Essays on Macroeconomic Policy and the Prices of Financial Assets

by

James A. Caldera

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2022

Doctoral Committee:

Professor Uday Rajan, Chair
Professor Jagadeesh Sivadasan
Professor Linda L. Tesar
Professor Toni M. Whited
Assistant Professor Emmanuel Yimfor

James A. Caldera

jcaldera@umich.edu

ORCID iD: 0000-0002-0359-2510

ACKNOWLEDGEMENTS

For their generous guidance and encouragement, I am indebted to my committee – Uday Rajan, Jagadeesh Sivadasan, Linda Tesar, Toni Whited, and Emmanuel Yimfor. I thank Ross Finance Area faculty members and fellow Ph.D. students for their helpful comments and suggestions. I am also grateful to Tammy Feldman and Brian Jones for helping to make my teaching experience at the University of Michigan as immensely rewarding as it was. Finally, I thank my family and friends for their unwavering support.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

ATM At-the-money

CAD Cumulative abnormal difference

CAPM Capital asset pricing model

CAR Cumulative abnormal return

CDS Credit default swap

CLO Collateralized loan obligation

CP Call-put

CPP Capital Purchase Program

DITM Deep in-the-money

DOTM Deep out-of-the-money

FABS Funding agreement-backed security

FF3 Fama-French three-factor model

FIG SNL Financial Institutions

FOMC Federal Open Market Committee

FSOC Financial Stability Oversight Council

ITM In-the-money

IV Option-implied volatility

LGD Loss given default

LM Lucca and Moench (2015)

MA Moving average

MR Modified restructuring

MTB Market-to-book

OTM Out-of-the-money

PC Property and casualty

PD Probability of default

P<*N*> Nth percentile

SIFI Systemically important financial institution

SMD Simulated minimum distance

TARP Troubled Asset Relief Program

TFP Total factor productivity

Abstract

My dissertation investigates what information embedded in financial prices reveals about questions relevant for macroeconomic policy. The first two chapters examine implicit government guarantees in the U.S. life insurance industry. My third chapter provides novel evidence on a phenomenon at the intersection of asset pricing and monetary policy.

To what extent do investors view life insurers as "too big to fail?" Chapter 1 provides new evidence on this question using a natural experiment from the 2008–09 financial crisis. The analysis examines market reactions to a U.S. Treasury announcement that raised expectations about government backstops for the life insurance industry. I find that a subset of large life insurers benefited from significant protection against bankruptcy, with implied risk neutral probabilities of a government rescue ranging from 21% to 37% at the one-year horizon. Rescue probabilities exhibit a broadly downward sloping term structure, suggesting that investors expected the protection to subside with time. Cross-sectional differences are also informative about which insurers would later be designated systemically important by regulators.

Chapter 2 builds on the findings presented in Chapter 1. In this chapter, I study the long-term dynamics of implicit government backstops for the U.S. life insurance industry and how this protection affects moral hazard. I structurally estimate a partial equilibrium model of life insurers protected in part by emergency bailouts. The estimates imply that for the 2001–20 period the investor expectation of the (physical) probability of a bailout is 9.1% for large insurers. Indicative of prominent time-variation, the standard deviation of these bailout probabilities is 10.3%. A counterfactual analysis reveals limited average levels of moral hazard in risk-taking practices. Structural estimates for small life insurers imply more modest support. Overall, the results are consistent with the presence of a significant and time-varying "too big to fail" subsidy for large life insurers.

In Chapter 3, I investigate a puzzle of asset pricing in relation to monetary policy communication. S&P 500 stocks earn extraordinary returns in the run-up to Federal Open Market Committee (FOMC) announcements. Recent literature documents that this pre-FOMC drift appears to have waned after early 2011. Using data extended through 2020, I show that pre-announcement returns have since rebounded in the most current period. What do these fluctuations reveal about the impetus behind this phenomenon? This chapter addresses this question with options-based methodology that disentangles variation in FOMC

policy impact from investor anticipation thereof. I decompose close-to-announcement returns into risk premia, price changes that reflect information about the announcement, and an unexplained (anomalous) component. I find that 30% to 55% of average pre-FOMC returns between 1996 and 2011 were due to an anomaly that has since vanished. During this time, risk premia account for 32% to 77% of this mean whereas informed trading in advance of positive news represents only between 1% and 12%. The cessation of anomalous returns and changes in risk premia appear most responsible for recent shifts in pre-FOMC returns.

With findings that are informative for macroprudential policy, the contributions of my first two chapters shed new light on the "too big to fail" problem in the insurance sector. My investigation of the pre-FOMC announcement drift in Chapter 3 further elucidates how capital markets are influenced by monetary policy decisions. The evidence I present highlights the usefulness of financial prices for informing policy.

Chapter 1

Implicit Government Guarantees in the U.S. Life Insurance Sector: Evidence from the TARP Program

1.1 Introduction

During the 2008–09 financial crisis AIG received $182.5 billion in U.S. government bailout funds. Additional tens of billions of dollars in total government support were provided to MetLife, Hartford Financial Services, and Lincoln National through emergency loans or Troubled Asset Relief Program (TARP) capital injections. A natural question raised by the unprecedented interventions of this time period asks whether, and to what extent, the U.S. life insurance industry has a "too big to fail" problem. That is, can the government commit to not bailing out the largest of these institutions should they face bankruptcy? This question accrues some urgency amid growing evidence that the scope for risk-taking in this industry extends well beyond the traditional model of life insurance. These non-traditional activities—which are incentivized by protracted low interest rate environments and frequently involve off-balance sheet financial engineering with a surprising breadth for regulatory arbitrage—can expose these institutions to a fragility not commonly seen outside of the banking sector (see, e.g., Koijen and Yogo, 2017; Foley-Fisher, Narajabad, and Verani, 2020; Koijen and Yogo, 2021). Recent extraordinary shifts in the post-crisis regulatory environment[1] also make it clear that answers to these questions are of great practical importance in the insurance industry and in policy circles alike.

This chapter provides evidence on the above question by asking what investor beliefs

[1]Under the government authority created by the 2010 Dodd-Frank Act, the Financial Stability Oversight Council (FSOC) voted to designate three life insurers (AIG, Prudential, and MetLife) as nonbank systemically important financial institutions (SIFIs) between 2013 and 2014. The designation carries with it enhanced regulatory oversight by FSOC and its member affiliates, which include the U.S. Treasury and the Federal Reserve. In a striking course reversal, all three of the insurers have since shed their SIFI status during the 2016 to 2018 period. Strictly speaking, MetLife was not officially a SIFI as its designation had been under appeal; for concise terminology, however, I refer to the company as a SIFI for the period between FSOC's December 2014 decision and MetLife's successful legal challenge in March 2016.

embedded in market prices reveal about implicit government guarantees for publicly-traded U.S. life insurers during the 2008–09 financial crisis.[2] The analysis examines market reactions to a U.S. Treasury Department announcement that introduced a shock to expectations about government backstops for the life insurance industry. The results indicate that a subset of large life insurance companies benefited from generous protection against bankruptcy. I place these results into perspective with credit default swap (CDS) implied probabilities of a government bailout. The approach asks how large the risk neutral probability of a bailout conditional on default must be if abnormal deviations in CDS rates at the time of the announcement were attributable to a change in protection of this kind.[3] I find that, depending on the insurer, implied bailout probabilities range from 21% to 37% at the one-year horizon. Cross-sectional differences are informative about which life insurers would later be designated systemically important by regulators, consistent with a "too big to fail" interpretation. I document two additional facts about market reactions to this event. First, credit spreads respond more strongly to news about the government support than equity prices. Second, implied bailout probabilities exhibit a generally downward-sloping term structure. This suggests that creditors expected the magnitude of the subsidy to retreat as economic conditions improve, highlighting the importance of dynamics.

The event at the heart of this analysis occurred during the latter stages of the crisis. On April 8, 2009 the U.S. Treasury Department issued a statement clarifying that life insurers may be eligible to receive aid through the TARP program. Quantification of market responses to this news follows from standard event study methodology, which I apply to a variety of financial claims issued by individual insurers. Asset price deviations defined over a brief event window containing the announcement serve as the basis for the analysis. The announcement's price impact is thus measured by event deviations that are abnormal with respect to a baseline factor pricing model designed to capture systematic financial and economic conditions. For each asset class, my evidence is robust to multiple baseline model specifications. The CDS findings, which are of particular importance for this chapter's contribution, are also robust across a wide range of contract tenors.

One potential concern is that the Treasury's announcement may have been well-anticipated.

[2]For results presented in this chapter, a life insurer refers to the publicly-traded parent company rather than a life insurance division within a broader organization. Although this complicates the distinction between life and non-life insurers, this unit of analysis is necessitated by the integral contribution of capital market prices in the methodology. Sample inclusion criteria are described in Section 1.3.

[3]It is likely that the government support manifests in a more nuanced fashion. For example, the protection may also lead to a decline in default probabilities as well. Viewing the implied support through the lens of bailout probabilities per se provides a simple way to interpret the magnitude of the results. Note that without knowledge of the pre-announcement bailout probabilities the event studies can recover only lower bounds for post-announcement bailout probabilities. If the ex ante bailout probabilities were zero (i.e., investors had no expectation of government protection), the lower bounds hold with equality.

As I discuss in Section 1.4.1, however, there is compelling anecdotal and empirical evidence that it came as a surprise to investors. An important consideration in this regard is that TARP was officially intended for bank and thrift-holding companies. From an ex ante perspective, efforts to expand eligibility to insurers could have run aground amid political pitfalls. I document that by the time of the announcement investors had no shortage of indicators to suspect that this was indeed the case. Placebo tests also serve to verify that the announcement presented a shock to beliefs about protection for life insurers in particular.

Even with anecdotal evidence about financial crisis aid to life insurers set aside, there exist a wealth of reasons why the U.S. government may be unable to avoid bailing out these institutions. Many of these reasons persist beyond the idiosyncrasies of the crisis and may continue to fuel an *ex ante* "too big to fail" problem with burgeoning force as of this writing. I explore this subject in detail in Section 1.2. At a high-level, however, life insurers play an important role in financial intermediation that extends well beyond the already apparent critical service provided to policyholders. Most notably, life insurers provide extensive credit provision in the U.S., both indirectly through corporate and mortgage-backed bond investment (Cummins and Weiss, 2013; FSB 2019) and increasingly in a direct fashion by way of private lending (Foley-Fisher, Heinrich, and Verani, 2020). These institutions are also riskier than the traditional model of a life insurance company might suggest. Life insurers' significant engagement in recent decades with non-traditional investment and borrowing activities, which are incentivized by low interest rate regimes, exposes the industry to runs (e.g., Foley-Fisher, Narajabad, and Verani, 2020) and other sources of fragility (see Peirce, 2014; Koijen and Yogo, 2021; among others). Significant engagement with regulatory arbitrage also makes the depth of life insurers' risk exposure difficult for authorities to measure, let alone regulate (Cetina et al., 2016; Koijen and Yogo, 2016).

The remainder of this chapter proceeds as follows. Section 1.2 summarizes the institutional details of the life insurance industry that set the stage for a potential "too big to fail" problem in the sector. In Section 1.3, I describe the data sources and sample construction. Section 1.4 presents my main findings from the TARP news event. I conclude in Section 1.5.

1.2 Institutional Background and the "Too Big to Fail" Problem

Beyond the AIG bailout, there exists a compelling prima facie case that large life insurers might be subject to implicit government guarantees. This case is twofold.

The first aspect concerns what might happen should a large life insurance carrier become severely distressed amid broader financial market turmoil. Life insurance serves an important role in the financial economy. Core products such as whole and term life insurance policies and annuity contracts are fundamental tools for households' labor, health, and retirement planning. U.S. life insurers' balance sheet liabilities were $8.4 trillion in 2020, which account for 7.5% of total liabilities in the domestic financial sector. The risk to policyholders posed by the failure of a large life insurance carrier may serve as a powerful motive for government authorities weighing the possibility of a rescue package.[4]

In addition to providing a fundamental service to households, life insurance companies represent a prominent source of credit in the U.S. In 2020, life insurance carriers held 22.0% of outstanding foreign and corporate bonds. These institutions also held 11.8% of outstanding commercial mortgages as of 2011 and 16% of outstanding collateralized loan obligations (CLOs) in 2018 (Cummins and Weiss, 2013; FSB 2019). This sizable participation in credit markets raises the possibility of fire sale externalities. Examining the prices of recently-downgraded bonds held by regulatory constrained life insurers, Ellul, Jotikasthira, and Lundblad (2011) provide evidence of insurers' capacity to catalyze fire sales in the corporate bond market. Evidence of fire sale risk from insurers in the mortgage-backed securities market is also documented by Merrill, Nadauld, Stulz, and Sherlund (2012).

Life insurers' role in credit provision extends to more direct means than the kind provided by bond market investing. In post-crisis years, these companies are increasingly involved in private debt origination. Foley-Fisher, Heinrich, and Verani (2020) report that of the $1 trillion in retail and institutional annuity capital added by U.S. life insurers between 2008 and 2018, private real estate and private credit account for approximately $500 billion and $200 billion, respectively. These private debt investments are held both on- and off-balance sheet, the study indicates, with a significant share securitized in an originate-to-distribute model. The authors document that gross CLO issuance by U.S. life insurers was $60 billion in 2018, with new issuance by these insurers comprising roughly 23% of total new CLO issuance that year. Together, the above facts suggest that the failure of a large life insurer could threaten significant disruption to credit provision in the U.S. economy.

The second half of the prima facie case is that life insurance carriers may be more fragile than the traditional model of a life insurer would suggest. Of particular consequence is the risk that the maturity or liquidity transformation employed by life insurers can subject these

[4]Explicit added protection for policyholders already exists for some life insurance products in the form of state guaranty funds. The Financial Stability Oversight Council speculates, however, that these associations may have inadequate funds to fully cover qualifying policies in the event of insolvency for some of the largest and most complex insurance companies (FSOC 2014). Moreover, if the bankruptcy of a large carrier exhausts state guaranty funds, this may transmit distress throughout the U.S. insurance industry even if the failing insurer's immediate policyholders are adequately covered. Concern about this possibility among official authorities may also motivate implicit government protection.

institutions to runs. Traditionally associated with the core functions of banking, maturity (liquidity) transformation arises when longer-term (less liquid) assets are funded with shorter-term (more liquid) liabilities. There exist multiple ways in which modern life insurance companies engage in this transformation. A range of life insurance products routinely embed deposit contract features, accumulating a cash value that is redeemable by policyholders on demand. For an example of the magnitude of this exposure in one large insurer, \$49 billion of MetLife's \$308 billion in general account liabilities in 2013 comprised insurance products eligible for withdrawal with little or no penalty (FSOC 2014).

Security lending provides another avenue for this transformation. In a security lending arrangement, the security borrower deposits collateral, often cash, with their counterparty. The security lender can invest cash collateral at its discretion, but remains liable for any shortfalls in the collateral account's value incurred by these investments upon return of the borrowed security. A mismatch arises when the security lender invests cash collateral in assets that are riskier or less liquid than the lent security. It is notable that although much of the common knowledge about AIG's near-collapse during the financial crisis centers on its CDS liabilities, a commensurate factor was its aggressive securities lending program (Peirce, 2014; McDonald and Paulson, 2015). Life insurers leverage their extensive financial asset portfolios to command a dominant role in the securities lending market; Koijen and Yogo (2017) report that in 2007 life insurance companies accounted for nearly 90% of the \$128 billion market. Insurers can also generate a liquidity mismatch with other borrowing sources such as funding agreement-backed securities (FABS)—debt instruments issued off-balance sheet through special purpose entities that may be treated *pari passu* with claims on insurance policies. The concern that short-term funding secured in this fashion could expose life insurers to runs has direct empirical support (Foley-Fisher, Narajabad, and Verani, 2020). From a general perspective, moderately liquid and highly liquid obligations account for approximately 54% of life insurance liabilities (Paulson et al., 2014). This is in contrast to the relatively illiquid nature of life insurance assets.

The provision of market insurance through guaranteed return products comes with added vulnerability for life insurers during economic contractions. Counter to the traditional practice of assuming largely diversifiable risks from their policyholders, insurers that provide guaranteed return products are, in essence, shorting market put options. The breadth of this exposure may be far from trivial. Koijen and Yogo (2021) document that variable annuities, one of multiple product lines through which life insurers provide return guarantees, accounted for 34% of life insurance liabilities in 2015.

Life insurers may also be riskier than an examination of their financial statements—or their compliance with state insurance regulations—suggest. As described in Foley-Fisher, Narajabad, and Verani (2020), security lending may be used to relax risk-based capital constraints. This is because high quality securities lent against cash collateral may remain on insurers' balance sheets, regardless of the assets in which this cash is reinvested. The authors note how FABS can supply proceeds for securities used in this type of capital charge mitigation. That these liabilities are also issued off-balance sheet warrants emphasis. Captive reinsurance provides an even more overt avenue for regulatory arbitrage. With captive reinsurance, a primary insurer cedes risks from policies it has underwritten to special purpose entities wholly-owned by the parent firm called captives. An otherwise binding regulatory constraint in one jurisdiction can thus be circumvented by transferring risks to captive reinsurers in jurisdictions where constraints are non-binding.[5] Among life insurers engaged in this practice in 2012, 25 cents of each dollar they insured were ceded to captive reinsurers (Koijen and Yogo, 2016). Summarizing their report on the subject in the U.S. Treasury's OFR Brief Series, Cetina, Fliegelman, Glicoes, and Leung (2016) state plainly, "Publicly available data are insufficient to analyze fully the risks from captives and the impact on insurers' financial condition," going on to contend, "Because life insurers are a material part of the financial system, these gaps may mask financial stability vulnerabilities."

