks on auction dates is larger than on non-auction dates. This pattern is consistent with auction results indeed influencing futures prices.

The table also reports moments for the zero lower bound (ZLB) and pre-ZLB periods. The variability of surprises for short maturities is considerably smaller during the ZLB period (December 2008 to the end of our sample) than outside the ZLB period (1995 to December 2008). For longer maturities, the volatility is similar for ZLB and pre-ZLB periods. However, these statistics mask important heterogeneity. As seen in Figure 2.6, during the Great Recession the volatility of surprises was elevated but then we observe strong compression for short maturities since the economy enters recovery. This finding is consistent with Swanson and Williams (2014) documenting that while the Fed's policies during the Great Recession compressed fluctuations of short-term rates, the behavior of long-term rates is still relatively normal.

Note that the shocks are in terms of futures (log) prices. Although futures contracts do not have a natural definition of yield, an approximate yield can be computed using the Treasury securities delivered at the end of the contract. Using this approximation, a one standard deviation change in the log price of each contract is equivalent to a 2.0 to 2.5 basis point change in yield for each contract.[9]

Additionally, Table 2.2 documents that price changes of Treasury futures strongly comove across maturities, with the strongest correlations between adjacent maturities. For example, on auction dates the correlation between $D_t^{(10Y)}$ and $D_t^{(30Y)}$ is 0.922 while the correlation between $D_t^{(2Y)}$ and $D_t^{(30Y)}$ is 0.672. Note that the correlations are generally stronger between short ($D_t^{(2Y)}$) and longer maturities during the non-ZLB period than during the ZLB period. At the same time, the comovement of $D_t^{(5Y)}$, $D_t^{(10Y)}$ and $D_t^{(30Y)}$ does not appear to be materially influenced by the binding ZLB. These correlations suggest that shocks to a given segment of the maturity spectrum generally affect not only prices of that particular segment but also prices in other parts of the spectrum, but there is heterogeneity across time in the strength of the correlation. This is a key result, which we explore in detail in Section 4.

2.3.2 Narrative Evidence

To provide a better understanding of what forces are behind these surprise movements, we plot the 30-year Treasury futures price during two 30-year Treasury bond auctions (Figure 2.7). The first is from an auction on August 11, 2011. Futures prices were relatively stable in the lead up to the close of the auction, but after the close and release of the auction results prices dropped sharply and immediately. The Financial Times wrote:

> "An auction of 30-year US Treasury bonds saw weak demand...bidders such as pension funds, insurers and foreign governments shied away. 'There's not too many ways you can slice this one, it was a very poorly bid auction.'"

The second is from December 9, 2010. This auction was a reopening of previously issued 30-year bonds from the month prior. Once again, the futures prices are relatively stable in the lead up to the close of the auction. After the auction closes and results are released, prices immediately spiked up. The Financial Times wrote:

> "Large domestic financial institutions and foreign central banks were big buyers at an auction of 30-year US Treasury bonds on Thursday. 'Investors weren't messing around...You don't get the opportunity to buy

[9] For details on how to convert between Treasury futures prices and the yield on the corresponding "cheapest-to-deliver" Treasury security, see Lauszweski et al. (2014).

large amounts of paper outside the auctions and 'real money' were aggressive buyers.'"

We interpret the two example auctions as follows. Before the auction closes, the market information set consists of all the supply information, both for outstanding securities as well as the amount on offer for the current 30-year auction. The 30-year futures prices reflect beliefs about the expected path of short-term interest rates, inflation expectations, and demand for long-term Treasury securities. After the auction closes and the results are released, the only update to the information set is the news regarding the bidding that took place in the auction, which solely reflects demand for Treasury debt. The change in the 30-year futures price reflects this unexpected shift in beliefs about Treasury demand. The contemporaneous articles in the financial press further suggest that the important driver of the demand shifts arise from foreign and domestic institutional investors. The last example also highlights why auctions can have important elements of price discovery: when investors have to purchase large amounts of Treasuries to meet their needs, they may prefer to use auctions rather than attempting to make substantial transactions on the secondary market. As a result, auctions reveal new information about demand.[10]

2.3.3 Demand Determinants

Our assumption is that $D_t^{(m)}$ captures unexpected shifts in the demand for Treasuries. We further hypothesize that these shocks are particularly driven by demand shifts arising from institutional investors. Figure 2.7 and the corresponding reporting in the financial press provided some narrative evidence in this direction. However, $D_t^{(m)}$ is a market-based measure and hence is an equilibrium response to the underlying shifts in demand. Because the mapping from shifts in demand to changes in futures prices may be complex, it is important to establish that the market interpretation of changes in demand is actually related to observable movements in demand.

One of the most commonly reported statistics in the financial press is the bid-to-cover ratio. It is a natural measure of the demand at a given auction (the higher is the bid-to-cover ratio, the higher is demand). The bin scatter plot in Figure 2.8 shows that the bid-to-cover ratio (after controlling for its four own lags) is a strong predictor of our measure of demand shocks. Table 2.3 presents more formal evidence by regressing our shocks on measures of demand reported at the auction:

$$D_t^{(m)} = \alpha^{(m)} + \beta^{(m)} X_t^{(m)} + \varepsilon_t^{(m)} \tag{2.2}$$

This specification is estimated separately for auctions corresponding to the Treasury futures maturity groups in columns (1)-(4). For example, column (1) restricts the

[10]We could not find any reference in the press about monetary policy (or leaked information about future monetary policy) being a source of unexpected movements. Consistent with this observation, we do not find any statistical power of surprise movements in Treasury futures around Treasury auctions to predict future monetary policy.

sample to include only auctions of 2-year notes and column (2) restricts the sample to include only auctions of notes with (2,5] year maturity. Column (5) reports results when we pool across maturities and impose that $\beta^{(m)}$ is the same across maturities m. To facilitate the comparison of the results, we standardize $D_t^{(m)}$ in these regressions to have zero mean and unit variance.

Panel A of Table 2.3 estimates equation (2.2) using the current bid-to-cover ratio as well as four lags of the bid-to-cover ratio (winsorized at the 1% level). These coefficients may be interpreted as the reaction of $D_t^{(m)}$ surprises to an innovation in the bid-to-cover ratio, and correspond to the slopes of the regression lines in Figure 2.8. The results show that the bid-to-cover ratio is positively associated with $D_t^{(m)}$ and the effect of an increase in the bid-to-cover ratio is statistically and economically significant. For example, a one standard deviation (0.5) increase in the bid-to-cover ratio (after controlling for its own four lags) in a Treasury auction for 30-year bonds raises the price of the 30-year Treasury futures by $2.119 \times 0.5 = 1.06$ standard deviations (this corresponds to a 0.26 log point increase in the price of the Treasury futures or an approximate change of 2.5 basis points in the yield).[11]

Panel B repeats the regressions from Panel A, but explicitly decomposes the bid-to-cover ratio into "expected" and "surprise" components. For these regressions, we first estimate a univariate AR(4) model of the bid-to-cover ratio, separately for each maturity group. We then construct the fitted (expected) and residual (surprise) values of the bid-to-cover ratio, and regress $D_t^{(m)}$ on these expected and surprise components. We find that the variation in our demand shocks is determined by the surprise component of the bid-to-cover ratio, and is unaffected by expected movements in the bid-to-cover ratio.

In order to assess sensitivity of futures prices to changes in demand by bidder type, Panel C reports estimates of equation (2.2) using the bid-to-cover ratio of Indirect Bidders, Direct Bidders, and Primary Dealers. The sensitivity of surprises D to unexpected demand of indirect bidders increases with maturity. For example, a unit increase in the bid-to-cover ratio for indirect bidders raises the price of 2-year Treasury futures by 2.7 standard deviations and the price of 30-year Treasury futures by 8.5 standard deviations. For direct bidders, the sensitivity is highest for short maturities. The sensitivity to changes in the bid-to-cover ratio coming from primary dealers for 2- and 5-year Treasury futures is smaller than the sensitivity for 10-year Treasury futures and greater than the sensitivity for 30-year Treasury futures. When we pool across maturities, demand of direct and especially indirect bidders generates *ceteris paribus* more variation in futures prices than demand of primary dealers.

Panel D uses additional investor allotment data from the Treasury to break down the amount accepted by types of bidders: Investment Funds, Foreign, Dealers, and Miscellaneous. Since the fractions by group add up to one, we set Dealers as the leave-out category. The estimated coefficients suggest that as the fraction accepted

[11] We found that controlling for other variables (e.g., policy uncertainty constructed in Baker et al. (2016)) in equation (2.2) does not materially change our estimates.

for investment funds and foreign buyers increases, $D_t^{(m)}$ increases too. The coefficients for the Miscellaneous category are generally smaller and less robust.

These results indicate that, indeed, a key determinant of $D_t^{(m)}$ surprises is movements in demand conditions as proxied by the bid-to-cover ratio. Furthermore, we observe that the demand from institutional investors is important in accounting for variation in $D_t^{(m)}$. In subsequent analyses, we will use this strong relationship to instrument $D_t^{(m)}$ with unexpected movements in the bid-to-cover ratio so that for our estimates we exploit variation due to changes in demand rather than variation due to fluctuations in market conditions.

2.3.4 Comovement Across Markets

We now turn to analyzing how our demand shocks for Treasuries propagate across other financial markets. Given the relatively high degree of correlation across our demand shocks, in the following analysis we will find it useful to compress $D_t^{(m)}$ into a single summary statistics: the first principal component of $D_t^{(m)}$. This time series captures the general movement of the yield curve in response to demand shocks for government debt with various maturities. The first principal component explains 88 percent of variation in our shock measures. We denote the first principal component by D_t, which has zero mean and unit variance.

We measure the impact of demand shocks on other asset prices by running simple bivariate regressions:

$$y_t = \gamma + \phi D_t + u_t \tag{2.3}$$

where y_t is the change in the price or yield of some asset on auction date t. Where available, we use intraday changes within the same time window as our shocks D_t. However we also examine changes at the daily frequency, partly due to data limitations but also because daily changes may pick up responses in other asset markets that don't occur immediately. A strong correlation between D_t and y_t signals either that D_t and y_t have a common determinant (e.g., changes in inflation expectations alter the behavior of bids in Treasury auctions and change prices of inflation swaps) or that y_t is a channel of propagation for D_t shocks (e.g., unexpected prices in an auction result in repricing of Treasuries in the secondary market). To preserve space, we focus on the OLS estimates of equation (2.3) and report very similar instrumental variable estimates (using unexpected changes in the bid-to-cover ratio as instrumental variables) in Appendix Table B1.

Panel A of Table 2.4 reports results for debt markets. The dependent variable in the first two rows are the intraday change in the Exchange Traded Funds (ETF) "TLT" and "SHY", which track Barclays Capital U.S. 20+ Year and 1-3 Year Treasury Bond indices, respectively. The third row of the panel is the intraday change in the ETF "LQD", which tracks the iBoxx Liquid Investment Grade Index. The coefficient should be interpreted as the impact in log points of a one standard deviation change in D_t. In all cases we observe a strong reaction to the Treasury demand shock, accounting for more than 50 percent of variation observed in these ETF prices during the short windows around the close and release of the Treasury auction results.

The final row reports the results for the daily change in corporate bond yields, as measured by Moody's Aaa corporate yields. Consistent with the intraday results, the negative coefficient implies that an increase in the price of Treasury futures (which means that the yield on Treasuries falls) is associated with a decrease in the Aaa bond yields. Specifically, a one standard deviation shock to D_t decreases the Aaa bond rate by 2.3 basis points. However, using daily rather than intraday changes as the dependent variable leads to a decline in R^2s, which underscore the benefits of using intraday data.

As expected, yields in the secondary market react strongly to the demand shock. Furthermore, this reaction is persistent in spite of the fact that our shocks are constructed from intraday movements. Figure 2.9 plots the contemporaneous reaction of 10-year Treasury spot rates (top panel) and the Aaa corporate bond yields (bottom panel) to our shocks D_t, as well as the reactions up to 60 days in the future. The reaction remains strongly statistically significant nearly 1 month later, while the point estimate is remarkably stable even 2 months later.

Panel B of Table 2.4 reports results for equities. Rows 1 and 2 report the results for the intraday change in ETFs tracking the S&P 500 and the Russell 2000 indices. Rows 3 and 4 are for the daily changes in these indices. Although the estimated slope is generally negative, the quantitative significance of shocks on equities is small: these shocks account for a tiny share of variation in equities.

Panel C of Table 2.4 presents results for inflation expectations and commodities. Rows 1 and 2 report the results for the daily change in inflation expectations implied by inflation swaps at the 10-year and 2-year horizon. We observe that demand shocks for Treasuries do not generate significant movements in inflation expectations. To explore the robustness of this finding, we examine price reactions of two additional assets which are often used to hedge against inflation. The dependent variable in row 3 is the intraday change in the ETF "GLD," which tracks the price of Gold Bullion. Row 4 reports results for the daily change in the S&P Total Commodity Index. For neither of these variables do we find a significant correlation with D_t. To further explore sensitivity of inflation expectations to demand shocks D_t, we plot reactions of inflation swap rates at all available maturities in Appendix Figure B1. We find that the change in the inflation expectation "yield" curve exhibits little reaction to D_t. We find a similar lack of sensitivity of inflation expectations when we use specific $D_t^{(m)}$ instead of summary series D_t. These results suggest that the demand shocks we measure at auctions are not driven by some underlying change in inflation expectations. Moreover, these shifts in demand do not propagate to changes in inflation expectations. This result is intuitive, as changes in demand of institutional investors or foreign monetary authorities are unlikely to generate future fluctuations in the rate of U.S. inflation.

Panel D of Table 2.4 reports results for various bond spreads and credit default swaps. The first row reports results for the daily change in the Moody's Baa-Aaa corporate yield spread. Rows 2 and 3 use daily changes in two CDS indices from Credit Market Analysis (CMA) that track the automotive industry (a highly cyclical industry) and banks (a proxy for the financial sector). These three measures proxy for expectations about future output and market conditions. We find that surprise movements D_t have no tangible effect on these measures, consistent with the view that D_t shocks do not capture superior information of Treasury auction bidders about future recessions and the like. In row 4, we document that D_t shocks are not associated with VIX (a measure of market perceptions about future volatility). Hence, it is unlikely that there is a common force that moves D_t and volatility or that D_t shocks propagate via volatility. Finally, we find (row 5) that D_t shocks are not significantly related to the 3-month LIBOR-Overnight Index Swap (OIS) spread (a measure of short-term liquidity risk) so that liquidity fluctuations are unlikely channels or determinants of D_t. In short, as with the case of inflation expectations, these null results suggest that our demand shocks are not being driven by changes in expectations regarding output, liquidity, default risk, or volatility

The results of Tables 2.3 and 2.4 allow for some broad observations. First, given our high-frequency approach and the structure of Treasury auctions, we know that our constructed shocks are only driven by new information regarding the demand side of the market. Second, these shifts are largely driven by shifts in the demand that arises from institutional investors. Third, as expected these demand shocks from the primary market propagate not only to the secondary market, but also to the corporate debt market. Finally, these demand shifts are not driven by some underlying shift in macroeconomic expectations (flight to quality, inflation expectations, etc.) that may move demand for Treasuries at all maturities. But it still remains the case that a variety of factors can generate movements in $D_t^{(m)}$ and one should not interpret $D_t^{(m)}$ as structural shocks.[12] Despite this limitation, the properties of $D_t^{(m)}$ shocks allow us to study how unexpected demand interventions at specific maturities propagate to other maturities.

2.4 Channels of Treasury Demand Shocks

Although demand shocks $D_t^{(m)}$ strongly comove with one another, the responses are not uniform across maturities. To see the heterogeneity in reactions, we use daily changes in zero-coupon spot rates as constructed in Gurkaynak et al. (2007), which

[12]To highlight this caveat, we plot sensitivities for select asset prices estimated over rolling windows in Figure B2. The sensitivity of LQD prices, Aaa interest rates, and Baa-Aaa spread is relatively stable over time. On the other hand, the sensitivity of S&P500 flips sign from positive in the late 1990s to generally negative since the early 2000s, which is consistent with Campbell et al. (2014). Although we do not have long-time series of inflation expectations (or assets used for hedging against inflation), we observe that during the Great Recession in the U.S., inflation expectations and D_t moved in opposite directions while in normal time these two series are approximately uncorrelated.

provide more granularity to the analysis (recall that we have only four maturities in Treasury futures contracts while the yield curve utilizes information for many more maturities but is available only at the daily frequency). Figure 2.10 plots "responses" of changes in the yield curve for auctions on two dates. On the first date (August 11, 2011), there was unexpectedly weak demand (as measured by changes in futures prices) during an auction of 30-year Treasuries. We observe that, although the whole yield curve shifted up, the strongest reaction was at long maturities. On the second date (February 6, 2007), there was an auction of 3-year government notes, and demand during this auction was unexpectedly strong. The whole yield curve shifted down, but the strongest reaction was at the 8-year maturity. These two cases illustrate that the "propagation" of demand shocks across maturities does not amount to simple upward or downward shifts.

This raises the question: to what extent do Treasury demand shocks have local effects? In other words, does the location of the demand shock in maturity space matter? And are the impacts state-dependent? The two auctions in Figure 2.10 provide suggestive evidence that the location can in fact matter. In order to better characterize the impact of these demand shocks, we now examine the impact on the term structure of Treasury rates through the lens of the *preferred habitat* model of investor demand. The key idea is the existence of "clientèle" investors who have idiosyncratic demand for Treasuries of specific maturities. The other side of the market are risk-averse arbitrageurs, who smooth out these demand shocks. Using a version of the model from Vayanos and Vila (2009), we create qualitative predictions of what happens to the term structure when hit with demand shocks to various parts of the maturity space during different economic regimes.

2.4.1 Preferred Habitat – Numerical Exercise

In our numerical exercises, we consider a "three-factor" version of the Vayanos and Vila (2009) model consisting of the instantaneous rate, and two demand factors that are otherwise equivalent, but are located in the "short" and "long" ends of the maturity space. We then solve the model and study the impact of each demand shock. A key element of the model is the level of risk aversion of the arbitrageurs, hence we study the reactions as risk aversion increases from very low to very high. We consider the case where the "short" shock is concentrated at the 3-year maturity, while the "long" shock is at 20 years (corresponding to the average length of short-term and long-term auctions in our empirical section). See Appendix A for details regarding the model and parameterization.

Figure 2.11 shows the change of the term structure in response to the short (top panel) and long (bottom panel) demand factors, as the risk aversion of arbitrageurs increases from low (lighter lines) to high (darker lines). In the case of low risk aversion, the impact is very similar: rates fall across the entire term structure, but the impact peaks at the short end of the yield curve, then drops off as the maturity increases.

The only difference between the impacts of the short and long shocks is that the long shock has a larger impact; but the shape of the response is nearly identical.

However, as risk aversion increases, the responses become quite different depending on the location of the demand shock. A shock to the short demand factor sharply decreases short-term rates; but this impact dies off quickly and even turns slightly positive at the long end of the term structure. On the other hand, the long demand factor decreases both short and longer term rates, but the impact is much stronger in the long end of the term structure.

These results confirm that some of the findings of Vayanos and Vila (2009) for the limiting cases of no risk aversion and infinite risk aversion also hold for intermediate cases of risk aversion. As they explain, the intuition for these results is as follows: when arbitrageurs are perfectly risk-neutral, demand shocks have no impact as the expected path of the instantaneous rate is the only determinant of the term structure. As arbitrageurs become somewhat risk averse, shocks to the instantaneous rate continue to be much more influential than demand shocks. But now arbitrageurs are concerned about instantaneous rate risk. Every bond is sensitive to instantaneous rate risk, and the market price of this source of risk is determined in equilibrium by the portfolio allocations of arbitrageurs. Because demand shocks from preferred habitat investors cause changes in these portfolios, even in cases of very low risk aversion these demand shocks do affect the term structure by altering the price of instantaneous rate risk.

What is the impact of the location of the demand shock? Consider an increase in preferred habitat demand. Regardless of the location of the shift, this causes arbitrageurs to sell bonds, reducing their exposure to instantaneous rate risk. Hence they require lower expected returns to hold bonds, pushing down rates. The location of the demand shock determines which bonds arbitrageurs sell; since these bonds have different sensitivity to instantaneous rate risk, this determines the magnitude of the overall reduction in rates. But regardless of the location of the demand shift, the bonds that respond the most are those most sensitive to instantaneous rate risk, which depends only on the stochastic properties of the instantaneous rate and demand shocks. In other words, when arbitrageurs have low risk aversion, the relative impact of short and long demand shocks to the term structure is roughly the same; only the overall size of the impact is affected by the location of the demand shock. In our calibration, for low values of risk aversion this leads to the peak impact occurring around $m \approx 4$. This could be lower or higher with different parameterizations, but remains largely independent of the location of the shock in maturity space.

As arbitrageur risk aversion increases, demand shocks become more prominent as additional sources of risk. Arbitrageurs try to limit their exposure to these sources of risk, leading to less propagation from the location of the demand shock to other parts of the term structure. This implies that the term structure response is more localized to each demand shock as arbitrageurs choose not to integrate bond markets across maturities. For example, following a demand shock for short-term bonds, all else equal arbitrageurs would like to buy longer term bonds to hedge the risk arising from the short demand shock, leading to upward pressure on prices (and downward pressure on spot rates). But this changes their exposure to the long demand shock as well, and this countervailing force leads arbitrageurs to sell sufficiently long-term bonds. As seen in the top panel of Figure 2.11, when risk aversion is sufficiently high this countervailing force can become strong enough to lead to an increase in rates for very long-term bonds.

In summary, this illustrative exercise indicates several qualitative predictions. When risk aversion is low, the impact of an increase in demand for either short-term or long-term debt causes a decrease in rates everywhere. Moreover, while the magnitude may differ, the response to both demand factors are similarly shaped, peaking at intermediate maturities and declining for very long-term maturities. Conversely, when risk aversion is high, demand factors have a stronger local component: increases in demand for short-term debt will have a maximal impact on shorter-term maturities, while long-term shocks will peak at long-term maturities. Additionally, the response of long-term (short-term) rates to long (short) demand shocks increases as risk aversion increases, respectively. Finally, although the magnitude of responses is more ambiguous, when risk aversion is high we expect the response of short-term rates to short demand shocks will be larger than to long demand shocks; and vice versa for responses of long-term rates.

2.4.2 Empirical Results

Comparing the theoretical results from Figure 2.11 with Figure 2.10 suggest that, at least during the auction in Panel A, the preferred habitat model with relatively high risk aversion does a good job explaining the response of the term structure; Panel B is more ambiguous. We now take a more rigorous approach to testing the predictions of our numerical exercise.

A key variable is a measure of risk aversion of arbitrageurs. We proxy this using the measure of financial crises in the United States from Romer and Romer (2017).[13] The crisis index is a continuous measure derived from narrative sources to identify periods of financial distress (higher values correspond to periods of more extreme financial crisis). Besides identifying financial distress during the recent financial crisis, the measure also identifies periods of distress in 1986, the early 1990s, and 1998 (Figure B3 in the Appendix).

[13]He and Krishnamurthy (2013), Kyle and Xiong (2001b) and others show how risk aversion can be endogenously higher in times of crises. Note that in contrast to other popular measures of financial stress (e.g., the Federal Reserve Board staff's Financial Stress Index), the Romer-Romer index does not use yields on Treasuries (outcome variables in our excercise) to identify stress/non-stress periods. Additionally, rather than using a narrative measure of financial distress, we used a market-based proxy for arbitrageur risk aversion. We define high risk aversion periods as those in which the "intermediary capital ratio" described in He et al. (2016) is low (Appendix Figure B4), and find similar results.

In order to measure the impact of demand shocks on the entire term structure, we estimate the following regression equations

$$\Delta R_t^{(m)} = C_t \left(\alpha^{(m,H)} + \beta^{(m,H)} D_t^{(m')}\right) + (1 - C_t)\left(\alpha^{(m,L)} + \beta^{(m,L)} D_t^{(m')}\right) + \varepsilon_t^{(m)} \quad (2.4)$$

for each maturity $m = 1, \ldots, 30$. $\Delta R_t^{(m)}$ is the daily change in the zero-coupon spot rate for the given maturity as measured by Gurkaynak et al. (2007). C_t is an indicator variable that is equal to 1 when the Romer and Romer (2017) measure of financial crisis is non-zero. The coefficients $\beta^{(m,L)}$ capture the impact of demand factor $D_t^{(m')}$ at maturity m' (our normalized intraday futures price shocks) during periods of low risk aversion; similarly $\beta^{(m,H)}$ capture the impact during periods of high risk aversion. While our shock measure is constructed at a higher intraday frequency, in order to capture the full extent of how markets absorb these shocks we prefer to use these daily estimates of the yield curve. To the extent that shocks are absorbed completely within smaller windows than a day, using daily changes as an outcome variable simply adds noise to our estimates, but shouldn't result in any bias.

A straightforward way in which to test the predictions of the preferred habitat model is to estimate equation (2.4) in two separate subsamples: *i*) days with short auctions; *ii*) days with long auctions. In our baseline regression we break up auctions into 2-7 years and 10-30 years. We choose the 10-year cutoff for long vs. short rather than 30-year in order to have a more balanced sample; the results are robust to choosing different cutoffs.[14] Breaking up the auctions in this manner allows us to more closely pinpoint the location of the demand shock in the maturity space, and ties closely with the numerical exercise above.

For our measure of demand shocks $D_t^{(m')}$ on the right-hand-side of equation (2.4), in our baseline results we take the same approach above and match each auction with the (normalized) futures surprises of closest maturity (e.g. for 5-year auctions, use the 5-year futures surprise). The β coefficients should be interpreted as the response of spot rates for maturity m to a one standard deviation demand shock at maturity m' on the day when maturity m' is auctioned.

Figure 2.12 plots the low and high coefficients from the two subsamples (Appendix Figure B5 plots p-values testing for equality of the coefficients). During periods of low risk aversion, the impact of short and long demand shocks on the term structure closely mirror one another. Both shocks decrease spot rates across the entire term structure, and are hump-shaped. But when risk aversion is high, the short and long demand shocks have differential impacts. Both shocks exhibit stronger local effects. For the long shock, the impact is no longer hump-shaped as the impact continues to remain large as the maturity increases. The magnitude is also considerably larger than the corresponding responses during periods of low risk aversion. The impact of the short demand shock peaks at intermediate rates and then begins declining;

[14]Consistent with our exercise in Section 4.1, the average maturities of "short" and "long" auctions are 3 and 20 years respectively.

further, the magnitude is larger than the corresponding response in periods of low risk aversion. Finally, when risk aversion is high, the response of short-term rates to short demand shocks is larger than their response to long demand shocks; and vice versa for the response of long-term rates.[15]

These results confirm the key predictions of our numerical exercise: during periods of low risk aversion, short and long demand shocks have relatively similar impacts; and these impacts peak at short to intermediate maturities. During periods of high risk aversion, the impacts are more localized: the impact of short demand shocks peak at short maturities, while the impact of long demand shocks peaks at the long end of the term structure. This local effect is particularly strong for long-term demand shocks. We also find that the peak responses for both short and long demand shocks are larger during periods of high risk aversion than during periods of low risk aversion.

As discussed above, the shocks $D_t^{(m')}$ are equilibrium reactions of market prices to changes in demand for Treasuries. As a result, these reactions may depend on market conditions (specifically on whether financial markets are disrupted) so that measured responses combine reactions to changes in demand and to how futures prices react to changes in demand. To isolate the effect of changes in demand for Treasuries, we employ unexpected changes in the bid-to-cover ratios for Treasury auctions as instruments for $D_t^{(m')}$ in equation (2.4). That is, we instrument $C_t \times D_t^{(m')}$ and $(1 - C_t) \times D_t^{(m')}$ with $C_t \times b_t^s$ and $(1 - C_t) \times b_t^s$, where b_t^s is the surprise movement in the bid-to-cover ratio for a Treasury auction at date t (as in Panel B of Table 2.3). This approach permits state-dependent mappings from demand shocks to futures prices. Consistent with our earlier results, we find that the unexpected changes in the bid-to-cover ratios are strong instruments for D_t (first-stage F-statistics above 10). The IV estimates of estimated reactions of spot rates (Figure 2.13) are similar to the OLS estimates (if anything, the responses during periods of financial distress are larger in magnitude, although confidence bands are wider) thus reassuring that our shocks D_t are good measures of the underlying demand shifts and hence the OLS estimates are capturing the response of the yield curve to these demand shifts.

As a robustness check, rather than use demand shocks as identified by intraday Treasury futures changes, we can also use the surprise component of the bid-to-cover ratio directly as a proxy of demand shocks from the auctions themselves. We re-estimate equation (2.4) using b_t^s in place of our demand shocks. Although the surprise component of the bid-to-cover ratio is not as clean a measure of demand shocks, this allows us to check the robustness of our results, as well as to expand the

[15] A downside of using the yield curve data from Gurkaynak et al. (2007) is that idiosyncratic changes along portions of the yield curve may be smoothed out. To address this concern, we repeat our empirical exercise but use security-level yield changes in the secondary market for Treasuries. In place of the Gurkaynak et al. (2007) rate changes, we use the daily changes in the yields for Treasury notes and bonds in the secondary market from CRSP. We can no longer trace out changes to zero-coupon rates at all points along the maturity space, but by grouping together Treasuries of similar maturities we can test for local demand shocks more directly. Appendix Table B3 confirms the findings from Figure 2.12.

sample period to 1979-2015. Appendix Figure B7 plots the results using the same sample 1995-2015 as in Figure 2.12, while Appendix Figure B9 uses the entire sample 1979-2015 (p-values are in Appendix Figures B8 and B10, respectively). As expected, the standard errors are a bit wider, but the qualitative responses are generally similar.

We also tried a number of additional empirical exercises, and find our results are robust to a variety of different specifications, including different cutoffs for short-term and long-term auctions and using different subsamples.[16]

2.5 Implications for QE

The responses of the yield curve to unexpected movements in demand during Treasury auctions offer several lessons for how one should understand the workings of QE programs implemented by the Fed and other central banks. For example, if the Fed is trying to decrease long-term Treasury rates relative to shorter-term rates, our results suggest that QE policies that directly purchase long-term Treasuries should be highly effective during financial crises. But if the Fed is trying to move the entire term structure of interest rates, during periods of high financial distress the Fed will have to be active in purchasing Treasuries throughout the yield curve. Thus, programs in spirit of "Operation Twist" may be an option because a central bank actively intervenes in multiple segments of the yield curve during a crisis.

As we move away from the most recent crisis, there have been discussions (see Blinder et al. (2016)) about whether central banks will continue to use unconventional policies in the future. Our results suggest that the impact of QE-style policies during non-crisis periods will likely differ greatly from those observed during the crisis. To the extent risk aversion is low and debt markets are more integrated, QE-type programs that attempt to move long-term rates relative to short-term rates may fail. During "normal" times of low risk aversion, the overall response of interest rates is less tied to the location of the shifts in demand. While we still expect targeted purchases of long-term Treasury debt from the Fed to reduce long-term rates, the largest declines may be for shorter term maturities that are not directly purchased by the Fed.

[16]Dropping auctions that occurred during the weeks of QE announcements leads to nearly identical results. In a handful of cases, auctions occur on the same days as FOMC announcements. Although our intraday windows do not overlap with the FOMC announcements, we also re-did our regressions dropping these dates, which leaves our results unchanged.

We additionally highlight one more robustness specification which more closely matches our numerical exercise. We take the first two principal components of our intraday shocks, D_t^ℓ and D_t^s, rotated such that D_t^s is uncorrelated with $D_t^{(30Y)}$ and normalized to have zero mean and unit variance. The first two principal components explain 97 percent of variation in our shocks. For long-term auctions the shock is D_t^ℓ; similarly short-term auctions use D_t^s. In this way we have two distinct "short" and "long" demand factors, which more closely matches our numerical exercise. Appendix Figure B11 plots these results (p-values in Appendix Figure B12), and finds very similar results as the baseline specification. The only difference is the response of long-term rates to short demand shocks falls more closely to zero.

Interestingly, our results suggest that the Fed may have a menu of options in terms of where it can intervene in the maturity space to hit the yield at a target maturity. For example, suppose the Fed wishes to decrease 30-year Treasury rates by purchasing $30 billion of notes and bonds. What is the benefit to targeting purchases in the longer end of the term structure? If the economy is not in a crisis and financial markets are healthy, the answer is not much. Given the size of a typical Treasury note auction in recent years, these $30 billion purchases represent a unit increase in the bid-to-cover. Our estimates imply that if the purchases were for short-term securities (2-7 years), we would expect to see an increase in our normalized demand shocks of 1.43 (standard error 0.19) and an ensuing decrease in secondary market 30-year rates of -1.84 basis points (0.36). If instead these purchases were of long-term securities (10-30 years), we would expect to see an increase in our normalized demand shocks of 1.90 (0.22) and a corresponding decrease in secondary market 30-year rates of -3.03 basis points (0.48). However, if these purchases took place in a period of financial turmoil, the purchases in the long end of the term structure relative to the short end become more effective. In this case, short-term security purchases lead to an increase in our normalized demand shocks of 1.39 (0.31) and a decrease in secondary market 30-year rates of -2.78 basis points (0.75). But long-term purchases lead to increase in our normalized demand shocks of 2.26 (0.38) and an ensuing decrease in secondary market 30-year rates of -7.22 basis points (0.86).

We can also use our results to assess what fraction of the market response to QE1 can be explained directly by shifts in demand for Treasury debt arising from the Fed.[17] To summarize the timeline of the Fed's actions, there were five announcements during QE1, four of which mentioned purchasing long-term Treasury securities. November 25, 2008: the Fed announced purchases of $100 billion in GSE debt and $500 billion in MBS. December 1, 2008: Chairman Bernanke stated that the Fed could purchase long-term Treasuries. December 16: the FOMC announced possible purchases of long-term Treasuries. January 28, 2009: the FOMC announced it is ready to expand agency debt and MBS purchases, and to begin purchasing long-term Treasuries. March 18, 2009: the FOMC announced it will purchase $300 billion in long-term Treasuries, along with an additional $750 billion in agency MBS and $100 billion in agency debt.

Using small intraday windows around the time of the four announcements which mentioned Treasury purchases, Chodorow-Reich (2014) estimates the 5-year Treasury rate reacted by -9.2, -16.8, 3.1, and -22.8 basis points, respectively. For the same dates but using larger 2-day windows to account for the possibility of slow responses due to liquidity effects, Krishnamurthy and Vissing-Jorgensen (2011) estimates the announcements moved the 5-year Treasury rate by -28, -15, 28, and -26 basis points respectively; additionally, they find the 5-year rate moved by -23 basis points after the initial November 25 announcement.[18] This gives a range of cumulative decline of

[17]We focus on QE1 since the surprise component of QE2 and QE3 was likely smaller than that of QE1.

[18]Chodorow-Reich (2014) drops the November 25, 2008, announcement because it occurred after

between 45 and 74 basis points. Note that, because QE1 set the stage for subsequent QE programs, this decline could combine the promise to purchase $300 billion in Treasuries in QE1 with the possibility of additional rounds of quantitative easing that would entail buying more government debt. In other words, the 45-74 basis point decline could overstate the response of the markets relative to a response one could have observed in the case when the Fed credibly committed to spend only $300 billion to purchase government bonds during the entirety of *all* its quantitative easing programs.

With this caveat in mind, we can carry out a back-of-the-envelope calculation to assess how much of the response of yields is due purely to the shift in demand for Treasuries from the Fed. The large majority of the Fed Treasury purchases during QE1 were concentrated in the 2-7 year range, and the magnitude of purchases would correspond to a ten-fold increase in the bid-to-cover ratio observed during auctions over the same period. During this period of financial distress, our estimates imply that we should expect a shock of this size to decrease 5-year secondary market spot rates by 44 basis points (95% confidence interval of 29–59 basis points). Our estimate is close to the estimates from **?**, reporting that Treasury purchases during QE1 reduced yields by about 30 basis points.

Although this exercise represents a very large out-of-sample forecast for our data, it shows that the actual market reaction to QE1 announcements is consistent with the predictions of a preferred habitat model and the behavior of the market in response to observed shifts in *private* demand for Treasuries. Since the mechanism for the market segmentation channel is the same regardless of the source of demand shifts (recall that market segmentation is about how private arbitrageurs absorb these shifts rather than the source of the demand shocks), this finding implies that the *net* effect of other channels of QE (e.g., inflation expectations, forward guidance, signaling) could be smaller than thought before. Consistent with this observation, Krishnamurthy and Vissing-Jorgensen (2012) document that there was little movement in 5-year inflation expectations in response to QE1 announcements.

2.6 Concluding Remarks

Quantitative easing (QE) was a massive policy experiment which likely influenced the economy via multiple channels. To understand how QE worked, we need to unbundle these channels so that future policy can be designed to maximize the effectiveness of QE-like tools in crisis and non-crisis times. In this paper, we focus on the "preferred habitat" channel, which posits that, because of market segmentation and limited arbitrage, interest rates for a given maturity range may be influenced by targeted buying or selling of assets within this range. We utilize Treasury auctions of government debt

trading hours. In addition, the positive response to the January 28, 2009, announcement seems to be because markets were expecting a concrete statement about purchases.

to identify Treasury demand shocks arising from changes in institutional investor demand to study how shocks in one maturity segment propagate to other segments, and how this propagation is affected by the condition of financial markets.

While these shocks do not have structural interpretation, they provide us with variation that is not related to some prominent theories of how QE works (inflation expectations, forward guidance, signaling) and instead allow us to focus attention on the role of preferred habitat mechanisms. Crucially, these mechanisms are dependent on how private agents in the market for Treasury debt absorb these demand shocks, regardless of the source of these shocks. Therefore, we can use this variation to examine whether preferred habitat theory can rationalize responses of interest rates to unexpected changes in demand for government debt with specific maturities during regular Treasury auctions and, by extension, QE rounds.

We find a strong local component of demand shocks (i.e., with some oversimplification, purchases of assets in a particular segment move prices more strongly in that segment), but the local concentration is decreasing in risk-bearing capacity. That is, local effects are stronger when markets are segmented (e.g. due to a crisis) than when markets are integrated. The magnitude of the responses during Treasury auctions is large enough to account for a large part of interest rate movements in response to QE announcements, consistent with the view that QE programs worked mainly via market segmentation. Our analysis suggests that QE can be an effective policy tool in crises, but will be less powerful in moving specific segments of the debt market in normal times. Finally, the net contribution of other hypothesized channels of QE propagation may be quantitatively less important than thought before.

Figure 2.1: Volume and composition of SOMA's holdings of U.S. Government Debt

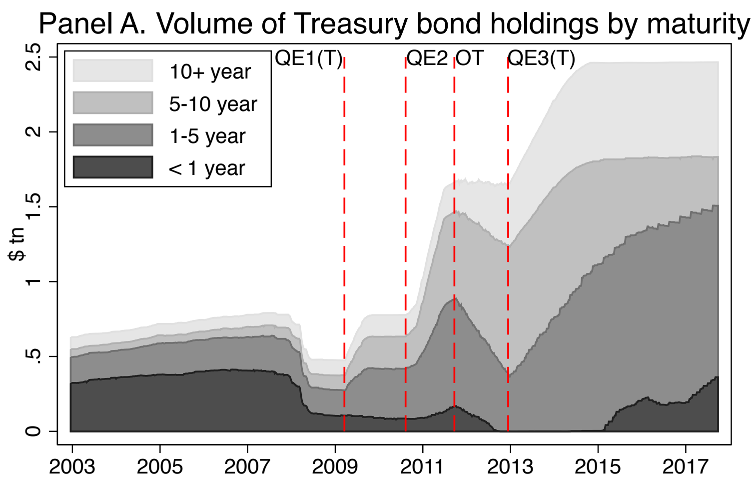

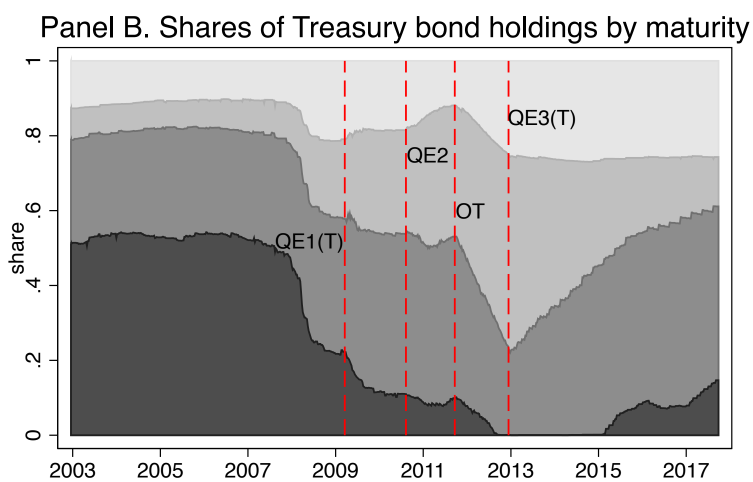

Notes: QE1(T) denotes time when the Fed announced its decision to buy U.S. Treasuries as a part of the first round of quantitative easing. QE2 denotes the announcement of the second round of quantitative easing. QE3(T) denotes the time when the Fed announced its purchases of U.S. Treasuries as a part of the third round of quantitative easing. OT denotes the announcement of "Operation Twist". Source: FRED database.

Figure 2.2: Auction Timing

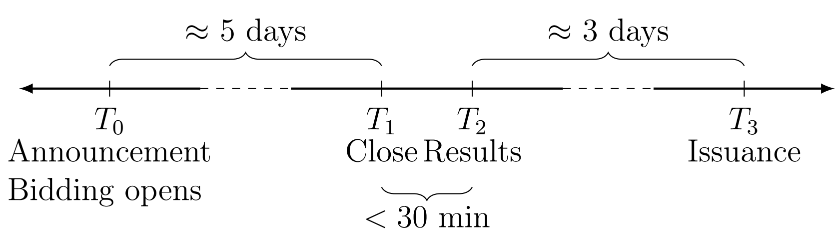

Figure 2.3: Example of an Auction Announcement

August 03, 2011 202-504-3550

TREASURY OFFERING ANNOUNCEMENT [1]

Term and Type of Security	30-Year Bond
Offering Amount	$16,000,000,000
Currently Outstanding	$0
CUSIP Number	912810QS0
Auction Date	August 11, 2011
Original Issue Date	August 15, 2011
Issue Date	August 15, 2011
Maturity Date	August 15, 2041
Dated Date	August 15, 2011
Series	Bonds of August 2041
Yield	Determined at Auction
Interest Rate	Determined at Auction
Interest Payment Dates	February 15 and August 15
Accrued Interest from 08/15/2011 to 08/15/2011	None
Premium or Discount	Determined at Auction
Minimum Amount Required for STRIPS	$100
Corpus CUSIP Number	912803DT7
Additional TINT(s) Due Date(s) and	August 15, 2041
CUSIP Number(s)	912834KP2
Maximum Award	$5,600,000,000
Maximum Recognized Bid at a Single Yield	$5,600,000,000
NLP Reporting Threshold	$5,600,000,000
NLP Exclusion Amount	$0
Minimum Bid Amount and Multiples	$100
Competitive Bid Yield Increments [2]	0.001%
Maximum Noncompetitive Award	$5,000,000
Eligible for Holding in Treasury Direct Systems	Yes
Eligible for Holding in Legacy Treasury Direct	No
Estimated Amount of Maturing Coupon Securities Held by the Public	$24,430,000,000
Maturing Date	August 15, 2011
SOMA Holdings Maturing	$2,205,000,000
SOMA Amounts Included in Offering Amount	No
FIMA Amounts Included in Offering Amount [3]	Yes
Noncompetitive Closing Time	12:00 Noon ET
Competitive Closing Time	1:00 p.m. ET

Figure 2.4: Example of an Auction Result Announcement

August 11, 2011 202-504-3550

TREASURY AUCTION RESULTS

Term and Type of Security	30-Year Bond
CUSIP Number	912810QS0
Series	Bonds of August 2041
Interest Rate	3-3/4%
High Yield[1]	3.750%
Allotted at High	41.74%
Price	100.000000
Accrued Interest per $1,000	None
Median Yield[2]	3.629%
Low Yield[3]	3.537%
Issue Date	August 15, 2011
Maturity Date	August 15, 2041
Original Issue Date	August 15, 2011
Dated Date	August 15, 2011

	Tendered	**Accepted**
Competitive	$33,305,800,000	$15,985,160,000
Noncompetitive	$14,855,600	$14,855,600
FIMA (Noncompetitive)	$0	$0
Subtotal[4]	**$33,320,655,600**	**$16,000,015,600**[5]
SOMA	$489,928,400	$489,928,400
Total	**$33,810,584,000**	**$16,489,944,000**

	Tendered	**Accepted**
Primary Dealer[6]	$23,734,000,000	$10,921,532,000
Direct Bidder[7]	$6,567,000,000	$3,119,654,000
Indirect Bidder[8]	$3,004,800,000	$1,943,974,000
Total Competitive	**$33,305,800,000**	**$15,985,160,000**

Figure 2.5: Number of Auctions per Year

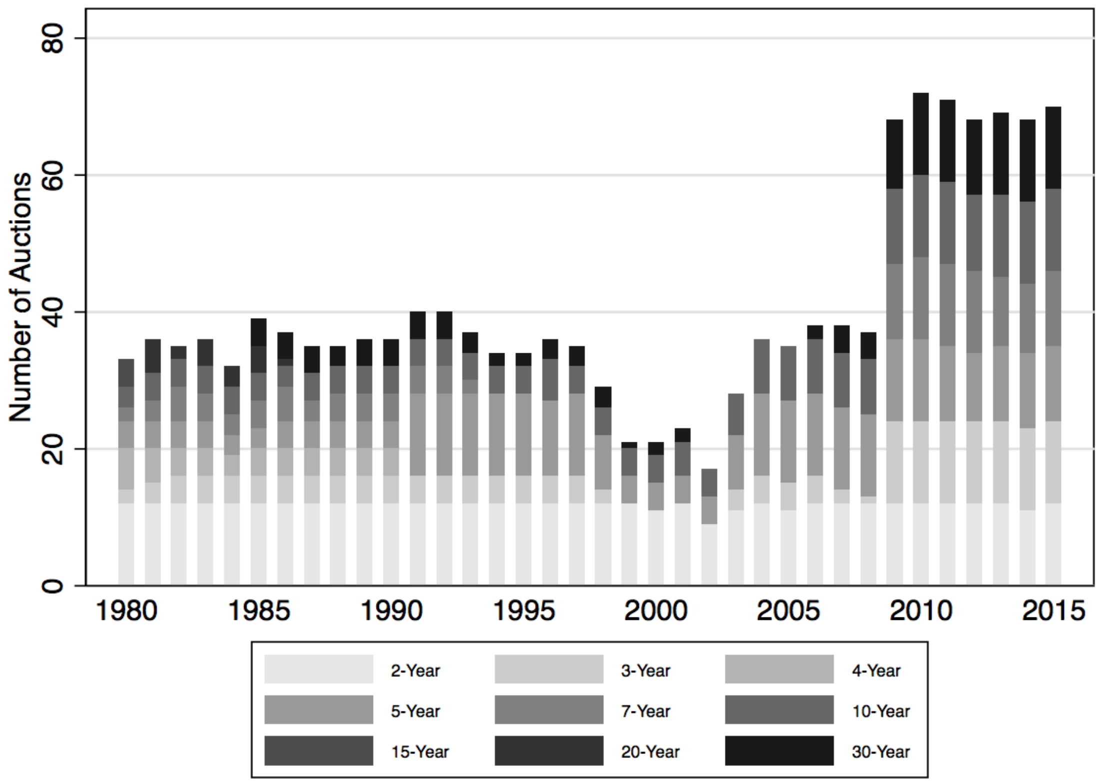

Notes: Number of note and bond Treasury auctions per year by term length in our dataset. The number of auctions temporarily fell in the late 1990s and early 2000s, before increasing sharply after the Great Recession.

Figure 2.6: Time series of surprises in Treasury futures

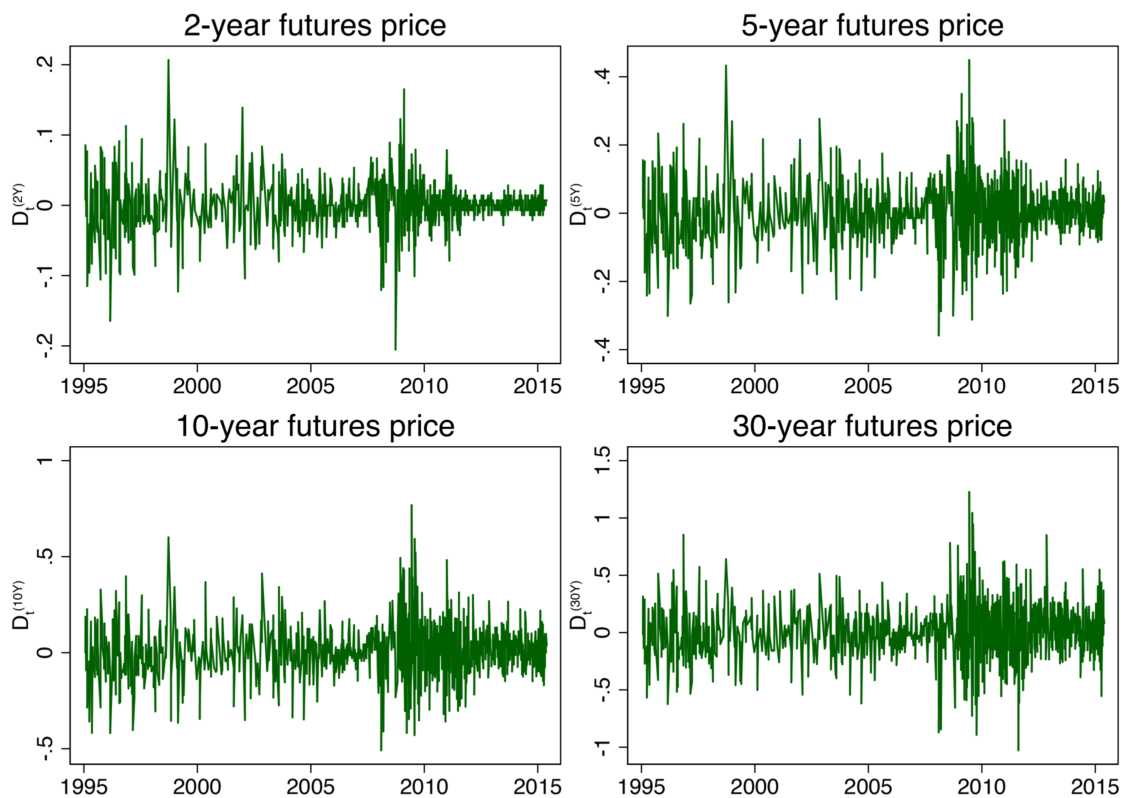

Notes: The figure plots times series of surprise movements in Treasury futures with maturities 2, 5, 10 and 30 years. The movements are reported in log points.

Figure 2.7: 30-year Auctions

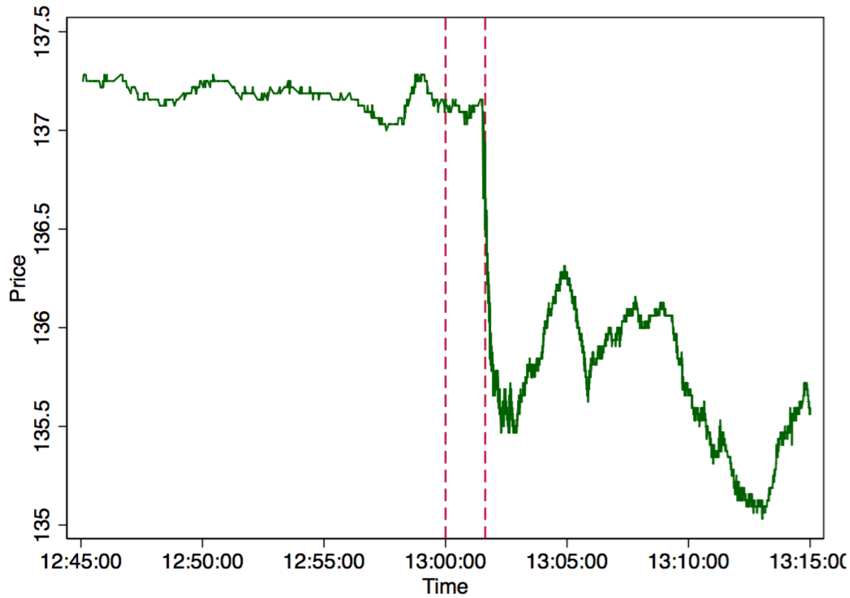

Notes: 30-year Treasury futures prices on August 11, 2011. An auction for 30-year Treasury bonds closed at 1:00pm (first vertical line), and results were released shortly after (second vertical line). Immediately following the release, Treasury futures prices dropped sharply.

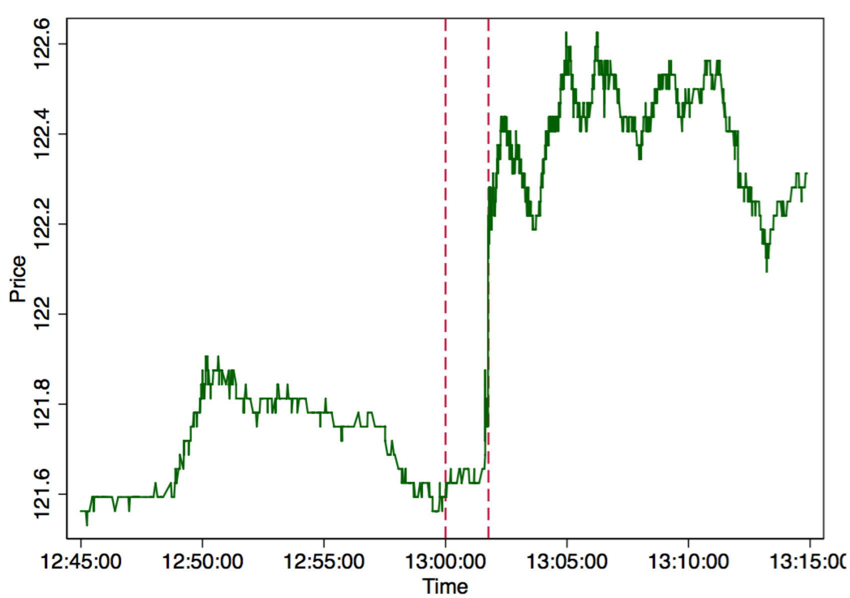

Notes: 30-year Treasury futures prices on December 9, 2010. An auction for 30-year Treasury bonds closed at 1:00pm (first vertical line), and results were released shortly after (second vertical line). Immediately following the release, Treasury futures prices rose sharply.

Figure 2.8: Demand Shocks and Bid-to-Cover

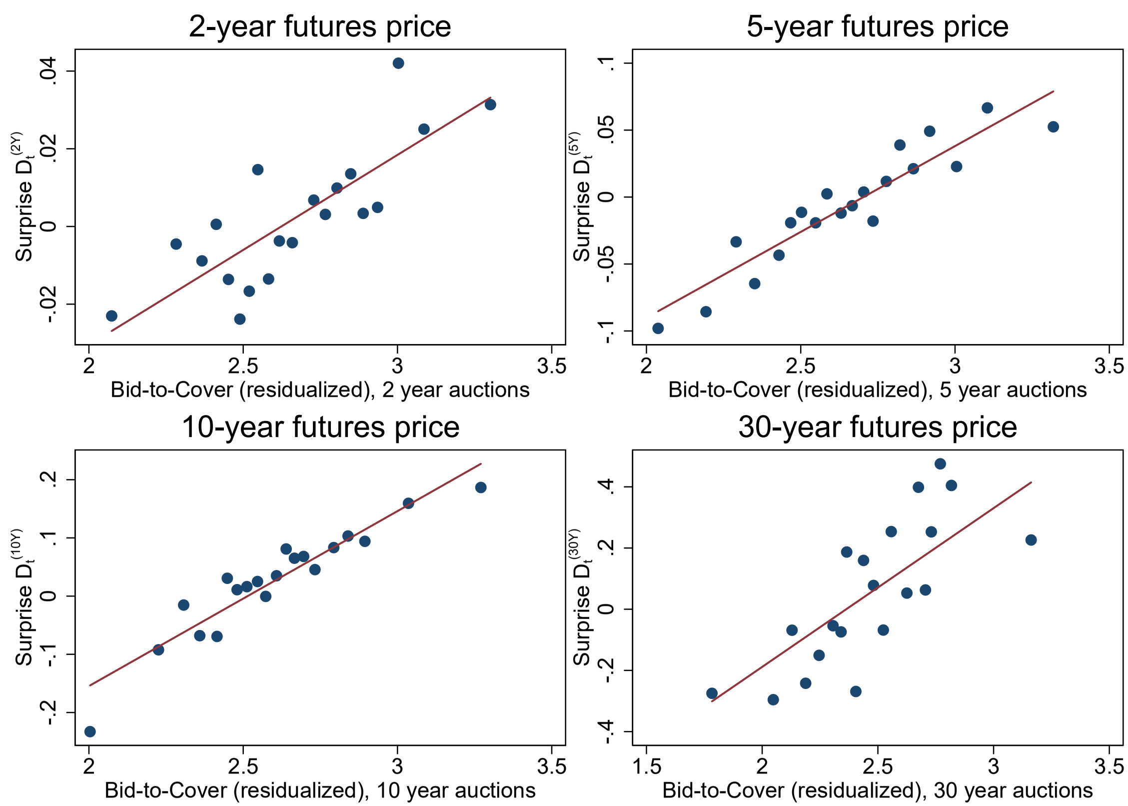

Notes: Bin scatter plot comparing demand shocks D_t and the bid-to-cover ratio from the auction (winsorized at 1% level). The bid-to-cover ratio is the fraction of dollar value of bids received to accepted at a given auction. Demand shocks D_t are reported in log points.

Figure 2.9: Long-Difference Responses to Shock D_t

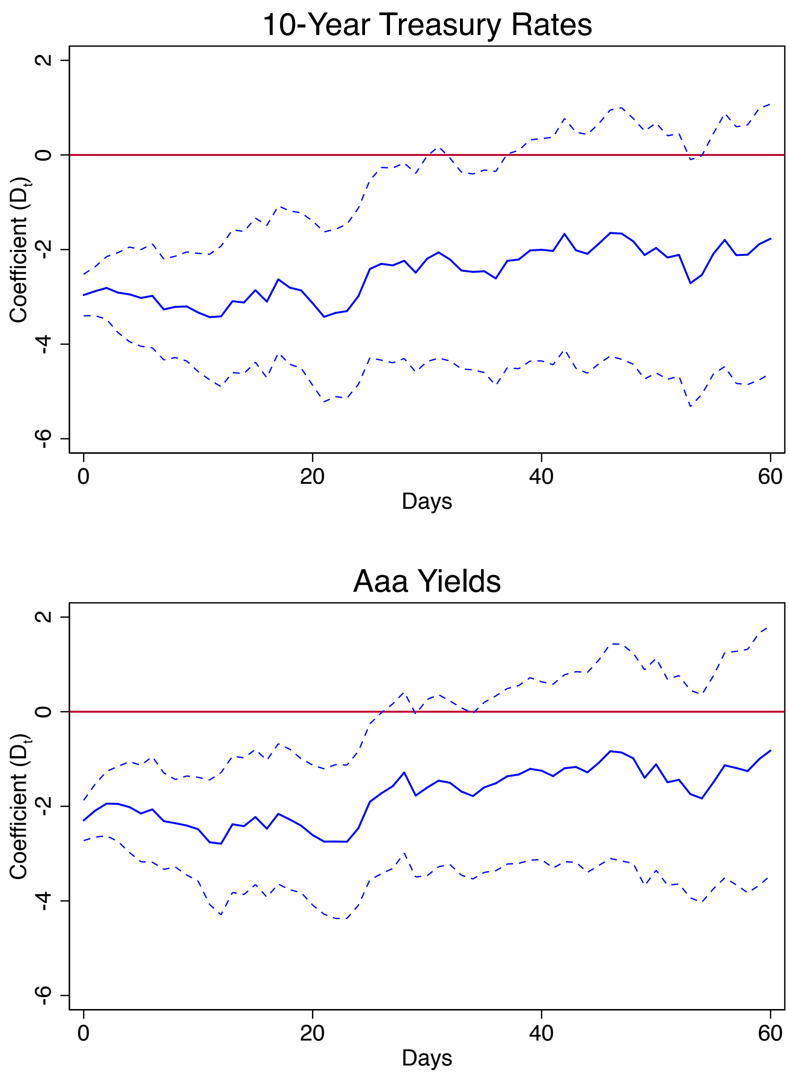

Notes: responses of 10-year Treasury spot rates (top panel) and Moody's Aaa yields (bottom panel) to a unit shock in the first principal component D_t. Spot rates come from Gurkaynak et al. (2007), estimated from daily prices from the secondary market for Treasuries. The regressions are "long-difference" regressions: on an auction date t, the dependent variable is $R_{t+h} - R_{t-1}$, i.e. the change (in terms of basis points) h days after the auction relative to the day before the auction. We plot the coefficients from regressions for $h = 0, \ldots, 60$. The solid line plots the point estimates, while dashed lines plot two-standard deviation (Newey-West) confidence bands.

Figure 2.10: Changes in yield curves on select Treasury auction days

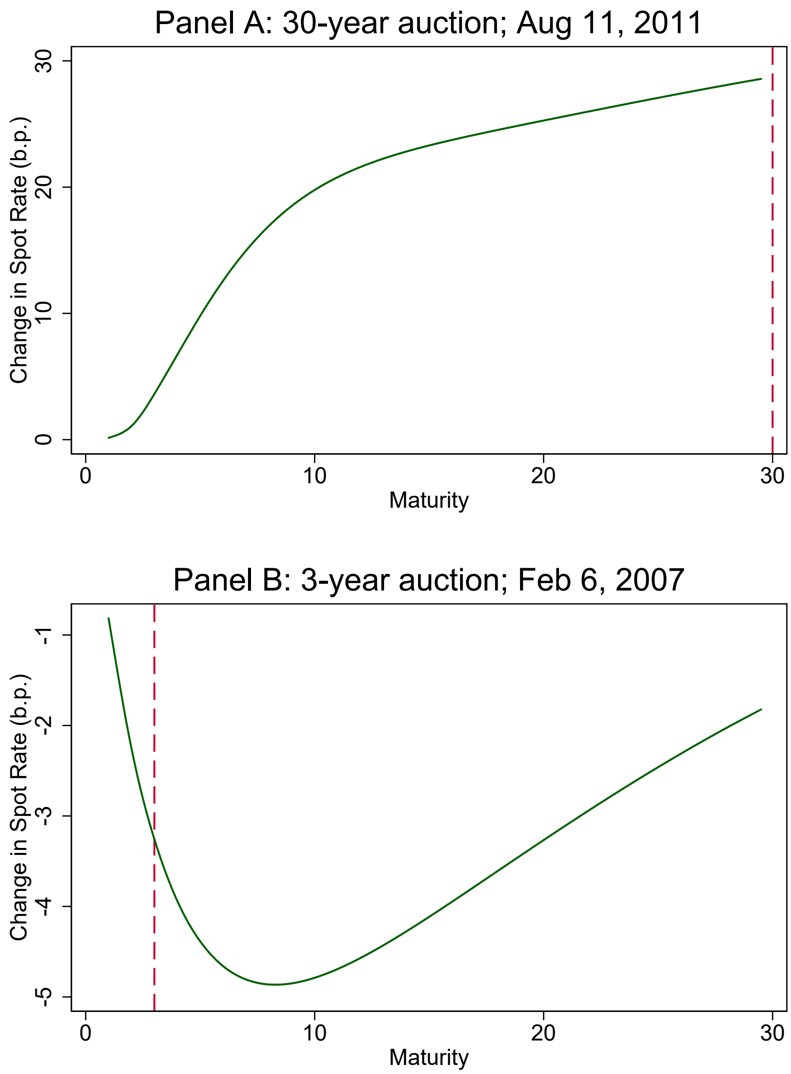

Notes: the figure plots changes in spot rates after 30-year auction on August 11, 2011 (top panel) and 3-year auction on February 6, 2007. The dashed vertical line shows the “location” of the auction in the maturity space.

Figure 2.11: Numerical Exercise

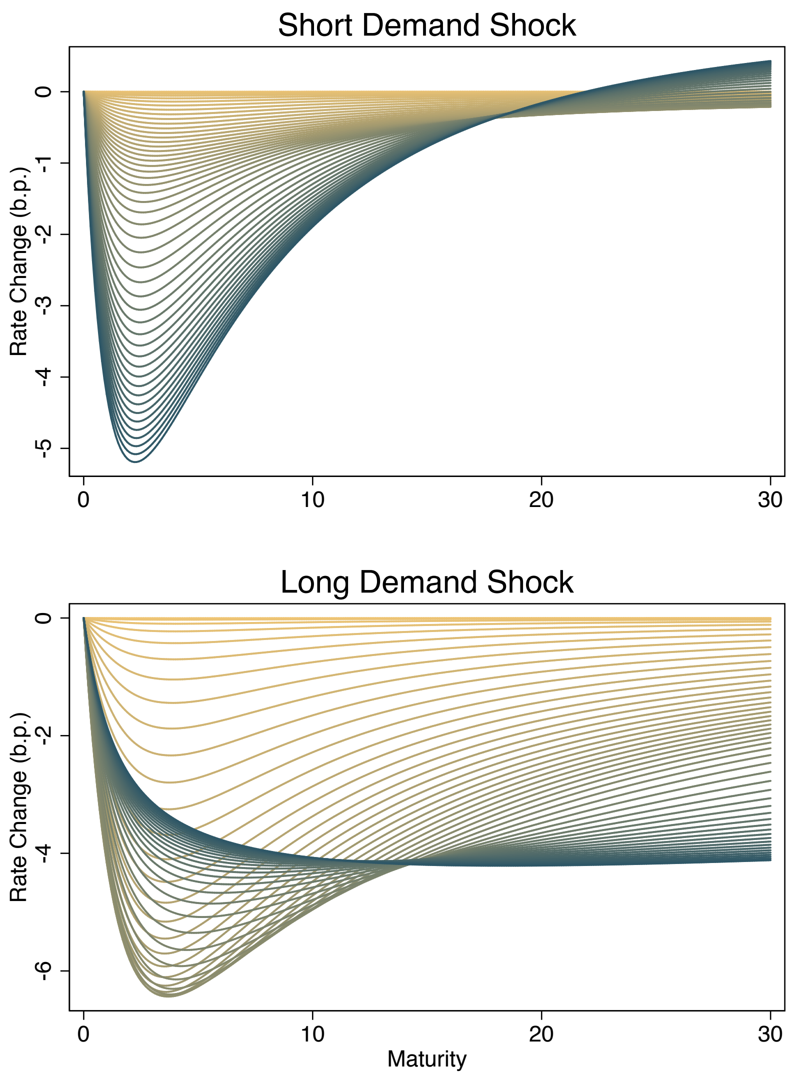

Notes: Numerical exercise studying the change in term structure of spot rates in response to one-standard deviation positive demand shocks, as risk aversion increase from low (lighter) to high (darker). The top panel is the impact of a short demand shock, and the bottom panel is the impact of a long demand shock.

Figure 2.12: Rate Responses (intraday Futures surprises)

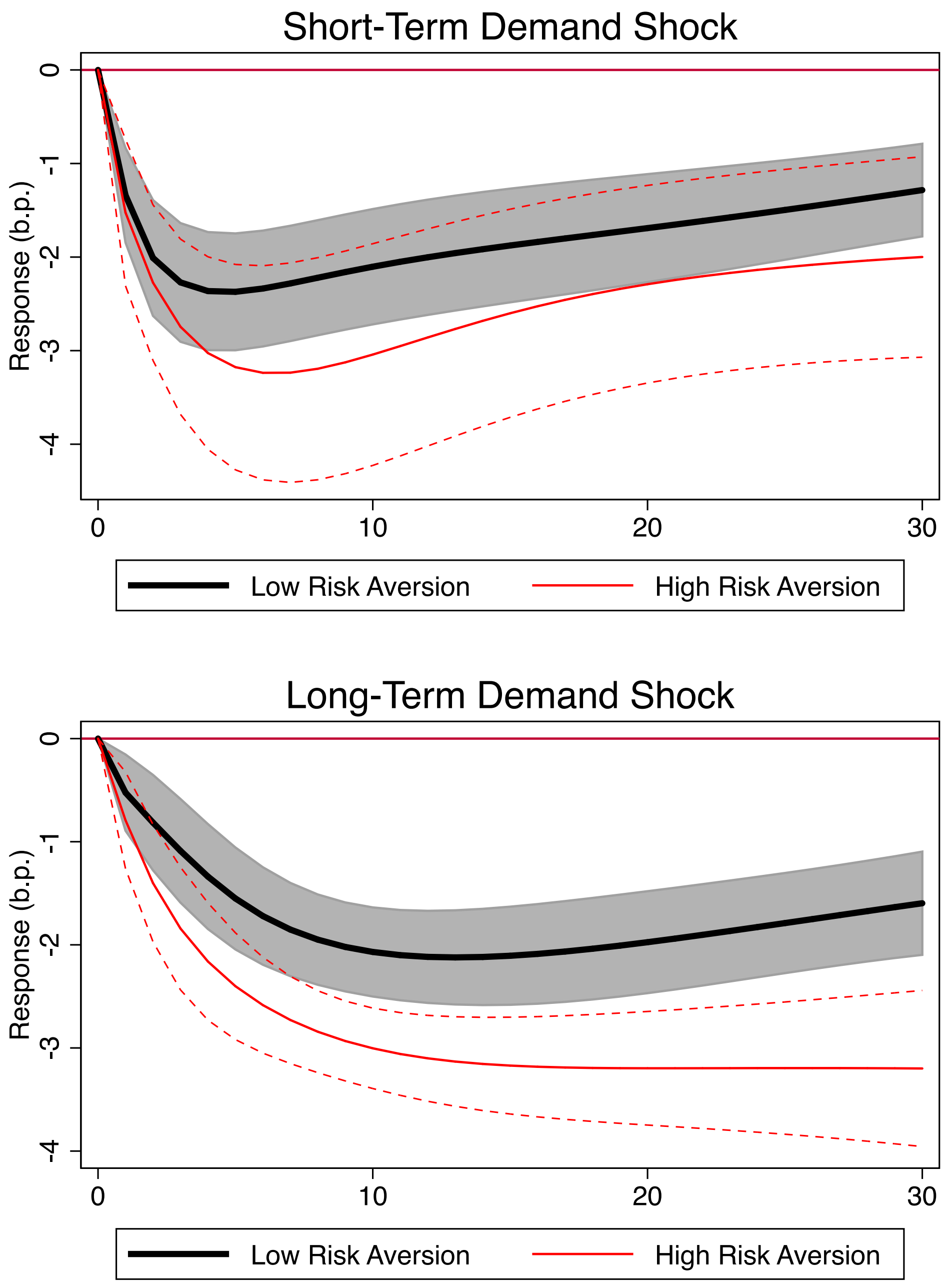

Notes: Plots of the regression coefficients on the demand shocks $D_t^{(m')}$ from regression equation (2.4). For each auction the demand shock $D_t^{(m')}$ is the normalized futures surprise that most closely corresponds to the maturity of the auction (e.g. a 5-year auction corresponds to $D_t^{(5Y)}$). Each curve is from the subsample combinations: short-term and long-term auctions; and periods of high and low risk aversion. 2 standard error (Newey-West) confidence intervals are included.

Figure 2.13: Rate Responses (IV specification)

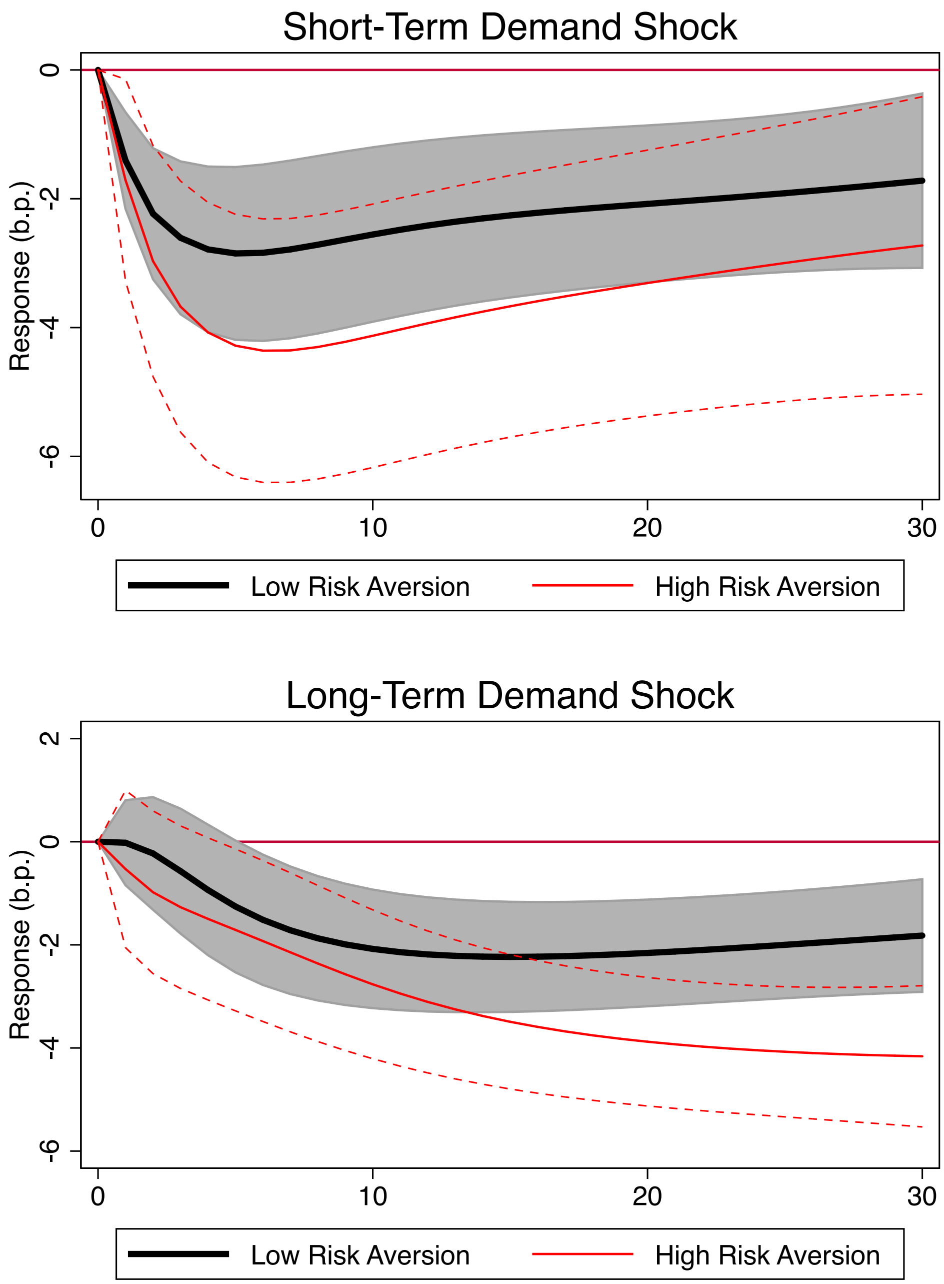

Notes: Plots of the regression coefficients on the demand shocks $D_t^{(m')}$ from regression equation (2.4), instrumented by the surprise component of the bid-to-cover ratio. Each curve is from the subsample combinations: short-term and long-term auctions; and periods of high and low risk aversion. The first-stage F-statistic for the short-term auctions is 10.16, while for the long-term auctions the F-statistic is 17.58. 2 standard error (Newey-West) confidence intervals are included.