Some context for these non-traditional activities casts light on one final point. Fundamentally, the role of a life insurer is to collect premiums on long-term liabilities and invest these premiums in a portfolio of largely fixed-income securities, earning a profit on the spread between investment income and net payouts to policyholders. Because life insurance policies can extend well beyond single-decade horizons, the market for bonds of comparable duration is limited (Hartley, Paulson, and Rosen, 2017). This provides insight into why it is challenging for life insurers to remain profitable amid low interest rates while satisfying policyholders' rate-sensitive preferences with traditional products backed by low risk investments (Hartley, Paulson, and Rosen, 2017; Foley-Fisher, Narajabad, and Verani, 2020). The secular decline of interest rates in recent decades suggests that non-traditional activities in the life insurance industry are here to stay. If their potential for fragility serves to motivate implicit government protection, the contemporary interest rate regime may speak to the persistence of this subsidy as well.

[5]Captive reinsurers are more concentrated in jurisdictions providing the most favorable regulatory or tax environments, with notable domiciles including South Carolina, Vermont, and offshore locales such as Bermuda, Barbados, and the Cayman Islands (Koijen and Yogo, 2016).

1.3 Data

This chapter draws upon multiple data sources to construct a view of the U.S. life insurance sector from both accounting and market perspectives. Company-level accounting data come from the SNL Financial Institutions (FIG) Database. These data provide a range of insurance-specific items on a quarterly basis. My sample supplements the SNL FIG data with fundamentals from Compustat North America, which include general accounting items at quarterly frequency and a limited set of insurance items at an annual frequency. For both of these databases, I limit attention to observations recorded on a GAAP accounting basis. I further restrict my sample to life insurance companies with publicly-traded equity.

CRSP and IvyDB OptionMetrics supply equity market data. Key variables contributed by CRSP include market capitalizations and stock returns. Option-implied volatility data from OptionMetrics provide insight into investors' expectations about the distribution of equity value. The database reports historical option implied volatility for at-the-money options in its standardized option price files. For reference entities with a sufficiently broad set of liquid options, OptionMetrics also constructs a term structure of implied volatility for call and put options at various levels of moneyness. These data are provided in the database's volatility surface files. Option-related variables in my sample draw upon data from both of these components.

A final source provides credit market prices. For these data, I use historical credit default swap (CDS) spreads from the IHS Markit database. In particular, my sample is composed of CDS par spreads for senior tier debt at tenors ranging from six months to ten years. I focus exclusively on pricing of contracts with modified restructuring (MR) document clauses, the most heavily traded CDS referencing insurers in my sample.

Market prices in the databases described above are almost exclusively for claims on the publicly-traded parent company, rather than on businesses within an insurance group. A minor challenge is that this unit of analysis does not readily lend itself to clean distinctions between life and non-life insurance. Most public U.S. insures engage in variety of businesses that range from life insurance to medical and property-causalty (PC) insurance, among other activities.[6] In this chapter, the sample selection errs on the side of inclusion. First, I consider a company to be an insurance carrier if the first four digits of its NAICS code is 5241, if the first two digits of its SIC code is 63, or if its average insurance obligations as a share of total liabilities exceed 50%. Agencies and brokerages are excluded by filtering out records in which

[6]There are theoretical reasons why this empirical challenge is less an accident of the data than it is a fundamental of the insurance business. For example, Koijen and Van Nieuwerburgh (2020) make the point that there are benefits to combining life insurance with medical insurance as the latter can reduce mortality risk.

the first four digits of a company's NAICS code is 5242.

To identify life insurers in particular, I begin by requiring that at least 50% of average insurance liabilities are categorized by SNL FIG as belonging to life and health insurance lines. To distinguish carriers from those chiefly engaged in medical insurance, I further require that less than half of average premiums and average claims are for managed care policies. Attention is restricted to public life insurers with at least two years of non-missing accounting data. I require that carriers have nonmissing stock return and market capitalization data for at least half of their total observations during the 2001 to 2020 time period. Lastly, I exclude a small number of companies with erroneous ownership type or industry classification data. Together, these selection criteria yield an unbalanced panel of 51 public U.S. life insurers over the 2001–2020 period. A life insurer in this chapter is therefore a public insurance company with significant engagement in the life insurance business. Companies in the sample are reported across Tables 1.1 and 1.2, along with basic information about the availability of market data for each institution. The sample shares broad agreement with existing literature on public life insurers (e.g., Hartley, Paulson, and Rosen, 2017 and Koijen and Yogo, 2021). It should be noted that a non-negligible share of the companies in the cross-section are only present for a fraction of the 2001–2020 time period. Sample exits in these data are largely driven by mergers and acquisitions.[7]

To help address the "too big to fail" question, I partition the data into large and small insurance carrier subsamples. Carriers are designated as large if their total asset value is in the sample's top quartile for at least 60% of insurer-quarter observations. All other carriers are categorized as small. Note that this categorization rule is static; insurers do not move from one subsample to the other as total assets change. The 12 large life insurers according to this definition are listed in Table 1.1. Table 1.2 reports the remaining 39 small insurers. In the absence of a clear precedence for size categories in this context, the sample split employed here is driven by two considerations. The first is practical. Of the 51 insurers in the sample, only 18 have CDS prices. This is a fundamental limitation of the data, as liquid single-name CDS markets in general only exist for relatively large reference entities. The size split employed here ensures that each subsample has roughly half of the total CDS observations in the data. The second consideration takes account of the financial crisis experience. As detailed in the next section, the U.S. Treasury determined in 2009 that six life insurers would be eligible to receive TARP funding at the time. If "too big to fail" is to be viewed through a binary lens, and if present, it is plausible that this event provides some guidance on where insurer size might attain critical mass in this regard. For this reason, the splitting rule is chosen so that

[7]One of the more high-profile examples is John Hancock Financial Services, which merged with the Canadian life insurance company ManuLife in 2004.

the smallest of the TARP-eligible insurers qualifies as large.

1.4 TARP Guidance for Life Insurers

The financial crisis yields of a wealth of suggestive evidence that a "too big to fail" problem may exist in the life insurance sector. While the bailout of AIG is perhaps the most salient example, this period also provides a natural experiment about the breadth and depth of government support for the industry. During the crisis, a number of life insurance companies applied for emergency funding through the U.S. government's TARP Capital Purchase Program (CPP). This program granted the U.S. Treasury Department authority to provide emergency relief to struggling banks or thrifts through the purchase of equity stakes. A number of life insurers that applied to the program—in some instances while racing to meet the bank or thrift regulatory requirement—were granted access in May 2009. In the intervening period before the program's formal extension to life insurers, market reaction to a news event in April about their eligibility yields detailed insight into investor beliefs about the government's support for these institutions.

1.4.1 April 8 Announcement and Analysis Methodology

Life insurance companies that sought aid through the TARP CPP program submitted their applications by the November 14, 2008 deadline. In the months that followed, the arrival of a new presidential administration and a dearth of guidance from the U.S. Treasury Department cast doubts on whether authorities were willing to fund life insurers through TARP under the broad interpretation of banks and thrifts set by the previous administration. On April 8, 2009, the Treasury Department signaled its inclination to support struggling life insurers with the following statement:

> [T]here are a number of life insurers who met the requirements for the Capital Purchase Program because of their thrift or bank-holding company status and applied within appropriate deadlines. These are among the hundreds of financial institutions in the CPP pipeline that will be reviewed and funded as appropriate on a rolling basis. (Patterson, 2009)

Not only was this announcement greeted with significant attention (e.g., Lawder, 2009; Moyer, 2009; Patterson, Solomon, and Scism, 2009), it was a headline news item of the day in the financial press.[8] The immediate optimistic reaction of credit rating agencies in particular

[8] For example, an April 8 business day wrap-up article in Bloomberg outlining U.S. stock market performance carried the title "U.S. Stocks Gain as Shares of Life Insurers, Centex Lead Rally," (Nazareth, 2009; see also

highlights the unambiguous nature of the signal for life insurers' liabilityholders (see Business Wire 2009; Patterson, 2009).[9]

Before turning our attention to the capital market data, it is helpful to consider the subsequent course of events. Figure 1.1 provides a basic illustration of the timeline. A little more than one month after providing the above guidance, the U.S. Treasury announced on May 14, 2009 that six life insurance companies would be eligible for the TARP program: Allstate, Ameriprise Financial, Hartford Financial Services, Lincoln National, Principal Financial, and Prudential Financial. Hartford accepted $3.4 billion in TARP aid. Lincoln National accepted $950 million. The remaining four companies ultimately declined TARP funding in the weeks following the May 14 determination.

A few features emerge on visual inspection of life insurers' capital market prices around the time of the April 8 announcement. Figures 1.2 and 1.3 plot daily CDS spreads for contracts referencing 5-year senior tier debt for large life insurers, including all five of the six TARP-eligible carriers in the sample.[10] For a subset of these companies, the April 8 news appears to be met with pronounced reductions in CDS spreads. This is most apparent for Hartford, Lincoln National, and Prudential. Another subset consisting of Allstate, AIG, and Principal Financial exhibits no visual evidence of a readily discernible response. These graphs hint at significant cross-sectional heterogeneity in market reaction sensitivity to the TARP guidance. Apparent reaction in CDS prices to the May 14 eligibility determination is generally muted by comparison. It is likely that the latter news was more heavily anticipated by investors.

Figures 1.4 and 1.5 chart daily time series of stock prices juxtaposed with 91-day option-implied volatility (IV) for the same set of insurers during this period. While most of the features apparent in the CDS data carry over to the equity side, one difference is noteworthy. Though the direction of price movements at the time of the April 8 news is consistent with rising stock value and falling volatility, responses appear less pronounced in these data than they do in CDS spreads. This contrast between CDS and equity prices is most striking

BusinessWeek 2009).

[9] An April 9 press release from Fitch described the possibility of TARP support as "positive development" for the financial strength of eligible insurers that "could temper future downgrades," (Business Wire 2009). Weighing in on the news, Joel Levine, a Senior Vice President at Moody's, stated, "The receipt of TARP capital, assuming it's a meaningful amount, could mitigate against the downward ratings pressure," (Patterson, 2009).

[10] CDS spreads for Ameriprise Financial are not included in Markit's insurance sector pricing data. Moreover, because life insurance constitutes a relatively modest component of its business, this company does not meet the sample inclusion criteria delineated in the previous section. Other large insurers not appearing in the spring 2009 charts do not exist in the sample during this period. Brighthouse Financial and Voya Financial had not been spun-off from their original parent companies until 2016 and 2014, respectively. Equitable Holdings, a U.S. carrier in which French insurance company AXA holds a sizable stake, completed its IPO in 2018.

for Hartford, for which evolution of both stock and IV data near the event are visually unremarkable.

To formalize the analysis, I conduct event studies of life insurers' capital market responses to the U.S. Treasury's TARP guidance. Due to the importance of government rescues for financial institutions' liabilityholders in particular, the analysis primarily attends to an investigation of the CDS data. However, the section concludes with a brief examination of the equity side using event studies of stock returns and option IV. All event studies are performed at the daily frequency.

The approach for IV and CDS event studies is modeled closely on established methodology for stock returns prevalent in the literature (see, e.g., MacKinlay, 1997). Here I begin by describing the process for CDS spreads. Below I explain the minor differences in implementation for the other asset classes. For each insurer i and for each CDS contract tenor j, I estimate the following baseline time series model of differenced CDS spreads, $\Delta y_{i,j,t}$:

$$\Delta y_{i,j,t} = \beta_{i,j}\mathbf{X}_t + \varepsilon_{i,j,t}, \tag{1.1}$$

where $\mathbf{X}_t$ is a vector of aggregate economic state variables and subscript t indexes time. Coefficient vector $\beta_{i,j}$ is the immediate analogue of loadings in stock return factor models such as CAPM or the Fama-French model. The regression serves to orthogonalize CDS rates with respect to the ordinary economic factors that govern their price evolution. With this perspective in mind, error term $\varepsilon_{i,j,t}$ can be viewed as representing abnormal deviations in CDS rates. Let the event window, defined here by the interval $[T_0, T_1]$, be a period of time containing the event date. The cumulative abnormal difference (CAD) in CDS spreads over the course of this event period is thus computed by

$$\sum_{t=T_0}^{T_1} \varepsilon_{i,j,t}. \tag{1.2}$$

I use three-day event windows centered on event date April 8, 2009. The start date for my model estimation window is set 90 calendar days prior to initial event date T_0. The end date is likewise set 90 calendar days after final event date T_1. As is customary, the model estimation period leaves out the event window and an additional carveout period of N days adjacent to this window. The aim is to mitigate contamination of model parameter estimates from the event itself. Following prior literature (e.g., Acemoglu et al., 2016), I use carveout periods of $N = 30$ calendar days. Though it is more common to restrict estimation windows to pre-event time frames, joint use of pre- and post-event periods for the regression window is not without precedent (e.g., Copeland and Mayers, 1982; Agrawal, Jaffe, and Mandelker, 1992;

Ahern, 2009). In this chapter, the approach helps to manage one of the challenges entailed by studying an event set amid the backdrop of a financial crisis. Centering the regression window on the event date provides sufficient degrees of freedom to estimate richer multi-factor models, without using data from fall 2008—when crisis pathologies in asset pricing were at their most extreme.

Contrary to the case of stock returns, extant literature provides little guidance on the composition of factor vector $\mathbf{X}_t$ for event studies of single-name CDS spreads. Therefore, I perform the analysis using three different types of baseline models: an intercept-only model, a single-factor model, and a multi-factor model. The single-factor models use first differences in the Intercontinental Exchange (ICE) Bank of America U.S. corporate bond effective yield indices. These indices are published for different maturity ranges, including the 1–3 year range, the 5–7 year range, and the 7–10 year range, among others. For each single-factor study, I use the ICE corporate bond index with the range nearest to the CDS tenor being analyzed. In case of ties, the index with the higher maturity range is chosen (e.g., single-factor event studies of 5-year CDS contracts use the 5–7 year range for the ICE corporate index rather than the 3–5 year range). All multi-factor models employ the same six covariates: the CRSP value-weighted total stock market excess return, U.S. Treasury yields for the 3-month and 10-year maturities, ICE corporate bond effective yield indices for the 1–3 and 7–10 year maturity ranges, and the Chicago Board Options Exchange's VIX index. The latter five of these regressors are expressed in first differences. The covariates in the multi-factor model emphasize interest rate and corporate credit risks, the two most prominent risk exposures faced by life insurers. In addition to capturing general uncertainty, the VIX index can help to serve as a proxy for liquidity risk (Nagel, 2012). Adjusting price changes for this risk exposure may be particularly appropriate in this context given life insurers' role in liquidity provision as well the widespread liquidity shortages characteristic of the crisis.

Before examining the results for the CDS prices, an additional consideration is in order. The possibility that investors anticipated the April 8 announcement presents a potential challenge for the analysis. As Borochin, Celik, Tian, and Whited (2022) demonstrate, standard event study methodology struggles to evaluate even cross-sectionally relative differences in the impact of heavily anticipated events. There are two main reasons why investor anticipation in the present context may not have been trivial. First, insurers needed be registered as thrifts or bank holding companies as a prerequisite for TARP eligibility. A number of life insurance companies, including Hartford, Lincoln National, and Genworth, acquired regulated savings and loan institutions in fall 2008 to meet this requirement. It seems unlikely that these insurers would have endeavored to adopt these regulatory classifications without a reasonable expectation that doing so would open a door to government aid. Second, the history of business news reporting on the subject suggests that it may have been on investors' minds well in advance of April 8 (Wutkowski and Rucker, 2009; Tibken, 2009). Of particular note is a leak about the Treasury's planned statement the day prior, which was first reported in the Wall Street Journal (Moyer, 2009).[11]

Yet, by April 2009 investors had no shortage of reasons to doubt that aid was forthcoming. Though insurers had compelling reasons to anticipate government support when they submitted their TARP applications in late 2008, a New York Times report in February 2009 documented the growing suspicion that the subsequent arrival of a new presidential administration may have heralded a course reversal (Williams Walsh, 2009). It is not hard to imagine why policymakers might have hesitated on the concern that providing insurers access to a bailout program officially for banks and thrifts could invite a political backlash. On whether the industry at the time received guidance from the Treasury, president of the American Council of Life Insurers (ACLI) Frank Keating is quoted in the article as saying, "We have not had our phone calls answered," prefacing this with the statement, "As we say in the monastic life, it's the magnum silencium — the great silence." While there is no question that the April 8 announcement may have been partially anticipated,[12] the preponderance of news evidence makes it clear that the expectation of government aid was tempered by well-reasoned skepticism. It remains an inherent limitation of the event study, however, that high levels of investor anticipation cannot be ruled out with the techniques presented here.[13]

1.4.2 Results

Event study results for CDS spreads are reported across Tables 1.3 to 1.4. Each table presents all CAD estimates corresponding to a specific CDS tenor. Individual companies lining the rows of each table are first divided into two main groups which categorize insurers into those that were subsequently determined eligible for TARP aid in May 2009 and those that were not. It is important to bear in mind that an ineligible company need not have been deliberately denied TARP funding by the U.S. Treasury. This group also includes insurers that had not applied to the program in the first place, with MetLife among the most prominent examples

[11] The leak, which was apparently published during non-trading hours, also precludes intraday analysis as much of the stock price reaction appears to have been priced in at market open on April 8.