Table 2.1: Auction Summary Statistics

Panel A: 1979-2015

	Mean	Median	Std. Dev.	Min	Max
Offering Amount (billions)	17.08	14.00	10.13	1.50	44.00
Total Tendered (billions)	47.44	36.38	31.97	2.37	160.96
Bid-to-Cover	2.60	2.57	0.52	1.22	5.88
Term (Years)	7.46	5.00	8.08	2.00	30.25
High Yield	5.39	4.77	3.66	0.22	16.28
High-Median Spread	0.03	0.03	0.02	0.00	0.14

Panel B: 1995-2015

	Mean	Median	Std. Dev.	Min	Max
Offering Amount (billions)	22.03	21.00	9.36	5.00	44.00
Total Tendered (billions)	61.46	52.98	32.04	11.35	160.96
Term (Years)	7.83	5.00	8.42	2.00	30.25
High Yield	3.26	3.20	1.91	0.22	7.79
High-Median Spread	0.03	0.03	0.02	0.00	0.13
Bid-to-Cover	2.62	2.60	0.49	1.22	4.07
Bid-to-Cover by type$^{\sharp}$					
Direct Bidders	0.24	0.25	0.18	0.00	0.84
Indirect Bidders	0.50	0.50	0.16	0.03	1.02
Primary Dealers	1.98	1.92	0.35	0.97	3.12
Fraction Accepted†					
Depository Institutions	0.01	0.00	0.02	0.00	0.32
Individuals	0.01	0.00	0.02	0.00	0.19
Dealers	0.58	0.58	0.14	0.20	0.98
Pensions	0.00	0.00	0.01	0.00	0.21
Investment Funds	0.20	0.18	0.13	0.00	0.64
Foreign	0.20	0.19	0.09	0.00	0.61
Other	0.00	0.00	0.00	0.00	0.03

Notes: Summary statistics for Treasury note and bond auctions. † indicates that the moments are computed for 2000-2015, the period for which these data are available. ♯ indicates that the moments are computed for 2003-2015, the period for which these data are available.

Table 2.2: Treasury Futures Shocks Summary Statistics

Mat.	Mean	Med.	SD.	N	Corr.			
					$D_t^{(2Y)}$	$D_t^{(5Y)}$	$D_t^{(10Y)}$	$D_t^{(30Y)}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Auction								
$D_t^{(2Y)}$	-0.000	0.000	0.034	871	1.000			
$D_t^{(5Y)}$	0.002	0.000	0.092	871	0.866	1.000		
$D_t^{(10Y)}$	0.007	0.007	0.143	871	0.782	0.958	1.000	
$D_t^{(30Y)}$	0.006	0.000	0.245	871	0.672	0.848	0.922	1.000
Pabel B. No auction								
$D_t^{(2Y)}$	-0.000	0.000	0.031	4031	1.000			
$D_t^{(5Y)}$	-0.001	0.000	0.072	4096	0.862	1.000		
$D_t^{(10Y)}$	-0.002	0.000	0.107	4100	0.794	0.945	1.000	
$D_t^{(30Y)}$	-0.005	0.000	0.172	4099	0.674	0.830	0.905	1.000
Panel C. Auction, non-ZLB period								
$D_t^{(2Y)}$	-0.002	0.000	0.043	424	1.000			
$D_t^{(5Y)}$	-0.005	-0.007	0.099	424	0.922	1.000		
$D_t^{(10Y)}$	-0.005	0.000	0.143	424	0.866	0.962	1.000	
$D_t^{(30Y)}$	-0.016	-0.026	0.223	424	0.778	0.878	0.933	1.000
Panel D. Auction, ZLB period								
$D_t^{(2Y)}$	0.002	0.000	0.022	447	1.000			
$D_t^{(5Y)}$	0.009	0.007	0.083	447	0.811	1.000		
$D_t^{(10Y)}$	0.018	0.014	0.143	447	0.736	0.960	1.000	
$D_t^{(30Y)}$	0.027	0.023	0.263	447	0.642	0.840	0.918	1.000

Notes: On auction dates, shocks $D_t^{(m)} = \log P_{t,post}^{(m)} - \log P_{t,pre}^{(m)}$ are the log intraday change in Treasury futures prices before and after the close of an auction, for each contract $m = 2, 5, 10, 30$ years. For non-auction dates, the shocks are the log intraday changes in Treasury futures prices using the same window. Binding zero lower bound (ZLB) period covers 2008M12-2015M12. Non-ZLB period covers 1995M1-2008M11. Statistics in columns (1)-(3) are reported in log points.

Table 2.3: Demand Shocks and Measures of Demand

Panel A: Total bid-to-cover ratio

	(1) $D_t^{(2Y)}$	(2) $D_t^{(5Y)}$	(3) $D_t^{(10Y)}$	(4) $D_t^{(30Y)}$	(5) Pool D_t
Bid-to-Cover	1.441***	1.399***	2.099***	2.119***	1.645***
	(0.240)	(0.230)	(0.216)	(0.565)	(0.142)
Observations	238	306	227	100	871
R^2	0.156	0.201	0.302	0.270	0.215

Panel B: Expected and Unexpected bid-to-cover ratio

	(1) $D_t^{(2Y)}$	(2) $D_t^{(5Y)}$	(3) $D_t^{(10Y)}$	(4) $D_t^{(30Y)}$	(5) Pool D_t
Bid-to-Cover (exp.)	0.031	-0.041	-0.454*	-1.374	-0.076
	(0.113)	(0.120)	(0.239)	(1.654)	(0.081)
Bid-to-Cover (unexp.)	1.382***	1.374***	2.113***	2.157***	1.645***
	(0.242)	(0.236)	(0.216)	(0.634)	(0.142)
Observations	238	306	227	100	871
R^2	0.124	0.189	0.294	0.215	0.198

Panel C: Total bid-to-cover ratio by bidder type

Indirect Bidder	2.716***	3.664***	4.528***	8.532***	4.451***
	(0.366)	(0.667)	(0.493)	(1.235)	(0.436)
Direct Bidder	2.236**	1.026	0.295	1.145	1.173***
	(1.034)	(0.702)	(0.956)	(0.951)	(0.448)
Primary Dealer	0.831**	0.762**	1.517***	0.057	0.887***
	(0.387)	(0.316)	(0.317)	(0.536)	(0.178)
Observations	138	228	187	80	633
R^2	0.350	0.309	0.383	0.650	0.370

Panel D: Fraction accepted by bidder type

Investment Funds	4.800***	3.401***	4.563***	6.436***	4.749***
	(0.908)	(0.854)	(0.902)	(1.462)	(0.494)
Foreign	2.797**	3.604***	5.173***	7.974***	4.393***
	(1.162)	(0.847)	(1.220)	(2.404)	(0.676)
Misc	4.815*	2.506**	0.034	0.853	2.353**
	(2.614)	(1.203)	(3.713)	(5.119)	(1.193)
Observations	174	241	201	84	700
R^2	0.214	0.128	0.287	0.391	0.191

Notes: Regressions of demand shocks $D_t^{(m)}$ on the bid-to-cover ratio, total and broken up by bidder type (winsorized at 1% level). Four lags of bid-to-cover ratios (or fractions accepted) are included but not reported. Column (1) restricts the sample to include only auctions of 2-year notes. Column (2) restricts the sample to include only auctions of notes with (2,5] year maturity. Column (3) restricts the sample to include only auctions of notes with [7,10] year maturity. Column (4) restricts the sample to include only auctions of bonds with (10,30] year maturity. Shocks $D_t^{(m)}$ are standardized to have zero mean and unit variance. Newey-West standard errors in parentheses.

Table 2.4: Reaction of market to surprises at Treasury auctions

Dep.variable: asset type	Estimate (s.e.)	N	R^2	Sample
	(1)	(2)	(3)	(4)
Panel A. Debt				
TLT	0.312***	662	0.679	2002-2015
	(0.016)			
SHY	0.022***	662	0.528	2002-2015
	(0.001)			
LQD	0.110***	662	0.544	2002-2015
	(0.008)			
Aaa†	-2.295***	871	0.173	1995-2015
	(0.212)			
Panel B. Equities				
SPY	-0.020	871	0.005	1995-2015
	(0.018)			
IWM	-0.081***	706	0.034	2000-2015
	(0.024)			
SP500†	-0.072	871	0.004	1995-2015
	(0.064)			
Russell 2000†	-0.169**	871	0.013	1995-2015
	(0.069)			
Panel C. Inflation expectations and commodities				
10Y Inflation Swap†	-0.172	618	0.003	2004-2015
	(0.131)			
2Y Inflation Swap†	0.044	618	0.000	2004-2015
	(0.229)			
GLD	0.021	595	0.004	2004-2015
	(0.015)			
GSCI†	0.008	871	0.000	1995-2015
	(0.056)			
Panel D. Spreads and credit default swaps				
Baa-Aaa†	-0.056	871	0.001	1995-2015
	(0.074)			
Auto CDS†	-3.254	627	0.000	2004-2015
	(5.796)			
Bank CDS†	0.426	627	0.004	2004-2015
	(0.450)			
3-month LIBOR-OIS†	-0.002	630	0.006	2003-2015
	(0.002)			
VIX†	0.058	871	0.001	1995-2015
	(0.082)			

Notes: The table reports estimates $\hat{\phi}$ from regression equation (2.3). Assets with † are measured at the daily frequency; other price changes are measured over the intrady window corresponding to what we use to construct surprises. The intraday changes are from ETFs that track various underlying securities or indices: TLT (long-term Treasuries); SHY (short-term Treasuries); LQD (corporate debt); SPY (S&P 500); IWM (Russell 2000); GLD (gold bullion). For the daily series: Aaa and Baa (Moody's corporate debt yields); SP500 (daily equity index); Russell 2000 (daily equity index); GSCI (S&P Total Commodity Index); Auto and Bank CDS (indsutry credit default swaps indices); LIBOR-OIS (3-month USD LIBOR-Overnight Index Swap spread); VIX (daily implied volatility index). Newey-West standard errors in parentheses.

Chapter 3

Polarized Expectations

3.1 Introduction

Political polarization is rising.[1] Much of the discussion in the political sphere has been concerned with how polarization undermines political discourse, leading to a lack of compromise and endless government gridlock.[2] But an equally important question is how polarization interacts with the economic beliefs and actions of households. Paradoxically, growing polarization has occurred contemporaneously with an increase in the accessibility of economic information. The proliferation of media and news sources makes economic information easier to consume than ever before, and yet polarization shows no signs of abating.

This paper explores theoretically and empirically how household beliefs and actions interact with political polarization. We first present a model where heterogeneous agents have imperfect information about the state of the economy, and must choose how to acquire costly information in order to inform their decision-making. We derive a "paradox of information" whereby declining information costs can actually increase ex-post disagreement about the economy; moreover, disagreement persists even with arbitrarily small information costs.

This "paradox of information" can help rationalize the secular increase in political polarization and the simultaneous increase in the ease of acquiring economic information. A naive model would suggest that increased access to information should reduce disagreement about economic fundamentals. Our model shows that, while households are able to learn about the macroeconomy more precisely, they process information in the manner which is most advantageous for satisfying their idiosyncratic preferences. When households are not identical, this implies that more information can actually exacerbate ex-post disagreement about the economy.

Empirically, we confirm that political polarization affects both household beliefs

[1] This chapter is based on my joint paper with Rupal Kamdar.

[2] Among others, Pew Research Center (2014) has documented a long-term increase in negative views of the opposing party or ideology within the U.S.

and actions. Using survey data on U.S. consumer beliefs, we find that households generally have persistent and stable economic beliefs and forecasts. However, there are a few striking exceptions: household belief persistence breaks down following elections where the White House changes party. During these periods, households that were optimistic about economic conditions become more likely to become pessimistic. Similarly, pessimistic households are more likely to become optimistic about economic outcomes following the election. This effect has been increasing since the 1980s, with the largest impact coming after the 2016 election.

Of course, changes in expectations should theoretically lead to changes in actions, such as shifts in consumption and savings patterns. But empirically, the evidence demonstrating this relationship is not as clear. We use the 2016 election to test whether these shifts in economic beliefs translate into differential consumption decisions. Using disaggregated geographical spending data, we find that polarization leads to differential changes in consumer spending: regions with a large Republican voteshare exhibited substantially higher consumption in the wake of the 2016 election.

We present a multi-dimensional rational inattention model in Section 3.2. Starting with a model of identical agents, we show that falling information costs eventually lead to ex-post agreement amongst agents regarding the realization of the state. However and more surprisingly, in general agents do not perceive the state accurately, even when information costs are arbitrarily small. We show analytically that misperceptions of the state persist as information costs fall to zero; this is because agents do not necessarily require perfect perception of the state in order to make optimal decisions. We then extend the analysis to the case of heterogeneous agents. Solving for the distribution of beliefs as a function of information costs, we derive two main analytical results: first, when information is very costly, falling information costs will actually increase the dispersion of beliefs across agents; and second, ex-post disagreement about the state may persist even in the limit as information costs fall to zero.

We then explore our theoretical results numerically with a simple two period model where agents differ in their level of risk aversion, and face uncertainty regarding income and prices. In response to a policy shock where the government levies a payroll tax to fund a retirement program, we find that highly risk-averse agents believe that prices will rise while less risk-averse agents believe prices will fall. Further, this disagreement persists even as information costs fall to zero. This stylized model shows that how agents perceive the economy in response to a policy change fundamentally depends on their preferences.

After presenting our general theoretical results, we turn to testing these stylized predictions empirically. In Section 3.3, we utilize survey data of U.S. consumers to show that political polarization affects macroeconomic beliefs. We show that beliefs tend to be persistent: people who think unemployment will fall tend to feel the same way when asked again in the future. A striking exception is following the 2016 presidential election, when this pattern reverses. Individuals who believed unemployment would rise before the election become more likely to say that unemployment will fall following the election. We observe an identical switch in beliefs for individuals who believed unemployment would fall before the election; after the election, these individuals believe that unemployment will rise. Similar patterns hold for other macroeconomic beliefs.

We then show that households' beliefs and forecasts about macroeconomic conditions as well as their own financial situation are almost entirely determined by a single component from a factor analysis. We show that this component is a measure of economic sentiment. There is widespread dispersion of sentiment across individuals at any point, but for a given individual we show that sentiment is highly persistent. Since 1980, the only exceptions are following presidential elections when the White House changed parties. During these periods, previously optimistic individuals are more likely to become pessimistic, and similarly pessimistic individuals are more likely to become optimistic. Moreover, this decline in persistence following elections has grown over time, with the most striking change in persistence coming after the 2016 presidential election.

There is a strong theoretical link between beliefs and economic actions, but it is an open question if this connection exists empirically. Section 3.4 demonstrates empirically that political polarization does lead to differential consumption responses. We combine disaggregated consumption data with voting data from California at the zip code level to study consumption responses in the weeks surrounding the 2016 election. We find that areas with a higher Republican voteshare exhibited substantially higher consumption in the weeks following the outcome of the 2016 election. Section 3.5 concludes and discusses avenues for future work.

This paper contributes to three literatures: (i) models of rationally inattentive agents that can choose the form of their signal, (ii) survey-based empirical investigations into macroeconomic beliefs, and (iii) the rise of polarization. First, we build upon the existing rational inattention literature to develop new theoretical results. The rational inattention framework, which was first proposed in Sims (2003), posits that agents do not observe economic variables perfectly, but instead must acquire costly information. Rational inattention can help explain the sluggish response to macroeconomic shocks, as in Maćkowiak and Wiederholt (2009). Kőszegi and Matějka (2018) generalizes the rational inattention methodology to analytically solve multi-dimensional rational inattention problems. They show that in general, agents choose to learn about combinations of variables rather than learning about economic variables independently. Learning about each economic variable independently is costly, and so optimal information-gathering implies that agents choose to learn about the economy in a manner which is most useful for deciding the agent's optimal course of action. Our model contributes to the rational inattention literature by studying what optimal information-processing implies for the dispersion of beliefs across agents. We derive new analytical results regarding the dispersion of beliefs across heterogenous agents, and how belief dispersion interacts with the cost of acquiring information.

Second, there is a growing literature that utilizes survey data to better understand the economic beliefs of individuals. Coibion et al. (2018) provide a summary of recent papers that demonstrate deviations from full-information rational expectations (FIRE). Our theoretical focus is on rational inattention, but other papers have also proposed other explanations for the failure of FIRE (e.g., Malmendier and Nagel 2016 and Kuchler and Zafar 2015 propose that lived experiences lead to deviations from rational expectations). But there is limited work on how consumer beliefs are tied to actions. Theoretical models in macroeconomics would suggest economic expectations (e.g., inflation and income expectations) should affect the today's actions (e.g., consumption and savings decisions). However, there is little empirical work done to establish this relationship. Furthermore, this empirical work has delivered mixed results. For example, survey-based confidence indices contain significant, but very small, predictive power for future aggregate consumer expenditure (e.g., Carroll et al. 1994, Bram and Ludvigson 1998, and Ludvigson 2004). Additionally, inflation expectations have been shown to affect household spending decisions, but the direction of the relationship has varied across environments and individuals studied (e.g., Bachmann et al. 2015, D'Acunto et al. 2016, and D'Acunto et al. 2018). In a political context (and empirically most similar to this paper) Mian et al. (2018) and Gillitzer and Prasad (2018) show that economic sentiments can change in the wake of elections in the U.S. and Australia, respectively. However, they have different results when it comes to consumption response to a positive political shock. Mian et al. (2018) find no effects on household spending, whereas Gillitzer and Prasad (2018) find a positive response in consumption. Our results support the view that households' macroeconomic expectations do play a role in their actual consumption decisions.

Third, there has been much discussion over the secular rise in polarization of the political discourse. In the context of U.S. politics, polarization has been particularly stark during and after the 2016 presidential election. Bartels (2002) demonstrates that party identification affects voter's interpretation of objective events. In addition to voters, this polarization exists amongst lawmakers. Andris et al. (2015) document a steady decline in cooperation and increase in partisanship amongst U.S. legislators in the post-war era. We add to this literature by demonstrating that both economic beliefs and actual consumption decisions are sensitive to the rise in political polarization.

3.2 A Rational Inattention Framework

This section explores theoretically how imperfect information-processing on the part of agents leads to disagreement about macroeconomic fundamentals. We start with a general multivariate rational inattention setup, as originally explained in Sims (2003) and generalized in Kőszegi and Matějka (2018). Suppose an individual has quadratic preferences over actions $\mathbf{y} \in \mathbb{R}^N$ and states $\mathbf{x} \in \mathbb{R}^J$:

$$U(\mathbf{x}, \mathbf{y}) = -\mathbf{y}^T \mathbf{C} \mathbf{y} + \mathbf{x}^\mathbf{T} \mathbf{B} \mathbf{y} \tag{3.1}$$

where $\mathbf{B} \in \mathbb{R}^{J \times N}$ and $\mathbf{C} \in \mathbb{R}^{N \times N}$, and is positive definite: $\mathbf{C} \succ 0$.

The individual does not observe $\mathbf{x}$ perfectly, and obtaining signals about the state is costly. The cost of information is the Shannon mutual information times a scaling parameter, $\lambda \geq 0$. The Shannon mutual information is the expected reduction of entropy (a measure of uncertainty) from the prior to the posterior. It is commonly used in the rational inattention literature as a cost of information. Intuitively, the more precise the posterior, the higher the Shannon mutual information. We scale the Shannon information by λ for flexibility. In particular, if $\lambda = 0$, information is free, and the individual can costlessly obtain perfect information about the realization of the state. If λ is very high, the agent decides to avoid gathering any new information. Accordingly, the agent's beliefs do not change from the prior. For intermediate values of $\lambda > 0$, the agent collects some information, but will not learn about the state perfectly.

For simplicity, assume a mean-zero and i.i.d. prior $\mathbf{x} \sim N(\mathbf{0}, \sigma_0^2 \mathbf{I})$. To briefly review known results, as shown in Kőszegi and Matějka (2018) the agent's optimization problem from Eq. (3.1) and the information problem given cost of information parameter $\lambda > 0$ can be rewritten as

$$\begin{aligned} &\max_{\mathbf{\Sigma}} - \mathbb{E}\left[(\tilde{\mathbf{x}} - \mathbf{x})^T \mathbf{\Omega} (\tilde{\mathbf{x}} - \mathbf{x})\right] + \frac{\lambda}{2} \log |\mathbf{\Sigma}| \\ &\text{s.t. } \sigma_0^2 \mathbf{I} - \mathbf{\Sigma} \succeq 0 \\ \text{where } \mathbf{\Omega} = &\frac{1}{4} \mathbf{B} \mathbf{C}^{-1} \mathbf{B}^T \end{aligned}$$

$\tilde{\mathbf{x}}$ is the agent's posterior mean and $\mathbf{\Sigma}$ is the posterior variance. The loss matrix $\mathbf{\Omega}$ characterizes how costly it is for the agent to misperceive the state. This characterization recasts the problem as one in which the agent optimally chooses to reduce uncertainty about the realization of the state by choosing the posterior variance $\mathbf{\Sigma}$. The restriction that $\sigma_0^2 \mathbf{I} - \mathbf{\Sigma}$ is positive semidefinite ensures that the agent cannot choose a posterior variance which is less precise than the prior; that is, the agent does not discard information. Note that when $\lambda > 0$, the agent will never learn about any aspects of the state perfectly (in which case the posterior variance is singular, and hence $|\mathbf{\Sigma}| = 0$) because the cost of acquiring such information grows without bound.

Lemma 3.1 (Optimal Signal Solution)**.** *Take the eigendecomposition of the loss matrix $\mathbf{\Omega} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$ where $\mathbf{V}$ is a matrix of orthonormal eigenvectors and $\mathbf{\Lambda}$ is a diagonal matrix of non-negative eigenvalues ordered $\Lambda_1 \geq \ldots \geq \Lambda_J$. Define*

$$S_{ii} = \begin{cases} \frac{\lambda}{2\Lambda_i} & \textit{if } 2\Lambda_i \sigma_0^2 > \lambda \\ \sigma_0^2 & \textit{otherwise} \end{cases}$$

$$\xi_i = 1 - \frac{S_{ii}}{\sigma_0^2}$$

and let $\mathbf{S} = \mathrm{diag}(S_{ii})$ *and* $\mathbf{\Xi} = \mathrm{diag}(\xi_i)$ *be diagonal matrices. Then the optimal choice of posterior variance is* $\mathbf{\Sigma} = \mathbf{V}\mathbf{S}\mathbf{V}^T$*. The posterior mean and optimal action is a function of the realization of the state* $\mathbf{x}$ *and independent standard normal noise* $\mathbf{e} \sim N(\mathbf{0}, \mathbf{I})$*:*

$$\tilde{\mathbf{x}} = \mathbf{W}_1\mathbf{x} + \mathbf{W}_2\mathbf{e} \tag{3.2}$$

$$\mathbf{y}^* = \mathbf{H}\tilde{\mathbf{x}} \tag{3.3}$$

where $\mathbf{W}_1 \equiv \mathbf{V}\mathbf{\Xi}\mathbf{V}^T$*,* $\mathbf{W}_2 \equiv \sigma_0\mathbf{V}\left(\mathbf{\Xi}(\mathbf{I} - \mathbf{\Xi})\right)^{1/2}$*, and* $\mathbf{H} \equiv \frac{1}{2}\mathbf{C}^{-1}\mathbf{B}^T$*.*

Lemma 3.1 characterizes the solution to this optimal signal problem. All proofs are in the Appendix. Essentially, the agent chooses to receive a signal that reduces uncertainty about the state in a manner which is most helpful for choosing an optimal action (subject to the cost of reducing uncertainty). The eigenvectors $\mathbf{V}$ govern which directions of the state are most important. The eigenvalues $\mathbf{\Lambda}$ govern how important this information is. Hence, each element of the diagonal matrix $\mathbf{\Xi}$ determines how much the agent reduces uncertainty about each direction of the state relative to the prior belief.

The matrices $\mathbf{W}_1$ and $\mathbf{W}_2$ in Eq. (3.2) characterize how the agent's beliefs respond to changes in the state and the noise in the signal, respectively. Perfect information implies that $\mathbf{W}_1 = \mathbf{I}$ and $\mathbf{W}_2 = \mathbf{0}$. When information is costly, the full-information signal structure is impossible to achieve. This leads to two key difference between beliefs under costly relative to full information. First, beliefs are less responsive to realizations of the state and sensitive to noise in the signals. Formally, this arises from the fact that the elements of $\mathbf{\Xi}$ are less than one. Second, beliefs about one state variable may react to realization of another state variable. As long as the loss matrix $\mathbf{\Omega}$ is not diagonal, this arises mathematically due to the fact that the eigenvector matrix $\mathbf{V} \neq \mathbf{I}$. Intuitively, this is because in general the agent chooses to learn about combinations of variables, which are more useful when it comes to choosing optimal actions, as this is less costly than learning about each element of the state independently form an information-acquisition perspective.

The optimal choice of action is given by Eq. (3.3). Note that the only difference between the optimal actions under full information and the solution with imperfect information is the difference in posterior beliefs. Because the agent has quadratic preferences, the optimal choice of action under perfect information takes the same form as Eq. (3.3). In this case $\tilde{\mathbf{x}} = \mathbf{x}$ and $\mathbf{y}^* = \mathbf{H}\mathbf{x}$.

One key insight is that an agent does not generally learn about each state variable equally. In fact, when the loss matrix $\mathbf{\Omega}$ is not full rank (so some eigenvalues $\Lambda_i = 0$), the agent never chooses to learn about some directions of the state. Moreover, the manner in which the agent learns about the state is directly tied to how fundamental preferences (the matrices $\mathbf{B}$ and $\mathbf{C}$) interact with the costs of information. Preferences therefore fundamentally affect how an agent views the economy, regardless of the cost of acquiring information.

3.2.1 Distribution of Beliefs

What does this optimal information processing imply about the distribution about beliefs? We characterize the distribution of misperceptions of the state across individuals with identical preferences following a realization of the state in the following Proposition. Further, we show how these moments vary with information costs λ.

Proposition 3.1 (Distribution of Belief Errors and Information Costs)**.** *Conditional on a realization of the state* $\mathbf{x}$*, the errors in posterior beliefs* $\tilde{\mathbf{x}} - \mathbf{x}$ *have mean and variance:*

$$\mathbb{E}\left[\tilde{\mathbf{x}} - \mathbf{x} \middle| \mathbf{x}\right] = \mathbf{V}(\boldsymbol{\Xi} - \mathbf{I})\mathbf{V}^T\mathbf{x} \tag{3.4}$$

$$\mathrm{Var}\left[\tilde{\mathbf{x}} - \mathbf{x} \middle| \mathbf{x}\right] = \mathbf{V}\boldsymbol{\Xi}\mathbf{S}\mathbf{V}^T \tag{3.5}$$

These moments are functions of the cost of information λ*:* $\frac{\partial}{\partial\lambda}\mathrm{Var}\left[\tilde{\mathbf{x}} - \mathbf{x} \middle| \mathbf{x}\right]$ *is negative semidefinite whenever* $2\Lambda_1\sigma_0^2 > \lambda > \Lambda_1\sigma_0^2$ *and positive semidefinite whenever* $\lambda < \Lambda_N\sigma_0^2$*.*

Further, as $\lambda \to 0$

$$\mathbb{E}\left[\tilde{\mathbf{x}} - \mathbf{x} \middle| \mathbf{x}\right] \to \mathbf{V}(\tilde{\mathbf{I}}_N - \mathbf{I})\mathbf{V}^T\mathbf{x}$$

$$\mathrm{Var}\left[\tilde{\mathbf{x}} - \mathbf{x} \middle| \mathbf{x}\right] \to \mathbf{0}$$

where $\tilde{\mathbf{I}}_N = \mathrm{diag}\begin{bmatrix}\mathbf{I}_N & \mathbf{0}_{J-N}\end{bmatrix}$*, with the zeros correspond to the zero eigenvalues of* Λ *(if they exist).*

Declining information costs implies that the individual's signal can be made more precise. That is, the noise in the signal becomes smaller and smaller. However, this does not mean that declining information costs always leads to a decline in the variance of belief errors. The first result in Prop. 3.1 shows that for high enough levels of information costs, declining information costs actually increases the (conditional) variance of belief errors.[3] This is because, when information costs are extremely high, the agent chooses to receive no signal (more precisely, the signal noise has infinite variance and so the agent completely ignores the signal). Once information costs fall below a certain threshold, the agent chooses to pay attention to the signal. But the noise in the error has very high variance, and hence this leads to large differences in signals received between agents. Individuals know that their signal is noisy and so do not update their prior too much, but this is still enough to increase the ex-post disagreement between individuals.

Eventually, once information costs are beneath another threshold (which depends on the smallest non-zero eigenvalue Λ_N), the conditional variance of belief errors

[3]Suppose the derivative of a variance-covariance Φ with respect to information cost λ is positive semidefinite ($\frac{\partial}{\partial\lambda}\Phi \succeq 0$). Then for two values of information cost $\lambda_1 < \lambda_2$, $\Phi_2 - \Phi_1 \succeq 0$. In other words, an increase in information costs implies an increase in the variance matrix Φ. The opposite logic applies for negative semidefinite derivative matrices.

begins to decline as information costs fall. Because the variance of the signal noise for each individual is now very small, signals are more similar across individuals. Hence, ex-post disagreement across individuals declines.

In the limit of zero information costs, the signal can be made arbitrarily precise. Since agents are identical, they choose the same signal structure. Unsurprisingly, when the noise in the signal falls to zero, there is no ex-post disagreement about the realization of the state. Hence, the conditional variance of belief errors approaches zero.

When $\lambda = 0$, the agent can learn the exact realization of the state costlessly. But surprisingly, it turns out that for arbitrarily small costs of acquiring information, the agent's beliefs do not necessarily approach the true state: in general, posterior belief errors do not vanish even as $\lambda \to 0$. Belief errors only vanish when the loss matrix $\mathbf{\Omega}$ is full rank, but this condition will fail if the number of actions is smaller than the number of states ($N < J$). When $\mathbf{\Omega}$ is not full rank, the agent will never choose to learn about certain directions of the state so long as information costs are not exactly zero. When information costs are actually zero, the agent is indifferent between learning perfectly the true state and learning about the subset of directions of the state that matter most to informing the optimal action.

The reason for this somewhat paradoxical result is that the agent is able to perform the optimal action without learning the state perfectly. By extracting information optimally, any additional information is not informative about the optimal action. As shown in Cor. 3.1.1, in the limit as information costs fall to zero, the agents' optimal choice (based on beliefs about the state) converges to the optimal choice under perfect information.

Corollary 3.1.1 (Optimal Actions and Information Costs)**.** *For any $\mathbf{\Omega}$, the agent's actions approach the optimal choice under full information. That is, as $\lambda \to 0$,*

$$\mathbf{y}^* \to \mathbf{H}\mathbf{x}$$

3.2.2 Belief Errors and Heterogeneous Agents

Now we suppose that there is a distribution of agents with heterogeneous preference. They all face the same information acquisition problem but have different preferences:

$$U^k(\mathbf{y}^k, \mathbf{x}) = -(\mathbf{y}^k)^T \mathbf{C}^k \mathbf{y}^k + \mathbf{x}^T \mathbf{B}^k \mathbf{y}^k$$

where superscript k indexes the individuals. These agents are distributed according to some distribution $k \sim F$.

Assume that individuals all have the same common prior, $\mathbf{x} \sim N(\mathbf{0}, \sigma_0^2 \mathbf{I})$. Then the individual k problem is the same as the previous section, and hence the posterior mean and variance of individual k are the same as from Lemma 3.1:

$$\tilde{\mathbf{x}}^k = \mathbf{W}_1^k \mathbf{x} + \mathbf{W}_2^k \mathbf{e}^k$$
$$\mathbf{\Sigma}^k = \mathbf{V}^k \mathbf{S}^k (\mathbf{V}^k)^T$$

However, heterogeneity in preferences will lead to differences in the distribution of beliefs relative to the case of homogeneous individuals.

Proposition 3.2 (Heterogeneous Belief Errors)**.** *Given a realization of the state* $\mathbf{x}$*, define* $\widetilde{\mathbf{W}}_1^k \equiv \mathbf{W}_1^k - \bar{\mathbf{W}}_1$ *and* $\bar{\mathbf{W}}_1 = \int_k \mathbf{W}_1^k \, \mathrm{d}F(k)$*, and*

$$\boldsymbol{\Sigma}_{\widetilde{\mathbf{W}}_1\widetilde{\mathbf{W}}_1|\mathbf{x}} \equiv \int_k \widetilde{\mathbf{W}}_1^k \mathbf{x}\mathbf{x}^T (\widetilde{\mathbf{W}}_1^k)^T \, \mathrm{d}F(k)$$
$$\boldsymbol{\Sigma}_{\mathbf{W}_2\mathbf{W}_2} \equiv \int_k \mathbf{W}_2^k (\mathbf{W}_2^k)^T \, \mathrm{d}F(k)$$

Then the errors in posterior beliefs $\tilde{\mathbf{x}}^k - \mathbf{x}$ *have conditional mean and variance:*

$$\mathbb{E}\left[\tilde{\mathbf{x}}^k - \mathbf{x} \middle| \mathbf{x}\right] = (\bar{\mathbf{W}}_1 - \mathbf{I})\mathbf{x} \tag{3.6}$$
$$\mathrm{Var}\left[\tilde{\mathbf{x}}^k - \mathbf{x} \middle| \mathbf{x}\right] = \boldsymbol{\Sigma}_{\widetilde{\mathbf{W}}_1\widetilde{\mathbf{W}}_1|\mathbf{x}} + \boldsymbol{\Sigma}_{\mathbf{W}_2\mathbf{W}_2} \tag{3.7}$$

Choose $\bar{\lambda}$ *such that*

$$\forall k, i: \ \bar{\lambda} \geq 2\Lambda_i^k \sigma_0^2$$
$$\exists k', i: \ \bar{\lambda} = 2\Lambda_i^{k'} \sigma_0^2$$

Then $\frac{\partial}{\partial \lambda}\mathrm{Var}\left[\tilde{\mathbf{x}}^k - \mathbf{x} \middle| \mathbf{x}\right]$ *is negative semidefinite at* $\bar{\lambda}$*. Further, as* $\lambda \to 0$*,*

$$\mathbb{E}\left[\tilde{\mathbf{x}}^k - \mathbf{x} \middle| \mathbf{x}\right] \to \int_k \mathbf{V}^k (\tilde{\mathbf{I}}_N - \mathbf{I})(\mathbf{V}^k)^T \, \mathrm{d}F(k)\, \mathbf{x}$$
$$\mathrm{Var}\left[\tilde{\mathbf{x}}^k - \mathbf{x} \middle| \mathbf{x}\right] \to \int_k \widetilde{\mathbf{W}}_1^{k*} \mathbf{x}\mathbf{x}^T (\widetilde{\mathbf{W}}_1^{k*})^T \, \mathrm{d}F(k)$$

where

$$\widetilde{\mathbf{W}}_1^{k*} = \mathbf{V}^k \tilde{\mathbf{I}}_N (\mathbf{V}^k)^T - \int_k \mathbf{V}^k \tilde{\mathbf{I}}_N (\mathbf{V}^k)^T \, \mathrm{d}F(k)$$

As in the case of homogeneous agents, at very high levels of information costs, declining information costs lead to a higher degree of dispersion of beliefs. The choice of $\bar{\lambda}$ in Prop. 3.2 is the smallest cost of information for which some agents are indifferent between reducing uncertainty about the state and not updating their prior. This is the point where falling information costs induce some set of agents to begin paying attention to their signal. Because the signal is imprecise, this increases disagreement amongst individuals about the realization of the state. The disagreement depends on how differently individuals respond to changes in the state relative to the average change in beliefs; $\bar{\mathbf{W}}_1$ is the average response of beliefs to a given change in the state $\mathbf{x}$ and $\widetilde{\mathbf{W}}_1^k$ is how individual k differentially responds compared to the average.

In contrast to the case of homogenous agents, this "paradox of information" continues even as information costs fall. Heterogeneity across preferences implies that large belief disagreements can persist even as information costs approach zero. As explained in the previous section, in general an individual will not choose to learn about all directions of the state regardless of how cheap information is. But with heterogeneous preferences, the directions that matter the most differ across individuals. Hence, even for very low information costs, different individuals receive different signals. These signals become arbitrarily precise, but beliefs of different agents do not approach the same value. Because agents have different preferences, they also choose different optimal actions; thus, they choose to learn about the state differently. Hence for arbitrarily small information costs, large ex-post disagreement still exists.

In the context of the secular changes observed in political polarization, this stylized model helps interpret how belief formation may have changed as the accessibility to information has increased. Over the past several decades, the ease of acquiring information regarding macroeconomic outcomes has increased substantially. This is due to the proliferation of print and televised media sources, as well as the creation and increased accessibility to the internet for most households across the developed world. It is now easier than ever to not only find information quickly, but also to find sources of information and news that are more tailored to an individual's tastes.

A naive model of information acquisition would suggest that increased access to information should reduce the degree of belief polarization across households. The rational inattention model described in this section shows that even though households are able to learn about the macroeconomy more precisely, this leads to increased disagreement about macroeconomic outcomes. Because households are not identical, households process information in the manner which is most advantageous for satisfying their idiosyncratic preferences.

3.2.3 A Two-Period Model

This section lays out a concrete partial equilibrium two-period model, where households choose how much to consume and save while young and old. The household problem is

$$\begin{aligned} U(C_1, C_2) &= u^k(C_1) + \beta u^k(C_2) \\ P_1 C_1 + S &\leq (1-\tau)Y \\ P_2 C_2 &\leq RS + \tau Y \end{aligned}$$

The agent chooses how much to consume of a single good while young and old (C_1 and C_2). Labor is supplied elastically, and they receive a nominal after-tax income $(1-\tau)Y$. Whatever is saved S earns a gross (nominal) return R.

Finally, the government may tax income in order to fund a retirement-savings program. The amount taxed, τY is returned to the agent while old (period 2). The return on the government savings program is smaller than the return on private savings; for simplicity, this government savings return is normalized to 1 while the private savings is $R = \frac{1}{\beta}$. The price of the good in both periods, P_1 and P_2, is uncertain, as is the period 1 income Y and net tax rate $1-\tau$.[4]

Individuals differ in their per-period utility function:

$$u^k(C) = \frac{C^{1-\eta^k}}{1-\eta^k}$$

Each agent has per-period CRRA utility, but differ across their relative risk aversion η^k.