[12] A Wall Street Journal article from March 12, 2009 reported that ACLI was expecting a decision from the Treasury later that month, though a representative for the industry group stressed that they did not have a "clear picture of which way that clarification would tend to go" (Patterson and Scism, 2009).

[13] Not otherwise reported is a structural estimate in the spirit of Borochin et al. (2022) of the risk neutral probability of positive TARP news during the two-week run-up to April 8. The estimation, which uses a Merton default model informed by market capitalization, option prices, and CDS spreads, places this implied probability at 37%. Details are available from the author upon request.

in this regard.[14] A sub-categorization within each eligibility group highlights the insurers that would later be designed by regulators as systemically important financial institutions (SIFIs).[15] Row categories also indicate which of the TARP-eligible insurers were takers. Note that the list of companies in each table varies slightly as a few of the smaller insurers have an adequate time series of daily data to estimate the baseline models for some CDS tenors but not others. Columns in each table correspond to different baseline models for CDS spreads. Cumulative abnormal difference values are reported in percentage points.

I begin discussion of the results with 5-year CDS as this is generally the most heavily-traded contract tenor. These results are reported in Table 1.4. Two features are immediately striking. First, CAD estimates are consistently negative for nearly all insurers. Notwithstanding which companies experienced significant price impacts, this is consistent with the idea that investors generally perceived the announcement as positive news for the creditworthiness of the industry. Second, consistently significant responders tend to conform to two categories: TARP-eligible insurers that would later accept the government aid (takers), and future SIFIs. Among takers, the multi-factor model 5-year CDS abnormal difference estimates are -2.45% for Hartford and -6.17% for Lincoln National. A clear pattern across all CDS event study tables is that the multi-factor baseline model generally provides the most conservative estimates. For future-SIFIs MetLife and Prudential, multi-factor CAD estimates are -1.83% and -0.80%, respectively. AIG, which does not exhibit significant responses, stands out as an exception in the SIFI group. This is unsurprising. By early April 2009 AIG had already received government bailout funds in excess of $150 billion. It is unlikely that the marginal impact of the April 8 statement carried meaningful information about the U.S. government's willingness to support the insurer.

Table 1.3 presents the event study results for the 1-year CDS prices. The results mirror what is observed in the longer-term contracts, with the same group of TARP-takers and SIFIs exhibiting significant responses across all models. One salient difference, however, is that the abnormal deviations tend to be of considerably greater magnitude. Most noticeable is the double-digit size of the effect for Lincoln National; abnormal difference estimates for this insurer range from -14.38% to -13.49%. The multi-factor model 1-year CAD estimate for TARP-taker Hartford is -4.44%. The multi-factor abnormal difference estimates at this horizon are -2.17% for MetLife and -6.21% for Prudential. One other point of departure

[14] Genworth Financial stakes out an ambiguous middle ground, as the insurer was unable to achieve regulatory approval for their registration as a savings and loan holding company in time to meet the TARP application deadline. This was not disclosed to the public until late in the day on April 9, 2009 (Kardos, 2009).

[15] In this chapter I use the SIFI label in the regulatory sense of the term. The set of financial institutions that are systemically important in the true economic sense need not be equivalent.

from the 5-year results is that Reinsurance Group of America, a company that was neither TARP-eligible nor a SIFI, is a consistently significant responder at the 1-year maturity. This responsiveness is not reflected in any of this insurer's longer term CDS rates. Results for 2-year, 3-year, and 4-year CDS spreads are qualitatively similar.

An obvious explanation for the responses of TARP-takers is that investors recognized that these insurers' creditworthiness stood to benefit most from the expected aid package. The results may also reflect information about government support for these institutions beyond the scope of the TARP program—a possibility that, if true, must shed light on cross-sectional differences in the "too big to fail" premium. This interpretation is perhaps most relevant for the two SIFI responders. What is interesting about these insurers in the context of the event studies is that not only did Prudential not accept TARP funds once they were available, MetLife had not even applied to the program. This suggests that U.S. Treasury's announcement may have presented a more general signal about the government's willingness to protect these institutions from failure.

A potential concern is that the heterogeneity in responses may be driven to a degree by cross-sectional differences in CDS liquidity. This cannot be ruled out with available CDS data alone. Limited supplemental evidence on CDS liquidity is available from the Depository Trust and Clearing Corporation (DTCC), which publishes on its website aggregated transaction data for CDS trades cleared by the organization. All insurers in the large subsample with the exception of Principal Financial were among the DTCC's top 1,000 single-name reference entities between June 2009 and March 2010, the earliest period for which these data are available. Smaller life insurers in this list include Genworth Financial and Unum Group. For trades cleared by the DTCC during this period, CDS referencing MetLife, for example, had an average of nine trades per day and an average daily notional transaction volume of $75 million.[16] Among insurers in the sample appearing in the top 1,000 data, the company with the lowest trade activity was Unum Group. CDS referencing this insurer had an average of two trades per day and an average daily notional transaction volume of $12.5 million. It is worth bearing in mind that not all CDS trades are cleared through the DTCC. Viewed in conjunction with this information, the lack of consistently significant responses in the event studies for relatively large insurers such as AIG and Genworth provides some evidence that the severity of this potential issue is within manageable limits.

To put the economic significance of the results into perspective, I back out from these estimates CDS-implied lower bounds on risk neutral government bailout probabilities. In the most high-level sense, the price of CDS is governed by the likelihood that the contract's

[16] The average daily notional transaction volume for CDS referencing Hartford, Lincoln National, and Prudential were $50 million, $17.5 million, and $22.5 million, respectively.

default settlement will be triggered and the expected payout to contract holders in the event that settlement occurs. To be more precise, let PD be the annualized risk neutral probability of default for debt insured by the contract. This is sometimes referred to as the CDS contract's intensity. Let $LGD \in (0, 1]$ denote the loss given default, expressed as a fraction of the notional amount, incurred by debtholders in the absence of default insurance. If we view these components as constant throughout the contract's tenor the per annum CDS spread, y, may be approximated by the relation $y \approx PD \times LGD$ (Houweling and Vorst, 2005; see also Giglio, 2016; Neuberg et al., 2016).[17] Implicit here is the assumption that contract holders bear no counterparty risk. Note that if default hazard rates are expected to change throughout the contract term we might interpret this static treatment of PD as a rough average of a sequence of default probabilities for different discrete time horizons. This risk neutral decomposition is also silent on the role of credit risk term premia. For example, if the price of credit risk is upward sloping, risk neutral default probabilities will be increasing in contract tenor even if physical default probabilities are flat. For these reasons approximations of this kind need to be interpreted with greater caution when applied to longer CDS maturities.

One natural way to extend this approximation to an environment with government backstops is to reformulate the contract intensity such that

$$y = LGD \approx PD \times (1 - PB),$$

where PB is the conditional risk neutral probability that the government will intervene to avert bankruptcy in the event of imminent default. In other words, the CDS default settlement is triggered only in the event that the reference entity defaults *and* the government chooses not to bail out the entity or its debtholders. In this representation, I assume for simplicity that government protection enters solely by way of conditional bailout probability PB. Though this comes at the cost of abstracting from more nuanced ways in which government support might reduce credit risk, it provides an intuitive lens for viewing the total magnitude of the support. An added benefit is that this conceptualization facilitates comparison with the structural estimates in companion research I present in Chapter 2. Because PB is a component of a more generalized notion of default probability, the above term structure and price of risk considerations for PD are analogous.

If the government bailout probability is the only price component above that changes

[17] For intuition, consider a simple two-period model as illustrated in Giglio (2016). Suppose a CDS contract in this framework fully insures a zero-coupon risky bond with face value $1. The risky bond defaults with risk neutral probability PD and incurs fractional loss LGD to debtholders in default. In the absence of counterparty risk the expected payoff for this insurance under the risk neutral measure is $PD \times LGD$. This is the competitive market premium of the CDS contract.

between times t and $t+1$, it is easy to see that

$$
\begin{aligned}
&-\frac{\Delta \mathrm{Pr}_{t+1}(\textit{No Bailout} \mid \textit{Default})}{\mathrm{Pr}_t(\textit{No Bailout} \mid \textit{Default})} \\
&\qquad = \frac{PB_{t+1} - PB_t}{1 - PB_t} = -\frac{\Delta y_{t+1}}{y_t} \leq PB_{t+1}.
\end{aligned}
\tag{1.3}
$$

The abnormal deviations in the event studies are estimates of the degree to which Δy is driven by the shock in government support that occurred on April 8, 2009. Using Equation 1.3, it is therefore straightforward to translate these results into CDS-implied government bailout probabilities.

Table 1.5 reports lower bounds for implied government bailout probabilities for each of the main responding institutions. Estimates are recovered from event studies of all available CDS tenors between 6-month and 10-year horizons. This table presents bailout probabilities using abnormal CDS difference estimates from the multi-factor baseline model. For added reference, the table also reports bailout probabilities using raw changes in CDS spreads. Because the multi-factor estimates are most conservative, however, I restrict discussion to results from this baseline model unless otherwise noted. Surveying the results through this lens makes it clear that the economic magnitude of government support provided on April 8 was substantial. At the 1-year horizon, lower bounds on implied bailout probabilities for TARP-takers Hartford and Lincoln National are 29.2% and 26.1% respectively. For SIFIs, lower bounds at this horizon are 20.5% for MetLife and 37.3% for Prudential.

The term structure of implied government support is easy to assess with Table 1.5. These term structures are broadly downward-sloping as hinted at in the foregoing CDS results. The shortest and most long-term tenors contribute some nuance to this observation, however. Between six months and one year, the bailout bounds for MetLife rise from a (statistically insignificant) 8.5% to a (significant) 20.5%. MetLife's results also shed their significance for longer term CDS tenors at 7-year and 10-year horizons, though it is clear from the raw differences estimates that the extent to which this is the case is sensitive to the choice of baseline model. In contrast to MetLife, bailout probabilities are largest at the six-month horizon for the TARP-eligible insurers. Most salient is Prudential with a 43.5% implied bailout probability for this horizon. The characterization of the bailout term structure as downward-sloping encounters minor caveats from Lincoln National and Prudential. For these two insurers, the magnitude of implied support is greater at the 10-year horizon than it is at the 7-year horizon.

It is worth bearing in mind that these estimates are based on a CDS pricing approximation most readily suited for shorter contract tenors. The risk neutral nature of these probabilities also warrants particular emphasis. Physical bailout probabilities may be less downward sloping than my findings suggest if the price of credit risk were downward sloping in April 2009. Though disentangling credit risk premia is outside the scope of the simple analysis presented here, the structural evidence provided in Chapter 2 addresses this limitation.

I now turn brief attention to the equity side of life insurers. The event study methodology of stock and options data is nearly identical to the approach employed with the CDS spreads. Examination of the options data is performed on first differences in option-implied volatility. As is typical in the literature, analysis of daily equity price data uses stock returns in excess of the risk-free rate. As with the CDS data, option IV use three different baseline models: the single-factor baseline model for option IV uses the CBOE VIX index. Multi-factor baselines for both the options and stock returns data use the same covariates as the CDS multi-factor model described earlier in this section. For stocks returns, the CAPM model and the Fama-French three-factor model (FF3) are used for the other two baseline specifications.

Event study results for daily stock returns are presented in Table 1.6. It is immediately clear from the table that these results are at variance with the CDS data. Among the four main responding institutions in the CDS results, only Lincoln National stands out with unequivocal evidence for a rise in stock price. Under the FF3 model—the most conservative of three baseline models for stock returns—Lincoln National exhibits a 37.8% cumulative abnormal return (CAR). There is mild evidence for Prudential, which has CARs significant at the 10% level for the CAPM and multi-factor baseline models. Neither Hartford nor MetLife respond significantly in stock returns, though the signs are consistent with an increase in equity value according to the majority of baseline models. Interestingly, Principal Financial, a TARP-eligible non-taker, exhibits a large and significant abnormal return. This is in contrast to both the insurer's CDS and options results, the latter of which is discussed below.

The view from the options data bears more similarity to stock returns than the CDS spreads. Table 1.7 reports results for 91-day at-the-money put option IV. Point estimates for Hartford and MetLife are consistent with a decrease in stock volatility, but abnormal differences in IV for these institutions are not statistically significant. Contrary to the stock return results, the evidence for a significant decline in volatility for Prudential is more clear. According to the multi-factor baseline model, the cumulative abnormal differences in 91-day ATM IV for Lincoln National and Prudential are -14.2% and -16.2%, respectively. Results for other maturities and for out-of-the-money (OTM) put options (not reported) are qualitatively similar.

A message from the event studies in aggregate is that debt responded more strongly to the news about TARP support than did equity claims. This is consistent with the idea that the receipt of a government bailout presents a more double-edged sword to shareholders. Both the direct support of averting bankruptcy and the indirect effect of subsidizing borrowing costs are positives for equity value. But with aid comes supervision. Restrictions imposed on an institution by a government stakeholder, such as limits on executive compensation schemes, may introduce costs borne by equityholders in particular. A more blunt view is that when government protection arrives it is directed at debtholders and counterparties; any immediate benefits to shareholders are incidental. The contrast in results between the equity side and the credit side provides some evidence that investor expectations comport with this perspective.

To ensure that the foregoing results reflect government support directed at life insurers in particular, I conduct placebo tests on non-life insurance carriers. The placebo tests repeat the event study analysis for public U.S. non-life insurers in the Markit historical CDS insurance data. These non-life insurers, which primarily operate in property-casualty insurance, are listed in Table 1.8. Placebo test results for 1-year and 5-year CDS spreads are presented in Table 1.9. Of the 14 non-life carriers with CDS data during the period, only one has abnormal differences that are consistently significant across baseline models. Moreover, the significant 5-year abnormal change in spread for this institution (W. R. Berkley) is *positive*. Results for stock returns and 91-day ATM IV are reported in Table 1.10. For stock returns, there is some evidence for significant responses for MGIC Investment Corporation and Markel Corporation. Cumulative abnormal return estimates for MGIC range from 19.3% to 30.5% across baseline models, though the former estimate is insignificant. For non-life insurers with options data, the signs for abnormal differences in 91-day ATM IV are generally consistent with a decline in stock volatility. In no instances, however, are these IV estimates significant. Non-life insurer results for other CDS tenors and for other option categories are qualitatively similar. In general, the placebo tests are consistent with the idea that life insurer liabilityholders were the foremost beneficiaries of the Treasury's April 8 TARP guidance.

1.5 Conclusion

The financial crisis raised important questions about the presence, nature, and extent of a "too big to fail" problem in the U.S. life insurance industry. Answers to these questions are of pressing concern for policymakers amid evidence of the sector's growing engagement with private debt extension and participation in other non-traditional activities notable for their fragility and their scope for regulatory arbitrage. The reduced-form evidence from the financial crisis presented in this chapter sheds new light on the matter. Not only does it help quantify implicit government protection at a time when the support was of great historical significance, the facts documented here provide additional context on the "too big to fail" problem in the industry. The event studies show that with the April 8 announcement the U.S. Treasury provided substantive support to a subset of large life insurers. Institutions that derived the most significant benefit fall into one of two categories: insurers that would go on to accept TARP funding, and insurers that would later be designated by regulators as SIFIs. The analysis reveals that, compared to equity, credit prices were particularly sensitive to the news. Finally, the largely downward-sloping term structure of CDS-implied bailout probabilities implies that investors expected the value of government support to decline over time.

Table 1.1: Public U.S. life insurers, large life insurer sample

			Nonmissing capital market data		
	Insurer	Ticker	CRSP?	OptionMetrics?	Markit?
1	Allstate Corporation	ALL	Y	Y	Y
2	American International Group	AIG	Y	Y	Y
3	Brighthouse Financial, Inc.	BHF	Y	Y	–
4	Equitable Holdings*	EQH	Y	Y	–
5	Hartford Financial Services	HIG	Y	Y	Y
6	John Hancock Financial Services	JHF	Y	Y	–
7	Lincoln National Corporation	LNC	Y	Y	Y
8	MetLife, Inc.	MET	Y	Y	Y
9	Nationwide Financial Services	NFS	Y	Y	Y
10	Principal Financial Group, Inc.	PFG	Y	Y	Y
11	Prudential Financial, Inc.	PRU	Y	Y	Y
12	Voya Financial, Inc.**	VOYA	Y	Y	–

* Formerly, AXA Equitable Holdings, Incorporated, ** Formerly, ING U.S.

Large public U.S. life insurance companies for the 2001–2020 period. An insurer is defined as large if its time t total asset value is in the life insurance sample's top time t quartile for at least 60% of its quarterly observations.

Table 1.2: Public U.S. life insurers, small life insurer sample

	Insurer	Ticker	*Nonmissing capital market data*		
			CRSP?	OptionMetrics?	Markit?
1	Aflac Incorporated	AFL	Y	Y	Y
2	Alfa Corporation	ALFA	Y	–	–
3	American Equity Investment Life Hldg. Co.	AEL	Y	Y	–
4	American Financial Group, Inc.	AFG	Y	Y	Y
5	American National Insurance Company	ANAT	Y	Y	–
6	Assurant, Inc.	AIZ	Y	Y	Y
7	Athene USA Corporation	AUSA	Y	Y	–
8	CNO Financial Group, Inc.	CNO	Y	Y	–
9	Citizens Financial Corporation	CFIN	Y	–	–
10	Citizens, Inc.	CIA	Y	Y	–
11	Cotton States Life Insurance Company	CSLI	Y	–	–
12	Delphi Financial Group, Inc.	DFG	Y	Y	–
13	Erie Family Life Insurance Company	ERIF	Y	–	–
14	FBL Financial Group, Inc.	FFG	Y	Y	–
15	Fidelity & Guaranty Life	FGL	Y	Y	–
16	Financial Industries Corp.	FNIN	Y	–	–
17	Genworth Financial, Inc.	GNW	Y	Y	Y
18	Globe Life*	GL/TMK	Y	Y	Y
19	Great American Financial Resources, Inc.	GFR	Y	–	–
20	Horace Mann Educators Corporation	HMN	Y	Y	Y
21	Independence Holding Company	IHC	Y	–	–
22	Investors Heritage Capital Corporation	IHRC	Y	–	–
23	Jefferson-Pilot Corporation	JP	Y	Y	
24	Kansas City Life Insurance Company	KCLI	Y	–	–
25	Kemper Corporation**	KMPR/UTR	Y	Y	Y
26	MONY Group Inc.	MNY	Y	Y	–
27	Nassau Companies of New York	NCNY	Y	Y	–
28	National Western Life Group, Inc.	NWLI	Y	–	–
29	Penn Treaty American Corporation	PTYA	Y	–	–
30	Presidential Life Corp.	PLFE	Y	Y	–
31	Primerica, Inc.	PRI	Y	Y	–
32	Protective Life Corporation	PL	Y	Y	Y
33	Reinsurance Group of America, Inc.	RGA	Y	Y	Y
34	Southern Security Life Insurance Co.	SSLI	Y	–	–
35	StanCorp Financial Group Inc.	SFG	Y	Y	–
36	Symetra Financial Corporation	SYA	Y	Y	–
37	UTG, Inc.	UTGN	Y	–	–
38	Unum Group	UNM	Y	Y	Y
39	Vesta Insurance Group, Inc.	VTAI	Y	Y	–

* Formerly Torchmark Corporation, ** Formerly Unitrin, Incorporated.

Small public U.S. life insurance companies for the 2001–2020 period. An insurer is defined as large if its time t total asset value is in the life insurance sample's top time t quartile for at least 60% of its quarterly observations. All remaining insurance companies in the sample are defined as small.

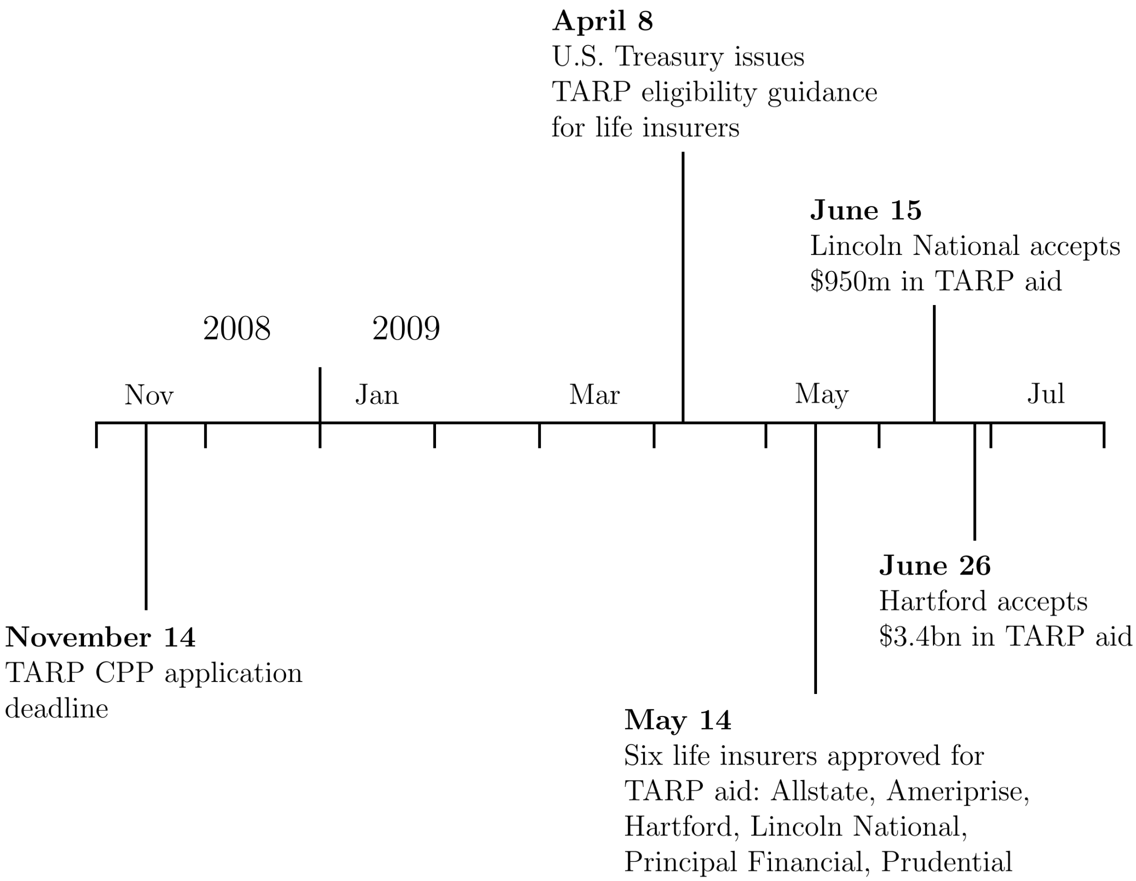

Figure 1.1: TARP aid for life insurers: main sequence of events. The deadline for TARP CPP program applications was November 14, 2008. On April 8, 2009, the U.S. Treasury Department issued a statement signaling that life insurer applications for TARP aid met the program's eligibility requirements and would be reviewed. The U.S. Treasury announced on May 14, 2009 preliminary approval for six life insurance companies: Allstate, Ameriprise Financial, Hartford Financial Services, Lincoln National, Principal Financial, and Prudential Financial. Hartford and Lincoln National accepted $3.4 billion and $950 million in TARP funding, respectively. The remaining four eligible companies declined the aid.

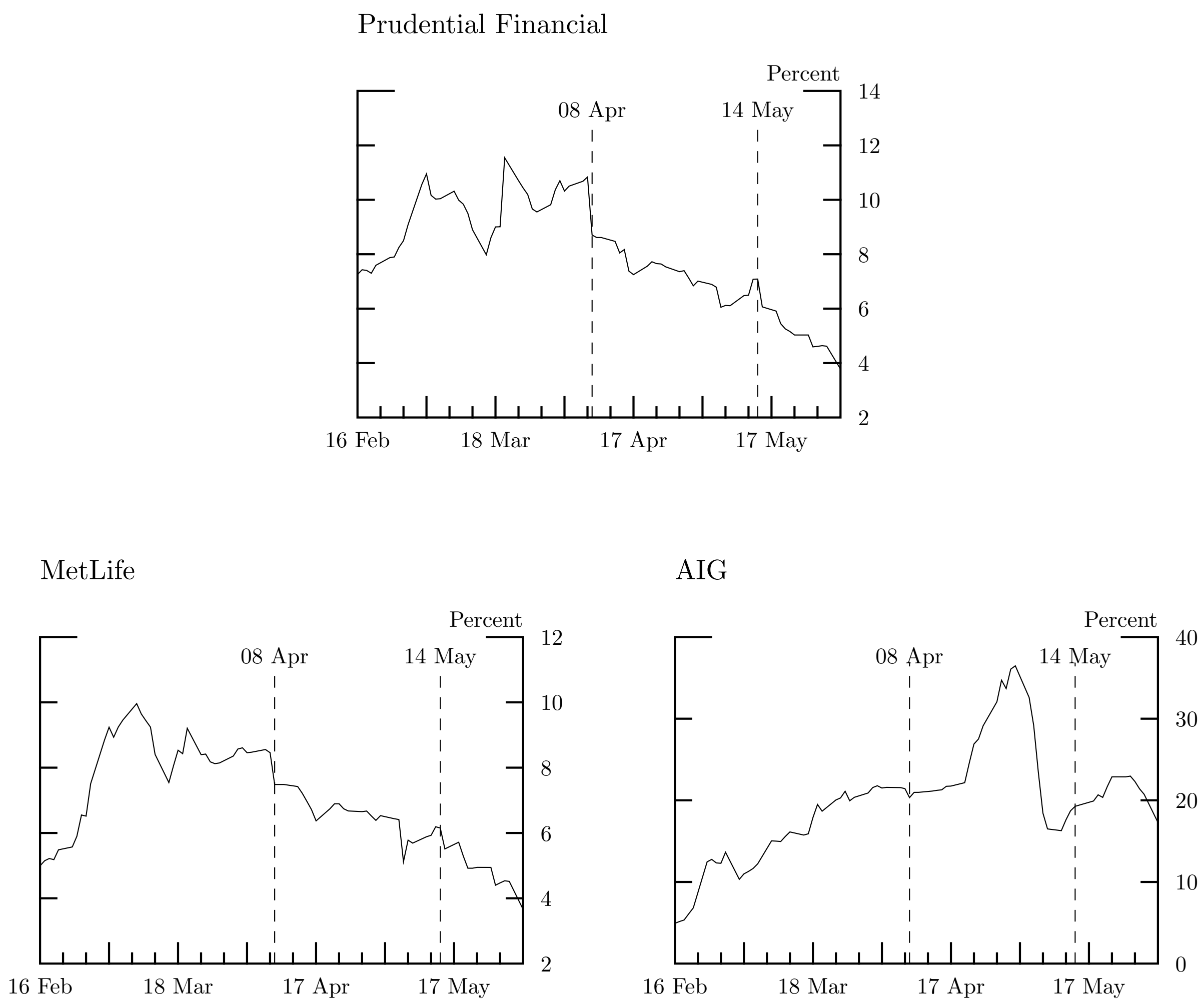

Figure 1.2: Daily 5-year senior tier CDS spreads, Feb. 2009 – May 2009, 1 of 2. The first date indicated by vertical line above, April 8, is the official statement date of the U.S. Treasury Department's guidance on TARP eligibility for life insurers. On May 14, the second date indicated above, the U.S. Treasury granted preliminary TARP funding approval to six life insurance companies. Values are expressed in percentage points.

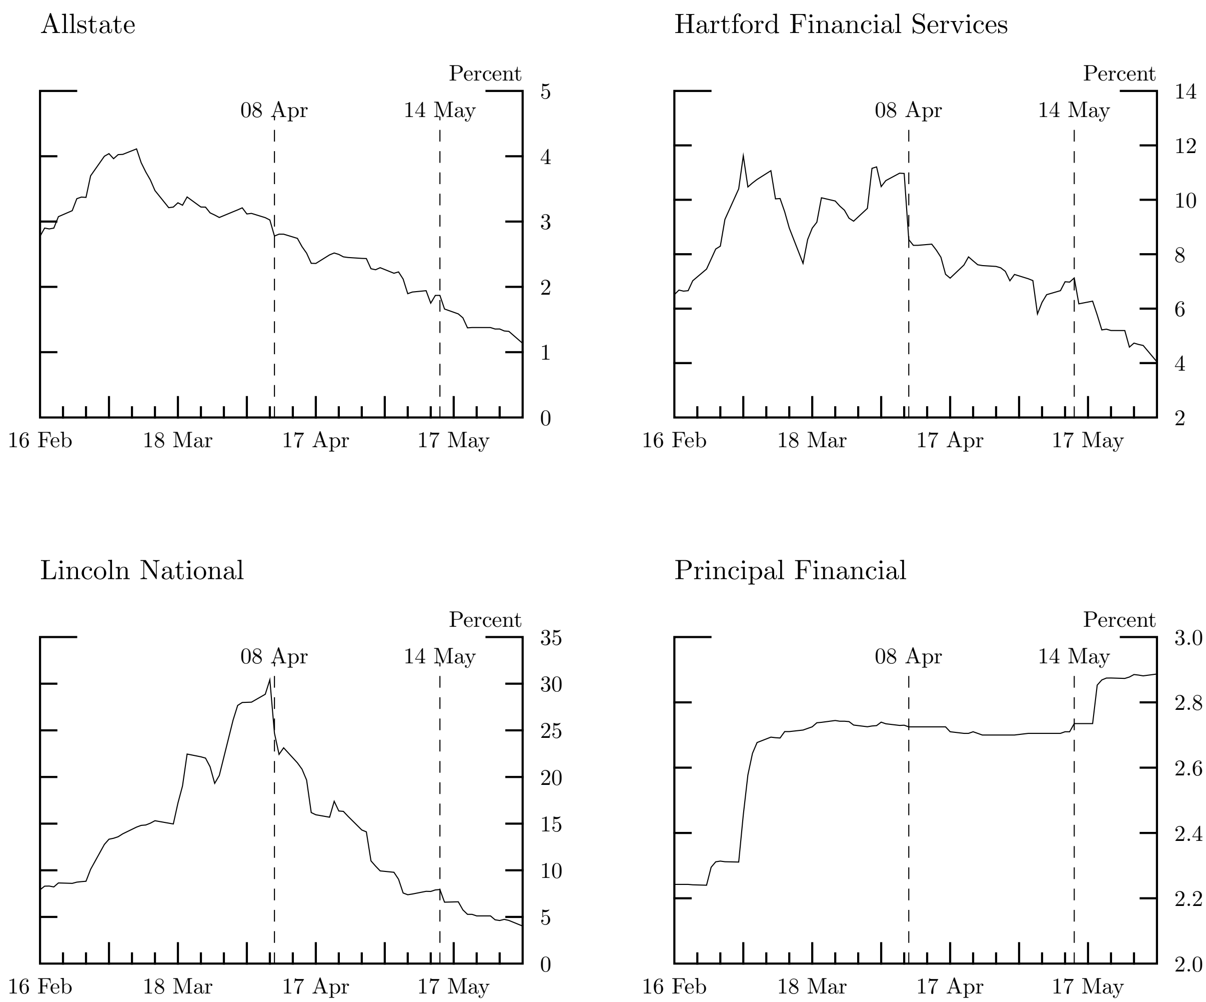

Figure 1.3: Daily 5-year senior tier CDS spreads, Feb. 2009 – May 2009, 2 of 2. The first date indicated by vertical line above, April 8, is the official statement date of the U.S. Treasury Department's guidance on TARP eligibility for life insurers. On May 14, the second date indicated above, the U.S. Treasury granted preliminary TARP funding approval to six life insurance companies. Values are expressed in percentage points.