In order to apply the rational inattention tools developed in Section 3.2.2, we take a log quadratic approximation of the household problem. Appendix C.2 shows that the approximation around the steady state means the problem can be re-written in the form of Eq. (3.1). The choices, states, and preference matrices are given by

$$y^k = c_1, \quad C^k = \frac{(1+\beta)^{\eta^k}\eta^k}{2\beta}$$

$$\mathbf{x} = \begin{bmatrix} y \\ 1-\tau \\ p_1 \\ p_2 \end{bmatrix}, \quad \mathbf{B}^k = \begin{bmatrix} \dfrac{(1+\beta)^{\eta^k}\eta^k}{\beta} \\ \dfrac{(1+\beta)^{\eta^k}\eta^k(1-\beta)}{\beta} \\ -\dfrac{(1+\beta)^{\eta^k}(\beta+\eta^k)}{\beta} \\ (1+\beta)^{\eta^k}(1-\eta^k) \end{bmatrix}$$

Note that the agent's budget is always exhausted (the budget constraint holds with equality), hence the agent only has one active choice. We write the model in terms of the choice of period-1 consumption. Since C^k is a scalar, this implies that the loss matrix is given by

$$\mathbf{\Omega}^k = \frac{1}{4C^k}\mathbf{B}^k(\mathbf{B}^k)^T$$

Hence $\mathbf{\Omega}^k$ is a rank-one matrix. Thus $\mathbf{B}^k$ is the only eigenvector (up to a normalization) associated with the non-zero eigenvalue:

$$\Lambda_1^k = \frac{1}{4C^k}(\mathbf{B}^k)^T\mathbf{B}^k$$

Further, the responsiveness of beliefs to changes in the state is governed by the matrix $\mathbf{W}_1^k$, which is proportional to $\mathbf{B}^k(\mathbf{B}^k)^T$. These objects, which govern how households optimally obtain and react to information, depend crucially on the risk aversion term η^k.

[4] Allowing for uncertainty regarding the return on savings does not affect the theoretical results, so to keep the size of the state smaller we assume that the return is known.

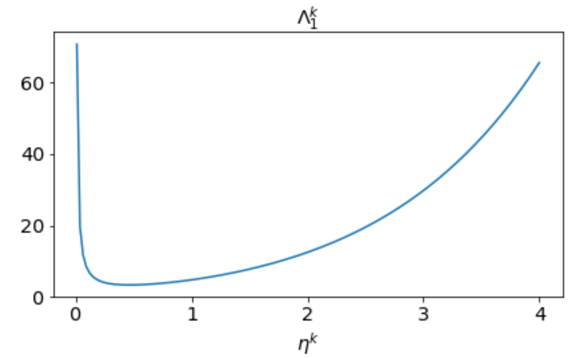

Figure 3.1: Importance of Information Across Agents

Notes: Plots of the non-zero eigenvalue Λ_1^k of the loss matrix Ω^k across individuals, who differ in their risk aversion parameter η^k (x-axis). This eigenvalue along with the prior variance σ_0^2 governs how much each individual is to willing to pay in order to reduce uncertainty.

Figure 3.1 shows how the single non-zero eigenvalue Λ_1^k changes with different values of risk aversion η^k. As $\eta^k \to \infty$ or $\eta^k \to 0$, this eigenvalue grows without bound. What this means is that individuals with either very high or very low levels of risk aversion are more willing to pay a high cost for obtaining information. However, there is a fundamental difference between the type of information individuals with high and low levels of risk aversion choose to obtain. Figure 3.2 plots the components of the associated eigenvector $\mathbf{V}^{k,1} \propto \mathbf{B}^k$ for different values of risk aversion. First, when $\eta^k \to 0$ so that the household has nearly linear preferences, the only information the agent cares about is (relative) prices. Second, there is a difference for agents with η^k above and below one. This determines whether the income or substitution effect dominates when considering a change in the real return.

What does this mean for how beliefs respond to changes in the shock? That is, how does $\frac{\partial \tilde{\mathbf{x}}^k}{\partial \mathbf{x}^T} \equiv \mathbf{W}_1^k$ vary across individuals? If agents observed the state perfectly, then of course $\frac{\partial \tilde{x}_i^k}{\partial x_j} = 0$ if $i \neq j$ and 1 otherwise for all agents. Figure 3.3, which plots $\mathbf{W}_1^k$ for different values of risk aversion, shows that the model deviates significantly from perfect information. Note that $\mathbf{W}_1^k$ is a function of household preferences as well as the cost of information λ. For simplicity, we plot the limit as $\lambda \to 0$; however, for $\lambda > 0$, the elements of $\mathbf{W}_1^k$ are proportional but smaller than in Figure 3.3).

It is clear that even in the limit of costless information, the response of beliefs to changes in the state do not approach the case of full information. First, beliefs about a given variable $\tilde{x}_i$ respond significantly less than one-for-one with actual changes in the state variable x_i. Second, beliefs regarding a given variable $\tilde{x}_i$ respond to changes in the realization of other state variables x_j. Finally, this imperfect information-processing differs across individuals. Individuals with nearly linear preferences ($\eta^k \to$

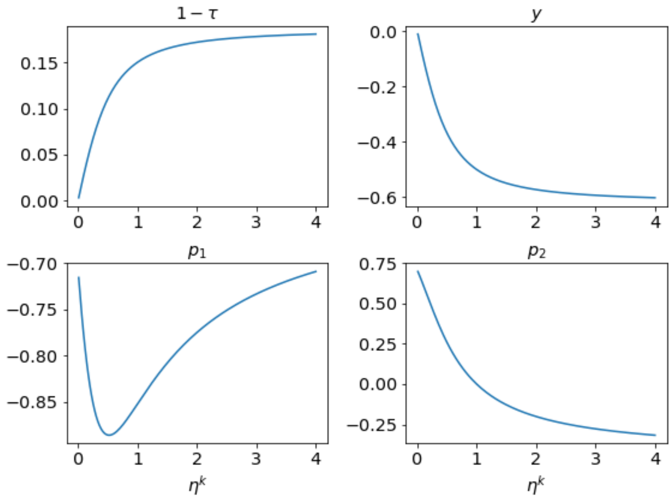

Figure 3.2: Information Directions Across Agents

Notes: Plots of the elements of the eigenvector associated with the non-zero eigenvalue of the loss matrix Ω^k across individuals, who differ in their risk aversion parameter η^k (x-axis). This eigenvector determines the direction in which each individual reduces uncertainty regarding the state. Each panel is the element of the eigenvector associated with the given state variable.

0) are unresponsive to changes in income or taxes, but instead only pay attention to changes in prices. Further, in response to an increase in period-2 prices, agents with risk aversion $\eta^k < 1$ reduce their posterior belief regarding period-1 prices; conversely, agents with $\eta^k > 1$ increase their period-1 price beliefs.

We now explore more explicitly how this heterogeneity in preferences affects the distribution of beliefs numerically. We assume that individual risk aversion parameters are distributed uniformly in $[0.5, 1.5]$. We set $\beta = 0.7$ and normalize $\sigma_0^2 = 1$. Note that with this range of η^k, agents with a higher degree of risk aversion will always have a higher willingness to pay for information. This can be seen in Figure 3.1: for $\eta^k > .5$, Λ_1^k is increasing in risk aversion.

The experiment we consider is a percentage point increase in the retirement tax τ (a decrease in the net tax rate $1-\tau$). Figure 3.4 plots the average error in beliefs across the population. When information costs are very high ($\lambda > 15$), no individual chooses to reduce uncertainty about the state at all, and hence beliefs stay at the prior. This means that all agents are mistaken about their beliefs regarding the tax rate, but in this case correctly believe that prices and income remain at their steady state.

As information costs fall (λ declines, moving from right to left on the x-axis), some agents with high risk aversion ($\eta^k > 1$) begin to learn about the shock. Unsurprisingly, these agents begin to reduce their misperceptions about the tax rate, and hence the avergage misperception about $1-\tau$ falls (upper left panel). But due to how agents with $\eta^k > 1$ optimally process information, these agents also begin to update their beliefs regarding income and prices, incorrectly inferring that prices and income have all increased. Hence, average misperceptions regarding income and prices also rise.

As information costs continue to fall, more and more agents find it optimal to begin learning about the state. In particular, for low enough information costs, agents with risk aversion $\eta^k < 1$ also begin to pay attention to their signal. Like agents with $\eta^k > 1$, agents with low risk aversion also incorrectly infer that income and period-1 prices have increased. Therefore, falling information costs increase the average misperceptions regarding these economic variables. However, agents with $\eta^k < 1$ incorrectly infer that period-2 prices have fallen. This means that on average, falling information costs begin to decrease average misperceptions. But for low enough information costs, the misperceptions of period-2 prices from agents with $\eta^k < 1$ begin to dominate. Hence, average misperceptions about period-2 prices eventually become negative when information costs $\lambda \to 0$.

The distribution of misperceptions across agents can be seen more clearly in Figure 3.5, which plots the variances and covariances of beliefs as a function of information costs. Falling information costs increase ex-post disagreement when information costs are high. Moreover, even in the limit of costless information, posterior beliefs do not approach the truth. Declining information costs allow agents to obtain precise signals, which eventually allows individuals to reduce their own uncertainty regarding the state. However, so long as information costs are not exactly zero, it is always optimal to learn about the state in a manner which is most informative for optimal actions (in this case, the optimal level of consumption). Hence, individual agents never learn about the state perfectly. Because agents with heterogeneous preferences make different choices even under perfect information, optimal information-processing implies that disagreements about the state persist for any level of information costs.

3.3 Empirical Evidence: Beliefs

The previous section showed how imperfect information and heterogeneity in preferences lead to ex-post disagreement about macroeconomic fundamentals following a policy shock. Moreover, falling information costs may actually exacerbate belief polarization. This section provides empirical evidence that political polarization does in fact play a role in how individuals form macroeconomic beliefs.

3.3.1 Data

We use the Michigan Survey of Consumers (MSC) to measure consumer beliefs. The MSC is a rotating panel survey of approximately 500 consumers per month. The survey began in 1978 and is still running today, which provides us with a long time series dimension to utilize. The MSC asks a variety of questions. For example, questions include: beliefs about current economic policy, expectations about future personal financial conditions, perceptions about past personal financial conditions, and inflation expectations.

For all questions that we use, the MSC solicits a categorical response from the consumer. The consumer has three choices of response, which fall broadly into an "optimistic" response, a "stay the same/neutral" response, and a "pessimistic" response. For example, in response to the question "As to the economic policy of the government – I mean steps taken to fight inflation or unemployment – would you say the government is doing a good job, only fair, or a poor job?" consumers may respond either "good job," "only fair," or "poor job." Similarly, for the question on 12 month unemployment expectations relative to today, the consumer chooses if they believe unemployment will "rise," "stay about the same," or "fall."

Based on these responses, the MSC also publishes a widely reported aggregate measure of consumer sentiment. Hence, an obvious first pass to assess whether political shocks affect consumer beliefs is to simply study the time variation in this measure. Figure 3.6 plots the MSC sentiment measure during the months before and after the recent 2016 presidential election.

At first glance, Figure 3.6 seems to suggest that consumer sentiment did not react strongly despite the polarizing presidential campaign and election. The consumer sentiment measure does appear to jump up slightly following the campaign. But this variation is well within normal fluctuations, and below particularly large changes (e.g. during recessions).

However, the theoretical results above show that focusing on average beliefs across heterogeneous agents is not sufficient to understand how polarization affects belief formation. Figure 3.7 bears this out. Here we plot the fraction of people whose answer regarding government policy was more optimistic than the answer they gave 6 months previously (green line). Similarly, the orange line is the fraction of people whose answer became more pessimistic. Usually, these two lines are negatively correlated: periods in which more people become optimistic about government policy are the same times in which fewer people are becoming pessimistic.

However, this pattern does not hold during the 6 months immediately following the 2016 election. During this period, the fraction of people becoming more pessimistic and more optimistic both increased. Note that during this 6 month period, individuals' previous answers in the MSC came before the outcome of the 2016 election.

3.3.2 Belief Responses to Political Shocks

The previous results are suggestive that political shocks differentially affect consumers' beliefs. However, government policy is directly related to the outcome of the presidential election, so it is perhaps not surprising to see sharp responses of beliefs regarding economic policy following a highly contentious and unexpected election. In this section we more formally explore whether the 2016 election had differential affects on beliefs. We also extend the analysis and compare the 2016 election with previous elections to see if these patterns have changed over time.

First, we more precisely estimate the transition probabilities surrounding the 2016 election. For a broader set of questions and using a multinomial logit estimation model, we examine the conditional probability that individuals answer optimistically or pessimistically, given that they previously answered optimistically, neutrally, or pessimistically.

Figure 3.8 plots these estimates for each month. The three columns correspond to estimated probabilities for three questions regarding government policy, expected business conditions, and expected unemployment. The top row plots the conditional probabilities for individuals answering optimistically, while the bottom row plots corresponds to pessimistic individuals.

The results show that usually, beliefs tend to be unchanged: people who had a positive view of government policy 6 months previously tend to think the same today. Similarly, people who previously believe business conditions will improve, or unemployment was declining, tend to feel the same way today. A handful of individuals who felt neutral will switch to positive, but very few people who felt negatively switch to a positive view. Similarly, those with pessimistic views of government policy, business conditions, or unemployment tend to continue to remain negative.

The major exception is the 6 months following the 2016 election. During this period, the pattern is exactly the opposite. Those who had a negative view of government policy 6 months ago (before the election) are now actually slightly more likely to have a positive view of government policy. Moreover, those who previously held a positive view of policy 6 months ago (before the election) are now unlikely to hold a positive view of policy. The results are the same for those who view policy negatively following the election. Additionally, the patterns for feelings about government policy are strikingly similar for other macroeconomic beliefs as well.

As argued in Kamdar (2018), when using consumer survey beliefs in a component analysis the first component explains the large majority of the variation in consumer beliefs. Furthermore this component can be viewed as a measure of optimism/pessimism. In order to more parsimoniously explore the behavior of expectations over time, we collapse the responses to a single component of a multiple correspondence analysis (MCA).

The results of the MCA estimation are in Figure 3.9.[5] The standardized coordinates of the first dimension are on the x-axis, while the second dimension corresponds to the y-axis. The first component explains over 87% of the inertia in the responses, while the second component only explains an additional 3%. Each different marker type corresponds to a question, while the colors correspond to the possible responses. Blue markers are for the pessimistic response, yellow is for the neutral response, and green is for the positive response.

[5]Note: the MCA is estimated on additional questions which include more than three categories; for readability, questions with more than three categorical answers are not included in the plot.

The estimated MCA corrdinates imply that the crucial first component acts like a measure of sentiment. Note that the coordinates of the first component are ordered such that pessimistic responses are less than neutral responses, which in turn are smaller than the optimistic responses (the blue scatter points fall to the left of the yellow points, which fall to the left of the green scatter points). The only slight exception is the question regarding interest rates (square markers). However, while the other responses can be unambiguously mapped to optimistic, neutral, and pessimistic responses, the change in interest rates is more ambiguous. Therefore, we interpret this component as a general measure of the sentiment of the consumer (see Kamdar (2018) for additional component analyses and robustness exercises).

Interestingly, the second dimension seems to correspond to a measure of change. The coordinates of the second dimension are ordered such that neutral responses are less than both optimistic and pessimistic responses (the yellow scatter points fall below both the blue and green scatter points). However, the first component explains the vast majority of inertia in the data, while the second component adds very little. Hence, we focus on the first component only.

We use the MCA to construct a measure of economic sentiment $f_{i,t}$ for each individual across time. From the results in Figure 3.9, we see that consumers who respond positively to any of the questions end up with a higher level of the first component. We use this sentiment measure to study how beliefs react over time to political shocks.

Figure 3.10 shows how the distribution of sentiment across individuals has changed over time. The solid line is our median economic sentiment measure across individuals in a given month. As expected, this measure is related to the business cycle, falling during recessions and increasing during booms. However, the dotted lines (plotting the 90-10 percent distribution) show that there is substantial variation across individuals regarding economic sentiment during booms and busts. For instance, even during the 2009 recession, more than 10% of individuals exhibited positive economic sentiment.

To explore how sentiment varies across individuals, we estimate the following regression

$$f_{i,t} = \alpha_t + \beta_t f_{i,t-6} + \varepsilon_{i,t} \tag{3.8}$$

The variable $f_{i,t}$ is the first component from the MCA analysis for an individual i at time t. We regress this on the individual's previous response 6 months prior, $f_{i,t-6}$. Hence, the coefficient β_t in Eq. (3.8) measures how persistent is individuals' sentiment measure over time. We estimate Eq. (3.8) period by period, pooling over individuals.

Figure 3.11 plots the estimated coefficients $\hat{\beta}_t$. The shaded regions correspond to the 6 month period following the outcome of presidential elections. In these periods, $f_{i,t}$ is determined by responses given after the outcome of the election, while $f_{i,t-6}$ is determined by responses from before the election. Hence, in the shaded regions Eq. (3.8) regresses responses following an election on the same individual's responses that were given before the election.

In general, sentiment is highly persistent: individuals who were optimistic 6 months prior tend to be optimistic today; similarly, pessimistic individuals remain pessimistic. Virtually the only exception to this is during certain presidential elections. Moreover, these elections where the persistence falls is exactly the elections where the White House changed party: 1980, when Reagan (R) replaced Carter (D); 1992, when Clinton (D) replaced Bush Sr. (R); 2000, when Bush Jr. (R) replaced Clinton; 2008, when Obama (D) replaced Bush Jr., and 2016, when Trump (R) replaced Obama. Further, the fall in persistence has been increasing over time. The 2016 election is the most striking, with an estimated negative coefficient. This confirms the individual question analysis from the previous section: individuals who were pessimistic before the 2016 election tend to become optimistic following the election, and vice versa for optimistic individuals.

3.4 Empirical Evidence: Actions

The theoretical results of Section 3.2 imply a tight connection between agent's beliefs and actions. Hence, the results in the previous section suggest that political shocks should lead to large changes in consumption. This section provides empirical evidence that political polarization during the 2016 presidential election played a crucial role in how individuals make spending decisions.

3.4.1 Data

We use the Nielsen Homescan data to study the response of consumption and spending to the outcome of the 2016 presidential election.[6] Nielsen Homescan is a panel dataset which measures U.S. consumer behavior. Panelists use scanners to record all purchases of products tracked by Nielsen in food and non-food categories. The dataset tracks when, where, and how much of each product a given panelist purchases across time. Nielsen also records demographic and geographical information on the panelists. Since 2007, the Homescan data includes roughly 60,000 households.

Using the Nielsen Homescan data, we construct a high-frequency measure of consumption spending at the 5-digit zip code level. In order to more precisely study the effects of the 2016 election, we aggregate spending to the weekly frequency; since presidential elections fall on Tuesdays, our weekly measure starts on Wednesday and runs to the following Tuesday.

[6]Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

We map our consumption data to voting data at the zip code level. Voting data is recorded at the precinct level, which is not observed in the Homescan dataset (and in general is not a common regional measure outside of elections). Instead, we use geographical shape data to allocate precinct-level voting data to zip codes. Precincts are smaller geographical areas than zip codes, and most precincts fall entirely within zip codes. For precincts that fall within multiple zip codes, we allocate votes proportionally to each zip code based on geographical size overlap.

Our voting and geographical precinct data is from the Statewide Database maintained by Berkeley Law; however, this data is restricted to the state of California. As an example of our zip code voting data, Figure 3.12 plots the vote shares by zip code in Orange County. Orange County voted for Clinton by a margin of 8%, but as the Figure shows this does not imply that votes were uniformly distributed across precincts.

We then merge our consumption panel data with the California zip code voting data. We focus on zip codes with at least 1,000 recorded votes. This leaves us with 1,046 zip codes. The median number of votes in our sample is roughly 12,000, with the largest made up of about 35,000 votes. The voteshare of the election in the median zip code median in our sample was 62% for Clinton. The largest and smallest zip code vote share for Clinton in our sample is 93% and 19%, respectively.

3.4.2 Event Study Results

In order to assess how spending reacted to the 2016 election, we estimate the following event study regression:

$$c_{z,t} = \alpha_z + \gamma_t + \sum_{k=-T}^{T} \beta_k v_z^{16} \cdot \mathbf{I}_{t=t^*+k} + \varepsilon_{z,t} \tag{3.9}$$

The outcome variable $c_{z,t}$ is (log) consumption in zip code z at week t. v_z^{16} is the Trump vote share in zip code z in the 2016 election. t^* is the week in which the 2016 presidential election took place; we set $T = 7$ in order to study consumption patterns in the 7 weeks before and after the election. The regression equation Eq. (3.9) also includes zip code and time fixed effects. Hence, the coefficient β_k represents the predicted percentage increase in consumption in a zip code with 1 percentage point higher Trump voteshare in the 2016 election.

Figure 3.13 plots the estimated $\hat{\beta}_k$ from the event study in Eq. (3.9). The increase in consumption is large and statistically significant in the weeks immediately following the election. In contrast, the estimates for the weeks preceding the election are small and not significantly different from zero; this shows there was no differential time variation in consumption patterns related to voting propensities in the lead-up to the election. The lack of pre-trends is reassuring that the change in consumption is in fact a response to the election.

Further, although the standard errors are large, the point estimate remains quite large for most of the weeks in the sample period following the election. Economically, our estimates are large: a zip code with 1 percentage point higher Republican voteshare is associated with a 0.5 percent increase in consumption over the weeks following the election. To put this in perspective, consider two hypothetical zip codes: the first voted 75% for Clinton while the second voted 75% for Trump. Then the Trump zip code is predicted to have 25% higher consumption relative to the Clinton zip code in response to the election outcome.

As a placebo, Figure 3.14 plots the same event study but using the equivalent timeframe in 2015. Since no election or other major political shock took place during this time period, we would not expect to see differential consumption patterns. The results in Figure 3.14 show that this is the case. None of the estimated coefficients are significantly different from zero. Further, the point estimates are never as large as the estimates we find following the 2016 election. This allows us to rule out any differential consumption patterns in zip codes that have a higher propensity to vote Republican for reasons unrelated to political shocks (e.g. different seasonal patterns in consumption).

Our empirical consumption results are in line with the results from a similar empirical design in Australia, as shown in Gillitzer and Prasad (2018). In the context of the U.S., our results contrast with the findings in Mian et al. (2018), who find little evidence that political shocks lead to differential consumption responses. This is likely due to two factors. First, we use voting data at the zip code level, which allows for a much tighter link between voting propensity of a given region and the consumption responses. Second, we use a higher-frequency measure of consumption. This allows us to pick out possibly shorter-lived consumption responses (note that while our estimates are economically large, they are only significantly different from zero at the 5% level for the two weeks following the election). However, one downside is the consumption data is largely composed of nondurable consumption. To the extent that durable consumption is a more important part of households' overall consumption spending, our results may miss a big part of the story.

3.5 Concluding Remarks

This paper argues that political polarization plays a large role in how individuals learn about the macroeconomy, and that these polarized expectations lead to differential consumption decisions. We first demonstrate theoretically that heterogeneous preferences lead to differential information-processing in a rational inattention model. Moreover, we show that declining information costs actually exacerbate the degree of ex-post disagreement about macroeconomic fundamentals under very general conditions. This "paradox of information" can help explain the secular increase in polarization: as information about the economy becomes easier to access, individuals rationally choose to seek out more and more precise information that is most useful to inform their actions. But when individuals differ in what is most important to them, this leads to individuals learning about the economy in fundamentally different ways. Hence, falling information costs may lead to larger ex-post disagreement about the macroeconomy.

We then use survey data of U.S. consumers to show that political shocks do lead to differential changes in expectations. We first construct an individual sentiment measure, which explains the vast majority of consumer's beliefs regarding not only current and future macroeconomic conditions, but also current and future personal financial conditions. This sentiment measure is highly persistent for a given individual, with the crucial exception of periods following presidential elections when the White House changed parties. During these periods, previously optimistic individuals are more likely to become pessimistic, and similarly pessimistic individuals are more likely to become optimistic. Moreover, the effect of presidential switches has been growing since the 1980s.

Finally, using the recent presidential election, we show that political shocks do lead to differential changes in consumption. Using a weekly measure of consumption spending at the zip code level tied to California voting records, we find that a 1 percentage point increase in the Republican voteshare of a zip code implies an increase of 0.5% in the weeks following the election.

The theoretical results in our model hold for general preferences, but our modeling framework still contain certain limitations. One important extension for future work is to allow for an infinite-horizon dynamic problem, along the lines of Maćkowiak et al. (2018). Empirically, extending the analysis to the entire U.S. is an important next step, though geographically precise precinct voting data is not readily available.

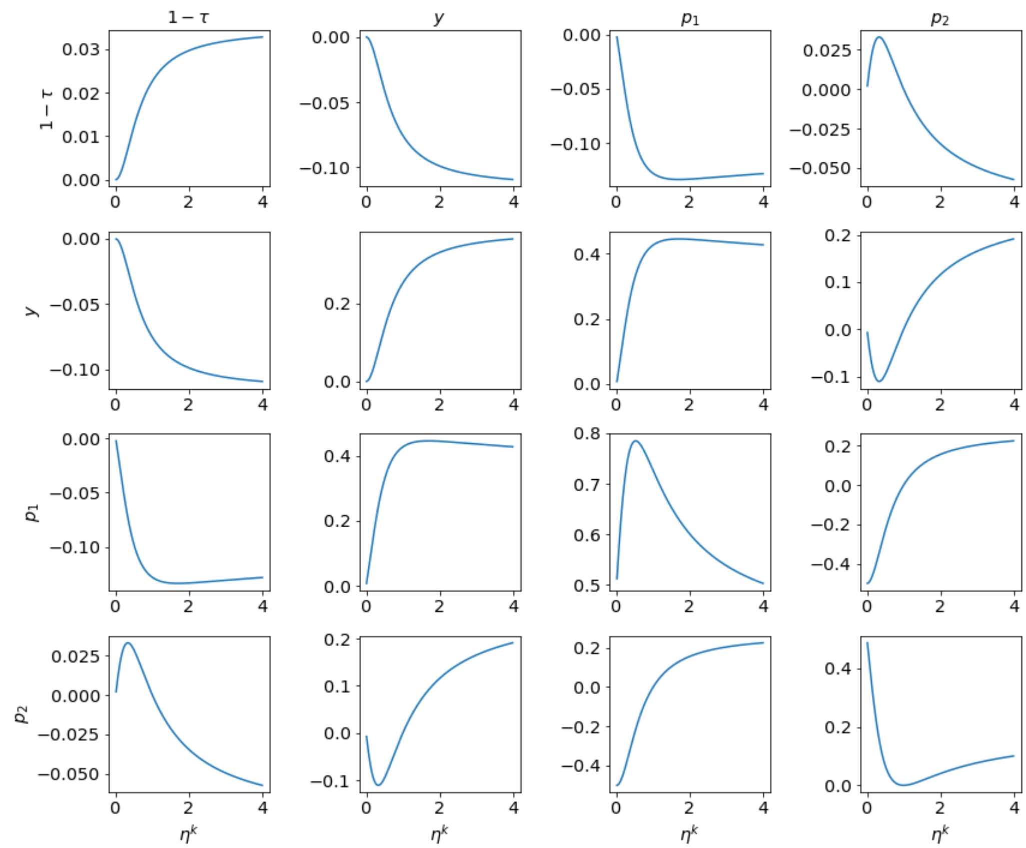

Figure 3.3: Belief Response Across Agents

Notes: Plots of the elements of the matrix $\mathbf{W}_1^k$ in the limit as information costs $\lambda \to 0$ across individuals, who differ in their risk aversion parameter η^k (x-axis). This matrix determines how each individual changes their belief regarding a variable (panel rows) in response to a unit change in the realization of another state variable (panel columns); note the matrix is symmetric. For instance, the third row and second column panel is the response of beliefs about period-1 prices p_1 in response to a unit increase in income y.

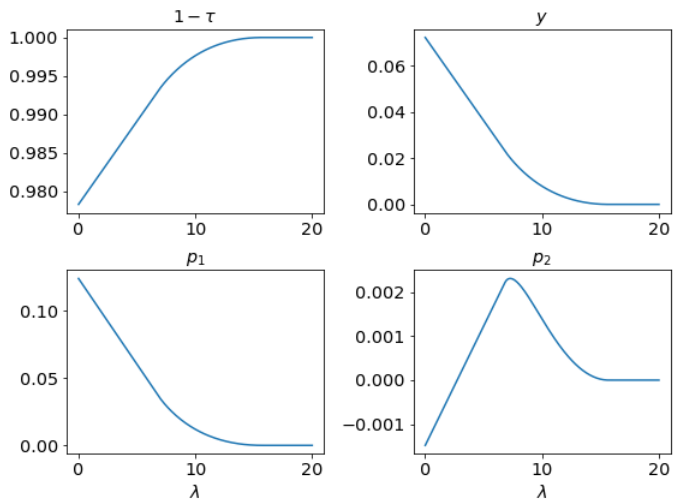

Figure 3.4: Average Beliefs and Information Costs: Tax Shock

Notes: plots of the mean error of beliefs across individuals as a function of information costs λ (x-axis). This experiment is conditional on a unit increase in the tax τ.

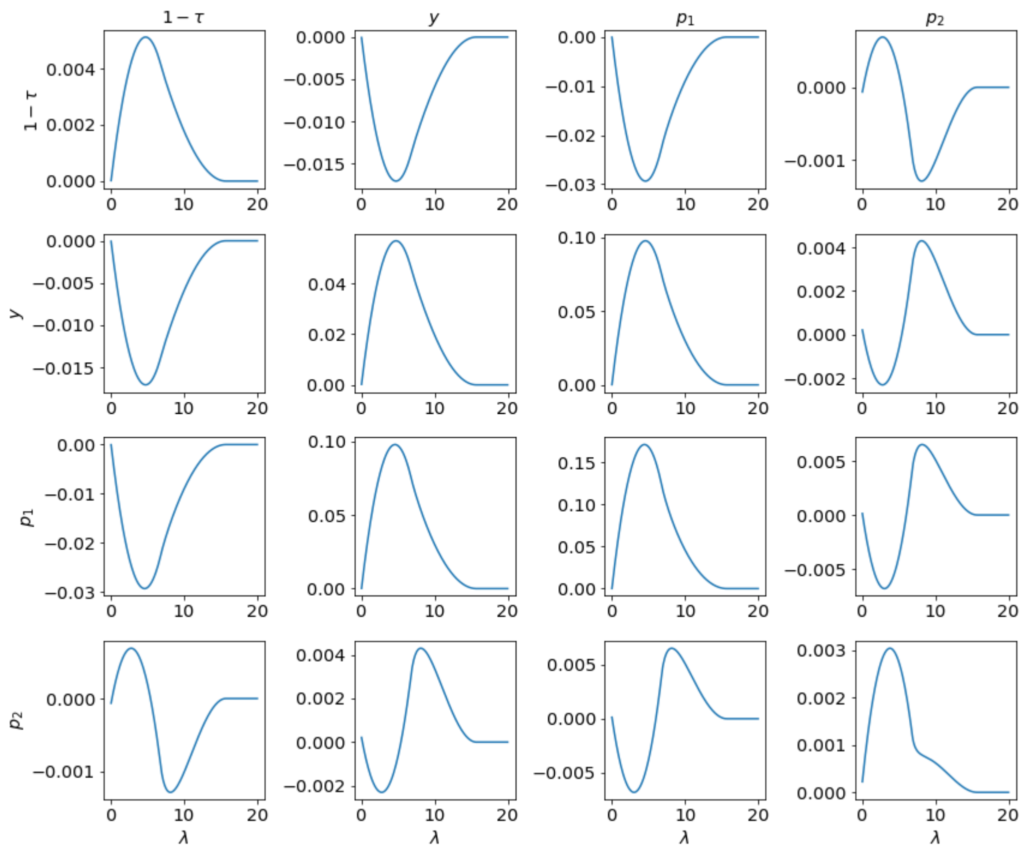

Figure 3.5: Belief Variances and Information Costs: Tax Shock

Notes: plots of the variance-covariance of belief errors across individuals as a function of information costs λ (x-axis). This experiment is conditional on a unit increase in the tax τ). Variances are on the diagonal panels; covariances are on the off-diagonal panels.

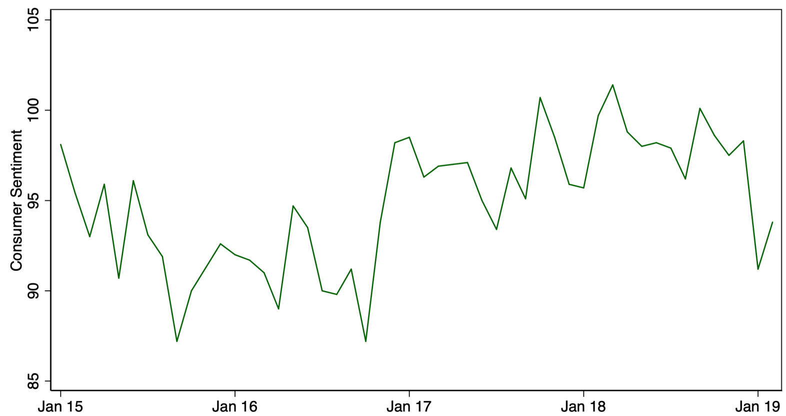

Figure 3.6: Consumer Sentiment

Notes: level of the Consumer Sentiment Index as published by the Michigan Survey of Consumer. Monthly data, from 2015-2019.

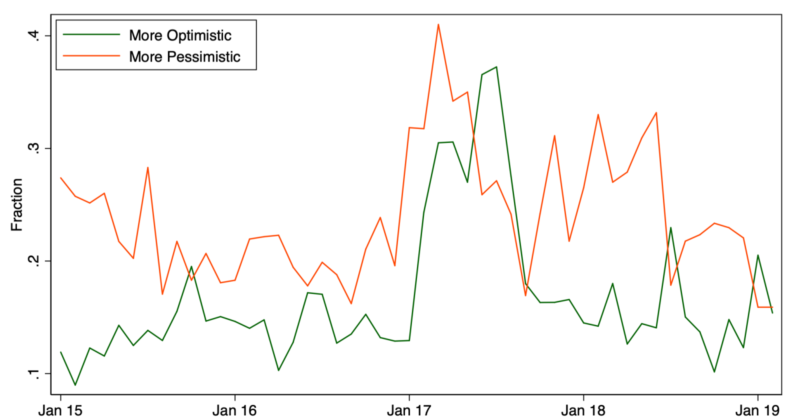

Figure 3.7: Government Policy Beliefs

Notes: the green line plots the fraction of individuals who become more optimistic about government policy relative to their previous response; the orange line plots the fraction of individuals who become more pessimistic. Monthly data, from 2015-2019.

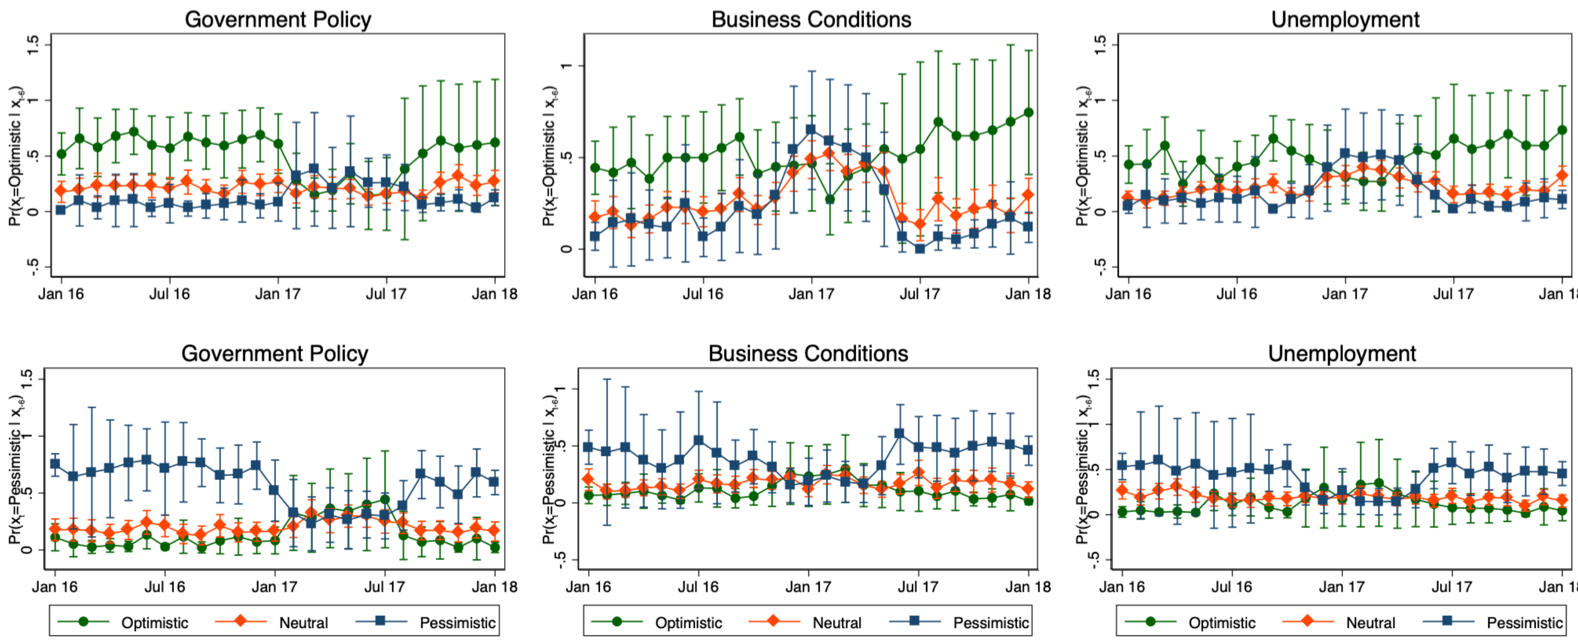

Figure 3.8: Switching Probabilities

Notes: probability of an optimistic (top row panels) or pessimistic response (bottom row panels) conditional on the individual's response from the previous survey 6 months ago. Survey questions are from the MSC regarding Government Policy (first column panels), Business Conditions (second column panels), or Unemployment (third column panels). The green circles are probabilities conditional on a previous optimistic response; the orange diamonds are probabilities conditional on a previous neutral response; and the blue squares are probabilities conditional on a previous pessimistic response. Estimates from a period-by-period multinomial logit model; vertical lines represent 95% confidence intervals.

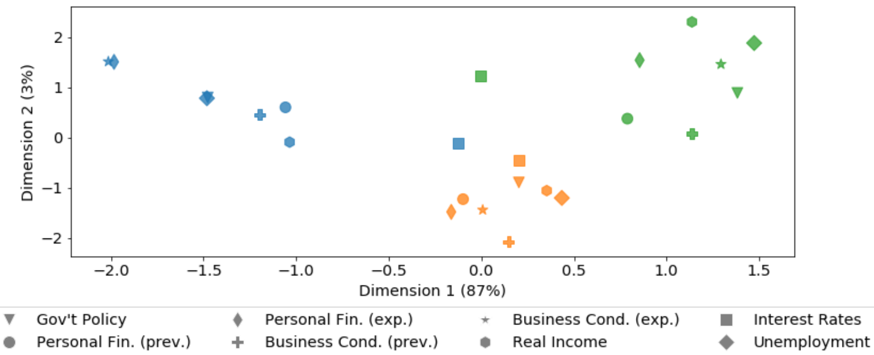

Figure 3.9: MCA Coordinates

Notes: selected coordinates for the first dimension (x-axis) and second dimension (y-axis) of a multiple correspondence analysis (MCA). Green points correspond to optimistic responses; orange points to neutral responses; and blue points to pessimistic responses. Each marker type corresponds to different questions from the MSC.

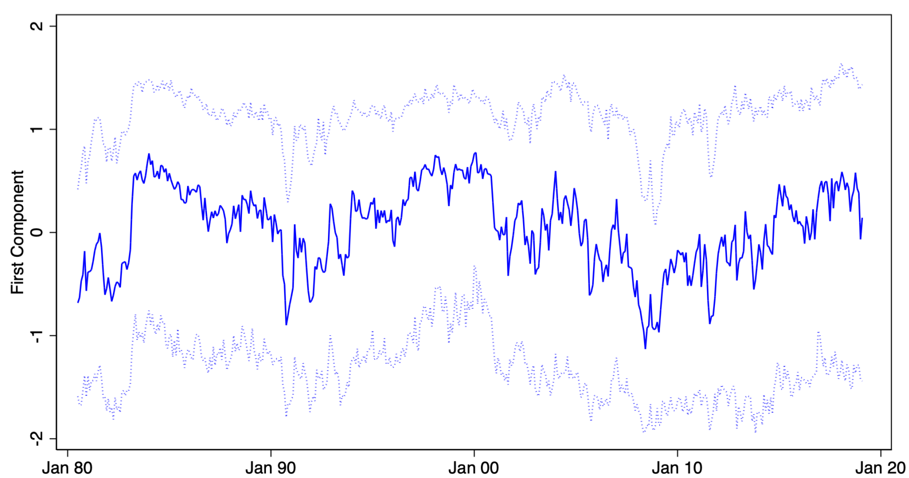

Figure 3.10: Sentiment Distribution Across Time

Notes: time series of the first component $f_{i,t}$ from an MCA analysis. The solid line is the median value of sentiment, while the dotted lines are the 90-10 percent distribution.

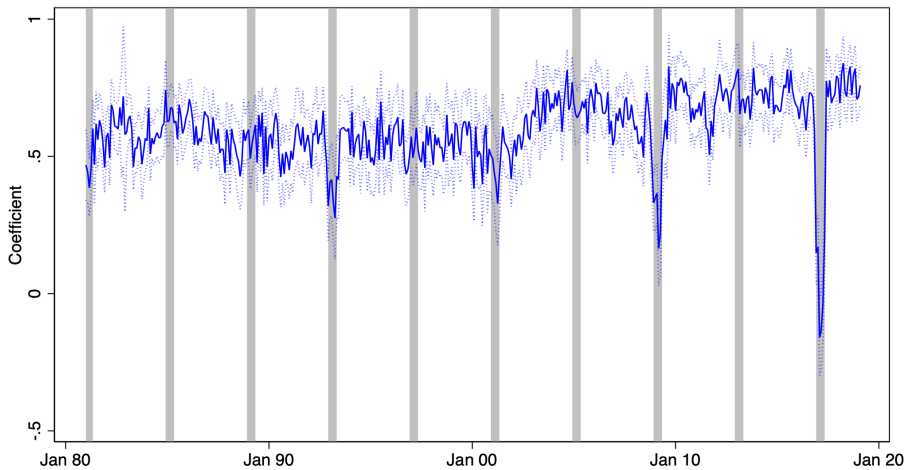

Figure 3.11: Sentiment Autocorrelation

Notes: coefficient from period-by-period regressions pooled across respondents $f_{i,t} = \alpha_t + \beta_t f_{i,t-6} + \varepsilon_{i,t}$. $f_{i,t}$ is the first component from an MCA analysis. Shaded regions correspond to 6-month periods following presidential elections. Dotted lines represent 2-standard error bands.

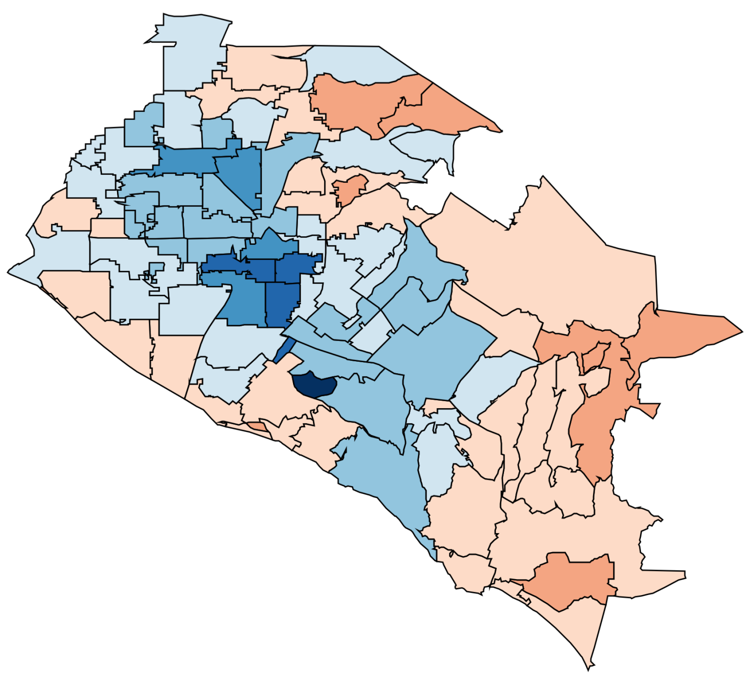

Figure 3.12: Orange County Votes by Zip Codes

Notes: vote shares in the 2016 presidential election in Orange County, California, at the zip code level. Precinct votes were allocated to the zip codes based on overlapping area.

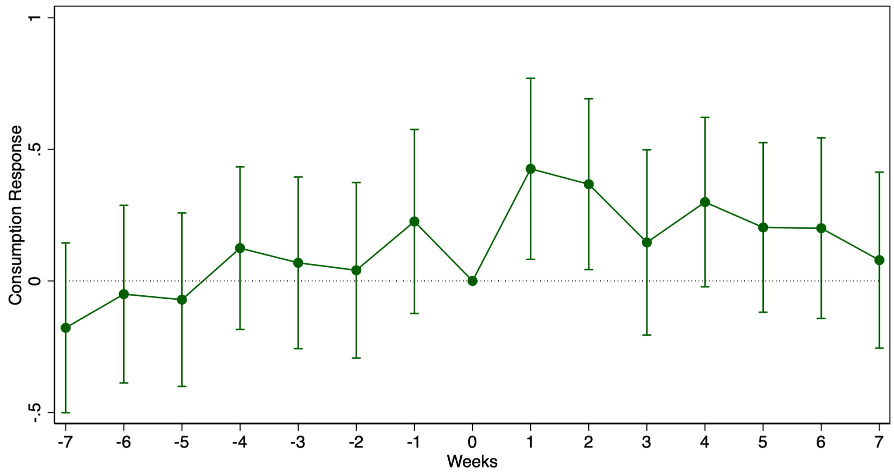

Figure 3.13: Event Study of the Consumption Response

Notes: $\hat{\beta}_k$ from event study described in Eq. (3.9) in the weeks preceding and following the 2016 presidential election. Vertical lines represent 95% confidence intervals. Standard errors are clustered at the zip code level.

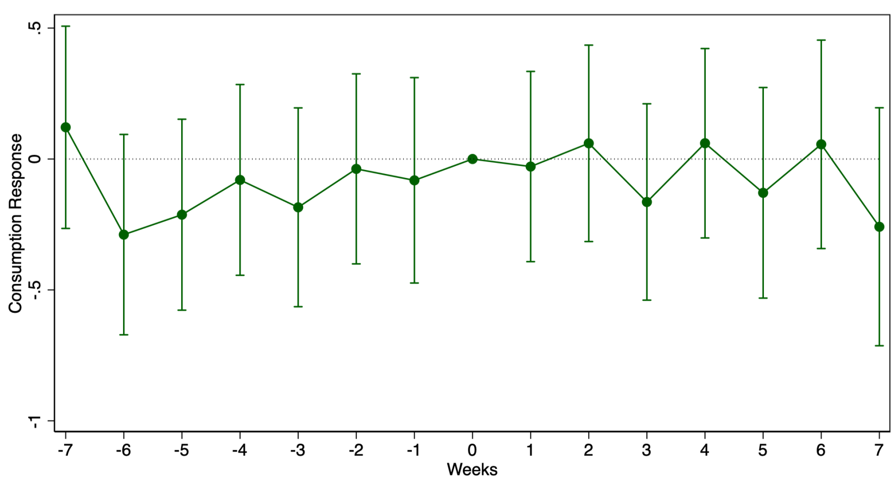

Figure 3.14: Placebo Event Study of the Consumption Response

Notes: $\hat{\beta}_k$ from event study described in Eq. (3.9) in the weeks preceding and following the week in 2015 one year prior to the 2016 presidential election. Vertical lines represent 95% confidence intervals. Standard errors are clustered at the zip code level.

Bibliography

Adrian, T. and Duarte, F. (2018). Financial Vulnerability and Monetary Policy. CEPR Discussion Papers 12680, C.E.P.R. Discussion Papers.

Alvarez, F., Atkeson, A., and Kehoe, P. J. (2002). Money, Interest Rates, and Exchange Rates with Endogenously Segmented Markets. *Journal of Political Economy*, 110(1):73–112.

Andris, C., Lee, D., Hamilton, M. J., Martino, M., Gunning, C. E., and Selden, J. A. (2015). The rise of partisanship and super-cooperators in the us house of representatives. *PloS one*, 10(4):e0123507.

Ang, A. and Piazzesi, M. (2003). A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables. *Journal of Monetary Economics*, 50(4):745–787.

Angeletos, G.-M. and Lian, C. (2018). Forward guidance without common knowledge. *American Economic Review*, 108(9):2477–2512.

Bachmann, R., Berg, T. O., and Sims, E. R. (2015). Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence. *American Economic Journal: Economic Policy*, 7(1):1–35.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Bartels, L. M. (2002). Beyond the running tally: Partisan bias in political perceptions. *Political Behavior*, 24(2):117–150.

Bauer, M. D. and Rudebusch, G. D. (2014). The Signaling Channel for Federal Reserve Bond Purchases. *International Journal of Central Banking*, 10(3):233–289.

Beraja, M., Fuster, A., Hurst, E., and Vavra, J. (2015). Regional Heterogeneity and Monetary Policy. Staff Reports 731, Federal Reserve Bank of New York.

Bernanke, B. S. and Kuttner, K. N. (2005). What Explains the Stock Market's Reaction to Federal Reserve Policy? *Journal of Finance*, 60(3):1221–1257.

Bhattarai, S., Eggertsson, G. B., and Gafarov, B. (2015). Time Consistency and the Duration of Government Debt: A Signalling Theory of Quantitative Easing. NBER Working Papers 21336, National Bureau of Economic Research, Inc.

Blanchard, O. J. and Kahn, C. M. (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, 48(5):1305–1311.

Blinder, A. S., Ehrmann, M., de Haan, J., and Jansen, D.-J. (2016). Necessity as the mother of invention: Monetary policy after the crisis. Working Paper 22735, National Bureau of Economic Research.

Bram, J. and Ludvigson, S. (1998). Does Consumer Confidence Forecast Household Expenditure? A Sentiment Index Horse Race. *Economic Policy Review*, (Jun):59–78.

Brunnermeier, M. K., Eisenbach, T. M., and Sannikov, Y. (2012). Macroeconomics with Financial Frictions: A Survey. NBER Working Papers 18102, National Bureau of Economic Research, Inc.

Buiter, W. H. (1984). Saddlepoint Problems in Continuous Time Rational Expectations Models: A General Method and Some Macroeconomic Examples. *Econometrica*, 52(3):665–680.

Campbell, J. R., Evans, C. L., Fisher, J. D., and Justiniano, A. (2012). Macroeconomic Effects of Federal Reserve Forward Guidance. *Brookings Papers on Economic Activity*, 43(1 (Spring):1–80.

Campbell, J. Y., Pflueger, C., and Viceira, L. M. (2014). Monetary Policy Drivers of Bond and Equity Risks. NBER Working Papers 20070, National Bureau of Economic Research, Inc.

Carlstrom, C. T., Fuerst, T. S., and Paustian, M. (2017). Targeting Long Rates in a Model with Segmented Markets. *American Economic Journal: Macroeconomics*, 9(1):205–242.

Carroll, C. D., Fuhrer, J. C., and Wilcox, D. W. (1994). Does Consumer Sentiment Forecast Household Spending? If So, Why? *American Economic Review*, 84(5):1397–1408.

Chen, H., Cúrdia, V., and Ferrero, A. (2012). The macroeconomic effects of large-scale asset purchase programmes. *The Economic Journal*, 122(564):F289–F315.

Chodorow-Reich, G. (2014). Effects of Unconventional Monetary Policy on Financial Institutions. *Brookings Papers on Economic Activity*, 45(1 (Spring)):155–227.

Christensen, J. H. E. and Rudebusch, G. D. (2012). The Response of Interest Rates to US and UK Quantitative Easing. *Economic Journal*, 122(564):385–414.

Clarida, R., Galí, J., and Gertler, M. (2000). Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *The Quarterly Journal of Economics*, 115(1):147–180.

Coibion, O. and Gorodnichenko, Y. (2012). Why are target interest rate changes so persistent? *American Economic Journal: Macroeconomics*, 4(4):126–62.

Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018). The Formation of Expectations, Inflation, and the Phillips Curve. *Journal of Economic Literature*, 56(4):1447–1491.

D'Acunto, F., Hoang, D., Paloviita, M., and Weber, M. (2018). Human Frictions in the Transmission of Economic Policy. Manuscript.

D'Acunto, F., Hoang, D., and Weber, M. (2016). The Effect of Unconventional Fiscal Policy on Consumption Expenditure. NBER Working Paper 22563.

D'Amico, S. and King, T. B. (2013). Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply. *Journal of Financial Economics*, 108(2):425–448.

Driessen, G. A. (2016). How treasury issues debt. *Congressional Research Service Report R40767*.

Farhi, E. and Werning, I. (2017). Monetary Policy, Bounded Rationality, and Incomplete Markets. NBER Working Papers 23281, National Bureau of Economic Research, Inc.

Fleming, M. J. (2007). Who buys treasury securities at auction? *Federal Reserve Bank of New York Current Issues in Economics and Finance*, 13(1).

Fleming, M. J. and Liu, W. (2016). Intraday pricing and liquidity effects of us treasury auctions.

Gabaix, X. (2016). A Behavioral New Keynesian Model. NBER Working Papers 22954, National Bureau of Economic Research, Inc.

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

Garbade, K. D. and Ingber, J. F. (2005). The Treasury auction process: objectives, structure, and recent acquisitions. *Current Issues in Economics and Finance*, 11(Feb).

Gertler, M. and Karadi, P. (2013). QE 1 vs. 2 vs. 3. . . : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool. *International Journal of Central Banking*, 9(1):5–53.

Gillitzer, C. and Prasad, N. (2018). The Effect of Consumer Sentiment on Consumption: Cross-Sectional Evidence from Elections. *American Economic Journal: Macroeconomics*, 10(4):234–269.

Gorodnichenko, Y. and Ray, W. (2017). The Effects of Quantitative Easing: Taking a Cue from Treasury Auctions. NBER Working Papers 24122, National Bureau of Economic Research, Inc.

Gorodnichenko, Y. and Weber, M. (2016). Are Sticky Prices Costly? Evidence from the Stock Market. *American Economic Review*, 106(1):165–199.

Greenwood, R., Hanson, S. G., Rudolph, J. S., and Summers, L. (2016). Debt Management Conflicts Between the US Treasury and the Federal Reserve. In Wessel, D., editor, *The $13 Trillion Question: How America Manages Its Debt*, chapter 2, pages 43–75. Brookings Institution Press.

Greenwood, R. and Vayanos, D. (2014). Bond Supply and Excess Bond Returns. *Review of Financial Studies*, 27(3):663–713.

Gurkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. Treasury Yield Curve: 1961 to the Present. *Journal of Monetary Economics*, 54(8):2291–2304.

Hamilton, J. D. and Wu, J. C. (2012a). The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment. *Journal of Money, Credit and Banking*, 44:3–46.

Hamilton, J. D. and Wu, J. C. (2012b). The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment. *Journal of Money, Credit and Banking*, 44:3–46.

He, Z., Kelly, B., and Manela, A. (2016). Intermediary asset pricing: New evidence from many asset classes. Working Paper 21920, National Bureau of Economic Research.

He, Z. and Krishnamurthy, A. (2013). Intermediary Asset Pricing. *American Economic Review*, 103(2):732–770.

Hördahl, P., Tristani, O., and Vestin, D. (2006). A Joint Econometric Model of Macroeconomic and Term-Structure Dynamics. *Journal of Econometrics*, 131(1-2):405–444.

Kamdar, R. (2018). The Inattentive Consumer: Sentiment and Expectations. Working paper.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4):1199–1239.

Kőszegi, B. and Matějka, F. (2018). An Attention-Based Theory of Mental Accounting. Manuscript.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2011). The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy. *Brookings Papers on Economic Activity*, (2 (Fall)):215–265.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The Aggregate Demand for Treasury Debt. *Journal of Political Economy*, 120(2):233–267.

Kuchler, T. and Zafar, B. (2015). Personal Experiences and Expectations About Aggregate Outcomes. Staff Reports 748, Federal Reserve Bank of New York.

Kuttner, K. (2006). Can central banks target bond prices? Technical Report 12454, National Bureau of Economic Research.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.

Kyle, A. S. and Xiong, W. (2001a). Contagion as a Wealth Effect. *Journal of Finance*, 56(4):1401–1440.

Kyle, A. S. and Xiong, W. (2001b). Contagion as a Wealth Effect. *Journal of Finance*, 56(4):1401–1440.

Lauszweski, J. W., Kamradt, M., and Gibbs, D. (2014). Understanding treasury futures. https://www.cmegroup.com/education/files/understanding-treasury-futures.pdf.

Lou, D., Yan, H., and Zhang, J. (2013). Anticipated and Repeated Shocks in Liquid Markets. *Review of Financial Studies*, 26(8):1891–1912.

Ludvigson, S. C. (2004). Consumer Confidence and Consumer Spending. *Journal of Economic Perspectives*, 18(2):29–50.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2018). Dynamic Rational Inattention: Analytical Results. *Journal of Economic Theory*, 176(C):650–692.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.

Malmendier, U. and Nagel, S. (2016). Learning from Inflation Experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Martin, C. and Milas, C. (2012). Quantitative easing: a sceptical survey. *Oxford Review of Economic Policy*, 28(4):750–764.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10):3133–3158.

Mian, A., Sufi, A., and Khoshkhou, N. (2018). Government Economic Policy, Sentiments, and Consumption. Fama-miller working paper, Chicago Booth Research Paper.

Pew Research Center (2014). Political Polarization in the American Public. Technical report.

Ray, W. (2017). Monetary policy and preferred habitat. Working paper.

Romer, C. D. and Romer, D. H. (2017). New evidence on the aftermath of financial crises in advanced countries. *American Economic Review*, 107(10):3072–3118.

Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.

Rudebusch, G. D. and Wu, T. (2008). A Macro-Finance Model of the Term Structure, Monetary Policy and the Economy. *Economic Journal*, 118(530):906–926.

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3):665–690.

Swanson, E. T. and Williams, J. C. (2014). Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates. *American Economic Review*, 104(10):3154–3185.

Vayanos, D. and Vila, J.-L. (2009). A Preferred-Habitat Model of the Term Structure of Interest Rates. NBER Working Papers 15487, National Bureau of Economic Research, Inc.

Wallace, N. (1981). A Modigliani-Miller Theorem for Open-Market Operations. *American Economic Review*, 71(3):267–274.

Werning, I. (2011). Managing a Liquidity Trap: Monetary and Fiscal Policy. NBER Working Papers 17344, National Bureau of Economic Research, Inc.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Appendix A

Monetary Policy and the Limits to Arbitrage

A.1 Proofs

Proof of Lemma 1.1. The results follow from analyzing the characteristic polynomial of Υ:

$$c(\lambda) = \lambda^2 - \kappa_r\lambda - \varsigma^{-1}\kappa_r\phi_x\hat{A}_r$$

1. Note that $c(\lambda)$ is increasing iff $\lambda > \frac{1}{2}\kappa_r$. Additionally, $c(\lambda) \to \infty$ whenever $\lambda \to \pm\infty$. Note

   $$c(0) = c(\kappa_r) = -\varsigma^{-1}\kappa_r\phi_x\hat{A}_r$$

   Hence if $\hat{A}_r > 0$, the expression above is negative which implies $c(\lambda)$ has two roots: $\lambda_1 > \kappa_r$ and $\lambda_2 < 0$.

2. Setting $c(\lambda)$ to zero and solving for $\hat{A}_r$ gives eq. (1.11). Note that $h(\lambda)$ is a positive quadratic with zeros at $\lambda = 0$ and $\lambda = \kappa_r$; for $\lambda \in (0, \kappa_r)$, $h(\lambda)$ is negative.

3. The eigenvector associated with λ_1 is

   $$q_1 \equiv \begin{bmatrix} -\frac{\lambda_1}{\varsigma^{-1}\hat{A}_r} \\ 1 \end{bmatrix}$$

   In this case Ω, the matrix governing the equilibrium dynamics of the jump variables in eq. (1.10), is simplify a scalar, given by $\omega_x \equiv \frac{q_{21}}{q_{11}}$. Substituting the solution for $\hat{A}_r$ gives the result.

□

Proof of Lemma 1.2. This result is the scalar counterpart of the optimality conditions derived in Prop. 1.7. The expected instantaneous return simplifies to

$$\mu_{t,\tau} = A_r'(\tau) r_t + C'(\tau) + \lambda(r_t - r^{SS}) A_r(\tau) + \frac{1}{2}\sigma_r^2 A_r(\tau)^2 \tag{A1}$$

□

Proof of Lemma 1.3. The functional forms of $A_r(\tau)$, $C(\tau)$, and ν are again the scalar analogues of the results in Prop. 1.7. In this case, $\mathbf{M} \equiv \nu$ and simplifies to eq. (1.16) so $C(\tau)$ simplifies to

$$C(\tau) = n_1(\tau) Z_C - \frac{1}{2}\sigma_r^2 n_2(\tau) \tag{A2}$$

where

$$\begin{aligned}
Z_C &= \frac{a\sigma_r^2 N_2 + \lambda r^{SS}}{1 + a\sigma_r^2 N_1} \\
n_1(\tau) &= \int_0^\tau A_r(u)\, \mathrm{d}u \\
n_2(\tau) &= \int_0^\tau A_r(u)^2\, \mathrm{d}u \\
N_1 &= \int_0^T \alpha(\tau) A_r(\tau) n_1(\tau)\, \mathrm{d}\tau \\
N_2 &= \int_0^T \alpha(\tau) A_r(\tau) \left(\bar{\beta}(\tau)\tau + \frac{1}{2}\sigma_r^2 n_2(\tau) \right) \mathrm{d}\tau
\end{aligned}$$

If $a = 0$ it immediately follows that $\nu = \lambda$, so throughout assume that $a > 0$. In this case, solving for ν is a fixed point problem. Since $\alpha(\tau) > 0\ \forall \tau$, this implies

$$a\sigma_r^2 \int_0^T \alpha(\tau)\tau^2 f(\nu\tau)^2\, \mathrm{d}\tau > 0$$

Hence, the right-hand-side of eq. (1.16) is strictly greater than λ.

Next, note

$$\frac{\partial f(\nu\tau)^2}{\partial \nu} = 2\tau f(\nu\tau) f'(\nu\tau)$$

$f(\cdot)$ is a strictly decreasing, convex function, approaching 0 as $x \to \infty$. So the above expression is negative. This implies the right-hand-side of eq. (1.16) is strictly decreasing in ν. Further, as $\nu \to \infty$, if $\tau > 0$ then $f(\nu\tau) \to 0$. Thus

$$\int_0^T \alpha(\tau)\tau^2 f(\nu\tau)^2\, \mathrm{d}\tau \to 0$$

which implies the right-hand-side of eq. (1.16) approaches λ as $\nu \to \infty$.

Hence there exists a unique $\nu \in (\lambda, \infty)$ such that eq. (1.16) is satisfied. Treating ν as a function of λ, I derive some additional properties. Recall $\nu(\lambda) > \lambda$. Further,

$$\begin{aligned}
\frac{\partial f(\nu\tau)^2}{\partial \lambda} &= 2\tau f(\nu\tau) f'(\nu\tau) \frac{\partial \nu}{\partial \lambda} \\
\implies \frac{\partial \nu}{\partial \lambda} &= \left(1 - 2a\sigma_r^2 \int_0^T \alpha(\tau)\tau^3 f(\nu\tau) f'(\nu\tau)\,\mathrm{d}\tau\right)^{-1}
\end{aligned}$$

Also note $\alpha(\tau)\tau^3 f(\nu\tau) \geq 0$ and $f'(\nu\tau) \leq 0$, so

$$0 < \frac{\partial \nu}{\partial \lambda} < 1$$

Fixing λ, note that

$$\frac{\partial \nu}{\partial a} = \frac{\sigma_r^2 \int_0^T \alpha(\tau)\tau^2 f(\nu\tau)^2\,\mathrm{d}\tau}{1 - 2a\sigma_r^2 \int_0^T \alpha(\tau)\tau^3 f(\nu\tau) f'(\nu\tau)\,\mathrm{d}\tau}$$

Hence ν is increasing in a.

Finally, $\hat{A}_r$ is given by

$$\hat{A}_r \equiv \int_0^T \frac{\eta(\tau)}{\tau} A_r(\tau)\,\mathrm{d}\tau$$

eq. (1.17) is obtained by substituting eq. (1.15). If $\eta(\tau)$ is the Dirac delta function, then $g(\lambda) = 1\ \forall\lambda$. Otherwise,

$$g'(\lambda) = \int_0^T \eta(\tau)\tau f'(\nu\tau)\,\mathrm{d}\tau\,\frac{\partial \nu}{\partial \lambda}$$

hence $g'(\lambda) < 0$ and additionally $g(\lambda) \to 0$ as $\lambda \to \infty$.

$\square$

Proof of Prop. 1.1. In general equilibrium, λ_1 and $\hat{A}_r$ are determined by the intersection of eqs. (1.17) and (1.11). Recall from the proof of Lemma 1.3, for $\lambda > 0$, $g(\lambda)$ is strictly positive, decreasing, and approaches 0 as $\lambda \to \infty$. From the proof of Lemma 1.1, $h(\lambda)$ is negative for $0 < \lambda < \kappa_r$, but is strictly increasing and grows without bound when $\lambda \geq \kappa_r$. Hence there exists a unique $\lambda_1 > \kappa_r > 0$ such that $g(\lambda_1) = h(\lambda_1)$. Figure 1.2 plots examples of the intersection of g and h.

Given any value of r^*, the steady state is

$$\begin{aligned}
r^{SS} &= \frac{\bar{r} - \hat{C}}{\hat{A}_r} \\
x^{SS} &= \frac{r^* - r^{SS}}{\phi_x}
\end{aligned}$$

$\square$

Proof of Corollary 1.1.1. Recall from Prop. 1.1 that $\hat{A}_r$ (and the associated eigenvalue λ_1) are determined by the intersection of a downward sloping curve $g(\lambda)$ and an upward sloping (in the neighborhood of the intersection λ_1) curve $h(\lambda)$. And while these curves depend on the parameters of the model, the parameter dependence is disjoint: $g(\cdot)$ only depends on the parameters governing the term structure side of the model, while $h(\cdot)$ depends on the macro parameters. This greatly simplifies studying comparative statics.

1. The proof of Lemma 1.3 showed that ν is increasing in a. Then

$$\frac{\partial g(\lambda)}{\partial a} = \int_0^T \eta(\tau) f'(\nu\tau)\,\mathrm{d}\tau\,\frac{\partial \nu}{\partial a}$$

   which implies that for any given λ, an increase in a leads to a downward shift in $g(\lambda)$. Since $h(\lambda)$ is unchanged, this implies that the point of intersection shifts downward and to the left.

   Further, note that as $a \to \infty$, $\nu \to \infty$. To see why, suppose instead that $\nu \to \nu^* < \infty$. Then $f(\nu\tau) \to f(\nu^*\tau) > 0$. But then

$$a\sigma_r^2 \int_0^T \alpha(\tau)\tau^2 f(\nu\tau)^2\,\mathrm{d}\tau \to \infty$$

   and therefore the right-hand side eq. (1.16) grows without bound, a contradiction. This immediately implies that $f(\nu\tau) \to 0$ and therefore so too does $g(\lambda)$ for $\lambda > 0$. As before, this says that the point of intersection of g and h continues shifting downward and to the left until $\hat{A}_r = 0$ and $\lambda_1 = \kappa_r$.

   Note that if $a = 0$, $g(\lambda)$ is independent of changes in σ_r. But when $a \neq 0$, the same arguments above show that

$$\frac{\partial \hat{A}_r}{\partial \sigma_r} < 0, \quad \frac{\partial \lambda_1}{\partial \sigma_r} < 0$$
$$\lim_{\sigma_r \to \infty} \hat{A}_r = 0, \quad \lim_{\sigma_r \to \infty} \lambda_1 = \kappa_r$$

2. Differentiating eq. (1.11) with respect to κ_r gives

$$-\frac{\lambda^2}{\varsigma^{-1}\kappa_r^2\phi_x} < 0 \ \ \forall\lambda$$

   Hence the intersection shifts down and to the right.

3. Differentiating eq. (1.11) with respect to ϕ_x gives

$$-\frac{\lambda(\lambda-\kappa_r)}{\varsigma^{-1}\kappa_r\phi_x^2}$$

   which is negative whenever $\lambda > \kappa_r$. Hence the intersection shifts down and to the right.

4. Note for arbitrary functions f and g, where f and g satisfy the following:

$$f(x) > 0, \quad f'(x) < 0$$
$$g(x) > 0 \iff x < x^*, \quad \int g(x)\,\mathrm{d}x = 0$$

Then $\int f(x)g(x)\,\mathrm{d}x > 0$. So set $f(\nu\tau)$ as f, and $\eta^s(\tau) - \eta^\ell(\tau)$ as g, which gives

$$\int_0^T (\nu\tau)\eta^s(\tau)\,\mathrm{d}\tau > \int_0^T (\nu\tau)\eta^\ell(\tau)\,\mathrm{d}\tau$$

which says that $g^s(\lambda) > g^\ell(\lambda)$ (using the notation from a previous Lemma). In other words, g shifts down moving from the η^s weights to the η^ℓ weights. Since h is unchanged, the equilibrium results follow.

□

Proof of Corollary 1.1.2. The expression for the optimal target eq. (1.18) follows from the steady state results derived in the proof of Prop. 1.1.

Next, consider demand functions $\bar{\beta}^h(\tau) \geq \bar{\beta}^h(\tau)$. Letting superscripts h and ℓ denote the equilibrium outcomes under $\bar{\beta}^h(\tau)$ and $\bar{\beta}^\ell(\tau)$, note that the coefficient function $A_r^h(\tau) = A_r^\ell(\tau)$, so the only difference comes about due to changes in the constant function. Then note that

$$\hat{C}^h - \hat{C}^\ell = -\hat{A}_r(r^{*h} - r^{*\ell})$$
$$\implies r^{*h} - r^{*\ell} = -\frac{a\sigma_r^2\hat{n}_1(N_2^h - N_2^\ell)}{\lambda\hat{n}_1 + \hat{A}_r(1 + a\sigma_r^2 N_1)}$$

where the second line follows from eq. (A2), and the expressions defined in the proof of Lemma 1.3. Then, since $\bar{\beta}^h(\tau) \geq \bar{\beta}^\ell(\tau)$,

$$N_2^h - N_2^\ell = \int_0^T \alpha(\tau)\tau A_r(\tau)\left(\bar{\beta}^h(\tau) - \bar{\beta}^\ell(\tau)\right)\mathrm{d}\tau > 0$$
$$\implies r^{h*} \leq r^{\ell*}$$

where the inequalities are strict whenever $a > 0$.

□

Proof of Prop. 1.2. To solve for x_0, for $t \leq t^\diamond$ the dynamics of the policy rate under the peg imply that

$$E_0 r_t = r^\diamond + e^{-\kappa_r^\diamond t}(r_0 - r^\diamond)$$
$$\implies E_0 r_{t^\diamond} = r^\diamond$$

In order to satisfy transversality conditions, once the central bank returns to a Taylor rule, the output gap must satisfy $x_{t^\diamond} = \omega_x(r_{t^\diamond} - r^{SS})$. Thus,

$$E_0 x_{t^\diamond} = \omega_x(r^\diamond - r^{SS})$$

Solving for $E_0 x_{t^\diamond}$ given x_0 gives

$$E_0 x_{t^\diamond} = x_0 + t^\diamond \varsigma^{-1}(\hat{A}_r^\diamond r^\diamond + \hat{C}^\diamond - \bar{r})$$

Hence,

$$\begin{aligned}\frac{\partial x_0}{\partial r^\diamond} &= \omega_x - t^\diamond \varsigma^{-1} \hat{A}_r^\diamond \\ \frac{\partial^2 x_0}{\partial r^\diamond \partial t^\diamond} &= -\varsigma^{-1} \hat{A}_r^\diamond\end{aligned}$$

Then the result follows since

$$\frac{\partial \hat{A}_r^\diamond}{\partial a} < 0, \quad \left|\frac{\partial \omega_x}{\partial a}\right| < 0$$

$\square$

Proof of Lemma 1.4. Following the same steps as in Lemma 1.1:

$$\hat{A}_r = \frac{\lambda_1(\lambda_1 - \kappa_r)}{\varsigma^{-1}\kappa_r \phi_x}$$

Substituting this into Υ, the eigenvalue decomposition is

$$\begin{aligned}\Lambda &= \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \kappa_\beta & 0 \\ 0 & 0 & -\lambda_1 + \kappa_r \end{bmatrix} \\ Q &= \begin{bmatrix} \frac{\kappa_r \phi_x}{-\lambda_1 + \kappa_r} & \frac{\kappa_r \phi_x}{\kappa_r - \kappa_\beta} & \frac{\kappa_r \phi_x}{\lambda_1} \\ 0 & \frac{\kappa_\beta^2 - \kappa_\beta \kappa_r - \lambda_1^2 + \lambda_1 \kappa_r}{(\kappa_r - \kappa_\beta)\varsigma^{-1}\hat{A}_\beta} & 0 \\ 1 & 1 & 1 \end{bmatrix}\end{aligned}$$

Solving for the rational expectations equilibrium gives

$$\begin{aligned}\Gamma &= \begin{bmatrix} \lambda_1 & \frac{\kappa_r \phi_x \varsigma^{-1} \hat{A}_\beta}{\lambda_1 - \kappa_r + \kappa_\beta} \\ 0 & \kappa_\beta \end{bmatrix} \\ \Omega &= \begin{bmatrix} \frac{-\lambda_1 + \kappa_r}{\kappa_r \phi_x} & \frac{-\varsigma^{-1}\hat{A}_\beta}{\lambda_1 - \kappa_r + \kappa_\beta} \end{bmatrix}\end{aligned}$$

$\square$

Proof of Lemma 1.5. This follows from the general case in Prop. 1.7, where

$$\Gamma = \begin{bmatrix} \gamma_1 & \gamma_{12} \\ 0 & \gamma_2 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & 0 \end{bmatrix} \implies \mathbf{M} = \begin{bmatrix} \nu & 0 \\ \nu_{12} & \gamma_2 \end{bmatrix}$$

and the elements of $\mathbf{M}$ are given by

$$\begin{aligned} \nu &= \gamma_1 + a\sigma_r^2 \int_0^T \alpha(\tau) A_r(\tau)^2 \, \mathrm{d}\tau \\ \nu_{12} &= \gamma_{12} + a\sigma_r^2 \int_0^T \alpha(\tau) \left[A_\beta(\tau) - \tau\theta(\tau)\right] A_r(\tau) \, \mathrm{d}\tau \end{aligned}$$

The eigenvector decomposition gives

$$\mathbf{D} = \begin{bmatrix} \nu & 0 \\ 0 & \gamma_2 \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} 1 & 0 \\ \frac{\nu_{12}}{\nu - \gamma_2} & \frac{\nu_{12}}{\gamma_2 - \nu} \end{bmatrix}$$

and thus

$$\begin{aligned} & \mathbf{A}(\tau) = \mathbf{G}\mathbf{D}^{-1} \left[\mathbf{I} - \exp(-\mathbf{D}\tau)\right] \mathbf{1} \\ \implies & A_r(\tau) = \tau f(\nu\tau) \\ & A_\beta(\tau) = \frac{\nu_{12}}{\nu - \gamma_2} \tau (f(\nu\tau) - f(\gamma_2\tau)) \end{aligned}$$

Hence, rewriting the expression for ν gives

$$\nu = \gamma_1 + a\sigma_r^2 \int_0^T \alpha(\tau) \tau^2 f(\nu\tau)^2 \, \mathrm{d}\tau$$

which is equivalent to eq. (1.16).