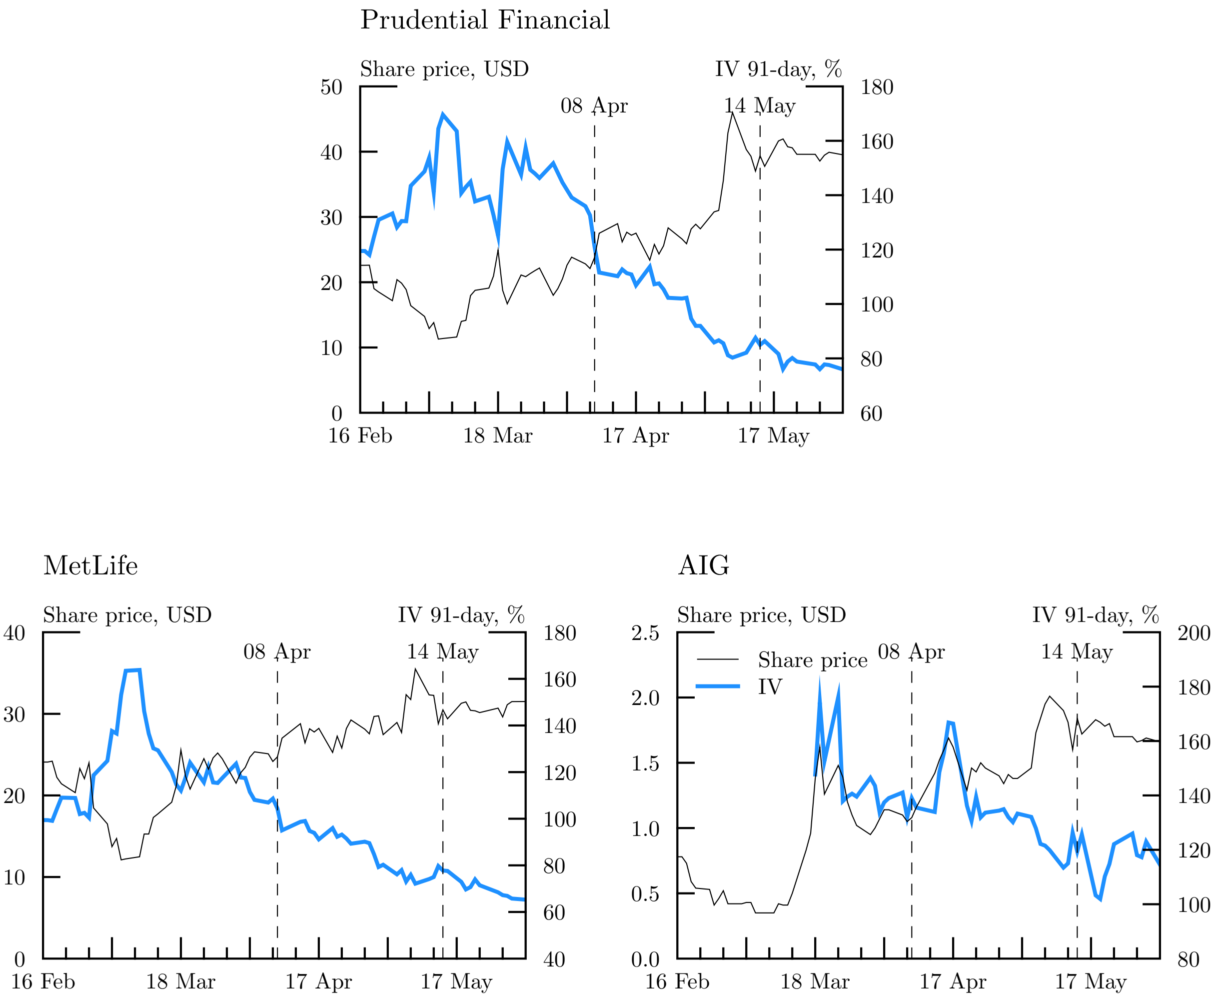

Figure 1.4: Stock price and option implied volatility (IV), Feb. 2009 – May 2009, 1 of 2. Stock prices are expressed in USD and represented by a thin black line. Option IV, expressed in percentage points and represented by a thick blue line, is for 91-day at-the-money options. The first date indicated by vertical line above, April 8, is the official statement date of the U.S. Treasury Department's guidance on TARP eligibility for life insurers. On May 14, the second date indicated above, the U.S. Treasury granted preliminary TARP funding approval to six life insurance companies.

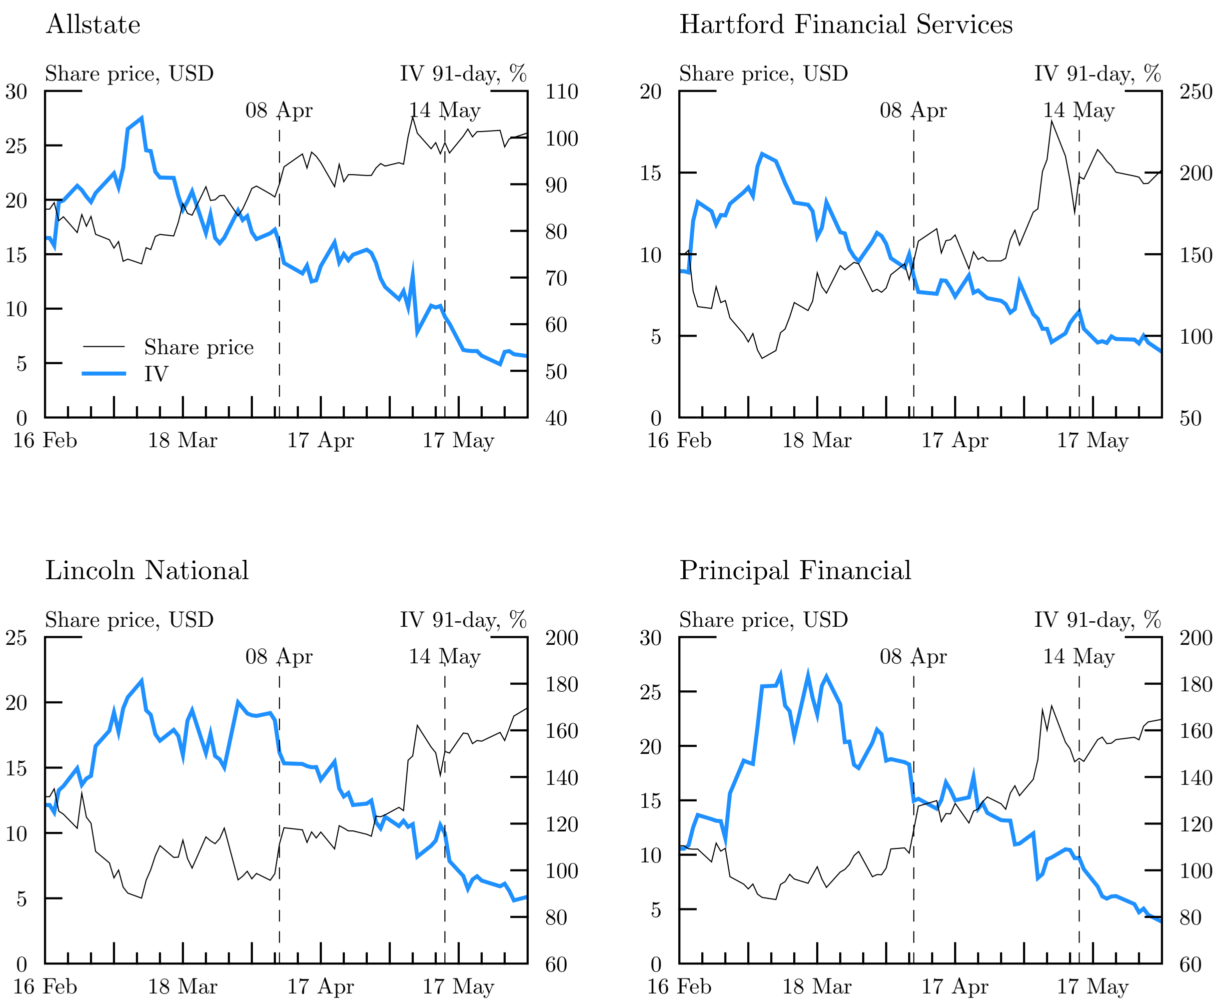

Figure 1.5: Stock price and option implied volatility (IV), Feb. 2009 – May 2009, 2 of 2. Stock prices are expressed in USD and represented by a thin black line. Option IV, expressed in percentage points and represented by a thick blue line, is for 91-day at-the-money options. The first date indicated by vertical line above, April 8, is the official statement date of the U.S. Treasury Department's guidance on TARP eligibility for life insurers. On May 14, the second date indicated above, the U.S. Treasury granted preliminary TARP funding approval to six life insurance companies.

Table 1.3: Cumulative abnormal first difference in 1-year CDS spread, April 7, 2009 – April 9, 2009

A. TARP eligible

	Insurer	(1)	(2)	(3)
SIFI				
Non-takers	PRU	-6.80***	-7.03***	-6.21***
Non-SIFI				
Takers	HIG	-4.98***	-4.97***	-4.44***
	LNC	-14.30***	-14.38***	-13.49***
Non-takers	ALL	-0.12	-0.04	0.11

B. TARP ineligible

	Insurer	(1)	(2)	(3)
SIFI	AIG	-2.85	-0.75	-0.24
	MET	-2.66***	-2.80***	-2.17**
Non-SIFI	AFG	-0.05	-0.08	-0.04
	AIZ	0.76	0.81	0.87
	GNW	-5.65	-4.54	-4.25
	RGA	-0.62***	-0.63***	-0.59***
	TMK	-0.17	-0.14	-0.27
	UNM	-0.28	-0.27	-0.14

This table reports 1-year CDS event study results for the April 8 TARP announcement. Cumulative abnormal first differences in CDS spreads, expressed above in percent, are computed over a three-day window centered on the event date. Insurers categorized as SIFIs would later be assigned the designation by regulatory authorities between 2013 and 2014. Takers are insurers officially determined TARP-eligible on May 14 that subsequently accepted the emergency funding. Results in column (1) are for the intercept-only baseline model, column (2) values are for the single factor model, and column (3) results are for the multi-factor baseline model. Statistical significance levels are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Table 1.4: Cumulative abnormal first difference in 5-year CDS spread, April 7, 2009 – April 9, 2009

A. TARP eligible

	Insurer	(1)	(2)	(3)
SIFI				
Non-takers	PRU	-2.19***	-2.23***	-1.83***
Non-SIFI				
Takers	HIG	-2.86***	-2.87***	-2.45***
	LNC	-6.64***	-6.66***	-6.17***
Non-takers	ALL	-0.31*	-0.31*	-0.15
	PFG	-0.03	-0.03	-0.02

B. TARP ineligible

	Insurer	(1)	(2)	(3)
SIFI	AIG	-0.78	-0.64	0.67
	MET	-1.24**	-1.25**	-0.80*
Non-SIFI	AFG	-0.00	-0.01	0.00
	AIZ	-0.27	-0.29	-0.23
	GNW	-4.16*	-3.92	-2.90
	HMN	-0.06	-0.05	-0.01
	RGA	-0.01	-0.01	-0.01
	TMK	-0.01	0.02	0.41
	UNM	-0.20	-0.20	-0.07

This table reports 5-year CDS event study results for the April 8 TARP announcement. Cumulative abnormal first differences in CDS spreads, expressed above in percent, are computed over a three-day window centered on the event date. Insurers categorized as SIFIs would later be assigned the designation by regulatory authorities between 2013 and 2014. Takers are insurers officially determined TARP-eligible on May 14 that subsequently accepted the emergency funding. Results in column (1) are for the intercept-only baseline model, column (2) values are for the single factor model, and column (3) results are for the multi-factor baseline model. Statistical significance levels are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Table 1.5: CDS-implied lower bounds on bailout probabilities from April 8 announcement

$$\Pr{}^{post}(Bailout \mid Default) \geq -\frac{\Delta \Pr{}^{post}(No\ Bailout \mid Default)}{\Pr{}^{pre}(No\ Bailout \mid Default)} = -\frac{\Delta y^{\ post}}{y^{\ pre}}$$

	Panel A: Multi-factor baseline model				Panel B: Raw differences (i.e., $\Delta y_t = \epsilon_t$)			
	TARP takers		*SIFI non-takers*		*TARP takers*		*SIFI non-takers*	
Term	HIG	LNC	MET	PRU	HIG	LNC	MET	PRU
6m	37.8***	29.1***	8.5	43.5***	38.7***	31.3***	12.1	45.1***
1y	29.2***	26.1***	20.5**	37.3***	30.4***	26.8***	23.0***	39.3***
2y	25.4***	25.4***	18.5***	24.5***	26.2***	26.2***	21.3***	26.5***
3y	25.9***	19.6***	12.4**	22.2***	27.2***	20.6***	14.8**	24.6***
4y	23.6***	19.6***	10.6*	18.8***	25.1***	20.7***	13.6**	20.8***
5y	22.3***	21.4***	9.3*	17.1***	24.2***	22.3***	12.5**	19.4***
7y	21.0***	21.5***	8.3	14.8***	22.7***	22.6***	12.2**	17.1***
10y	19.5***	26.8***	8.2	16.0***	21.0***	28.4***	12.1*	18.5***

Table 1.5 reports implied lower bounds for the probability of a government bailout recovered from CDS event studies of the U.S. Treasury's April 8 announcement. Holding the probability of default and expected debt recovery value fixed, the abnormal deviation in a CDS spread, y, observed over the event window may be interpreted as capturing the change in the probability of a government rescue conditional on default. Under these assumptions, $-\Delta y_{t+1}/y_t$ establishes a lower bound for this bailout probability. Results in Panel A use cumulative abnormal CDS differences from the multi-factor baseline model. Bailout probabilities in Panel B are recovered from raw differences in CDS spreads. Tickers HIG and LNC reference Hartford Financial Services and Lincoln National Corporation, respectively. The tickers MET and PRU correspond to MetLife, Inc. and Prudential Financial, Inc., respectively. TARP takers are insurance companies that accepted TARP funding after receiving approval from the U.S. Treasury on May 14, 2009. SIFI non-takers are insurers that did not accept TARP aid but were later designated systemically important financial institutions (SIFIs) by government authorities between 2013 and 2014. Values are reported in percentage points. Significance levels of the corresponding event studies are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Table 1.6: Cumulative abnormal stock returns, April 7, 2009 – April 9, 2009

A. TARP eligible				
	Insurer	(1)	(2)	(3)
SIFI				
Non-takers	PRU	11.62*	2.95	11.72*
Non-SIFI				
Takers	HIG	9.26	-1.39	11.52
	LNC	45.46***	37.81***	47.68***
Non-takers	ALL	7.13	3.51	6.22
	PFG	27.16***	22.17***	29.66***

B. TARP ineligible				
	Insurer	(1)	(2)	(3)
SIFI	AIG	10.19	0.64	7.73
	MET	0.86	-8.18	1.23
Non-SIFI	AEL	-1.10	-3.58	-2.38
	AFG	-0.42	-4.10	-1.38
	AFL	16.51	7.27	14.61
	AIZ	5.74	2.48	6.76
	AMIC	-4.87	-4.76	-7.97
	ANAT	4.03	-0.22	3.85
	CIA	-0.43	-6.98	-0.46
	CNO	-2.24	-13.55	4.20
	DFG	5.48	-3.12	5.30
	FFG	3.28	-5.53	3.09
	GNW	16.97*	10.05	15.08
	HMN	-0.28	-5.61	-0.55
	IHC	-2.08	-3.00	0.13
	KCLI	-10.46	-15.71**	-11.49
	NCNY	-4.43	-10.12	-4.14
	NWLI	7.22	2.85	6.30
	PL	5.69	0.08	9.47
	PLFE	6.51	0.98	5.98
	RGA	0.99	-1.78	0.85
	SFG	3.46	-0.77	4.25
	SNFCA	-6.39	-4.28	-14.87
	TMK	3.83	-1.41	3.69
	UFCS	-11.29***	-12.40***	-11.96***
	UNM	1.40	-5.04	2.61
	UTR	1.04	-1.97	0.08

This table reports stock return event study results for the April 8 TARP announcement. Cumulative abnormal stock returns, expressed in percent, are computed over a three-day window centered on the event date. Insurers categorized as SIFIs would later be assigned the designation by regulatory authorities between 2013 and 2014. Takers are insurers officially determined TARP-eligible on May 14 that subsequently accepted the emergency funding. Results are for the CAPM baseline model in column (1), the Fama French three-factor model in column (2), and the multi-factor model in column (3). Statistical significance levels are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Table 1.7: Cumulative abnormal first difference in 91-day at-the-money option-implied volatility, April 7, 2009 – April 9, 2009

A. TARP eligible

	Insurer	(1)	(2)	(3)
SIFI				
Non-takers	PRU	-26.67**	-19.87**	-16.18**
Non-SIFI				
Takers	HIG	-18.66	-12.27	-7.62
	LNC	-23.78**	-16.30*	-14.24*
Non-takers	ALL	-7.45	-3.93	-1.31
	PFG	-17.71	-13.13	-10.26

B. TARP ineligible

	Insurer	(1)	(2)	(3)
SIFI	AIG	-7.63	-6.65	-8.33
	MET	-14.34	-8.98	-4.85
Non-SIFI	AEL	-10.18	-11.89	-14.13
	AFG	3.30	1.09	1.81
	AFL	-10.93	-7.31	-2.79
	AIZ	-5.16	-3.87	-2.56
	CIA	1.27	1.83	1.34
	DFG	-3.33	-4.19	-5.71
	GNW	2.02	3.24	7.19
	HMN	-33.09	-42.59*	-46.00**
	PL	-4.89	-4.50	-1.58
	RGA	-1.84	-1.75	-1.24
	SFG	-1.03	-1.90	0.43
	TMK	-2.41	-2.05	-0.81
	UFCS	0.86	-3.33	-2.57
	UNM	-8.15	-3.58	-2.30
	UTR	-0.06	-4.17	-2.60

This table reports option implied volatility (IV) event study results for the April 8 TARP announcement. Cumulative abnormal first differences in 91-day at-the-money option IV, expressed in percent, are computed over a three-day window centered on the event date. Insurers categorized as SIFIs would later be assigned the designation by regulatory authorities between 2013 and 2014. Takers are insurers officially determined TARP-eligible on May 14 that subsequently accepted the emergency funding. Results in column (1) are for the intercept-only baseline model, column (2) values are for the single factor model, and column (3) results are for the multi-factor baseline model. Statistical significance levels are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Table 1.8: Public U.S. non-life insurers

			Industry code	
	Insurer	Ticker	NAICS	SIC
1	Ambac Financial Group, Inc.	AMBC	524126	6351
2	Berkshire Hathaway Inc.	BRK	999977	9997
3	Chubb Limited	CB	524126	6331
4	Cincinnati Financial Corporation	CINF	524126	6331
5	CNA Financial Corporation	CNA	524126	6331
6	HCC Insurance Holdings Inc.	HCC	524126	6331
7	Loews Corporation	L	524126	6331
8	Markel Corporation	MKL	524126	6331
9	MBIA Inc.	MBI	524126	6351
10	MGIC Investment Corporation	MTG	524126	6351
11	PMI Group, Inc.	PMI	524126	6351
12	Progressive Corporation	PGR	524126	6331
13	Radian Group, Inc.	RDN	524126	6351
14	Travelers Companies, Inc.	TRV	524126	6331
15	W. R. Berkley Corporation	WRB	524126	6331

Public U.S. non-life insurance companies in the Markit insurance CDS database during the 2001–2020 period. The non-life insurer sample is used strictly for the purposes of the event study placebo tests.