Substituting the solution for $A_\beta(\tau)$ into the expression for ν_{12} and solving gives

$$\nu_{12} = \frac{\gamma_{12} - a\sigma_r^2 \int_0^T \alpha(\tau)\tau^2 f(\nu\tau)\theta(\tau)}{1 - \frac{a\sigma_r^2}{\nu - \gamma_2} \int_0^T \alpha(\tau)\tau^2 (f(\nu\tau) - f(\gamma_2\tau)) \, \mathrm{d}\tau} \tag{A3}$$

Thus, integrating and weighting by $\frac{\eta(\tau)}{\tau}$ gives

$$\begin{aligned} \hat{A}_r &= \int_0^T \eta(\tau) f(\nu\tau) \, \mathrm{d}\tau \\ \hat{A}_\beta &= \frac{\nu_{12}}{\nu - \gamma_2} \int_0^T \eta(\tau)(f(\nu\tau) - f(\gamma_2\tau)) \, \mathrm{d}\tau \end{aligned}$$

and hence, since the fixed point problem that solves ν is equivalent to the baseline rigid price model, so is the expression for $\hat{A}_r$.

□

Proof of Prop. 1.3. 1. Combining the coefficients used in Lemmas 1.4 and 1.5:

$$\begin{aligned}
\gamma_1 &= \lambda_1 \\
\gamma_2 &= \kappa_\beta \\
\gamma_{12} &= \frac{\kappa_r \phi_x \varsigma^{-1} \hat{A}_\beta}{\lambda_1 - \kappa_r + \kappa_\beta} \\
&= \frac{\kappa_r \phi_x \varsigma^{-1}}{\lambda_1 - \kappa_r + \kappa_\beta} \left(\frac{\nu_{12}}{\nu - \gamma_2} \int_0^T \eta(\tau)(f(\nu\tau) - f(\gamma_2\tau)) \,\mathrm{d}\tau \right)
\end{aligned}$$

Substituting these expressions into eq. (A3) and rearranging gives

$$\nu_{12} = \frac{-a\sigma_r^2 N_1}{1 - \frac{\phi_x \kappa_r \varsigma^{-1}}{(\lambda_1 - \kappa_r + \kappa_\beta)(\nu - \kappa_\beta)} N_2 - \frac{a\sigma_r^2}{\nu - \kappa_\beta} N_3} \tag{A4}$$

where N_1, N_2, N_3 are integral expressions:

$$\begin{aligned}
N_1 &= \int_0^T \alpha(\tau)\tau^2 f(\nu\tau)\theta(\tau) \,\mathrm{d}\tau \\
N_2 &= \int_0^T \eta(\tau)(f(\nu\tau) - f(\kappa_\beta\tau)) \,\mathrm{d}\tau \\
N_3 &= \int_0^T \alpha(\tau)\tau^2(f(\nu\tau) - f(\kappa_\beta\tau)) \,\mathrm{d}\tau
\end{aligned}$$

First, focusing on the denominator of eq. (A4),

$$\nu > \kappa_\beta \iff f(\nu\tau) < f(\kappa_\beta\tau)$$

hence the denominator is strictly positive. Then, since by assumption $\theta(\tau) \geq 0$, the numerator of eq. (A4) is negative (strictly if $a > 0$). Therefore, $\nu_{12} \leq 0$ and $\hat{A}_\beta \geq 0$ (with strict inequalities when $a > 0$).

2. Taking the limit as $\kappa_\beta \to \infty$, from eq. (A4) the limiting value of ν_{12} is

$$\nu_{12} \to -a\sigma_r^2 \int_0^T \alpha(\tau)\tau^2 f(\nu\tau)\theta(\tau)$$

Since ν is independent of κ_β, this limit is bounded. Further, $f(\kappa_\beta\tau) \to 0$ as well. Hence, taking the limit of eq. (1.22) gives

$$\frac{\nu_{12}}{\nu - \kappa_\beta} \int_0^T \eta(\tau)(f(\nu\tau) - f(\kappa_\beta\tau)) \,\mathrm{d}\tau \to 0$$

□

Proof of Lemma 1.6. 1. The characteristic polynomial of Υ is

$$c(\lambda) \equiv \lambda^3 - (\kappa_r - \rho)\lambda^2 - \varsigma^{-1}(\hat{A}_r\kappa_r\phi_x + \delta + \kappa_r\rho\varsigma)\lambda - \varsigma^{-1}\kappa_r(\hat{A}_r(\delta\phi_\pi + \rho\phi_x) - \delta)$$

Hence:

$$\begin{aligned}\lim_{\lambda\to+\infty} c(\lambda) &= +\infty\\ \lim_{\lambda\to-\infty} c(\lambda) &= -\infty\\ c(\kappa_r) &= -\varsigma^{-1}\kappa_r\hat{A}_r(\kappa_r\phi_x + \delta\phi_\pi + \rho\phi_x) < 0\end{aligned}$$

Hence there is some (real) $\lambda_1 > \kappa_r > 0$ such that $c(\lambda_1) = 0$. For the other two roots, note

$$\begin{aligned}c'(\lambda) &= 3\lambda^2 - 2(\kappa_r - \rho)\lambda - \varsigma^{-1}(\hat{A}_r\kappa_r\phi_x + \delta + \kappa_r\rho\varsigma)\\ c(0) &= -\varsigma^{-1}\kappa_r(\hat{A}_r(\delta\phi_\pi + \rho\phi_x) - \delta)\\ c'(0) &= -\varsigma^{-1}(\hat{A}_r\kappa_r\phi_x + \delta + \kappa_r\rho\varsigma)\end{aligned}$$

So $c'(0) < 0$. If $c(0) > 0$, since $c(\kappa_r) < 0$ there will be another value $\lambda_2 \in (0, \kappa_r)$ such that $c(\lambda_2) = 0$. Hence $c(0) < 0$ is necessary for determinacy. This is satisfied iff eq. (1.24) holds.

To see that this condition is sufficient, since $c'(0) < 0$, $c(\lambda)$ has a local maximum for some $\lambda < 0$, denoted $c(z) = y$. If $y > 0$ then there are two real negative zeros, one less than z and one greater than z. If $y = 0$ then there are two duplicated real negative zeros (equal to z). Finally, if $y < 0$ then there are two conjugate complex zeros. The line that intersects $(\lambda_1, 0)$ and is tangent to $c(\lambda)$ is upward sloping, and is tangent to $c(\lambda)$ at a point (M, H) where $M < z$. The real components of the complex zeros is given by M.

2. Setting $c(\lambda)$ to zero and solving for $\hat{A}_r$ gives eq. (1.25). The fact that $h(\lambda) \to \infty$ as $\lambda \to \infty$ follows for the sign restrictions on the parameters.

Note that the zeros of $h(\lambda)$ are

$$\left\{\kappa_r,\ \frac{1}{2}\left(-\rho \pm \sqrt{\rho^2 + 4\delta\varsigma^{-1}}\right)\right\}$$

so there are two positive (real) zeros, and one negative (real) zero. Order these $z_1 > z_2 > 0 > z_3$. Then $z_2 + z_3 \le -\rho < 0$ (equality holds if $z_1 = \kappa_r$). Write

$$h(\lambda) = \frac{(\lambda - z_1)(\lambda - z_2)(\lambda - z_3)}{\varsigma^{-1}\kappa_r\left(\delta\phi_\pi + \rho\phi_x + \lambda\phi_x\right)}$$

Differentiating with respect to λ gives

$$h'(\lambda) = h(\lambda)\left[\frac{1}{\lambda - z_1} + \frac{1}{\lambda - z_2} + \frac{1}{\lambda - z_3} + \frac{-\phi_x}{\delta\phi_\pi + \rho\phi_x + \lambda\phi_x}\right]$$

If $\lambda \in [0, z_2)$, then $h(\lambda) > 0$ while the first and fourth terms in brackets are negative. Then

$$\frac{1}{\lambda - z_2} + \frac{1}{\lambda - z_3} = \frac{2\lambda - (z_2 + z_3)}{(\lambda - z_2)(\lambda - z_3)} < 0$$

since $z_3 < 0 < \lambda < z_2$ by assumption, and $z_2 + z_3 < 0$. Hence for $\lambda \in [0, z_2)$, $h'(\lambda) < 0$ and hence $h(0)$ is the maximum in this range.

Finally, $h(\lambda) > 0$ whenever $\lambda > z_1$. The second and third in brackets are positive, while

$$\frac{1}{\lambda - z_1} - \frac{\phi_x}{\delta\phi_\pi + \rho\phi_x + \lambda\phi_x} = \frac{\delta\phi_\pi + \rho\phi_x + z_1\phi_x}{(\lambda - z_1)(\delta\phi_\pi + \rho\phi_x + \lambda\phi_x)} > 0$$

hence for $\lambda > z_1$, $h(\lambda)$ is a positive strictly increasing function.

Also, note that the condition for determinacy is equivalent to

$$\hat{A}_r > h(0)$$

3. Substituting the expression for $\hat{A}_r$ into Υ and carrying out the eigenvalue decomposition and solving for Ω gives the result

□

Proof of Prop. 1.4. 1. In general equilibrium, $\hat{A}_r$ is determined by the intersection of eqs. (1.17) and (1.25). Recall that for $\lambda \geq z_1$ where z_1 is the largest root of h, h is strictly increasing and grows without bound. Further, g is a positive, strictly decreasing function approaching 0. Hence there exists a unique $\lambda_1 \geq z_1$ such that

$$\hat{A}_r = g(\lambda_1) = h(\lambda_1)$$

2. For uniqueness, since g is defined for positive values only I show that there is no other $\lambda' \in [0, z_1]$ such that $g(\lambda') = h(\lambda')$. From Lemma 1.6, $h(0) > h(\lambda')$ in this range. Additionally, Lemma 1.3 says that $g(\lambda)$ is decreasing, so that $g(\lambda_1) > g(\lambda')$. If the model is determinate, then the condition from Lemma 1.6 gives $g(\lambda_1) > h(0)$. So $g(\lambda') > h(\lambda')$ for all $\lambda' \in [0, z_1]$. Hence $\hat{A}_r$ is unique. Figure 1.5 plots examples of intersections of g and h.

□

Proof of Corollary 1.4.1. The proofs are similar to Cor. 1.1.1.

□

Proof of Corollary 1.4.2. Since $\delta > 0$, the right hand side of the inequality from eq. (1.24) is strictly greater than zero. Then since $\hat{A}_r \to 0$ as $a \to \infty$, there is some $\overline{a}$ for which $a > \overline{a}$ implies

$$\hat{A}_r < \frac{\delta}{\delta\phi_\pi + \rho\phi_x}$$

□

Proof of Prop. 1.5. When the determinacy condition eq. (1.24) is satisfied, the proof is the same as in Prop. 1.2. The result follows since

$$\frac{\partial \hat{A}_r^{\diamond}}{\partial a} < 0, \quad \left|\frac{\partial \omega_\pi}{\partial a}\right| < 0, \quad \left|\frac{\partial \omega_x}{\partial a}\right| < 0$$

□

Proof of Lemma 1.7. The proof is the same as Lemma 1.4, except that the rational expectations equilibrium matrices are now more complicated. The rational expectations dynamics matrices are given by

$$\Gamma = \begin{bmatrix} \lambda_1 & \gamma_{12} \\ 0 & \kappa_\beta \end{bmatrix} \tag{A5}$$

$$\Omega = \begin{bmatrix} \frac{\delta(-\lambda_1+\kappa_r)}{\kappa_r(\phi_\pi\delta+\phi_x(\lambda_1+\rho)} & \omega_{2,1} \\ \frac{(\lambda_1+\rho)(-\lambda_1+\kappa_r)}{\kappa_r(\phi_\pi\delta+\phi_x(\lambda_1+\rho)} & \omega_{2,2} \end{bmatrix} \tag{A6}$$

where γ_{12}, $\omega_{2,1}$ and $\omega_{2,2}$ are rational functions of the eigenvalue λ_1.

□

Proof of Prop. 1.6. Given the determinacy condition is satisfied, the the proof is similar to Prop. 1.3.

□

Proof of Prop. 1.7. Given the affine functional form assumption, I use Ito's Lemma to compute instantaneous returns. Write $P_{t,\tau}$ explicitly as a function of time t and variables $\mathbf{y}_t$:

$$P(t, \mathbf{y}) = \exp\left\{-\mathbf{y}^T\mathbf{A}(\tau(t)) - C(\tau(t))\right\}$$

Note that the dependence on the first argument t comes through τ in the coefficient functions $\mathbf{A}(\tau)$ and $C(\tau)$. Of course $\frac{\mathrm{d}\tau}{\mathrm{d}t} = -1$, which implies

$$\frac{\partial P}{\partial t} = P_{t,\tau}\left(\mathbf{y}_t^T\mathbf{A}'(\tau) + C'(\tau)\right)$$

The gradient and Hessian with respect to $\mathbf{y}$ are

$$\begin{aligned} \boldsymbol{\nabla}_{\mathbf{y}} P &= -P_{t,\tau}\mathbf{A}(\tau) \\ \mathrm{H}_{\mathbf{y}} P &= P_{t,\tau}\left(\mathbf{A}(\tau)\mathbf{A}(\tau)^T\right) \end{aligned}$$

Therefore, Ito's Lemma implies the instantaneous return is

$$\frac{\mathrm{d}P_{t,\tau}}{P_{t,\tau}} = \mu_{t,\tau}\,\mathrm{d}t - \mathbf{A}(\tau)^T\mathbf{S}\,\mathrm{d}\mathbf{B}_t \tag{A7}$$

$$\mu_{t,\tau} = \mathbf{y}_t^T\mathbf{A}'(\tau) + C'(\tau) + \left[\Gamma(\mathbf{y}_t - \overline{\mathbf{y}})\right]^T\mathbf{A}(\tau) + \frac{1}{2}Tr\left[\boldsymbol{\Sigma}\mathbf{A}(\tau)\mathbf{A}(\tau)^T\right] \tag{A8}$$

The arbitrageur's optimality conditions are given by:

$$\frac{\partial E_t \, \mathrm{d}W_t}{\partial b_{t,\tau}} = \frac{a}{2} \frac{\partial Var_t \, \mathrm{d}W_t}{\partial b_{t,\tau}}$$

Use eq. (A7) to compute the expectation and variance of the change in arbitrageur wealth:

$$\begin{aligned} E_t \, \mathrm{d}W_t &= \left[\left(W_t - \int_0^T b_{t,\tau} \, \mathrm{d}\tau \right) r_t + \int_0^T b_{t,\tau}\mu_{t,\tau} \, \mathrm{d}\tau \right] \mathrm{d}t \\ \implies \frac{\partial E_t \, \mathrm{d}W_t}{\partial b_{t,\tau}} &= (\mu_{t,\tau} - r_t) \, \mathrm{d}t \end{aligned}$$

and the variance is

$$\begin{aligned} Var_t \, \mathrm{d}W_t &= \left(\int_0^T b_{t,\tau} \mathbf{A}(\tau)^T \, \mathrm{d}\tau \right) \mathbf{\Sigma} \left(\int_0^T b_{t,\tau} \mathbf{A}(\tau) \, \mathrm{d}\tau \right) \mathrm{d}t \\ \implies \frac{\partial Var_t \, \mathrm{d}W_t}{\partial b_{t,\tau}} &= 2 \left(\int_0^T b_{t,\tau} \mathbf{A}(\tau)^T \, \mathrm{d}\tau \right) \mathbf{\Sigma} \mathbf{A}(\tau) \, \mathrm{d}t \end{aligned}$$

Note that

$$b_{t,\tau} = -\tilde{b}_{t,\tau} = \alpha(\tau) \left[\mathbf{y}_t^T \left(\Theta(\tau)\tau - \mathbf{A}(\tau) \right) + \bar{\beta}(\tau)\tau - C(\tau) \right] \tag{A9}$$

Substitute eq. (A9) into the optimality conditions derived above. Equating the coefficients on $\mathbf{y}_t$ terms (and assuming that r_t is ordered first) gives:

$$\mathbf{A}'(\tau) + \mathbf{M}\mathbf{A}(\tau) - \mathbf{e}_1 = 0$$

where $\mathbf{M}$ is defined by eq. (1.30). Note that since $\mathbf{M}$ is itself a function of integral terms involving $\mathbf{A}(\tau)$ this is a fixed point problem with no simple solution. However, treating $\mathbf{M}$ as fixed, with initial conditions $\mathbf{A}(\tau) = 0$ the general solution is given by eq. (1.31).

Suppose $\mathbf{M}$ is diagonalizable, with $\mathbf{M} = \mathbf{G}\mathbf{D}\mathbf{G}^{-1}$. Normalize $\mathbf{G}$ such that $\mathbf{G}\mathbf{1} = \mathbf{e}_1$ (that is, its first row sum is 1, and all other rows sum to 0). Then

$$\begin{aligned} \int_0^\tau \exp(-\mathbf{M}s) \, \mathrm{d}s \;\; \mathbf{e_1} &= \mathbf{G} \int_0^\tau \exp(-\mathbf{D}s) \, \mathrm{d}s \;\; \mathbf{G}^{-1}\mathbf{e_1} \\ &= \mathbf{G} \int_0^\tau \exp(-\mathbf{D}s) \, \mathrm{d}s \;\; \mathbf{1} \end{aligned}$$

and if $\mathbf{M}$ is invertible then eq. (1.31) is obtained.

Again substitute eq. (A9) into the optimality conditions. Equating constant coefficients gives:

$$\begin{aligned} C'(\tau) - (\Gamma\overline{\mathbf{y}})^T \mathbf{A}(\tau) + \frac{1}{2} Tr \left[\mathbf{\Sigma} \mathbf{A}(\tau) \mathbf{A}(\tau)^T \right] & \\ = a\mathbf{A}(\tau)^T \mathbf{\Sigma} \int_0^T \alpha(\tau) \left(\tau\bar{\beta}(\tau) - C(\tau) \right) \mathbf{A}(\tau) \, \mathrm{d}\tau & \end{aligned}$$

Define the vector

$$\mathbf{Z}_C \equiv a\mathbf{\Sigma}\int_0^T \alpha(\tau)\mathbf{A}(\tau)\left[\tau\bar{\beta}(\tau) - C(\tau)\right] \mathrm{d}\tau + \Gamma\overline{\mathbf{y}} \tag{A10}$$

Imposing the initial condition $C(0) = 0$ and integrating, the solution for $C(\tau)$ given by

$$C(\tau) = \mathbf{n}_1(\tau)^T\mathbf{Z}_C - \frac{1}{2}Tr\left[\mathbf{\Sigma}\mathbf{n}_2(\tau)\right] \tag{A11}$$

$$\mathbf{n}_1(\tau) = \int_0^\tau \mathbf{A}(u)\,\mathrm{d}u \tag{A12}$$

$$\mathbf{n}_2(\tau) = \int_0^\tau \mathbf{A}(u)\mathbf{A}(u)^T\,\mathrm{d}u \tag{A13}$$

Substitute $C(\tau)$ from (A11) into (A10) and solve for $\mathbf{Z}_C$:

$$\mathbf{Z}_C = \left[\mathbf{I} + a\mathbf{\Sigma}\mathbf{N}_1\right]^{-1}\left[a\mathbf{\Sigma}\mathbf{N}_2 + \Gamma\overline{\mathbf{y}}\right] \tag{A14}$$

where $\mathbf{I}$ is the identity matrix and

$$\mathbf{N}_1 = \int_0^T \alpha(\tau)\mathbf{A}(\tau)\mathbf{n}_1(\tau)^T\,\mathrm{d}\tau \tag{A15}$$

$$\mathbf{N}_2 = \int_0^T \alpha(\tau)\mathbf{A}(\tau)\left(\bar{\beta}(\tau)\tau + \frac{1}{2}Tr\left[\mathbf{\Sigma}\mathbf{n}_2(\tau)\right]\right)\mathrm{d}\tau \tag{A16}$$

□

A.2 Numerical Solution Algorithm

This section describes the numerical solution method used to solve the generalized model. This also requires obtaining closed-form solutions to a number of integral expressions.

A.2.1 Closed-Form Integral Expressions

Define the (scalar) functions:

$$\begin{aligned}
\phi_0(\nu,\tau) &\equiv \int_0^\tau u f(\nu u)\,\mathrm{d}u \\
&= \frac{\tau}{\nu}\left(1 - f(\nu\tau)\right) \\
\phi_1(\nu,\tau) &\equiv \int_0^\tau u^2 f(\nu u)\,\mathrm{d}u \\
&= \frac{\tau}{\nu^2}\left(\exp(-\nu\tau) - f(\nu\tau) + \frac{1}{2}\nu\tau\right) \\
\phi_2(\nu_i,\nu_j,\tau) &\equiv \int_0^\tau u^2 f(\nu_i u) f(\nu_j u)\,\mathrm{d}u \\
&= \frac{\tau}{\nu_i\nu_j}\left(1 - f(\nu_i\tau) - f(\nu_j\tau) + f((\nu_i+\nu_j)\tau)\right)
\end{aligned}$$

For any function $F(\tau)$, define

$$\tilde{F} \equiv \int_0^T \eta(\tau)F(\tau)\,\mathrm{d}\tau\,,\quad \hat{F} \equiv \int_0^T \frac{\eta(\tau)}{\tau}F(\tau)\,\mathrm{d}\tau$$

In particular, define

$$\tilde{f}(\nu) \equiv \int_0^T \eta(\tau)f(\nu\tau)\,\mathrm{d}\tau$$

Recall $\int_0^T \eta(\tau)\,\mathrm{d}\tau \equiv 1$. Therefore,

$$\begin{aligned}
\hat{\phi}_0(\nu,\tau) &= \frac{1}{\nu}(1 - \tilde{f}(\nu)) \\
\hat{\phi}_2(\nu_i,\nu_j,\tau) &= \frac{1}{\nu_i\nu_j}(1 - \tilde{f}(\nu_i) - \tilde{f}(\nu_j) + \tilde{f}(\nu_i+\nu_j))
\end{aligned}$$

For the integral terms involving $\alpha(\tau)$, define

$$\begin{aligned}
\phi_0^\alpha(\nu,\tau) &\equiv \int_0^\tau \alpha(u) u f(\nu u)\,\mathrm{d}u \\
\phi_1^\alpha(\nu,\tau) &\equiv \int_0^\tau \alpha(u) u^2 f(\nu u)\,\mathrm{d}u \\
\phi_2^\alpha(\nu_i,\nu_j,\tau) &\equiv \int_0^\tau \alpha(u) u^2 f(\nu_i u) f(\nu_j u)\,\mathrm{d}u
\end{aligned}$$

For the integral terms involving $\theta_k(\tau)$, define

$$\phi_1^{\alpha,\theta_k}(\nu,\tau) \equiv \int_0^\tau \alpha(u) u^2 \theta_k(u) f(\nu u)\,\mathrm{d}u$$

Since $\eta(\tau)$ is proportional to the pdf of a truncated Gamma distribution, this term can be written as

$$\tilde{f}(\nu) = \left(\int_0^T \tau \exp(-\eta_1 \tau) \, \mathrm{d}\tau \right)^{-1} \frac{(\nu \exp(\eta_1 T) - \eta_1 - \nu + \exp(-\nu T)\eta_1) \exp(-\eta_1 T)}{(\nu \eta_1 (\eta_1 + \nu))}$$

Further, I consider a more general form of $\alpha(\tau)$, given as follows:

$$\alpha(\tau) = \alpha_0 e^{-\alpha_1 \tau}$$

In the estimation, I simplify by setting $\alpha_1 = 0$ and $\alpha_0 = 1$.

Note that this implies

$$\alpha(u) f(\nu u) = \frac{\alpha_0}{\nu} \left[(\alpha_1 + \nu) f((\alpha_1 + \nu) u) - \alpha_1 f(\alpha_1 u) \right]$$

Therefore the integral expressions become

$$\begin{aligned} \phi_0^{\alpha}(\nu, \tau) &= \frac{\alpha_0}{\nu} \left[(\alpha_1 + \nu) \phi_0(\alpha_1 + \nu, \tau) - \alpha_1 \phi_0(\alpha_1, \tau) \right] \\ \phi_1^{\alpha}(\nu, \tau) &= \frac{\alpha_0}{\nu} \left[(\alpha_1 + \nu) \phi_1(\alpha_1 + \nu, \tau) - \alpha_1 \phi_1(\alpha_1, \tau) \right] \\ \phi_2^{\alpha}(\nu_i, \nu_j, \tau) &= \frac{\alpha_0}{\nu_i} \left[(\alpha_1 + \nu_i) \phi_2(\alpha_1 + \nu_i, \nu_j, \tau) - \alpha_1 \phi_2(\alpha_1, \nu_j, \tau) \right] \end{aligned}$$

For $\theta_k(\tau)$, I assume

$$\theta_k(\tau) = \delta(T_k - \tau)$$

where $T_k \in [0, T]$, and where δ is the Dirac delta function. Therefore, when $\tau > T_k$,

$$\phi_1^{\alpha, \theta_k}(\nu, \tau) = \alpha(T_k) T_k^2 f(\nu T_k)$$

Also, note that $\tau^2 f(\nu \tau) \to 0$ as $\tau \to 0$. Hence if $T_k = 0$, $\phi_1^{\alpha, \theta_k}(\nu, \tau) = 0$.

Next, extend these scalar functions to their multi-dimensional components.

$$\mathbf{F}(\mathbf{x}) = \mathbf{x}^{-1} \left(\mathbf{I} - e^{-\mathbf{x}} \right) \mathbf{1}$$

This implies the coefficient function can be written

$$\mathbf{A}(\tau) = \mathbf{G} \tau \mathbf{F}(\mathbf{D} \tau)$$

where

$$\left[\mathbf{F}(\mathbf{D}\tau) \right]_i = f(\nu_i \tau)$$

since $\mathbf{D} = diag\left[\nu_1, \ldots, \nu_J \right]$.

$$\boldsymbol{\Phi}_0(T) = \int_0^T \tau \mathbf{F}(\mathbf{D}\tau)\,\mathrm{d}\tau$$
$$\boldsymbol{\Phi}_1(T) = \int_0^T \tau^2 \mathbf{F}(\mathbf{D}\tau)\,\mathrm{d}\tau$$
$$\boldsymbol{\Phi}_2(T) = \int_0^T \tau^2 \mathbf{F}(\mathbf{D}\tau)\mathbf{F}(\mathbf{D}\tau)^T\,\mathrm{d}\tau$$

For integral terms involving the function $\alpha(\tau)$, define

$$\boldsymbol{\Phi}_1^{\alpha}(T) = \int_0^T \alpha(\tau)\tau^2 \mathbf{F}(\mathbf{D}\tau)\,\mathrm{d}\tau$$
$$\boldsymbol{\Phi}_2^{\alpha}(T) = \int_0^T \alpha(\tau)\tau^2 \mathbf{F}(\mathbf{D}\tau)\mathbf{F}(\mathbf{D}\tau)^T\,\mathrm{d}\tau$$

For integral terms involving the function $\boldsymbol{\Theta}(\tau)$, define

$$\boldsymbol{\Phi}_1^{\alpha,\theta}(T) = \int_0^T \alpha(\tau)\tau^2 \boldsymbol{\Theta}(\tau)\mathbf{F}(\mathbf{D}\tau)^T\,\mathrm{d}\tau$$

For integral terms involving the function $\eta(\tau)$, define

$$\tilde{\mathbf{F}} = \int_0^T \eta(\tau)\mathbf{F}(\mathbf{D}\tau)\,\mathrm{d}\tau$$
$$\hat{\boldsymbol{\Phi}}_0 = \int_0^T \frac{\eta(\tau)}{\tau}\boldsymbol{\Phi}_0(\tau)\,\mathrm{d}\tau$$
$$\hat{\boldsymbol{\Phi}}_2 = \int_0^T \frac{\eta(\tau)}{\tau}\boldsymbol{\Phi}_2(\tau)\,\mathrm{d}\tau$$

Therefore, all the vector functions can be written as

$$\mathbf{n}_1(\tau) = \mathbf{G}\boldsymbol{\Phi}_0(\tau)$$
$$\mathbf{n}_2(\tau) = \mathbf{G}\boldsymbol{\Phi}_2(\tau)\mathbf{G}^T$$
$$\mathbf{M} = \Gamma^T - a\left[\boldsymbol{\Phi}_1^{\alpha,\theta} - \mathbf{G}\boldsymbol{\Phi}_2^{\alpha}\right]\mathbf{G}^T\boldsymbol{\Sigma}$$
$$\mathbf{N}_1 = \mathbf{G}\left[\boldsymbol{\Phi}_1^{\alpha}\mathbf{1} - \boldsymbol{\Phi}_2^{\alpha}\right]\mathbf{D}^{-1}\mathbf{G}^T$$
$$\mathbf{N}_2 = \mathbf{G}\left[\boldsymbol{\Phi}_1^{\alpha,\beta} + \frac{1}{2}\boldsymbol{\Phi}_3^{\alpha} vec\left[\mathbf{G}^T\boldsymbol{\Sigma}\mathbf{G}\right]\right]$$
$$\hat{\mathbf{A}} = \mathbf{G}\tilde{\mathbf{F}}$$
$$\hat{\mathbf{n}}_1 = \mathbf{G}\hat{\boldsymbol{\Phi}}_0$$
$$\hat{\mathbf{n}}_2 = \mathbf{G}\hat{\boldsymbol{\Phi}}_2\mathbf{G}^T$$

I compute these integral vectors and matrices element by element, using the scalar functions above.

$$[\mathbf{\Phi}_0]_i = \phi_0(\nu_i, T)$$
$$[\mathbf{\Phi}_1]_i = \phi_1(\nu_i, T)$$
$$[\mathbf{\Phi}_2]_{i,j} = \phi_2(\nu_i, \nu_j, T)$$

For $\alpha(\tau)$ integrals:

$$[\mathbf{\Phi}_1^\alpha]_i = \phi_1^\alpha(\nu_i, T)$$
$$[\mathbf{\Phi}_2^\alpha]_{i,j} = \phi_2^\alpha(\nu_i, \nu_j, T)$$

For $\eta(\tau)$ integrals:

$$\left[\tilde{\mathbf{F}}\right]_i = \tilde{f}(\nu_i)$$
$$\left[\hat{\mathbf{\Phi}}_0\right]_i = \hat{\phi}_0(\nu_i, T)$$

For $\theta_k(\tau)$ integrals:

$$\left[\mathbf{\Phi}_1^{\alpha,\theta}\right]_{k,i} = \phi_1^{\alpha,\theta_k}(\nu_i, \tau)$$

A.2.2 Equilibrium Algorithm

This section describes a solution method for solving the fixed point problem described in Prop. 1.7. Define the function

$$F(\mathbf{M}) = \Gamma(\hat{\mathbf{A}})^T - a\left[\int_0^T \alpha(\tau)\left(\tau\Theta(\tau) - \mathbf{A}(\tau)\right)\mathbf{A}(\tau)^T\,\mathrm{d}\tau\right]\mathbf{\Sigma} - \mathbf{M}$$

Equilibrium is defined as a root of this function: $F(\mathbf{M}) = 0$.

The following algorithm is used to compute F. Given some initial value of the matrix $\mathbf{M}$:

1. Solve the eigen-decomposition to get $\mathbf{D}$, $\mathbf{G}$, which gives the implied $\mathbf{A}(\tau)$ and $\hat{\mathbf{A}}$.

2. Construct the implied dynamics matrix $\Upsilon(\hat{\mathbf{A}})$ and solve for the rational expectations equilibrium matrix $\Gamma(\hat{\mathbf{A}})$ (assuming stability conditions are met).

3. Using the integral expressions derived above, solve for $F(\mathbf{M})$.

Using the above algorithm, standard root-finding algorithms can be used to solve for the equilibrium (note that while $\mathbf{D}$ and $\mathbf{G}$ may contain complex values, $\mathbf{M}$ and F will always be real-valued).

For the dynamics matrix functions $\Upsilon(\hat{\mathbf{A}})$ considered in Section 1.5, when $a = 0$ the equilibrium is unique. However, for generic dynamics matrix functions $\Upsilon(\hat{\mathbf{A}})$ and when $a > 0$, multiple equilibria are possible. In order to rule out pathological equilbiria, I focus on the equilibrium which approaches the $a = 0$ equilibrium as $a \to 0$.

A.3 Microfoundations

This section describes the model from first principles. Time is continuous. The model consists of households, firms, arbitrageurs, and a government. The government conducts monetary policy (changes in policy rate), passive fiscal policy (changes in taxes, which play no role), and QE (changes in holdings of long-term bonds).

Households face radically incomplete markets: they can only borrow through a passive mutual fund offering an instantaneous return which is a weighted average of all yields. Infinitesimally-lived arbitrageurs are the marginal investors in financial markets.

The short rate is the main policy tool, and as in Woodford (2003) I consider a "cashless limit" economy.

A.3.1 Households and Firms

An infinitely-lived representative household seeks to maximize expected utility by choosing consumption, labor, and savings. The household problem is standard, except they are restricted to borrowing at an effective rate $\tilde{r}_t$. The household problem is

$$\max_{\{C_t, N_t\}} E_0 \int_0^\infty e^{-\rho t} \left(\frac{C_t^{1-\varsigma}}{1-\varsigma} - \frac{N_t^{1+\xi}}{1+\xi} \right) \mathrm{d}t \tag{C1}$$

$$\text{s.t. } \mathrm{d}W_t^H \leq \left(\tilde{r}_t W_t^H - P_t C_t + w_t N_t - T_t^H \right) \mathrm{d}t \tag{C2}$$

$$\lim_{T\to\infty} E_t[Q_{t,T}^H W_T^H] = 0 \tag{C3}$$

eq. (C2) is the household's flow budget constraint; eq. (C3) is a transversality condition. The parameter ς is the coefficient of relative risk aversion; ξ is the labor supply elasticity; ρ is the discount rate. W_t^H is household nominal wealth; C_t is consumption of composite good; N_t is labor; w_t is nominal wage; T_t^H is the net taxes and transfers to households.

$Q_{t,T}^H$ is the household's discount factor, given by

$$Q_{t,T}^H = \exp\left[-\int_t^T (\tilde{r}_s - \pi_s)\, \mathrm{d}s \right] \tag{C4}$$

Note that markets are not complete (and so eq. (C4) is not the relevant discount factor for pricing securities), but eqs. (C2) and (C3) can be replaced with an equivalent single intertemporal budget constraint

$$W_0^H = E_0 \int_0^\infty Q_{0,t}^H \left(P_t C_t - w_t N_t + T_t^H\right) \mathrm{d}t \tag{C5}$$

where W_0^H is given.

Then the intra-temporal optimality condition determining labor supply is

$$C_t^\varsigma N_t^\xi = \frac{w_t}{P_t} \tag{C6}$$

and linearized intertemporal optimality conditions give

$$\mathrm{d}c_t = \varsigma^{-1} \left(\tilde{r}_t - \pi_t - \rho\right) \mathrm{d}t \tag{C7}$$

Firms face Rotemberg (1982) pricing adjustment costs, and their problem is unchanged relative to benchmark New Keynesian models.

A.3.2 Arbitrageurs and Preferred Habitat Investors

Arbitrageurs face a mean-variance trade-off in their wealth:

$$\max_{b_{t,\tau}} E_t \,\mathrm{d}W_t^A - \frac{a}{2} Var_t \,\mathrm{d}W_t^A$$

subject to their flow budget constraint

$$\mathrm{d}W_t^A = \left(W_t^A - T_t^A - M_t^A - \int_0^T b_{t,\tau}\,\mathrm{d}\tau\right) r_t \,\mathrm{d}t + \int_0^T b_{t,\tau} \frac{\mathrm{d}P_{t,\tau}}{P_{t,\tau}}\,\mathrm{d}\tau + M_t^A r_t^M \,\mathrm{d}t \tag{C8}$$

This is the same as the one considered in the main text, with the exception that arbitrageurs can also hold monetary balances M_t^A (with a return r_t^m, which must satisfy $r_t \geq r_t^m$ with equality if money supply is non-zero), and they face taxes/transfers T_t^A.

The "preferred habitat" investors are aggregated into a single passive mutual fund, which offers households the effective borrowing rate, and invests the rest of its wealth in long-term bonds. The flow budget constraint is thus

$$\mathrm{d}W_t^F = -W_t^H \tilde{r}_t \,\mathrm{d}t + \int_0^T \tilde{b}_{t,\tau} \frac{\mathrm{d}P_{t,\tau}}{P_{t,\tau}}\,\mathrm{d}\tau + \left(W_t^F - W_t^H - T_t^G - \int_0^T \tilde{b}_{t,\tau}\,\mathrm{d}\tau\right) r_t \,\mathrm{d}t \tag{C9}$$

and also face taxes/transfers.

The demand for bonds $\tilde{b}_t$ is given by eq. (1.4). Thus the demand from the mutual fund for long-term bonds is reduced-form. Demand of this form could additionally be derived by adding another set of finite-lived, infinitely risk-averse investors as in the appendix of Vayanos and Vila (2009). When studying QE, I assume that the demand shocks operate through the preferred habitat investors. In this case, it may be useful to view this mutual fund as a type of government sponsored entereprise.

A.3.3 Rational Expectations Equilibrium

Assuming that all of the profits from the arbitrageurs and mutual funds are transfered lump-sum to the households, the log-linearized solution is the same as a benchmark model but with the effective rate $\tilde{r}_t$ in place of the short rate r_t in the IS curve.