Table 1.9: CDS event study placebo test, non-life insurers, April 7, 2009 – April 9, 2009

	1-year CDS			5-year CDS		
Insurer	(1)	(2)	(3)	(4)	(5)	(6)
AMBC	-1.52	0.94	1.58	-0.35	-0.16	1.05
BRK	-0.28	-0.21	-0.10	-0.26	-0.25	-0.06
CB	-0.02	-0.01	0.02	-0.05	-0.05	-0.02
CNA	-0.24	-0.18	-0.06	0.00	0.01	0.11
HCC	–	–	–	0.01	0.01	-0.03
L	-0.04	-0.04	-0.01	-0.04	-0.04	-0.03
MBI	0.06	0.81	1.62	-0.68	-0.66	0.15
MKL	–	–	–	0.03	0.03	0.00
MTG	-0.04	0.37	1.35	0.24	0.23	0.77
PGR	-0.16	-0.15	-0.09	-0.05	-0.05	-0.01
PMI	2.82	3.15	4.29**	0.97	0.98	1.66*
RDN	1.67	2.11	3.21	0.18	0.16	0.97
TRV	0.06	0.08	0.10	-0.05	-0.05	-0.01
WRB	0.04	0.08	0.09	0.66***	0.66***	0.65***

This table reports event study placebo test results for the April 8 TARP announcement using CDS data from public non-life insurance companies. Cumulative abnormal first differences in CDS spreads, expressed above in percent, are computed over a three-day window centered on the event date. Columns (1)–(3) report results for 1-year CDS spreads and columns (4)–(6) provide results for 5-year CDS spreads. Results in columns (1) and (4) are for the intercept-only baseline model, values in columns (2) and (5) are for the single factor model, and results in columns (3) and (6) are for the multi-factor baseline model. Statistical significance levels are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Table 1.10: Stock return and option IV event study placebo test, non-life insurers, April 7, 2009 – April 9, 2009

	Stock returns			91-day ATM IV		
Insurer	(1)	(2)	(3)	(4)	(5)	(6)
AMBC	1.38	-8.51	-0.28	–	–	–
BRK	-0.13	-2.10	0.16	–	–	–
CB	-0.36	-2.22	-0.16	-4.01	-2.00	-1.22
CINF	2.99	-0.35	2.61	-1.08	-0.32	0.44
CNA	1.41	-2.37	2.64	-4.68	-4.66	-4.43
HCC	-0.17	-2.66	0.23	3.02	-0.45	-0.22
L	3.00	1.38	3.60	-1.54	-2.20	-1.39
MBI	-8.60	-12.45	-7.74	1.57	3.84	6.46
MKL	6.10*	5.07	6.08*	–	–	–
MTG	30.50**	19.34	28.72**	4.93	6.74	9.24
PGR	2.54	0.32	2.43	-3.03	-1.60	-0.85
PMI	9.11	2.60	2.50	–	–	–
RDN	3.62	-9.30	1.03	-9.23	-12.31	-10.20
TRV	-1.83	-3.30	-2.31	-4.24	-2.25	-1.85
WRB	1.30	-1.49	0.85	-3.05	-3.41	-3.84

This table reports event study placebo test results for the April 8 TARP announcement using stock and option data from public non-life insurance companies. Cumulative abnormal responses are computed over a three-day window centered on the event date. Columns (1)–(3) report abnormal stock returns and columns (4)–(6) provide cumulative abnormal first differences in 91-day option implied volatility (IV) of at-the-money options. Stock return results are for the CAPM baseline model in column (1), the Fama French three-factor model in column (2), and the multi-factor model in column (3). Option IV results are for the intercept-only model in column (4), the single-factor model in column (5), and the multi-factor model in column (6). Statistical significance levels are indicated by *, **, and *** for 10%, 5%, and 1%, respectively.

Chapter 2

Risk-taking and "Too Big to Fail" Subsidies in the U.S. Life Insurance Industry

2.1 Introduction

In companion research presented in Chapter 1, I document evidence of implicit government guarantees in the U.S. life insurance industry during the 2008–09 financial crisis. The findings, which come from a natural experiment about life insurers' eligibility for the TARP program, provide a snapshot of government support for these institutions at a specific point in time. In this chapter I expand upon these results by addressing two questions left open by the companion analysis. First, is the implicit protection afforded to large life insurers sufficient to impact capital market prices over significant time frames? Second, to what extent are government backstops for the industry associated with moral hazard in risk-taking practices?

To address the above questions, I structurally estimate the time series properties of implicit bailout probabilities for public U.S. life insurance companies between 2001 and 2020. The estimates imply that large life insurers benefit from significant implicit guarantees, which are associated with moral hazard in risk-taking practices. Subsample analysis of large and small life insurers supports the interpretation that this represents a "too big to fail" subsidy. My findings also reveal that the implicit protection exhibits strong dynamic properties, consistent with the view that the government's propensity to bail out life insurance companies changes as economic and political conditions evolve.

The structural estimation employs a partial equilibrium model of a life insurance sector with emergency bailouts. In the model, agents observe time-varying government rescue probabilities, which are reflected in the market prices of life insurers' equity and liability obligations. Rescues take the form of equity injections sufficient to prevent an otherwise imminent bankruptcy. Shareholder value-maximizing insurers make decisions about leverage and risk exposure. The estimation recovers the time series properties of bailout probabilities by mapping this model to a rich array of capital market prices and insurance-specific accounting data. I find that for large life insurers the unconditional physical probability that the government will bail out a failing institution is 9.1% between 2001 and 2020. Corroborating the importance of dynamics hinted at in Chapter 1, the unconditional standard deviation of this rescue probability is 10.3%.

This sizable dispersion in support is important because consideration of the time-invariant aspect of implied bailout probabilities in isolation understates the depth of protection investors expect—particularly given that default risk is highest in states of the economy when the support is likely most elevated. Conditional on a bailout, the expected aid package for large insurers is 1.9% of general account assets. For perspective, 1.9% of Prudential Financial's general account assets as of 2020 is $11.7 billion. Bailout quantities also exhibit considerable variation. A conditional two standard deviation bailout for Prudential, for example, would amount to $42.3 billion in 2020.

An important advantage of the structural estimation presented here is that it can account for the moral hazard entailed by the implicit backstop. In a counterfactual analysis, I find evidence of modest excessive risk-taking practices on average. Among large insurers the per annum probability of severe financial distress falls from 2.3% under baseline parameter estimates to 2.1% in the absence of implicit guarantees.[1] Not only is a quantification of moral hazard informative in itself from a policy perspective, it renders the bailout probability estimates more robust to an important form of misspecification. To see why, consider what happens to the price of a five-year CDS contract when bailout probabilities rise. *Ceteris paribus*, the five-year CDS spread will fall because the insurer is less likely to default during the period in which the contract is in effect. This *ceteris paribus* qualifier is critical, however. Shareholders might respond to the protection by taking on greater risk. If this additional risk can outlive the shock to bailout probability, the net effect on a five-year CDS price is less obvious. By taking this into account, the estimated bailout probabilities are less contaminated by the muddying effect that moral hazard introduces.

An additional advantage is that the model can facilitate cost-benefit analysis of regulatory policy intended to address the "too big to fail" problem. For example, consider the taxpayer burden of MetLife's implicit government subsidy in comparison with the insurer's compliance costs for macroprudential regulation under the purview of Dodd-Frank. Assuming a 2%

[1] The limited scope for moral hazard is somewhat surprising. This is not due to the transient nature of the implicit protection per se; moral hazard is at least as limited if bailout probabilities are fixed at their unconditional expectation. Demand for life insurance liabilities may provide an explanation. Interest rates are of particular importance for life insurer liability holders (e.g., Hartley, Paulson, and Rosen, 2017; Foley-Fisher, Narajabad, and Verani, 2020). From the perspective of policyholders and lenders, it is possible that interest rate considerations often dwarf bankruptcy concerns. The muted effect of implicit guarantees on life insurer incentives may therefore reflect a second order status for credit risk in their borrowing costs.

annual discount rate and abstracting from asset growth, the expected value of government support over a twenty year horizon for the insurer as of mid-2016 is \$2.9 billion according to the model. This puts into perspective the \$1.4 billion Naubert and Tesar (2019) estimate for the compliance cost of the enhanced oversight borne by MetLife's shareholders should the regulation take full effect.[2] In general, the results suggest that cost-effective macroprudential regulation for life insurers is within policymakers' reach if it achieves sufficient reduction in the likelihood of future bailouts.[3]

A number of features of the reduced-form TARP evidence from the 2008–09 financial crisis presented in Chapter 1 are in line with the structural findings. Not only are the magnitudes of government support in broad agreement between the two analyses, the downward-sloping term structure of the implicit subsidies is consistent with the prominent dynamics found in the structural results. The TARP evidence also supports the model implication that debt is more sensitive to the implicit protection than equity—a feature that serves an important role in identifying structural parameters. Additional external validity is provided by analysis of a small life insurance carrier subsample. If the structural methodology offers a sound glimpse at "too big to fail" subsidies, it follows that parameter estimates should imply much more modest support for small life insurers. Indeed, the results indicate that for small life insurance companies the probability of a bailout in the event of default has an unconditional mean and standard deviation of 4.1% and 2.3%, respectively.

A strand of related literature investigates implicit government guarantees of banks and other firms in the financial services industry. Tsesmelidakis and Merton (2012) use a Merton default model calibrated on pre-crisis data to estimate the value of implicit guarantees for financial intermediaries throughout the crisis and post-crisis period. They find that the resulting wealth transfers to the entire financial sector, including insurers, amount to \$129.2 billion for equityholders and \$236.1 billion for creditors. Their findings also indicate that the implicit government support for banks was more extensive and more persistent compared to insurance companies. A limitation of this study is that it does not account moral hazard. As noted above, this can understate the extent of implicit support—particularly for instances in

[2]Though expected bailout expenditures in perpetuity present a more direct analogue to this compliance cost, the twenty year horizon value is less sensitive to long-run interest rate and asset growth assumptions. The approach also errs on the side of a more conservative expected bailout estimate.

[3]It could be argued to the contrary that the value of bailouts attributable to moral hazard per se is insufficient to warrant macroprudential oversight. For example, consider a counterfactual in which bailouts conform to the baseline distribution and liabilities are priced accordingly, yet life insurers act under the erroneous belief that the government committed to a no-bailout regime. In this scenario, assuming a 2% risk-free interest rate and no asset growth, the 20-year expected cost of bailouts for all large life insurers falls by only \$247 million compared to the baseline estimate. Nonetheless, even if bailout costs unrelated to moral hazard are excluded from public interest, the results suggest that there are at least some states of the world in which regulation-guided risk mitigation may be reasonable and cost-effective.

which these backstops are heavily dynamic. Kim (2016) structurally estimates time-invariant probabilities of government bailouts for large and small banks from Call Report data. The paper further investigates the associated moral hazard in terms of leverage and bank lending decisions. Other noteworthy contributions in this line of literature include Acharya, Anginer, and Warburton (2016), Atkeson et al. (2019), and Berndt, Duffie, and Zhu (2019).

The remainder of this chapter is organized as follows. In Section 2.2 I describe the model. Section 2.3 discusses the estimation methodology and parameter identification. In Section 2.4, I present the results of the structural estimation. I conclude in Section 2.5.

2.2 Model

2.2.1 Environment and Overview

The model that serves as the basis for the structural estimation that follows is of an insurance sector in partial equilibrium. Time is discrete. Life insurance carriers maximize the discounted present value of distributions to shareholders over an infinite horizon. Capital markets are competitive and efficient. Equityholders, creditors, and policyholders are risk averse and share the same pricing kernel.

Each period, insurers choose their leverage and their level of exposure to an insurance cash flow shock. Leverage governs exposure to two types of obligations: policyholder liabilities and ordinary debt, the latter of which is subordinate to the former. The cash flow shock perturbs what this chapter refers to as the core life insurance spread, defined as the difference between investment income and net payouts to policyholders. Composed of idiosyncratic and systematic components, this shock enters insurer cash flows as an asymmetric distribution to capture the higher moment risks associated with the insurance business.

Liabilities issued by insurance carriers are subject to liquidity default risk. If the net worth default threshold is crossed at time t, the government intervenes for the distressed insurer with probability $\phi_t \geq 0$. In the absence of an intervention, equityholders are wiped out and creditors recover the assets and internal funds that remain after servicing insurance obligations. Should a government intervention occur, an equity injection is provided to ensure that insurance claimants and debtholders are made whole, thereby allowing the insurance carrier to remain a going concern. As with the cash flow shock, government bailout probabilities follow an exogenous process consisting of idiosyncratic and systematic components. Insurers, capital market participants, and policyholders share rational expectations about these probabilities.

2.2.2 Life Insurance Liabilities and Income

The model abstracts from insurer size with the assumption that general account assets are fixed at one. Stock and flow variables are thus expressed as shares of general account assets.[4] Except where otherwise noted, all variables in the model are for carrier i at time t. For ease of notation, I suppress i subscripts.

Insurance policies are modeled as one-period obligations. Let λ_t denote the accounting value of insurance liabilities, expressed as a fraction of general account assets, in effect at time t. Policies generate premium revenue and incur benefit and claim expenses during the period in which they are in effect. The "float" from collected premiums is invested in assets which earn investment returns. The difference between investment returns and net payouts to policyholders is the fundamental source of income in the life insurance business. I refer to this net accounting item as the core spread, where the ad hoc prefacing term emphasizes distinction with the CDS spreads referenced extensively in my description of the estimation that follows.[5] In the model, the core spread at time t, ψ_t, before capital regulation costs is

$$\psi_t(.) = \big([\sigma_t i_t(\epsilon_t, z_t) - \gamma_t(\sigma_t, \sigma_{t+1})] - 1 \big)\lambda_t + p_t(.)\lambda_{t+1}, \tag{2.1}$$

where $i_t(\epsilon_t, z_t)$ is the insurance income risk function, $\gamma_t(\sigma_t, \sigma_{t+1})$ denotes risk management costs, and p_t is the fully informative time t market price of insurance liabilities maturing at time $t+1$. As with insurance liabilities, note that the core spread is represented as a fraction of general account assets. Exogenous state variables ϵ_t and z_t represent the insurance cash flow shock and standardized aggregate productivity, respectively. Aggregate productivity is modeled separately for the purposes of the asset pricing kernel described in Section 2.2.4 below. Endogenous state variable $\sigma_t > 0$ governs the insurance risk exposure. Both policy liabilities, λ_t, and risk exposure, σ_t, in effect at time t are governed by the carrier's decision-making in the previous period.

The variable σ_t may be interpreted in a broad sense as the beta on a life insurance portfolio that generates return i_t. This may also be construed to limited degree as capturing the "shadow leverage" that accumulates off-balance sheet. In general, however, the model is

[4] The general account is in line with more traditional notions of a firm's assets, whereas separate accounts are similar to the assets under management measure used in the asset management industry. Assets in the general account can be invested and used as collateral at the insurer's the discretion, within regulatory constraints. Separate account assets represent investments managed by insurance carriers on the behalf of clients. An important example of assets grouped under this secondary category in the life insurance sector are those belonging to variable annuity accounts.

[5] The core spread is a flow item in the model. Note, however, that because general account assets are fixed at 1 all accounting items are normalized accordingly. Therefore the core spread may also be viewed as a net interest rate which is expressed in terms of general account assets.

agnostic about the source of risk represented by an increase in σ_t. This is motivated in part by data limitations. It is also in the service of maintaining a tractable state space given life insurers' extensive range of activities, of which Chapter 1 provides an overview.

Before unpacking the components of the core spread further, note that Equation 2.1 is arranged to emphasize the liability nature of insurance policies. Functions i_t and γ_t in the bracketed term above are what distinguish liability λ_t as insurance technology. In the case that i_t and γ_t are fixed at zero, λ_t can be interpreted as a one-period debt obligation and Equation 2.1 is net financing cash flow from debt. To expand upon the debt analogy, insurance policies in general could be viewed as debt contracts that pay a stochastic interest rate to their liabilityholders. In unconditional expectation, this interest rate is lower than the risk-free rate and often negative as risk averse policyholders are willing to accept negative expected returns for investments that pay out in states where their marginal utility is highest.