After linearizing around the steady state and imposing affine functional forms to the term structure, the aggregate dynamics can be written in matrix form as in eq. (1.8). The solution to the rational expectations equilibrium is found by following Buiter (1984), the continuous time analogue of Blanchard and Kahn (1980). In general, let $\mathbf{Y}_t = [\mathbf{y}_t \;\; \mathbf{x}_t]^T$ where $\mathbf{x}_t$ are the "jump" variables and $\mathbf{y}_t$ are the state variables. Assuming that Υ is diagonalizable, partition the eigenvalues and eigenvectors as follows:

$$\Upsilon = Q\Lambda Q^{-1}$$
$$\Lambda = \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{bmatrix}, \quad Q = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix}$$

where the partitions correspond to the state and jump variables. If the number of "stable" eigenvalues (non-negative real parts) equals the number of state variables, then given some transversality conditions the rational expectations equilibrium dynamics are given by eq. (1.10), with

$$\Gamma = Q_{11}\Lambda_1 Q_{11}^{-1} \tag{C10}$$
$$\Omega = Q_{21}Q_{11}^{-1} \tag{C11}$$

A.3.4 Conditional and Unconditional Distributions

The conditional distribution of the state variables $\mathbf{y}_t$ given initial state $\mathbf{y}_0$ is normal:

$$\mathbf{y}_t|\mathbf{y}_0 \sim N(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$$

The conditional mean is

$$\boldsymbol{\mu}_t = \mathbf{y}^{SS} + e^{-\Gamma t}(\mathbf{y}_0 - \mathbf{y}^{SS})$$

The conditional variance-covariance is

$$\boldsymbol{\Sigma}_t = \int_0^t e^{\Gamma(u-t)}\boldsymbol{\Sigma}e^{\Gamma^T(u-t)}\,\mathrm{d}u$$

where again $\boldsymbol{\Sigma} = \mathbf{S}\mathbf{S}^T$. Then this simplifies to

$$\begin{aligned} \operatorname{vec}\boldsymbol{\Sigma}_t &= \left(\int_0^T e^{(\Gamma\oplus\Gamma)(u-t)}\,\mathrm{d}u\right)\operatorname{vec}\boldsymbol{\Sigma} \\ &= (\Gamma\oplus\Gamma)^{-1}\left(\mathbf{I} - e^{-(\Gamma\oplus\Gamma)t}\right)\operatorname{vec}\boldsymbol{\Sigma} \end{aligned}$$

where $\oplus$ is the Kronecker sum. Taking the limit as $t \to \infty$, the unconditional variance-covariance is given by

$$\operatorname{vec} \mathbf{\Sigma}_\infty = (\Gamma \oplus \Gamma)^{-1} \operatorname{vec} \mathbf{\Sigma}$$

and the present discounted value is computed as

$$\begin{aligned}
\widetilde{\mathbf{\Sigma}}_\infty &\equiv \int_0^\infty e^{-\rho t} \mathbf{\Sigma}_t \, \mathrm{d}t \\
\operatorname{vec} \widetilde{\mathbf{\Sigma}}_\infty &= (\Gamma \oplus \Gamma)^{-1} (\rho \mathbf{I} + \Gamma \oplus \Gamma)^{-1} \operatorname{vec} \mathbf{\Sigma}
\end{aligned}$$

Then from eq. (1.10), the conditional distribution of the jump variables $\mathbf{x}_t$ given initial state $\mathbf{y}_0$ (recall the initial values $\mathbf{x}_0$ are endogenous) is normal:

$$\mathbf{x}_t | \mathbf{y}_0 \sim N \left(\Omega(\boldsymbol{\mu}_t - \mathbf{y}^{SS}) + \mathbf{x}^{SS}, \mathbf{\Sigma}_t^{\mathbf{x}} \right)$$

and the equivalently defined covariances for the jump variables are easily computed as

$$\begin{aligned}
\mathbf{\Sigma}_t^{\mathbf{x}} &= \Omega \mathbf{\Sigma}_t \Omega^T \\
\mathbf{\Sigma}_\infty^{\mathbf{x}} &= \Omega \mathbf{\Sigma}_\infty \Omega^T \\
\widetilde{\mathbf{\Sigma}}_\infty^{\mathbf{x}} &= \Omega \widetilde{\mathbf{\Sigma}}_\infty \Omega^T
\end{aligned}$$

Hence if $\mathbf{x}^{SS} = 0$, $E_0 \left[\mathbf{x}_t \mathbf{x}_t^T \right] = \Omega \mathbf{\Sigma}_t \Omega^T$.

Appendix B

Unbundling Quantitative Easing

B.1 Numerical Exercise Details

In this section we briefly describe the model and calibration of our numerical exercise. For more details regarding the model setup, see Vayanos and Vila (2009).

B.1.1 Numerical Exercise Model

There is a continuum of zero-coupon bonds with maturities $m \in (0, T]$ in zero net supply. A bond with maturity m has a time t price of $P_{t,m}$ and pays \$1 at time $t+m$. The spot rate is $R_{t,m}$ which is given by

$$R_{t,m} = -\frac{\log P_{t,m}}{m}$$

There are two types of investors: idiosyncratic/clientèle investors and arbitrageurs. By assumption idiosyncratic demand takes the following form:

$$y_{t,m} = \alpha(m)m(R_{t,m} - \beta_{t,m})$$

where $\beta_{t,m}$ is a demand shifter which responds to K demand factors:

$$\beta_{t,m} = \bar{\beta} + \sum_{k=1}^{K} \theta_k(m)\beta_{k,m}$$

Arbitrageurs choose how much of each bond to hold (denoted by $x_{t,m}$). Their budget constraint is:

$$dW_t = \left(W_t - \int_0^T x_{t,m}dm\right) r_t dt + \int_0^T x_{t,m}\frac{dP_{t,m}}{P_{t,m}}dm$$

where r_t is the instantaneous rate: $\lim_{m\to 0} R_{t,m} = r_t$. Arbitrageurs maximize an instantaneous mean-variance trade-off:

$$\max_x E_t dW_t - \frac{a}{2} Var_t dW_t$$

where the parameter a governs the level of risk aversion. In equilibrium, we have $y_{t,m} = -x_{t,m}$.

We assume the instantaneous rate and demand factors are stacked in a $K+1$ vector $\mathbf{Y}$ which follows an Ornstein-Uhlenbeck process:

$$d\mathbf{Y}_t = -\Gamma(\mathbf{Y}_t - \overline{\mathbf{Y}})dt + \mathbf{S}d\mathbf{B}_t$$

where $\mathbf{B}_t$ is a vector of Brownian motions.

It turns out that the above is consistent with bond prices that are affine in the $\mathbf{Y}$ factors:

$$-\log P_{t,m} = \mathbf{Y}_t^T \mathbf{A}(m) + C(m)$$

We are interested in the response of the term structure to shocks to the demand factors, and hence need to solve the model for the coefficient functions $\mathbf{A}(m)$. Using the arbitrageur FOCs and taking into account the zero net supply condition, these functions must satisfy the system of differential equations

$$\mathbf{A}'(m) + \Gamma^T \mathbf{A}(m) - \mathbf{e}_1 = a\mathbf{M}\mathbf{A}(m)$$

where $\mathbf{e}_1$ is the first coordinate vector (assuming r_t is ordered first in $\mathbf{Y}$), and

$$\mathbf{M} = \left(\int_0^T \alpha(m) \left[m\Theta(m) - \mathbf{A}(m) \right] \mathbf{A}(m)^T dm \right) \mathbf{S}\mathbf{S}^T$$

Solving the above differential equation is made more difficult by the presence of the integral terms in $\mathbf{M}$. Vayanos and Vila (2009) solves the model for the limiting case when the risk aversion parameter $a \to 0$ or $a \to \infty$ (and the particular case when Γ and $\mathbf{S}$ are diagonal), but for intermediate values a solution must be found numerically. For details regarding the solution algorithm, see Ray (2017).

B.1.2 Numerical Exercise Calibration

For our numerical exercise, we take the number of demand factors to be $K = 2$. We will interpret the first demand factor as a "short" demand factor denoted by $\beta_{t,s}$. The second factor is taken to be a "long" demand factor denoted by $\beta_{t,\ell}$. We assume

$$\mathbf{\Gamma} = \begin{bmatrix} \kappa_r & 0 & 0 \\ 0 & \kappa_s & 0 \\ 0 & 0 & \kappa_\ell \end{bmatrix}$$

and set these mean reversion parameters to imply that shocks to the instantaneous rate have a half-life of approximately one year, while shocks to the demand factors have a half-life of 2.5 years.

For simplicity we also assume uncorrelated shocks, and that the size of the innovations for each factor is equal, i.e. $\mathbf{S} = \sigma \mathbf{I}$. We set $\sigma = .01$.

$\theta_s(m)$ and $\theta_\ell(m)$ govern where the demand shocks are located in maturity space. Although not realistic, we set these as Dirac delta functions, so that the short demand shock is entirely concentrated at idiosyncratic investors whose habitat is at $m = 3$ years; similarly for the long demand shock we choose $m = 20$ years. These maturities roughly correspond to the average maturity of the short-term and long-term auctions in the empirical counterpart. We could have instead assumed these functions have non-zero values for a continuum of bonds, but still concentrated at the long and short end of the maturity space. This complicates the numerical solution algorithm, but leads to largely similar results.

$\alpha(m)$ governs the sensitivity of idiosyncratic investors to changes in the price of bonds within their habitat. We don't have priors for this parameter, and for simplicity assume the function is constant. We set this value to match the following empirical counterpart: a standard deviation increase in our demand shock D_t is associated with an increase of 0.15 in the bid-to-cover ratio during short-term auctions. Given our parameterization above, a one-standard deviation positive short demand shock increases idiosyncratic demand by $3\alpha\sigma$. Equating these values implies $\alpha = 5$.

Finally, we let the risk aversion parameter vary from 0 to 500. The upper limit is ad-hoc; the value was chosen as the response of spot rates at this point begins to stabilize.

Appendix Table B4 summarizes the parameter calibration. The spirit of the numerical exercise is not to match the data perfectly, but rather to gain some qualitative predictions for intermediate levels of risk aversion. Finally, it is important to note that the parameters a, α, and σ in this specification enter multiplicatively. Hence an appropriate rescaling of these values will give numerically identical responses.

B.2 Additional Figures and Tables

Figure B1: Response of Inflation Swap Rates to Shock D_t

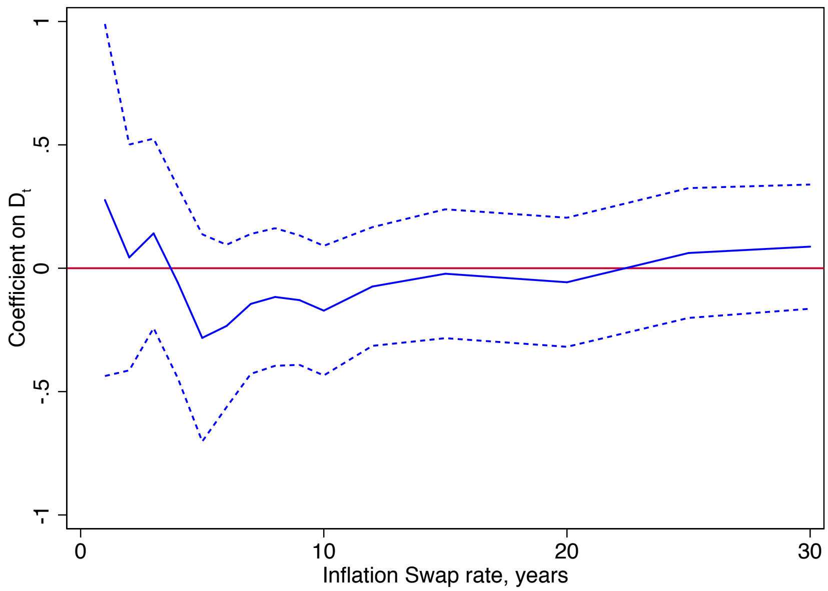

Notes: the figure plots responses of inflation swap rates across different maturities to a shock in the first principal component D_t. The solid line plots the point estimates, while dashed lines plot two-standard deviation (Newey-West) confidence bands.

Figure B2: Rolling Regressions

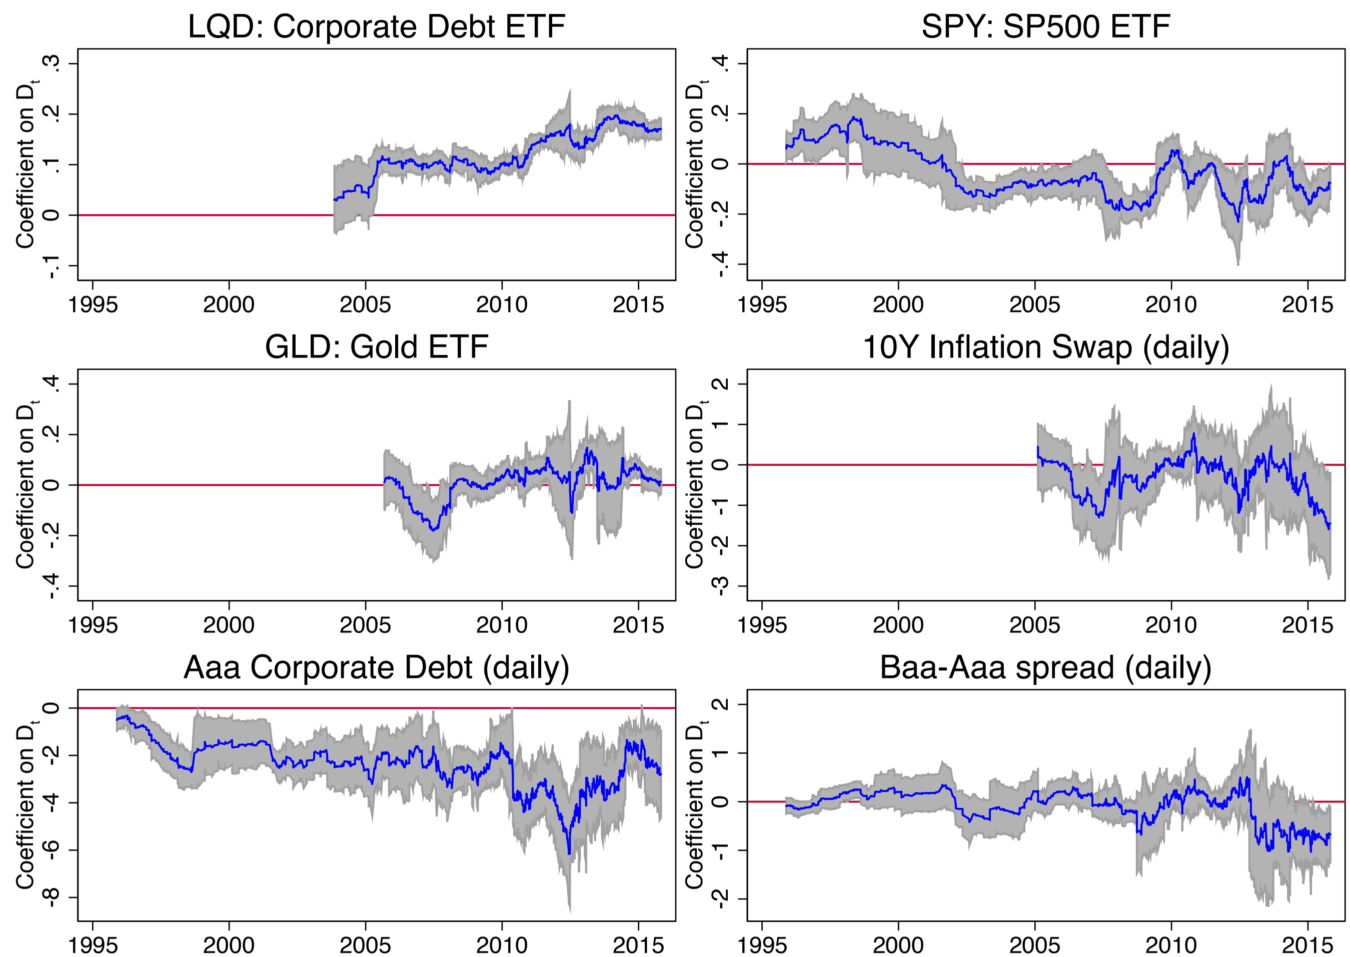

Notes: Coefficient estimates of rolling regressions; each data point is from estimating (2.3) for using the 60 most recent auctions. LQD, SPY, and GLD are intraday changes; 10-year inflation swap is daily change. Aaa Corporate Debt is the daily change in the interest rate for Aaa corporate bonds. Baa-Aaa Spread is the daily change in the spread between Baa and Aaa corporate bonds. The regressor is always intraday D_t, the first principal component in $D_t^{(2Y)}$, $D_t^{(5Y)}$, $D_t^{(10Y)}$, and $D_t^{(30Y)}$.

Figure B3: U.S. Financial Crises

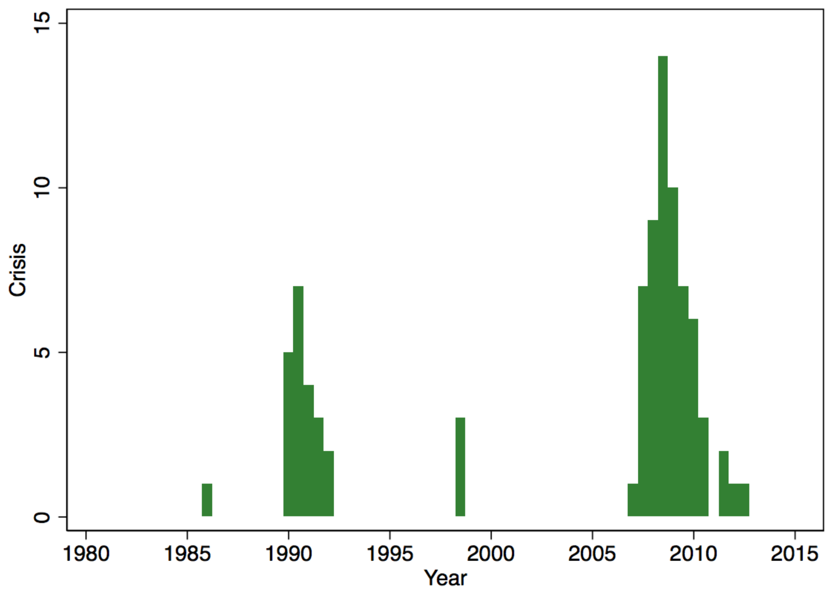

Notes: Financial Crisis indicator for the United States from Romer and Romer (2017).

Figure B4: Intermediary Capital Ratio

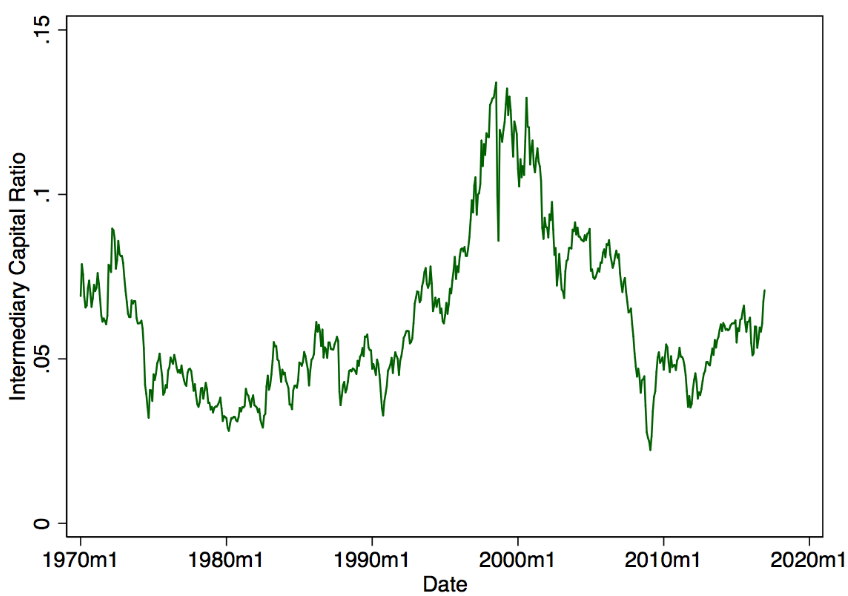

Notes: Intermediary capital ratio from He et al. (2016).

Figure B5: Rate Response P-Values

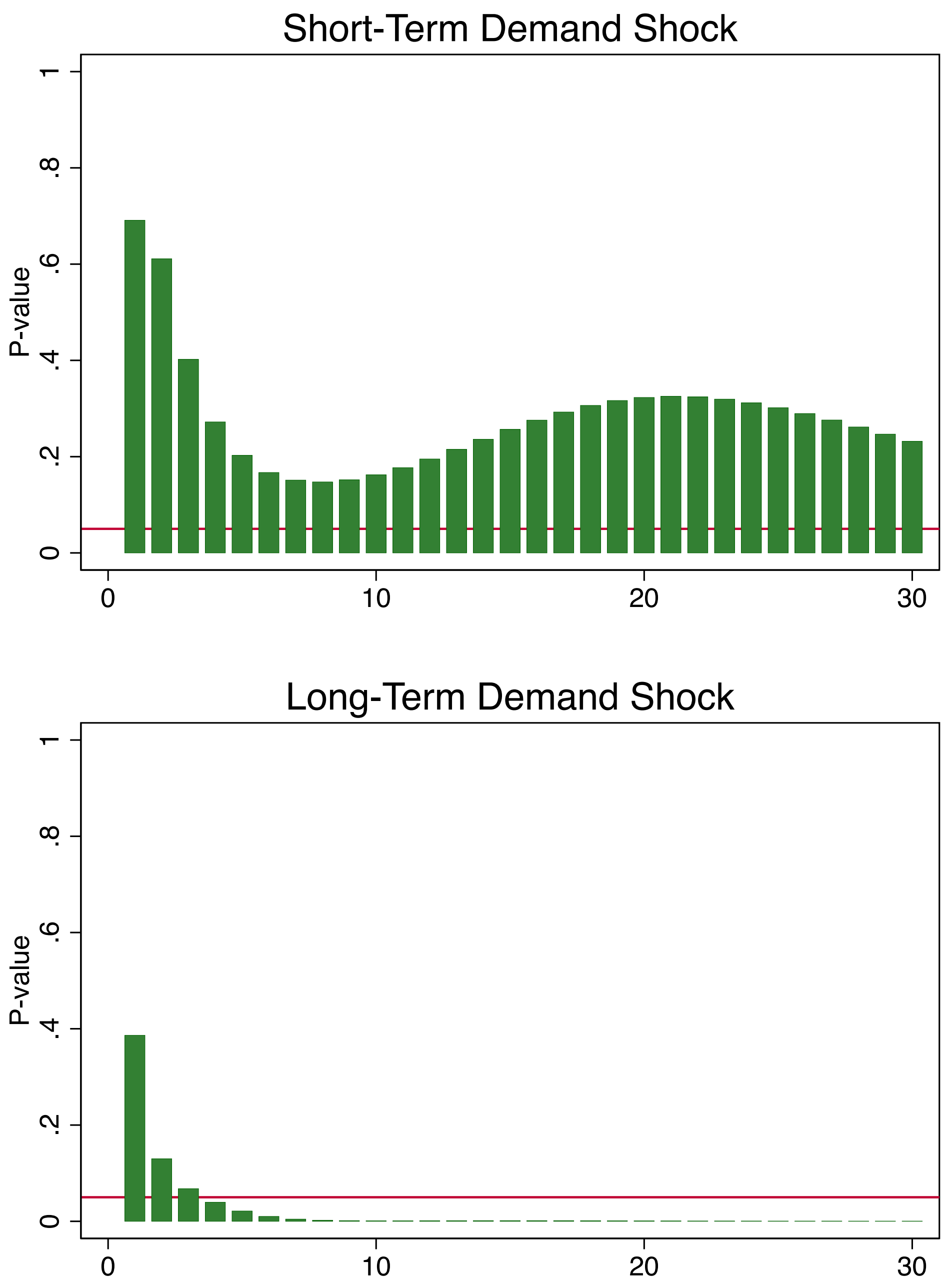

Notes: p-values testing equality of coefficients from Figure 2.12.

Figure B6: Rate Response P-Values

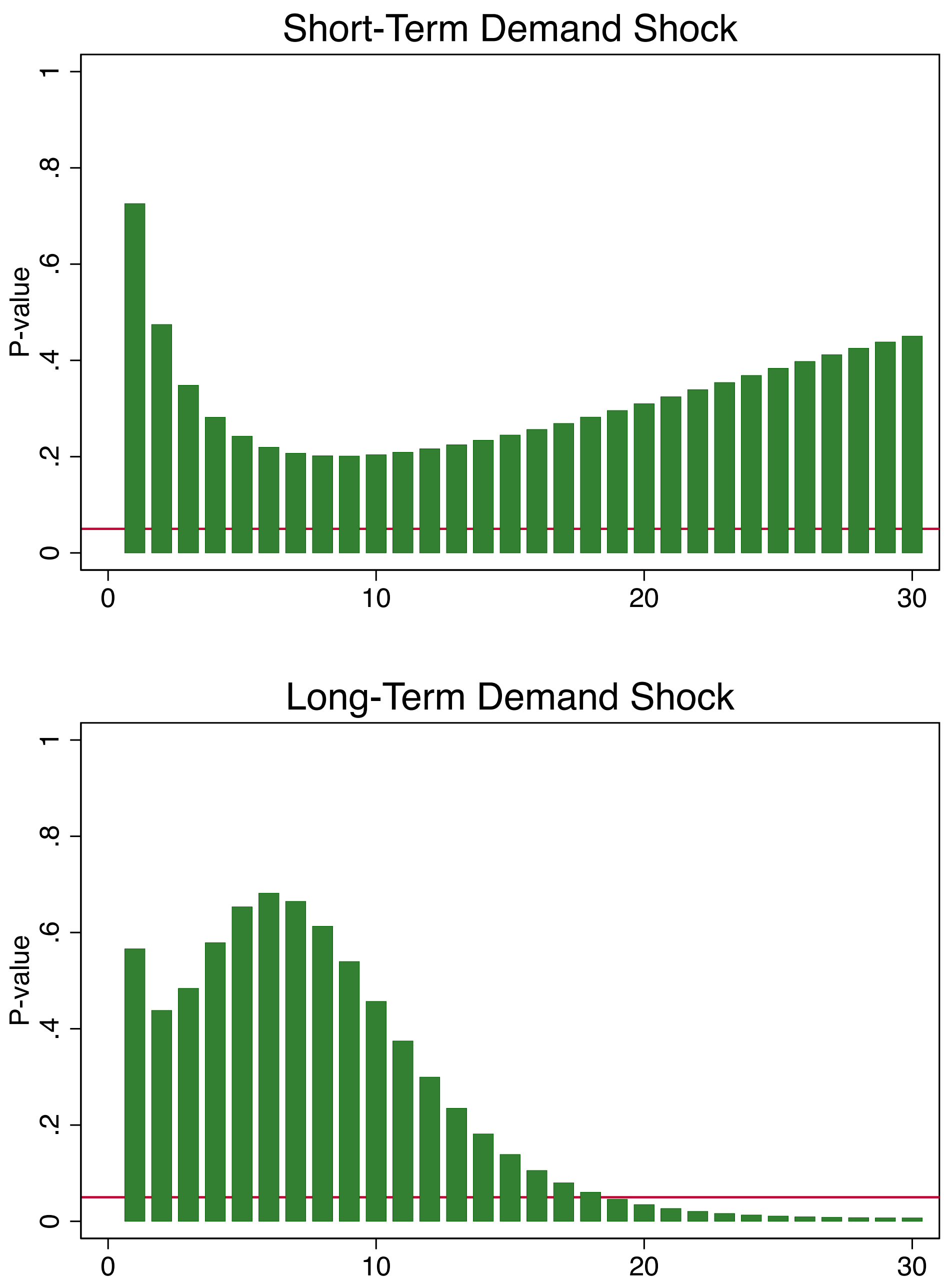

Notes: p-values testing equality of coefficients from Figure 2.13.

Figure B7: Rate Responses (Bid-to-Cover, 1995-2015)

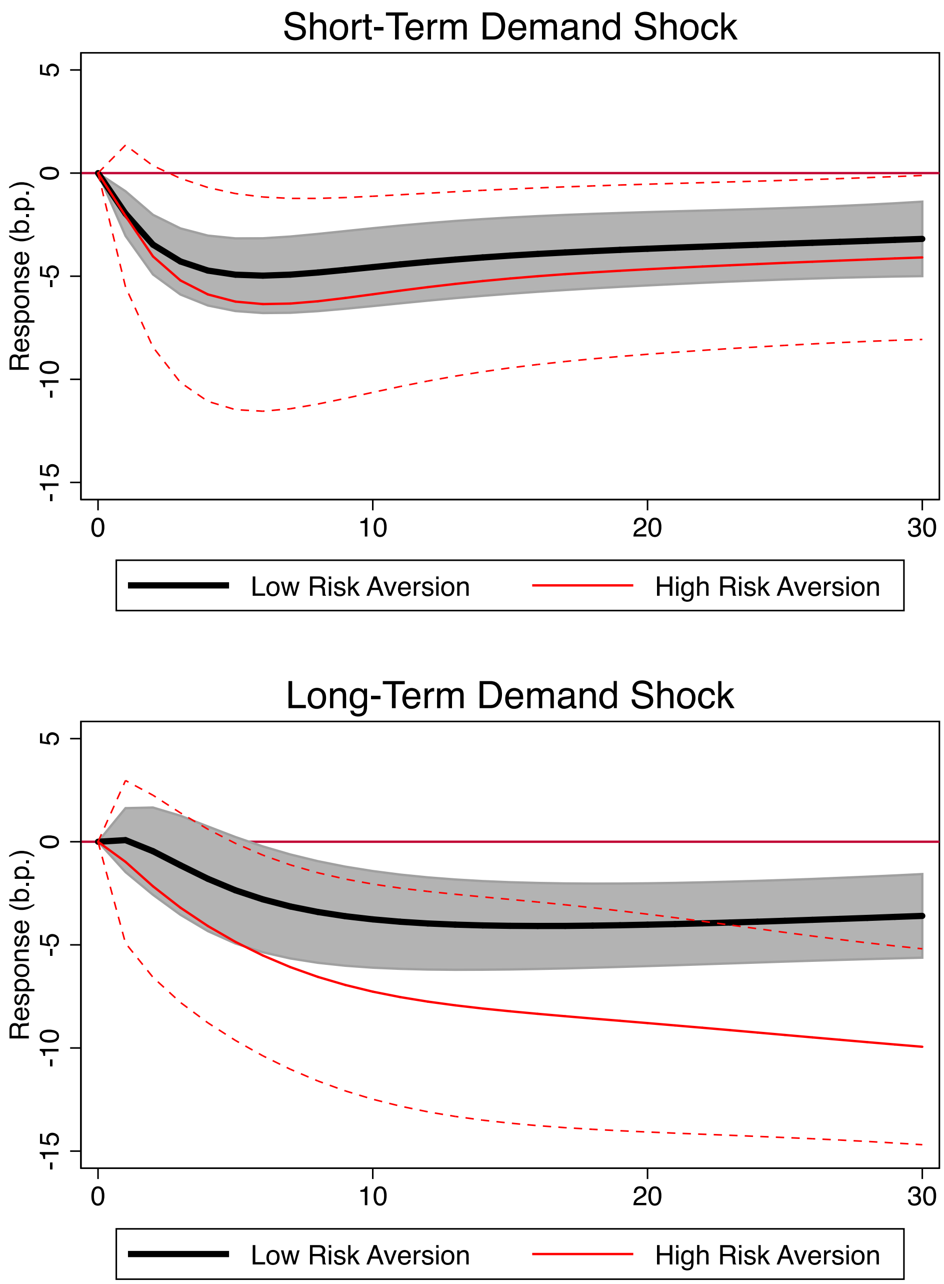

Notes: Plots of the regression coefficient on the surprise component of the bid-to-cover ratio from regression equation (2.4) for the sample 1995-2015. Each curve is from the subsample combinations: short-term and long-term auctions; and periods of high and low risk aversion as measured. 2 standard error (Newey-West) confidence intervals are included.

Figure B8: Rate Response P-Values (Bid-to-Cover, 1995-2015)

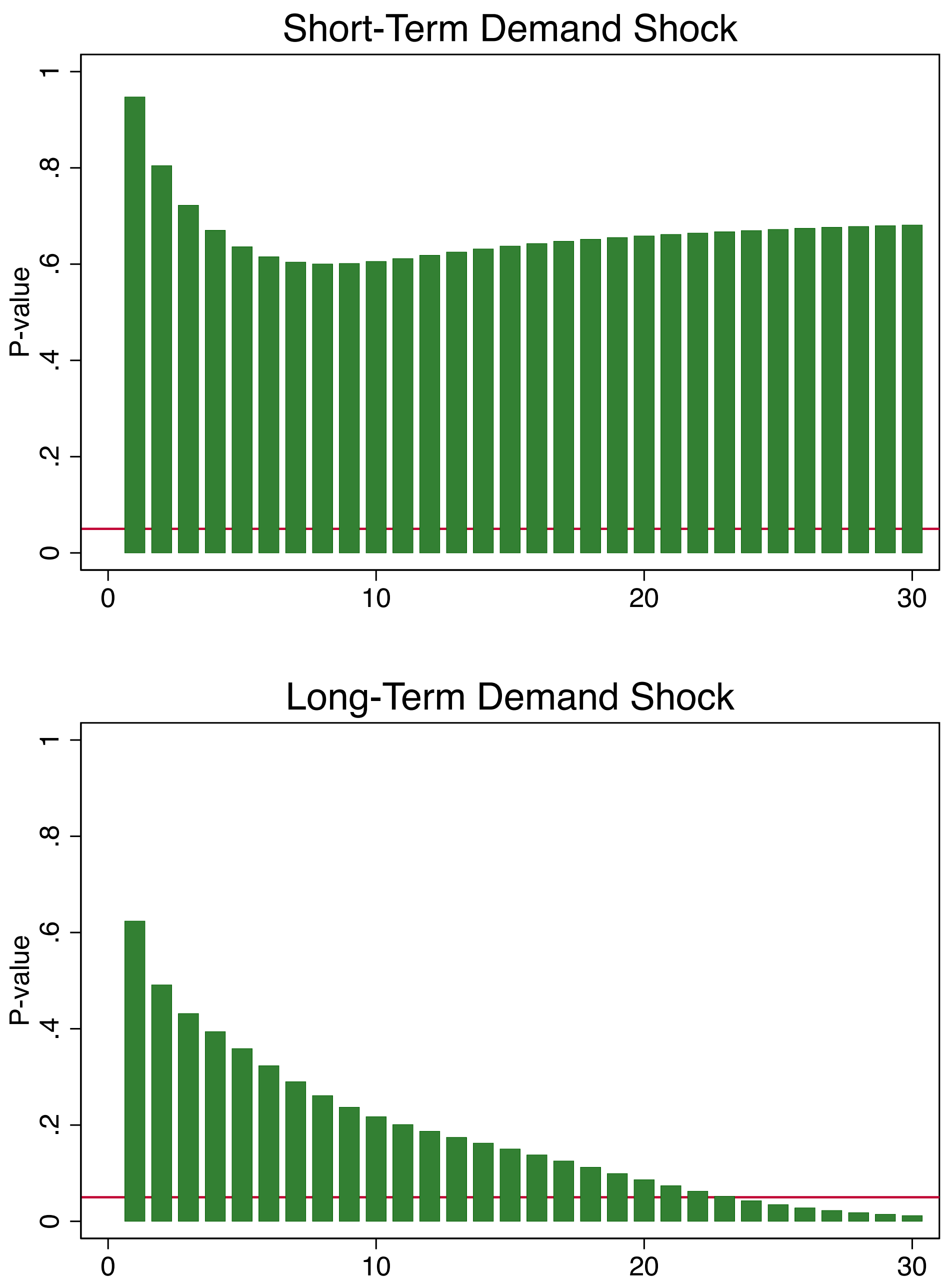

Notes: p-values testing equality of coefficients from Figure B7.

Figure B9: Rate Responses (Bid-to-Cover, 1979-2015)

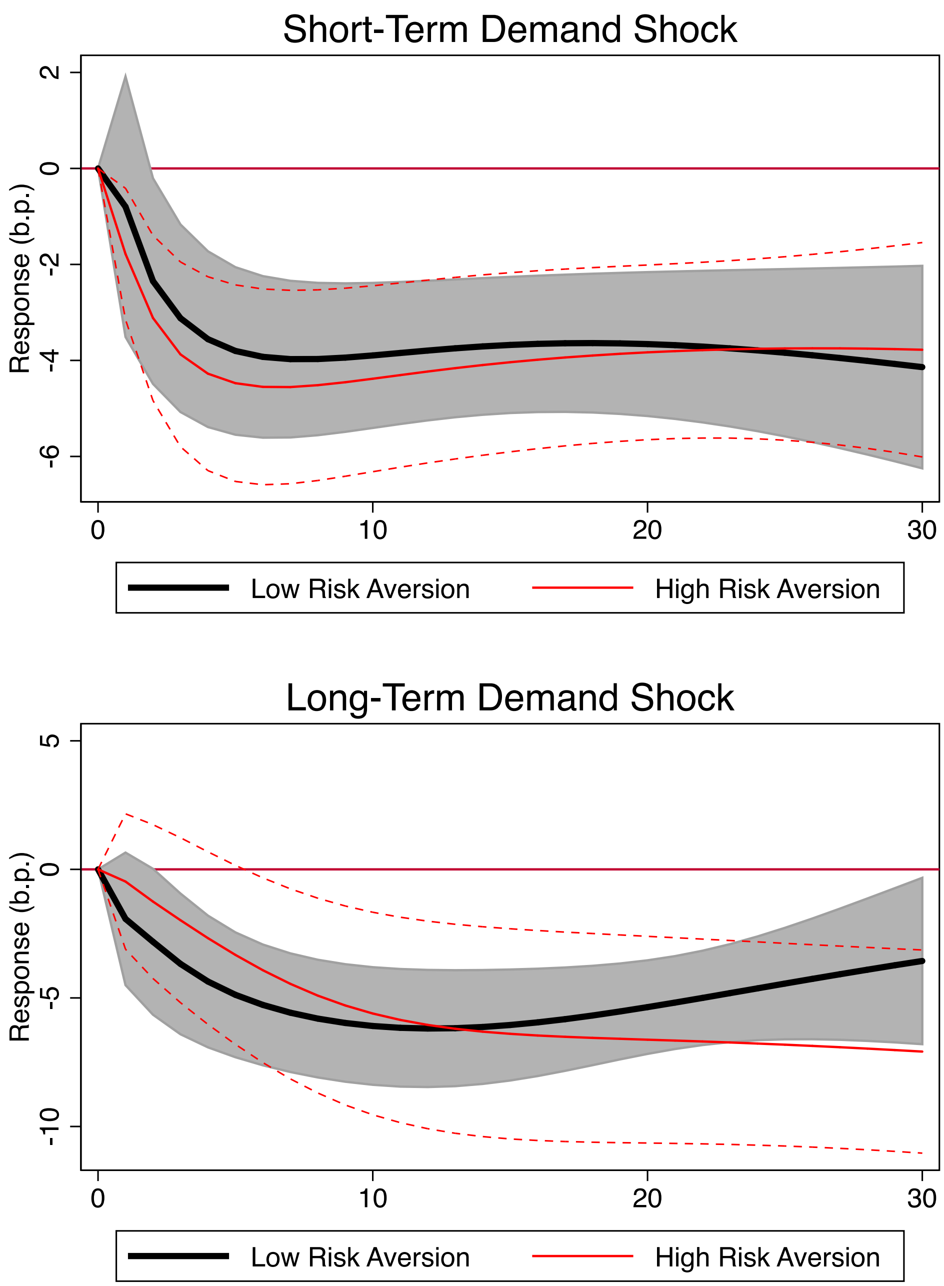

Notes: Plots of the regression coefficient on the surprise component of the bid-to-cover ratio from regression equation (2.4) for the sample 1979-2015. Each curve is from the subsample combinations: short-term and long-term auctions; and periods of high and low risk aversion as measured. 2 standard error (Newey-West) confidence intervals are included.

Figure B10: Rate Response P-Values (Bid-to-Cover, 1979-2015)

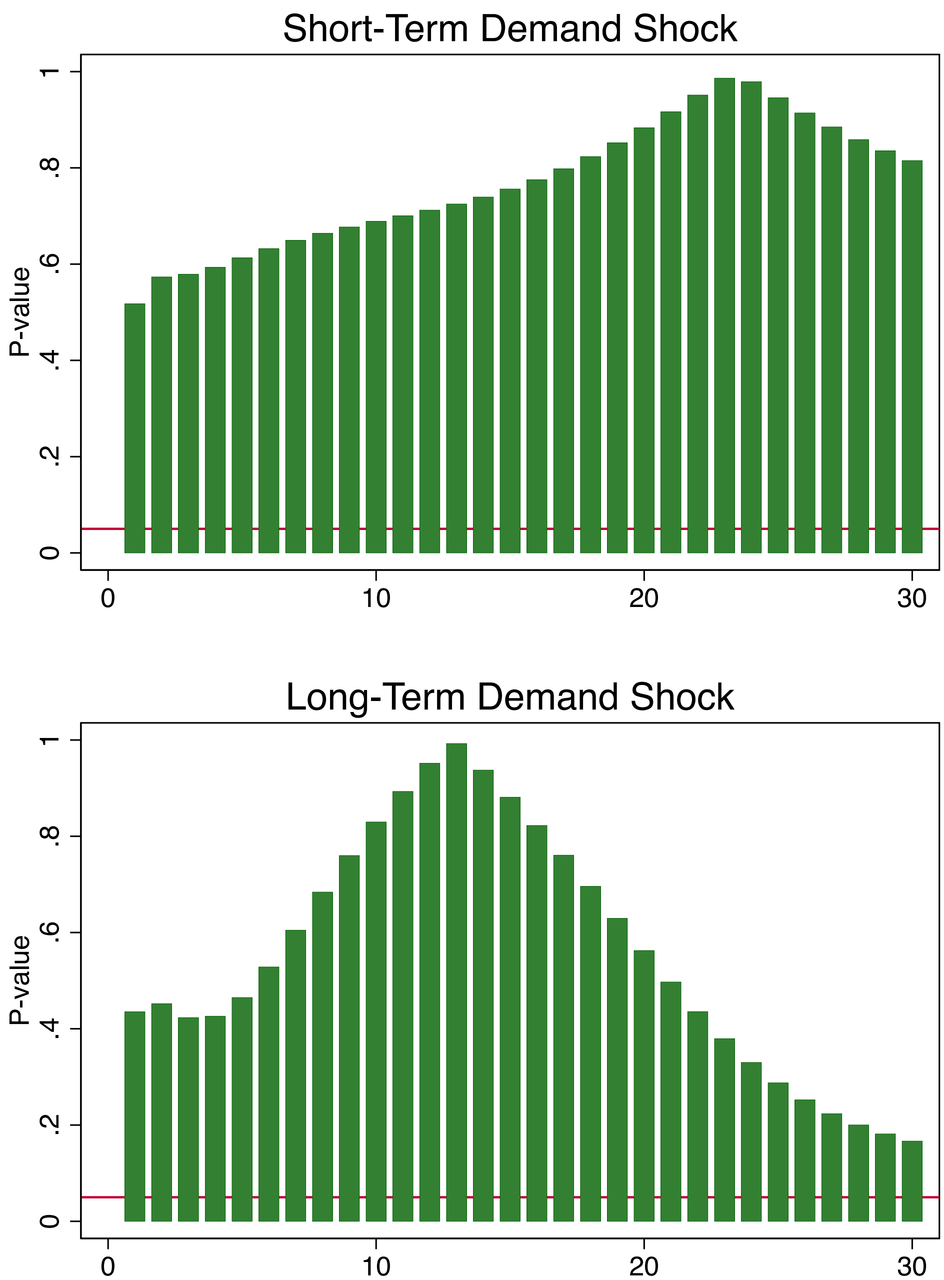

Notes: p-values testing equality of coefficients from Figure B9.

Figure B11: Rate Responses (rotated intraday Futures surprises)

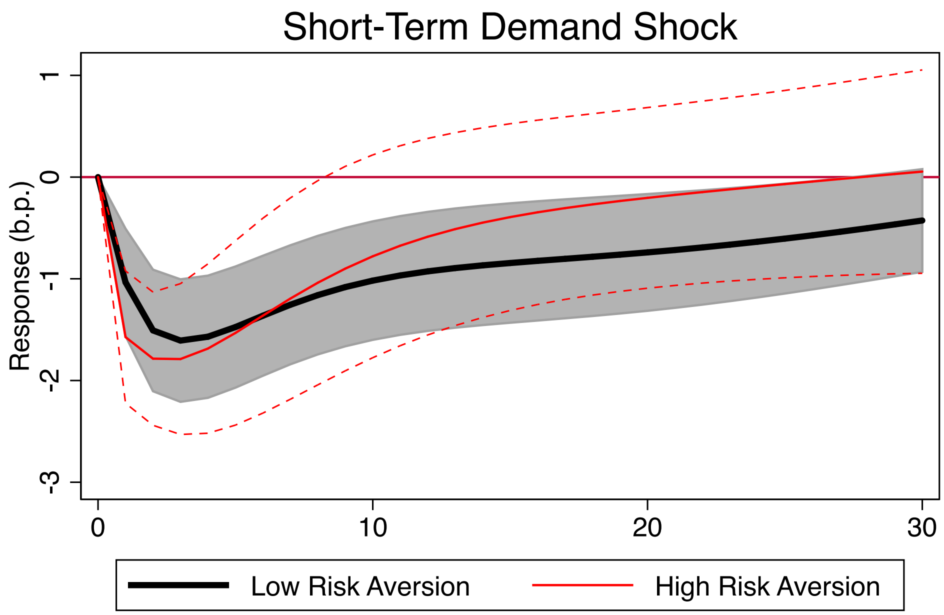

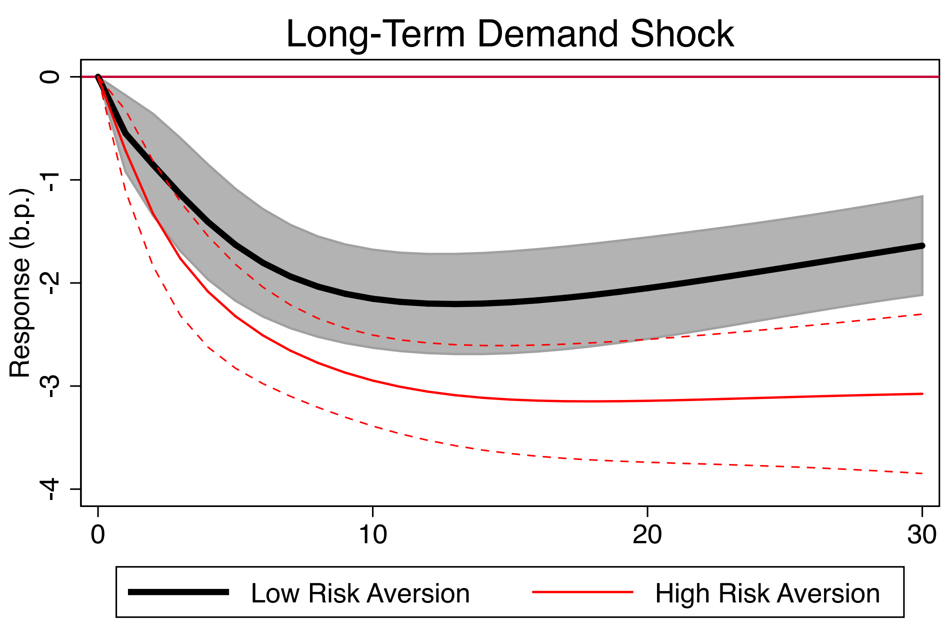

Notes: Plots of the regression coefficients on the demand shocks D_t^i from regression equation (2.4). The shocks are the first two principal components of our intraday shocks, D_t^ℓ and D_t^s, rotated such that D_t^s is uncorrelated with $D_t^{(30Y)}$. For long-term auctions the shock is D_t^ℓ; similarly short-term auctions use D_t^s. Each curve is from the subsample combinations: short-term and long-term auctions; and periods of high and low risk aversion. 2 standard error (Newey-West) confidence intervals are included.

Figure B12: Rate Response P-Values (rotated intraday Futures surprises)

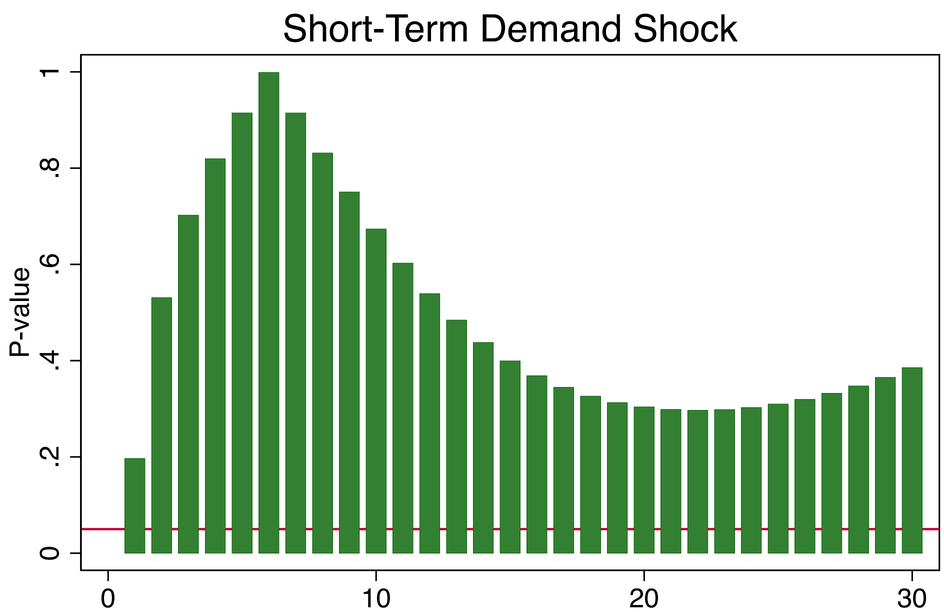

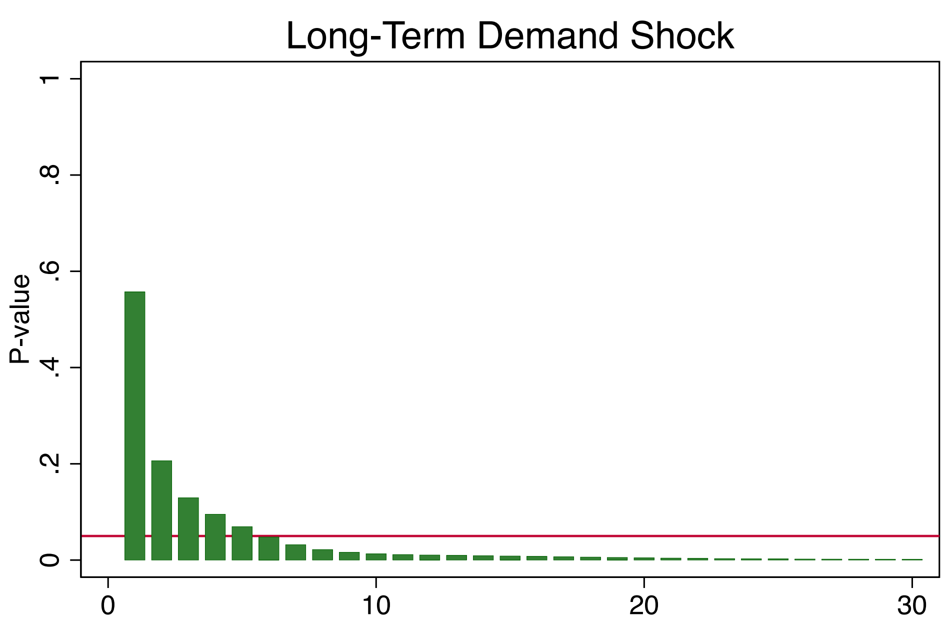

Notes: p-values testing equality of coefficients from Figure B11.

Table B1: Reaction of market to surprises at Treasury auctions (IV specification)

Dep.variable: asset type	Estimate (s.e.)	N	F-stat	Sample
	(1)	(2)	(3)	(4)
Panel A. Debt				
TLT	0.359***	662	78.9	2002-2015
	(0.032)			
SHY	0.024***	662	78.9	2002-2015
	(0.002)			
LQD	0.121***	662	78.9	2002-2015
	(0.015)			
Aaa†	-2.666***	871	126.6	1995-2015
	(0.406)			
Panel B. Equities				
SPY	0.016	871	126.6	1995-2015
	(0.027)			
IWM	0.039	706	91.4	2000-2015
	(0.060)			
SP500†	-0.111	871	126.6	1995-2015
	(0.113)			
Russell 2000†	-0.191	871	126.6	1995-2015
	(0.119)			

Notes: The table repeats the regressions from Table 2.4, but instruments $D_t^{(m')}$ with the surprise component of the bid-to-cover ratio. First-stage F-statistics are reported in column (3). Newey-West standard errors in parentheses.

Table B2: Reaction of market to surprises at Treasury auctions (IV specification)

Dep.variable: asset type	Estimate (s.e.)	N	F-stat	Sample
	(1)	(2)	(3)	(4)
Panel C. Inflation expectations and commodities				
10Y Inflation Swap†	-0.290	618	74.2	2004-2015
	(0.331)			
2Y Inflation Swap†	-0.001	618	74.2	2004-2015
	(0.669)			
GLD	0.041	595	72.3	2004-2015
	(0.030)			
GSCI†	0.023	871	126.6	1995-2015
	(0.107)			
Panel D. Spreads and credit default swaps				
Baa-Aaa†	-0.169	871	126.6	1995-2015
	(0.146)			
3-month LIBOR-OIS†	-0.006	630	77.3	2003-2015
	(0.004)			
Auto CDS†	-9.793	627	77.0	2004-2015
	(7.458)			
Bank CDS†	1.269	627	77.0	2004-2015
	(0.800)			
VIX†	-0.063	871	126.6	1995-2015
	(0.148)			

Notes: The table repeats the regressions from Table 2.4, but instruments $D_t^{(m')}$ with the surprise component of the bid-to-cover ratio. First-stage F-statistics are reported in column (3). Newey-West standard errors in parentheses.

Table B3: Secondary Market Rate Responses

Panel A: Short-Term (2-7 year) auctions

	(1) 0-2	(2) 2-5	(3) 5-8	(4) 8-11	(5) 11-20	(6) 20-30
C_t=0 × D_t	-1.32***	-2.16***	-2.48***	-1.97***	-2.46***	-1.79***
	(0.22)	(0.30)	(0.33)	(0.30)	(0.29)	(0.24)
C_t=1 × D_t	-1.44***	-2.61***	-3.09***	-3.02***	-2.72***	-2.09***
	(0.25)	(0.39)	(0.46)	(0.45)	(0.44)	(0.45)
Observations	35363	37304	15644	7813	9733	11085
Clusters	615	615	615	615	615	615
R^2	0.136	0.217	0.224	0.205	0.206	0.153
P-Value	0.712	0.366	0.274	0.053	0.625	0.561

Panel B: Long-Term (10-30 year) auctions

	(1) 0-2	(2) 2-5	(3) 5-8	(4) 8-11	(5) 11-20	(6) 20-30
C_t=0 × D_t	-0.56***	-1.07***	-1.55***	-1.83***	-2.21***	-2.13***
	(0.20)	(0.30)	(0.40)	(0.33)	(0.39)	(0.26)
C_t=1 × D_t	-1.04***	-1.96***	-2.49***	-2.95***	-3.00***	-3.15***
	(0.25)	(0.29)	(0.30)	(0.33)	(0.34)	(0.38)
Observations	15139	16525	7264	3385	4161	4559
Clusters	255	255	255	255	255	255
R^2	0.152	0.216	0.277	0.329	0.349	0.369
P-Value	0.144	0.032	0.064	0.017	0.136	0.028

Notes: The table reports results estimating equation (2.4), but using security-level changes in yields as the dependent variable. Panel A reports the results for short-term auction dates (2-7 years), while Panel B reports the results for long term auction dates (10-30 years). The columns break up the securities into different baskets based on the remaining maturity: column (1) contains all note and bonds with less than 2 years remaining before maturity; column (2) is 2-5 years; column, column (3) is 5-8 years; column (4) is 8-11 years; column (5) is 11-20 years; and column (6) is 20-30 years. P-values testing equality of coefficients are reported in the final row. Standard errors clustered at the auction level are in parentheses.

Table B4: Numerical Exercise Calibration

Parameter	Value
T	30
σ	.01
κ_r	0.7
κ_s	0.3
κ_ℓ	0.3
α	5
a	$(0, 500)$
$\theta_s(m)$	$\delta(m-3)$
$\theta_\ell(m)$	$\delta(m-20)$

Appendix C

Polarized Expectations

C.1 Proofs

Proof of Lemma 3.1. The characterization of posterior variance matrices $\mathbf{S}$, $\mathbf{\Xi}$, $\mathbf{\Sigma}$, and the optimal action $\mathbf{y}^*$ follow from Kőszegi and Matějka (2018).

To derive the form of Eq. (3.2), first consider the simple case of a noisy signal s for a single variable x.

$$
\begin{aligned}
s &= x + \sigma_e e \\
x &\sim N(\mu, \sigma_0^2) \text{ (prior)} \\
e &\sim N(0, 1) \text{ (known)}
\end{aligned}
$$

Then given a realization of the signal s, the posterior mean and variance is:

$$
\begin{aligned}
\tilde{x} = \mathbb{E}\left[x \middle| s\right] &= \mu + \frac{\sigma_0^2}{\sigma_0^2 + \sigma_e^2}(s - \mu) \\
\sigma^2 = \mathrm{Var}\left[x \middle| s\right] &= \frac{\sigma_0^2}{\sigma_0^2 + \sigma_e^2}\sigma_\varepsilon^2
\end{aligned}
$$

which can be written using $\xi \equiv 1 - \frac{\sigma^2}{\sigma_0^2} = \frac{\sigma_0^2}{\sigma_0^2+\sigma_e^2}$ as follows:

$$
\begin{aligned}
\tilde{x} &= \mu + \xi(s - \mu) = (1 - \xi)\mu + \xi s \\
\sigma^2 &= \xi\sigma_e^2 = (1 - \xi)\sigma_0^2
\end{aligned}
$$

Hence if the prior mean is zero ($\mu = 0$), then $\tilde{x} = \xi s$

In the general multivariate case with prior $\mathbf{x} \sim N(\mathbf{0}, \sigma_0^2\mathbf{I})$, the orthonormal eigenvector matrix $\mathbf{V}$ transforms the problem in a set of independent signals. If $\mathbf{v}^i$ is the

i^{th} element of $\mathbf{V}$, then

$$\begin{aligned}
(\mathbf{v}^i \cdot \tilde{\mathbf{x}}) &= \xi_i(\mathbf{v}^i \cdot \mathbf{x}) + \xi_i \sigma_{e_i} e_i \\
\implies \mathbf{V}^T\tilde{\mathbf{x}} &= \mathbf{\Xi}\mathbf{V}^T\mathbf{x} + \mathbf{\Xi}\Sigma_e^{1/2}\mathbf{e} \\
&= \mathbf{\Xi}\mathbf{V}^T\mathbf{x} + \sigma_0 \left(\mathbf{\Xi}(\mathbf{I}-\mathbf{\Xi})\right)^{1/2}\mathbf{e} \\
\implies \tilde{\mathbf{x}} &= \mathbf{V}\mathbf{V}^T\tilde{\mathbf{x}} \\
&= \mathbf{V}\mathbf{\Xi}\mathbf{V}^T\mathbf{x} + \sigma_0\mathbf{V}\left(\mathbf{\Xi}(\mathbf{I}-\mathbf{\Xi})\right)^{1/2}\mathbf{e} \\
&\equiv \mathbf{W}_1\mathbf{x} + \mathbf{W}_2\mathbf{e}
\end{aligned}$$

□

Proof of Prop. 3.1. Taking the conditional expectation of Eq. (3.2) gives

$$\begin{aligned}
\mathbb{E}\left[\tilde{\mathbf{x}}\middle|\mathbf{x}\right] &= \mathbf{W}_1\mathbf{x} + \mathbf{W}_2\mathbb{E}\left[\mathbf{e}\right] \\
&= \mathbf{W}_1\mathbf{x} \\
\implies \mathbb{E}\left[\tilde{\mathbf{x}} - \mathbf{x}\middle|\mathbf{x}\right] &= (\mathbf{W}_1 - \mathbf{I})\mathbf{x} \\
&= \mathbf{V}(\mathbf{\Xi} - \mathbf{I})\mathbf{V}^T\mathbf{x}
\end{aligned}$$

since $\mathbf{e}$ is mean-zero and independent of $\mathbf{x}$.

Note that $\mathbf{W}_1$ is symmetric and

$$\begin{aligned}
\mathbf{W}_1\mathbf{W}_1^T &= \mathbf{V}\mathbf{\Xi}\mathbf{V}^T\mathbf{V}\mathbf{\Xi}\mathbf{V}^T = \mathbf{V}\mathbf{\Xi}^2\mathbf{V}^T \\
\implies (\mathbf{W}_1)^n &= \mathbf{V}\mathbf{\Xi}^n\mathbf{V}^T \\
(\mathbf{I} - \mathbf{W}_1)^n &= \mathbf{V}(\mathbf{I} - \mathbf{\Xi})^n\mathbf{V}^T \\
\mathbf{W}_2\mathbf{W}_2^T &= \sigma_0^2\mathbf{V}\mathbf{\Xi}(\mathbf{I} - \mathbf{\Xi})\mathbf{V}^T = \mathbf{V}\mathbf{\Xi}\mathbf{S}\mathbf{V}^T \\
\implies \left(\mathbf{W}_2\mathbf{W}_2^T\right)^n &= \mathbf{V}(\mathbf{\Xi}\mathbf{S})^n\mathbf{V}^T
\end{aligned}$$

Hence the conditional variance is given by Eq. (3.5).

Taking derivatives with respect to information costs λ gives

$$\begin{aligned}
\frac{\partial \xi_i}{\partial \lambda} &= \begin{cases} -\frac{1}{2}(\sigma_0^2\Lambda_i)^{-1} & \text{if } 2\Lambda_i\sigma_0^2 \geq \lambda \\ 0 & \text{otherwise} \end{cases} \\
\frac{\partial \xi_i S_{ii}}{\partial \lambda} &= \begin{cases} \frac{1}{2}\frac{\Lambda_i\sigma_0^2 - \lambda}{\Lambda_i^2\sigma_0^2} & \text{if } 2\Lambda_i\sigma_0^2 > \lambda \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

Then the derivatives of the respective diagonal matrices are

$$\frac{\partial}{\partial \lambda}\mathbf{\Xi} = \operatorname{diag}\left(\frac{\partial \xi_i}{\partial \lambda}\right), \; \frac{\partial}{\partial \lambda}\mathbf{\Xi}\mathbf{S} = \operatorname{diag}\left(\frac{\partial \xi_i S_{ii}}{\partial \lambda}\right)$$

Since the eigenvectors $\mathbf{V}$ are independent of information costs,

$$\frac{\partial}{\partial\lambda}\mathrm{Var}\left[\tilde{\mathbf{x}} - \mathbf{x}\middle|\mathbf{x}\right] = \mathbf{V}\left[\frac{\partial}{\partial\lambda}\mathbf{\Xi S}\right]\mathbf{V}^T \tag{A1}$$

Note

$$\mathrm{sign}\,\frac{\partial\xi_i}{\partial\lambda} = \begin{cases} 0 & \text{if } \lambda > 2\Lambda_i\sigma_0^2 \\ -1 & \text{if } \lambda \leq 2\Lambda_i\sigma_0^2 \end{cases}$$

$$\mathrm{sign}\,\frac{\partial\xi_i S_{ii}}{\partial\lambda} = \begin{cases} 0 & \text{if } \lambda > 2\Lambda_i\sigma_0^2 \\ -1 & \text{if } 2\Lambda_i\sigma_0^2 \geq \lambda > \Lambda_i\sigma_0^2 \\ 0 & \text{if } \lambda = \Lambda_i\sigma_0^2 \\ 1 & \text{if } \lambda < \Lambda_i\sigma_0^2 \end{cases}$$

Recall that the eigenvalues are ordered $\Lambda_1 \geq \ldots \geq \Lambda_J$ and that some eigenvalues may be zero:

$$\mathbf{\Lambda} = \mathrm{diag}\begin{bmatrix}\Lambda_1 & \ldots & \Lambda_N & 0 & \ldots & 0\end{bmatrix}$$

Hence, if $2\Lambda_1\sigma_0^2 > \lambda > \sigma_0^2\Lambda_1$ then $\lambda > sigma_0\Lambda_i\ \forall i$ and hence

$$\frac{\partial\xi_1 S_{11}}{\partial\lambda} < 0$$
$$\frac{\partial\xi_i S_{ii}}{\partial\lambda} \leq 0$$

Thus, the matrix $\frac{\partial}{\partial\lambda}\mathbf{\Xi S}$ in Eq. (A1) is a diagonal matrix with nonpositive diagonal elements. Since these are the eigenvalues of $\frac{\partial}{\partial\lambda}\mathrm{Var}\left[\tilde{\mathbf{x}} - \mathbf{x}\middle|\mathbf{x}\right]$, this matrix is negative semidefinite.

A similar argument shows that $\lambda < \Lambda_N\sigma_0^2$ implies $\frac{\partial}{\partial\lambda}\mathrm{Var}\left[\tilde{\mathbf{x}} - \mathbf{x}\middle|\mathbf{x}\right]$ is positive semidefinite.

Note that as information costs approach zero,

$$\lim_{\lambda=0}\xi_i = \begin{cases} 1 & \text{if } \Lambda_i > 0 \\ 0 & \text{otherwise} \end{cases}$$

Hence as the cost of information $\lambda \to 0$,

$$\mathbf{\Xi} \to \tilde{\mathbf{I}}_N$$
$$\mathbf{S} \to \sigma_0^2(\mathbf{I} - \tilde{\mathbf{I}}_N)$$
$$\mathbf{\Xi S} \to \mathbf{0}$$

Hence

$$\begin{aligned} \mathbf{W}_1 &\to \mathbf{V}\tilde{\mathbf{I}}_N\mathbf{V}^T \\ \mathbf{W}_2\mathbf{W}_2^T &\to \mathbf{0} \end{aligned}$$

and the results follow.

$\square$

Proof of Cor. 3.1.1. Consider the matrix $\mathbf{V}(\tilde{\mathbf{I}}_N - \mathbf{I})$, which is made up of columns of zeros and columns of eigenvectors associated with zero eigenvalues (if they exist). For any such eigenvalue $\mathbf{v}$, we have

$$\mathbf{\Omega v} = \mathbf{0} \implies \mathbf{v}^T\mathbf{B}\mathbf{C}^{-1}\mathbf{B}^T\mathbf{v} = 0$$

Since $\mathbf{C}$ and hence $\mathbf{C}^{-1}$ are positive definite,

$$\begin{aligned} \mathbf{z}^T\mathbf{C}^{-1}\mathbf{z} = 0 &\iff \mathbf{z} = \mathbf{0} \\ &\implies \mathbf{B}^T\mathbf{v} = \mathbf{0} \\ &\implies \mathbf{B}^T\mathbf{V}(\tilde{\mathbf{I}}_N - \mathbf{I}) = \mathbf{0} \\ &\implies \mathbf{H}\mathbf{V}(\tilde{\mathbf{I}}_N - \mathbf{I})\mathbf{V}^T\mathbf{x} = \mathbf{0} \end{aligned}$$

Then note as $\lambda \to 0$:

$$\mathbf{H}(\tilde{\mathbf{x}} - \mathbf{x}) \to \mathbf{H}\mathbf{V}(\tilde{\mathbf{I}}_N - \mathbf{I})\mathbf{V}^T\mathbf{x} = \mathbf{0}$$

and hence $\mathbf{y}^* \to \mathbf{H}\mathbf{x}$.

$\square$

Proof of Prop. 3.2. Conditional on an individual k, the solution is identical to the problem considered in the case of homogenous agents. Hence

$$\mathbb{E}\left[\tilde{\mathbf{x}}^k \middle| \mathbf{x}, k\right] = \mathbf{W}_1^k\mathbf{x}$$

Then the law of iterated expectations implies

$$\begin{aligned} \mathbb{E}\left[\tilde{\mathbf{x}}^k \middle| \mathbf{x}\right] &= \mathbb{E}\left[\mathbb{E}\left[\tilde{\mathbf{x}}^k \middle| \mathbf{x}, k\right] \middle| \mathbf{x}\right] \\ &= \mathbb{E}\left[\mathbf{W}_1^k\right]\mathbf{x} \\ &= \int_k \mathbf{W}_1^k \, \mathrm{d}F(k)\, \mathbf{x} \equiv \bar{\mathbf{W}}_1\mathbf{x} \end{aligned}$$

hence the expression for the conditional misperception in Eq. (3.6) follows. Further,

$$(\tilde{\mathbf{x}}^k - \mathbf{x}) - \mathbb{E}\left[\tilde{\mathbf{x}}^k - \mathbf{x} \middle| \mathbf{x}\right] = \widetilde{\mathbf{W}}_1^k\mathbf{x} + \mathbf{W}_2^k\mathbf{e}^k$$

and

$$\left[\widetilde{\mathbf{W}}_1^k\mathbf{x}+\mathbf{W}_2^k\mathbf{e}^k\right]\left[\widetilde{\mathbf{W}}_1^k\mathbf{x}+\mathbf{W}_2^k\mathbf{e}^k\right]^T=\widetilde{\mathbf{W}}_1^k\mathbf{x}\mathbf{x}^T(\widetilde{\mathbf{W}}_1^k)^T$$
$$+\widetilde{\mathbf{W}}_1^k\mathbf{x}(\mathbf{e}^k)^T(\mathbf{W}_2^k)^T+\mathbf{W}_2^k\mathbf{e}^k\mathbf{x}^T(\widetilde{\mathbf{W}}_1^k)^T$$
$$+\mathbf{W}_2^k\mathbf{e}^k(\mathbf{e}^k)^T(\mathbf{W}_2^k)^T$$

Taking expectations of the above expression conditional on $\mathbf{x}$ and k gives

$$\widetilde{\mathbf{W}}_1^k\mathbf{x}\mathbf{x}^T(\widetilde{\mathbf{W}}_1^k)^T+\sigma_0^2\mathbf{W}_2^k(\mathbf{W}_2^k)^T$$

Thus, the law of iterated expectations implies that the conditional variance of posterior belief misperceptions is given by Eq. (3.7).

Note that from how $\bar{\lambda}$ is defined, $\Xi^k=\mathbf{0}\ \forall k$ at $\bar{\lambda}$. Hence $\widetilde{\mathbf{W}}_1^k=\mathbf{0}$ as well. Then note, taking derivatives with respect to information costs:

$$\frac{\partial}{\partial\lambda}\mathbf{\Sigma}_{\widetilde{\mathbf{W}}_1\widetilde{\mathbf{W}}_1|\mathbf{x}}=\int_k\left[\frac{\partial}{\partial\lambda}\widetilde{\mathbf{W}_1}^k\right]\mathbf{x}\mathbf{x}^T(\widetilde{\mathbf{W}_1}^k)^T\,\mathrm{d}F(k)$$
$$+\int_k\widetilde{\mathbf{W}_1}^k\mathbf{x}\mathbf{x}^T\left[\frac{\partial}{\partial\lambda}\widetilde{\mathbf{W}_1}^k\right]^T\mathrm{d}F(k)$$

The above expression is also equal to $\mathbf{0}$ at $\bar{\lambda}$.

Additionally,

$$\frac{\partial}{\partial\lambda}\mathbf{\Sigma}_{\mathbf{W}_2\mathbf{W}_2}=\int_k\mathbf{V}^k\left[\frac{\partial}{\partial\lambda}\mathbf{\Xi}^k\mathbf{S}^k\right](\mathbf{V}^k)^T\,\mathrm{d}F(k)$$

From the proof of Prop. 3.2, the integrand of the above expression is negative semidefinite at $\bar{\lambda}$. Hence, as the sum of negative semidefinite matrices, the above expression is negative semidefinite at $\bar{\lambda}$. Thus $\frac{\partial}{\partial\lambda}\mathrm{Var}\left[\tilde{\mathbf{x}}^k-\mathbf{x}\middle|\mathbf{x}\right]$ is negative semidefinite at $\bar{\lambda}$.

The results regarding the limit as $\lambda\to 0$ follow from the proof in Prop. 3.2, which shows that

$$\lim_{\lambda\to 0}\mathbf{W}_1^k=\mathbf{V}^k\tilde{\mathbf{I}}_N(\mathbf{V}^k)^T$$
$$\lim_{\lambda\to 0}\mathbf{W}_2^k(\mathbf{W}_2^k)^T=\mathbf{0}$$

□

C.2 Log-Quadratic Approximation

This section derives the log-quadratic approximation around the steady state, and applies the approximation to the two period model described in Section 3.2.3.

Suppose the individual has arbitrary preferences $U(\mathbf{Y}, \mathbf{X})$, where $(\mathbf{Y}, \mathbf{X})$ are the choices and state variables, respectively. The steady state of the state variables is known: $\bar{\mathbf{X}}$. Then the steady state of the choice variables $\bar{\mathbf{Y}}$ is implicitly defined by

$$\left.\frac{\partial U(\mathbf{Y}, \bar{\mathbf{X}})}{\partial \mathbf{Y}^T}\right|_{\mathbf{Y}=\bar{\mathbf{Y}}} = \mathbf{0}$$

Let lower case variables denote log deviations from steady state: $z = \log Z - \log \bar{Z}$. Then the utility function can be written equivalently as

$$\hat{u}(\mathbf{y}, \mathbf{x}) \equiv U(\bar{\mathbf{Y}} \circ \exp(\mathbf{y}), \bar{\mathbf{X}} \circ \exp(\mathbf{x})) \tag{B1}$$

where $\circ$ is Hadamard (elementwise) multiplication, and with a slight abuse of notation such that $\exp(\mathbf{y})$ is the element-wise exponential.

Define the first- and second-order partial derivatives evaluated at steady state:

$$\begin{aligned}
\hat{u}_{\mathbf{y}} &\equiv \left.\frac{\partial \hat{u}}{\partial \mathbf{y}}\right|_{\mathbf{x},\mathbf{y}=\mathbf{0}} \\
\hat{u}_{\mathbf{y},\mathbf{y}} &\equiv \left.\frac{\partial^2 \hat{u}}{\partial \mathbf{y} \partial \mathbf{y}^T}\right|_{\mathbf{x},\mathbf{y}=\mathbf{0}} \\
\hat{u}_{\mathbf{y},\mathbf{x}} &\equiv \left.\frac{\partial^2 \hat{u}}{\partial \mathbf{y} \partial \mathbf{x}^T}\right|_{\mathbf{x},\mathbf{y}=\mathbf{0}}
\end{aligned}$$

Then the second-order (log) approximation of the utility function around the steady state is

$$\begin{aligned}
U(\mathbf{y}, \mathbf{x}) &\approx -\mathbf{y}^T \mathbf{C} \mathbf{y} + \mathbf{x}^{\mathbf{T}} \mathbf{B} \mathbf{y} + \text{h.o.t.} + \text{t.i.c.} \\
\mathbf{C} &= -\frac{1}{2} \hat{u}_{\mathbf{y},\mathbf{y}} \\
\mathbf{B} &= \hat{u}_{\mathbf{y},\mathbf{x}}
\end{aligned}$$

Note the above ignores higher order terms, and terms independent of choice.

C.2.1 Two-Period Model

Now consider the model considered in Section 3.2.3. The budget constraint always holds with equality, so the utility function is equivalent to the concentrated utility function with

$$C_2 = \frac{1}{P_2} \left(R((1-\tau)Y - P_1 C_1) + \tau Y \right)$$

In steady state we have $\bar{Y} = 1$, $\bar{P}_1 = \bar{P}_2 = 1$, and $\overline{1-\tau} = 1$. Then steady state period-1 consumption is $\bar{C}_1 = \frac{1}{1+\beta}$. Re-writing the utility function in terms of deviations from steady state, as in Eq. (B1), and evaluating the first- and second-order partial derivatives involving period-1 log consumption c_1 gives

$$\begin{aligned}
\hat{u}_{c_1} &= 0 \\
\hat{u}_{c_1,c_1} &= -\frac{(1+\beta)^{\eta^k}\eta^k}{\beta} \\
\hat{u}_{c_1,y} &= \frac{(1+\beta)^{\eta^k}\eta^k}{\beta} \\
\hat{u}_{c_1,1-\tau} &= \frac{(1+\beta)^{\eta^k}\eta^k(1-\beta)}{\beta} \\
\hat{u}_{c_1,p_1} &= -\frac{(1+\beta)^{\eta^k}(\beta+\eta^k)}{\beta} \\
\hat{u}_{c_1,p_2} &= (1+\beta)^{\eta^k}(1-\eta^k)
\end{aligned}$$

Using these results gives the expressions for the quadratic preference matrices C^k (in this case a scalar) and $\mathbf{B}^k$ (in this case a column vector).