The insurance cash flow shock consists of a systematic industry component, $\epsilon_{s,t}$, common across all insurance carriers at time t, and an idiosyncratic component, $\epsilon_{i,t}$. In particular, let $\epsilon_t = \epsilon_{i,t} + \epsilon_{s,t}$, where the laws of motion for the idiosyncratic and industry components are

$$\begin{aligned} \epsilon_{i,t+1} &= \rho_{\epsilon,i}\epsilon_{i,t} + u_{\epsilon,i,t+1}, \quad u_{\epsilon,i,t} \sim \mathcal{N}\left(0, \iota[1-\rho_{\epsilon,i}^2]\right), \\ \epsilon_{s,t+1} &= \rho_{\epsilon,s}\epsilon_{s,t} + u_{\epsilon,s,t+1}, \quad u_{\epsilon,s,t} \sim \mathcal{N}\left(0, [1-\iota][1-\rho_{\epsilon,s}^2]\right), \end{aligned} \tag{2.2}$$

where $\iota \in [0, 1]$ is a parameter that governs the relative weight of risk idiosyncratic to the insurer. Disturbance terms $u_{\epsilon,i,t}$ and $u_{\epsilon,s,t}$ are independently distributed. Under the additional assumption that the idiosyncratic and systematic processes above share the same level of persistence, ϵ_t can be expressed as

$$\epsilon_{t+1} = \rho_\epsilon \epsilon_t + u_{\epsilon,t+1}, \quad u_{\epsilon,t} \sim \mathcal{N}\left(0, 1-\rho_\epsilon^2\right), \tag{2.3}$$

where $\rho_\epsilon = \rho_{\epsilon,i} = \rho_{\epsilon,s}$ and $u_{\epsilon,t} = u_{\epsilon,i,t} + u_{\epsilon,s,t}$. For the purposes of the life insurer's decision-making, state variable ϵ_t is a sufficient for both idiosyncratic and industry components of the institution's cash flow risk. In line with the production-based asset pricing literature, aggregate productivity also follows an AR(1) process with normal disturbance terms. The specification of the process is detailed in Section 2.2.4. Note that because z_t is the standardized version of aggregate productivity, its unconditional distribution is $\mathcal{N}(0, 1)$.

The insurance income risk function $i_t(.)$ transforms standard normal shocks ϵ_t and z_t into an asymmetric perturbation of the core life insurance spread. This is in line with existing theory established in Froot (2007), which emphasizes the asymmetry of payoffs in the insurance industry. If x_t is the scaled convex combination defined by $x_t \equiv [\omega\epsilon_t - (1-\omega)z_t]/\bar{s}$, with $\omega \in [0,1]$, income risk may be parameterized succinctly in the model by

$$i_t(\epsilon_t, z_t) = i_t(x_t) = \theta - \frac{e^{\eta x_t} - \mathbb{E}e^{\eta x_t}}{\text{Var}\left(e^{\eta x_t}\right)^{1/2}}. \tag{2.4}$$

That is, $i_t(x_t)$ is a shifted (negative) lognormal distribution with unconditional mean $\theta \geq 0$ and unconditional variance one. Shape parameter $\eta > 0$ governs the level of higher moment risk. The extent to which the exogenous component of the core spread is driven by aggregate productivity is controlled by parameter ω. For notational convenience, the scaling term $\bar{s}$ is chosen so that the variance of x_t is one; that is, $\bar{s} = \sqrt{\omega^2 + (1-\omega)^2}$.

Risk management costs have two components. The first is risk carrying expenses which make it costlier for insurers to maintain higher levels of risk exposure. In particular, I assume that risk carrying costs are convex and increasing in σ_t. The second component is risk adjustment costs which make it costly for insurers to change their level of exposure. I parameterize risk management costs in the model as follows:

$$\gamma(\sigma_t, \sigma_{t+1}) = \zeta\sigma_t^2 + \alpha_\sigma \left(\sigma_{t+1} - \sigma_t\right)^2. \tag{2.5}$$

The first quadratic term above represents risks carrying costs, which are governed by parameter $\zeta > 0$. The second term represents quadratic risk management costs with corresponding parameter $\alpha_\sigma > 0$. The risk management cost function captures in reduced-form a number of expenditures related to the management of risks among policies and investment assets. One interpretation of risk carrying costs is that they follow from state regulatory limits on risk-taking. Another, not mutually exclusive, interpretation is that they represent unmodeled nuances of policyholder demand. Risk adjustment expense can be viewed as arising in part from the costs entailed by re-balancing a large portfolio of often illiquid securities.

The insurer's operating income is the difference between the core spread, ψ_t, and underwriting costs, c_t. Underwriting expenditure, as a fraction of general account assets, is modeled as a quadratic adjustment cost for insurance liabilities,

$$c_t(\lambda_t, \lambda_{t+1}) = \alpha_\lambda \left(\lambda_{t+1} - \lambda_t\right)^2, \quad \alpha_\lambda \geq 0. \tag{2.6}$$

The model emphasizes the core spread at the expense of a richer treatment of underwriting expense. Due to the long-term nature of most life insurance products, costs arising from underwriting activities tend to be of second order importance in the life insurance industry. In the data, policy liabilities evolve in an incremental fashion over time that is largely devoid of "lumpy" changes. This is line with optimal policy subject to convex adjustment costs such as the kind assumed in Equation 2.6.

2.2.3 Liability Pricing and Government Intervention

In addition to policy liabilities, insurers issue debt. Debt is one-period in the model, with d_t denoting principal value of debt maturing at time t expressed as a fraction of general account assets. Let q_t represent the time t price of debt maturing the following period. Debt is subordinate to insurance liabilities. Note that this implies that the data analogue to debt in the model is ordinary balance sheet debt. Off-balance sheet borrowings such as funding agreement backed securities are unmodeled, though captured indirectly by insurance liabilities, λ_t, and shadow leverage entailed by σ_t.

Before continuing on to examine the pricing of debt and insurance liabilities, note that distributions to shareholders, e_t, before issuance costs is

$$e_t(.) = \psi_t(.) - c_t(.) - d_t + q_t(.)d_{t+1}. \tag{2.7}$$

In words, distributions before issuance costs are the core spread less underwriting expenditure plus net financing cash flow from debt. It follows from the composition above that e_t is modeled as a fraction of general account assets. Equity issuance is not costless. I assume that raising new equity capital incurs linear issuance cost $\chi > 0$. With issuance costs, net distributions to shareholders is

$$e(.)\Big(1 + \chi \mathbf{1}_{\{e(.)<0\}}\Big),$$

where indicator function $\mathbf{1}_{\{e(.)<0\}}$ above is one if pre-issuance cost distributions, e_t, are negative and zero otherwise.

In the model, bankruptcy occurs when insurers have insufficient net worth to back policyholder and credit obligations. Let $\widetilde{\psi}_t$ denote the core spread *before* adjustment costs for risk exposure. Likewise, call underwriting expenditure before policy adjustment costs $\widetilde{c}_t$. Absent a government rescue, the life insurer is in default if the net worth condition below is violated:

$$1 + \widetilde{\psi}_t(.) - \widetilde{c}_t(.) \equiv h_t(.) \geq \lambda_t + d_t. \tag{2.8}$$

That is, the bankruptcy is triggered in the event that general account assets, 1, and internal funds before adjustment costs, $\widetilde{\psi}_t - \widetilde{c}_t$, are not sufficient to cover general account liabilities $\lambda_t + d_t$. In default, $h_{\mathcal{F},t}$ resources, defined by $h_{\mathcal{F},t}(.) \equiv \xi + \widetilde{\psi}_t(.) - \widetilde{c}_t(.)$, are recoverable by liabilityholders. Parameter $\xi \in [0, 1]$ represents general account assets after fire sale losses and other financial distress costs.

The default condition established above means that the insurer cannot put its continuation value up as collateral. Implicit in this setup is the assumption that bankruptcy is driven by illiquidity rather than insolvency. This approach draws motivation from the financial crisis, which was associated with defaults that arose largely from liquidity shortages. It is important to note, however, that among life insurers the dissonance between strategic value default and liquidity default is more narrow than usual. Market-to-book ratios of these institutions tend to trade close to one, as the value of life insurers primarily reflects the value of their financial asset portfolios. This insight is helpful in the present context because it explains why the estimation is less exposed to an important misspecification problem. In particular, the concern is that the structural methodology cannot disentangle market expectations of government bailouts from beliefs about observationally similar rescues from other agents such as white knight investors. Because liquidity default often implies insolvency in the life insurance industry, however, the extent to which intervention of white knights resembles government bailouts is limited.

Violation of the net worth constraint in Equation 2.8 casts the insurance carrier into a state of financial distress. Without immediate outside help, bankruptcy is imminent. In the event that the insurer becomes financially distressed at time t, the government intervenes with probability $\widetilde{\phi}_t$. A government bailout takes the form of a costly equity injection equivalent to the shortfall amount $h_{\mathcal{F},t}(.) - (\lambda_t + d_t)$. Conditional bailout probabilities follow a censored normal distribution. As with insurance cash flow shocks, these probabilities are a function in part of aggregate productivity. Specifically, the time t probability of a bailout conditional on default is

$$\widetilde{\phi}_t = \max\left\{\min\left\{\mu + \nu\frac{\omega\phi_t - (1-\omega)z_t}{\sqrt{\omega^2 + (1-\omega)^2}}, 1\right\}, 0\right\}, \tag{2.9}$$

where $\mu \in [0,1)$, $\nu \geq 0$, and ϕ_t evolves according to the AR(1) process

$$\phi_{t+1} = \rho_\phi \phi_t + u_{\phi,t+1}, \quad u_{\phi,t} \sim \mathcal{N}\left(0, 1-\rho_\phi^2\right). \tag{2.10}$$

By construction, the fraction term in Equation 2.9 has a standard normal unconditional distribution. The minimization and maximization functions censor the distribution to ensure that probability $\widetilde{\phi}_t$ lies in the unit interval. Implicit above is the simplifying assumption that process weight parameter ω is the same for perturbations to the core spread as it is for bailout probabilities. That is, aggregate productivity matters for implicit government guarantees no more and no less than it does for life insurance fundamentals. Note that this assumption relates only to co-movement with aggregate productivity; it has no bearing on the dispersion of bailout probabilities. For example, $\widetilde{\phi}_t$ has a degenerate distribution at μ if dispersion parameter ν is zero. The AR(1) process for ϕ_t consists of idiosyncratic and industry components related in a fashion analogous to the decomposition of ϵ_t in Equation 2.2. As with ω, I make the simplifying assumption that idiosyncratic weight parameter ι is the same for bailout process ϕ_t as it is for insurance cash flow shock ϵ_t. The estimation is not sensitive to this assumption as none of the moments that are important for the identification of bailout parameters lean disproportionately on cross-sectional information.

In the event of bankruptcy, the resources available to liabilityholders are distributed according to absolute priority rules. Resources are paid out to policyholders first; creditors recover only what remains after policyholders are made whole. The market value of liabilities are priced competitively under symmetric information. The price of insurance liabilities, p_t, satisfies:

$$\begin{aligned} p_t(.)\lambda_{t+1} = \beta \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} & m(\widetilde{z}_t, \widetilde{z}_{t+1})\Big(\lambda_{t+1}\mathbf{1}_{\{h_{t+1}\geq\lambda_{t+1}+d_{t+1}\}} \\ & + \Big[\widetilde{\phi}_{t+1}\lambda_{t+1} + (1-\widetilde{\phi}_{t+1})\min\{h_{\mathcal{F},t+1}, \lambda_{t+1}\}\Big] \\ & \times \mathbf{1}_{\{h_{t+1}<\lambda_{t+1}+d_{t+1}\}}\Big) dF(\epsilon_{t+1}, z_{t+1}, \phi_{t+1}|\epsilon_t, z_t, \phi_t), \end{aligned} \tag{2.11}$$

where β is the time-preference parameter and $m(\widetilde{z}_t, \widetilde{z}_{t+1})$ is the pricing kernel, which is expressed in directly in terms of non-standardized aggregate productivity, $\widetilde{z}_t$, in accordance with usual practice in the production-based asset pricing literature. Pricing kernel $m(\widetilde{z}_t, \widetilde{z}_{t+1})$ is shared by equityholders, creditors, and policyholders. The first indicator function in the integral above is one if the insurer is not in financial distress at time $t+1$ and zero if it is not. Likewise, the second indicator function is one if the carrier is financially distressed and zero otherwise. Price p_t is therefore the expected present value of distributions to policyholders. The bracketed term in Equation 2.11 details what policyholders receive if the insurer enters financial distress. With probability $\widetilde{\phi}_{t+1}$ the carrier is bailed out and policyholders are made whole with the full repayment the insurance liability λ_{t+1}.[6] In the event that the insurer is not bailed out, which occurs with probability $1-\widetilde{\phi}_{t+1}$, policyholders receive distressed resources, $h_{F,t+1}$, up to the promised insurance obligation, λ_{t+1}. Note that the recovery value of policies is required to be non-negative. Implicitly, the $h_{\mathcal{F},t+1}$ in the minimization function above is defined as the maximum of $h_{\mathcal{F},t+1}$ and zero.

As with policy liabilities, the capital market price of debt, q_t, is the discounted expectation

[6] This assessment of default risk faced by policyholders abstracts from benefits incurred at time $t+1$, which are captured indirectly by the $h_{F,t+1}$ term. Implicitly, this immediate component of policy liabilities is risk-free in the model. Due to the long-term nature of life insurance liabilities in the data, however, expected benefits one quarter ahead represent a small share of total policy benefits at risk.

of distributions to creditors. The price of debt satisfies:

$$
\begin{aligned}
q_t(.)d_{t+1} = \beta \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} & m(\widetilde{z}_t, \widetilde{z}_{t+1})\Big(d_{t+1}\mathbf{1}_{\{h_{t+1}\geq\lambda_{t+1}+d_{t+1}\}} \\
& + \Big[\widetilde{\phi}_{t+1}d_{t+1} + (1-\widetilde{\phi}_{t+1})\min\{h_{\mathcal{F},t+1} - \lambda_{t+1}, d_{t+1}\}\Big] \\
& \times \mathbf{1}_{\{h_{t+1}<\lambda_{t+1}+d_{t+1}\}}\Big) dF(\epsilon_{t+1}, z_{t+1}, \phi_{t+1}|\epsilon_t, z_t, \phi_t).
\end{aligned} \tag{2.12}
$$

Creditors are made whole, receiving the principal value of debt d_{t+1}, if the insurer is not financially distressed or if the government intervenes to guarantee the insurer's liabilities in the event of distress. If the life insurer enters bankruptcy, debtholders receive any resources that remain after fulfilling the insurance obligation, $h_{\mathcal{F},t+1} - \lambda_{t+1}$, but no more than the face value of debt, d_{t+1}. As with policies, the recovery value of debt is non-negative. The $h_{\mathcal{F},t+1} - \lambda_{t+1}$ term enters Equation 2.12 as zero in the event that it is negative.

The model's treatment of default risk faced by policyholders abstracts from official government support provided by state guaranty funds. There are a number of factors to suggest, however, that the approach employed here is a better approximation to real world data than appears on first glance. First, not all policies are fully covered by state guaranty protections. Similar to FDIC deposit insurance, these guarantees are often specified with maximum coverage limits which vary by state and by contract type. Second, a related consideration is that a fraction of insurance obligations, particularly among large life insurers, bear greater resemblance to debt than traditional life insurance policies. Most notable in this regard are funding agreements, coverage of which by guaranty funds remains unclear.[7] Third, it remains an open question whether state guaranty associations are adequately funded to meet their coverage commitments in the event that the largest and most complex life insurers enters bankruptcy. The Financial Stability Oversight Council notes, for example, that state guaranty funds may be unable to fully cover qualifying policies for MetLife should the large insurer default on its obligations (FSOC 2014).[8] Lastly, extant empirical evidence suggests that policyholders are more sensitive to insurer insolvency risk than would be expected if their preferences conformed with typical risk averse capital market investors (Phillips, Cummins, and Allen, 1998). In the interest of parsimony, the modeling approach adopted here splits the difference by assuming that policyholders are no more risk averse than capital market investors but they receive no special protection from guaranty funds.

[7] Protection for funding agreement-backed securities would appear to be outside the spirit of guaranty funds' purview. As of this writing, however, this is an untested question.

[8] When AIG requested additional funding from the U.S. Treasury in March 2009, which it received, the insurer argued that its failure could overwhelm state guaranty funds (Christie and Frye, 2009).

2.2.4 Pricing Kernel and Capital Regulation

The asset pricing kernel shared by capital market investors and policyholders is formulated in line with the production-based asset pricing literature. In particular, I adopt the pricing kernel from Jones and Tuzel (2013) which has the specification:

$$\log(m_{t+1}) = -u_{z,t+1}e^{m_0+m_1\widetilde{z}_t} - \frac{1}{2}\left(s_z e^{m_0+m_1\widetilde{z}_t}\right)^2, \tag{2.13}$$

where $m_0 > 0$, $m_1 < 0$, and the law of motion for aggregate productivity, $\widetilde{z}_t$, is

$$\widetilde{z}_{t+1} = \rho_z \widetilde{z}_t + u_{z,t+1}, \quad u_{z,t} \sim \mathcal{N}\left(0, s_z^2\right). \tag{2.14}$$

Notice that for any level of productivity, $\widetilde{z}_t$, the conditional expectation $\mathbb{E}_t[m_{t+1}|\widetilde{z}_t] = 1$. Therefore, by construction, the risk-free rate generated by this pricing kernel is nonstochastic. Note that the time preference parameter for policyholders and creditors is $\beta = (1 + r_f)^{-1}$, where r_f is the nonstochastic risk-free rate expressed on a quarterly basis. The time preference parameter for shareholders, denoted by β_e, satisfies the inequality $\beta_e < \beta$. As noted in Michaels, Page, and Whited (2019), this captures in reduced-form the tax benefits of distributions to liabilityholders.

Life insurers face capital regulation from state regulators. The model abstracts from the trade-offs of carrying different types of insurance risks in different types of jurisdictions. It is necessary, however, to capture the role that regulatory constraints have on decision-making about leverage. Therefore, I assume that general account liabilities incur regulatory cost $g(\lambda_t, d_t) = g(\lambda_t + d_t)$, where $g(.)$ is increasing and convex in liabilities. These regulatory costs are modeled as an additional expense in core spread, ψ_t, with $g(.)$ modeled as a fraction of general account assets and parameterized as

$$g(\lambda_t, d_t) = g(\lambda_t + d_t) = \frac{\kappa}{1 - (\lambda_t + d_t)}, \quad \kappa \geq 0. \tag{2.15}$$

With the addition of capital regulation costs, the core spread becomes

$$\psi_t(.) = \Big([\sigma_t i_t(\epsilon_t, z_t) - \gamma_t(\sigma_t, \sigma_{t+1})] - 1\Big)\lambda_t + p_t(.)\lambda_{t+1} - g(\lambda_t, d_t). \tag{2.16}$$

A treatment of regulatory capital requirements is necessary to generate realistic quantities for leverage. More specific rationale behind the soft regulatory constraint above is as follows. As required capital approaches its constraint in one jurisdiction, the opportunity cost of forgoing additional insurance opportunities from that area or the expenditure of utilizing reinsurance (captive or otherwise) for such opportunities grows prohibitively large. The regulatory expense function also captures more direct costs. These include the human capital cost entailed by regulatory compliance with low margins of error as well as the expected fees and legal costs for noncompliance.

2.2.5 Life Insurer Problem

This section describes the recursive formulation of the life insurer problem. For ease of notation I now switch to prime notation. As noted above, the model abstracts from decision-making about the balance between insurance liabilities and balance sheet debt. I assume that the debt share of liabilities, $\bar{d}$, is fixed. The life insurer manages directly general account liabilities, $\lambda + d$. In conjunction with parameter $\bar{d}$, this decision pins down the accounting value of insurance liabilities, $\lambda = (1 - \bar{d})(\lambda + d)$, and the face value of debt, $d = \bar{d}(\lambda + d)$. Therefore, the life insurer problem has two endogenous state variables, $(\lambda + d, \sigma)$, and three exogenous state variables, (ϵ, z, ϕ).

Due to the important role of bailouts in the model, the life insurer value function, $v(\lambda + d, \sigma, \epsilon, z, \phi)$, lends itself to a piecewise formulation:

$$v(\lambda + d, \sigma, \epsilon, z, \phi) = \begin{cases} v^{\mathcal{H}}(\lambda + d, \sigma, \epsilon, z, \phi), & h(.) \geq \lambda + d \\ v^{\mathcal{F}}(\lambda + d, \sigma, \epsilon, z, \phi), & h(.) < \lambda + d \end{cases} \tag{2.17}$$

Above, $v^{\mathcal{F}}(.)$ denotes the value function of an insurer that is financially distressed, and $v^{\mathcal{H}}(.)$ denotes the value function of one that is not. The Bellman equation for the financially healthy life insurance carrier is

$$\begin{aligned} v^{\mathcal{H}}(\lambda + d, \sigma, \epsilon, z, \phi) = \max_{\lambda' + d', \sigma'} \Big\{ & e(.) \left(1 + \chi \mathbf{1}_{\{e(.)<0\}}\right) \\ & + \beta_e \mathbb{E}\left[m(\widetilde{z}, \widetilde{z}') v(\lambda' + d', \sigma', \epsilon', z', \phi')\right] \Big\}, \end{aligned} \tag{2.18}$$

subject to:

$$\lambda = (1 - \bar{d})(\lambda + d) \text{ and } d = \bar{d}(\lambda + d), \tag{2.19}$$

$$e(.) = \psi(.) - c(.) - d + q(.)d', \tag{2.20}$$

$$\psi(.) \text{ and } c(.) \text{ satisfy Equations 2.6 and 2.16,} \tag{2.21}$$

$$(\epsilon, z, \phi) \text{ evolve according to Equations 2.3, 2.10, and 2.14,} \tag{2.22}$$

$$p(.) \text{ and } q(.) \text{ satisfy Equations 2.11 and 2.12.} \tag{2.23}$$

Note that the expectation in the Bellman equation is conditional on the current state, i.e. $(\lambda + d, \sigma, \epsilon, z, \phi)$.

The Bellman equation for the life insurer that begins the current period in a state of financial distress is

$$\begin{aligned} v^{\mathcal{F}}(\lambda + d, \sigma, z, \epsilon, \phi) = \widetilde{\phi} \times \max_{\lambda' + d', \sigma'} \Big\{ & [e(.) + \chi \min\{e(.), h_{\mathcal{F}} - (\lambda + d)\}] \\ & + \beta \mathbb{E}\left[m(\widetilde{z}, \widetilde{z}')v(\lambda' + d', \sigma', z', \epsilon', \phi')\right] \Big\}, \end{aligned} \tag{2.24}$$

subject to the same conditions as above. To summarize in words, shareholders are wiped out in the event of bankruptcy. If the distressed carrier is bailed out, which occurs with probability $\widetilde{\phi}$, the insurer remains a going concern upon receipt of a costly equity injection equal to the net worth shortfall $h_{\mathcal{F}} - (\lambda + d)$. Note that the minimum function within the Bellman equation prevents double counting of issuance costs in the event that shareholders of the bailed-out insurer elect to raise capital in excess of the bailout amount.

2.3 Estimation Methodology and Identification

2.3.1 Model Solution and Estimation

The model is solved using value function iteration. I employ a Markov chain approximation for exogenous variables (ϵ, z, ϕ) according to the Tauchen (1986) method for order-one vector autoregressions. Real world data help to inform the construction of the grid for general account liabilities, $\lambda + d$. The upper and lower endpoints for this grid are set at the 99th and 1st percentile of this variable in the data, respectively. A Gaussian-quadrature based method is used to set the spaces between liability grid points. This method generates finer spacing near the center, which is selected to correspond with the unconditional sample mean of insurance liabilities. Because the risk exposure level, σ, is unobservable in the data, the sample is less immediately expressive about how its grid should be arranged under sensible parameterizations. I therefore use equal-spacing for this grid. Its lower endpoint is set to $1\mathrm{e}-10$, as the model requires that $\sigma > 0$. The upper endpoint is chosen so that its corresponding risk carrying costs, when λ is at its sample mean, are implausibly large—i.e., triple the first percentile of the core spread in the data.

I estimate the model on samples of large and small publicly-traded U.S. life insurance companies according to the methodology described below. The estimation targets features of insurance-specific accounting data from the SNL Financial Institutions (FIG) database as well as Compustat. Also targeted are moments related to stock, stock option, and CDS data from CRSP, IvyDB OptionMetrics, and IHS Markit, respectively. The data sources, sample construction, and sample partitioning criteria are detailed in companion research presented in Chapter 1.

The majority of parameters are estimated using simulated minimum distance estimation (SMD). The estimation proceeds as follows. For a given parameterization, the model is solved using the numerical solution method described above. Data are then simulated for multiple sets of panels, each with the same cross-sectional and time characteristics as the real world sample. For large life insurers, I simulate 200 panels of 12 carriers with 20 years of quarterly data. The small life insurer estimation uses 60 panels of simulated data for 39 carriers with 20 years of quarterly data. The numbers of panels are chosen so that each sample utilizes roughly the same quantity of simulated data, both cases entailing ratios of simulated data to real world data in excess of ten (see Michaelides and Ng, 2000). Moments and functions of moments computed from the simulated data are compared to their real world counterparts.[9] Distances between them are aggregated into a scalar summary value by the SMD objective function. This aggregation is performed according to the optimal weight matrix estimated from the empirical moment covariance matrix. Intuitively, moments are weighed in the objective function according to how precisely they are measured in the data. Model parameters are thus estimated by minimizing this aggregate distance. I use simulated annealing in conjunction with generalized pattern search to locate the objective-minimizing set of parameters.

Not all model parameters are estimated via SMD. Some of these external parameters are estimable directly from the data. A set of remaining parameters, most notably those related to the pricing kernel, are calibrated. Table 2.1 reports all external parameters for both the

[9] Note that because functions of moments are themselves moments, this chapter uses the latter term to refer to the former except in cases where the role of an auxiliary model warrants emphasis. With this perspective, the methodology used here could also be described as a simulated method of moments (SMM) estimation.

large and small life insurer samples. Estimated separately from the SMD estimation is the debt share, $\bar{d}$, the idiosyncratic shock weight parameter, ι, and the risk-free rate, r_f. The debt share is set at the sample average, which is 7.52% for large insurers and 5.47% for small insurers.

Parameter ι is estimated from a panel regression of a transformation of the core spread on its corresponding industry aggregate. To describe the process in more detail, first note that for convenience of terminology I refer to the core spread as adjusted if it is expressed as a share of insurance liabilities. To estimate the idiosyncratic weight parameter, adjusted core spreads are first orthogonalized with respect to aggregate productivity. This orthogonalized measure is the residual series from a regression of the adjusted core spread on the level of standardized (i.e., expressed as a z-score) TFP and a third degree polynomial of first differences in standardized TFP. Parameter ι is then set as one minus the OLS coefficient from a pooled panel regression of this measure on the value-weighted industry average, where both dependent and independent variables are standardized and the cross-sectional industry means are weighted by general account assets. The resulting estimates for ι are 0.576 and 0.730 for large and small carrier samples, respectively.

The nonstochastic risk-free rate is estimated as the mean 3-month U.S. Treasury bill rate for the 2001–2020 sample period. This yields a risk-free rate of 1.34% in annual terms. Note that the risk-free rate pins down both the time preference parameter for policyholders and creditors, β, and the time preference parameter for shareholders, β_e. Following Michaels, Page, and Whited (2019), I set the rate for equityholders 20% higher than r_f. This corresponds with an effective corporate income tax rate of 20%.

It is not possible to identify jointly financial distress costs and bailout parameters using the estimation methodology in this chapter. I therefore calibrate general account assets in distress, ξ, to one. By assuming the best-case asset recovery value for liabilityholders the estimation errs on the side of conservative estimates for government bailout probabilities. The assumption is also less restrictive in this model than it appears: the higher moment risks governed by parameter η likely capture some of the distress costs ordinarily approximated with $\xi < 1$. Equity issuance cost parameter χ is set at 4.46% using the estimate reported by Michaels, Page, and Whited (2019). Note that this estimate comes from a sample of industrial firms. It is plausible that equity issuance costs are systematically different for life insurers. For realistic levels of $\chi > 0$, however, the model presented here is not sensitive to this parameter.

The remaining auxiliary parameters govern aggregate productivity and the asset pricing kernel. In line with the baseline calibration of Jones and Tuzel (2013), I set the constant price of risk parameter, m_0, to 3.22 and time-varying price of risk parameter, m_1, to -14.75. Following convention, the law of motion parameters for aggregate productivity are drawn from

King and Rebelo (1999).[10] Their quarterly estimates for the persistence parameter and the conditional standard deviation of productivity are $\rho_z = 0.979$ and $s_z = 0.0072$, respectively.

2.3.2 Identification

Having addressed the external parameters as described above, 11 parameters remain: $\{\theta$, η, ρ_ϵ, ζ, α_σ, α_λ, κ, ω, μ, ν, $\rho_\phi\}$. These parameters are estimated jointly by SMD. In this section, I explain how these parameters are identified from variation in the data. The basic idea is to target moments or auxiliary models that are sensitive to changes in parameters. In the process of selecting moments it is important not to "cherry pick," lest the econometrician exert influence on the parameter estimates or obscure important areas where the model struggles to describe the data. For this reason my approach to moment selection errs in the direction of inclusion.

The SMD estimation targets 38 moments. The full list of moments is provided across Tables 2.2 and 2.3. These tables include a note column to detail how a moment is computed in instances where the abbreviated description is not self-explanatory. Table 2.2 and Table 2.3 also group moments into type categories based on the model parameters for which they are most informative. The model is heavily overidentified. Of these 38 moment restrictions, the first 16 are uncategorized (general) moments because they describe basic features of the data without contributing indispensable identifying information for any parameter in particular. The rationale for the sizable number of overidentifying restrictions is twofold. First, as noted above, I avoid singling out moments that are unrepresentative of the overall message from the data by emphasizing broad inclusion. In light of this approach, the high number moments is a direct consequence of the wealth of capital market pricing variables supplied by the options and CDS data. Second, the success of the estimation in the present context hinges on its capacity to extract investor expectations from financial prices. To interpret these model-implied expectations with confidence, it is incumbent upon the model to describe observable features of the capital market data as accurately as practical.

I now focus attention on how parameters are identified by the moment restrictions. It is important to bear in mind that no direct mapping exists between individual parameters and individual moments. The SMD parameters are identified jointly. Understanding which moments are most informative for individual parameters, however, is helpful for developing intuition about the identification. As θ governs the central tendency of the core spread, it follows that the unconditional mean of this variable is highly informative for this parameter. Three moments aid in the identification of η. Because this parameter controls the asymmetry

[10] A replication of the quarterly productivity regression in King and Rebelo (1999) on an updated time series yields estimates that are statistically no different from the original reported values.

of insurance cash flows, features of the IV distributions for put and call options lend vital identifying information. The first moment is the 95th percentile of ATM put IV. Note that all moment restrictions related to option IV are for the 91-day time horizon. The second moment for η is the variance of OTM call IV interacted with an indicator that is one if the core spread is below its mean and zero otherwise. This type of moment is referred to here as a quasi-conditional variance; it is not a true conditional variance as the moment is computed on observations where the indicator function is zero. A third moment for η is the OLS coefficient from a regression where the dependent variable is the difference between OTM put IV and OTM call IV, which I refer to as the call-put OTM IV gap. The independent variable is the standardized (i.e., expressed as a z-score) adjusted core spread squared and interacted with its sign. Also included in this regression, but not directly targeted, is an intercept.

To identify ρ_ϵ, I target the persistence parameter from an AR(1) autoregression of the core spread. Note that targeted OLS coefficients are for regressions without intercepts or controls unless otherwise indicated. Three moments are most informative for ω: the correlation of first differences in standardized total factor productivity (TFP) and the core spread, the correlation of first differences in TFP and excess stock returns, and the correlation of standardized TFP and the 5-year CDS spread. Two moments help to identify risk carrying cost parameter ζ: the 95th percentile of core spreads and the 95th percentile of excess stock returns. Risk carrying costs most directly impact observables by way of their influence on insurer decision-making about the upper end of their risk exposure level. This is most limiting when fundamentals are favorable. Accordingly, the right tails of the core spread and return distributions shrink as ζ increases.

Three moments are informative for risk adjustment cost parameter α_σ. The first is the variance of ATM put IV. The second moment is the variance of the core spread interacted with an indicator function that is one when the core spread is within its interquartile range and zero otherwise. By emphasizing dispersion among moderate values of the core spread, this quasi-conditional variance offers insight into α_σ with less contaminating force from η and ζ, parameters that most directly govern the tails. The third moment for α_σ is the correlation of 6-month CDS spreads with 5-year CDS spreads. This correlation is increasing in α_σ; long-term credit risk and short-term credit risk grow more disconnected when the cost of adjusting risk falls.

Identifying policy adjustment cost parameter α_λ is less challenging than α_σ because insurance liabilities are observable in the data. Naturally, α_λ is pinned down by the variance of insurance liabilities. For capital regulation cost parameter κ, the most informative moment is the mean of general account liabilities. This is intuitive. Generating realistic quantities for leverage is the raison d'être of the soft regulatory capital constraint. Similar to models in the structural banking literature, the theoretical framework presented here implies unrealistically high levels of leverage for life insurers in the absence of an opposing regulatory force.

This leaves the three parameters that govern the law of motion for implicit guarantee probabilities: $\{\mu, \nu, \rho_\phi\}$. A class of auxiliary models assist in the identification of these parameters. In particular, useful identifying information comes from the distance between realized CDS rates and Merton default model (i.e., the auxiliary model) implied credit spreads. In this chapter, I refer to this distance as the Merton gap. Using stock option IV, the market value of equity, book liabilities data, and the risk-free rate, Merton models provide a computationally tractable way of estimating the price of debt within a simple risk neutral environment that abstracts from strategic decision-making about leverage and asset risk. Variation in this Merton model gap serves as a rough proxy for movement in bailout probabilities. The reason is that investor expectations about these probabilities manifest in capital market data primarily through a wedge between equity and debt pricing. To see why, first note that a government rescue is designed to avert imminent default. Because credit spreads concern only the likelihood and severity of bankruptcy scenarios, a reduction in default probability via government protection has a first-order effect on these spreads, even when credit risk is low. The direct effect on equity value is muted by comparison. By the point that an insurer's default is imminent without a rescue package, most of the market value of equity has already been erased—often recently and in precipitous fashion. In most states of the world, therefore, the bailout put option constitutes a much larger share of the CDS value than it does equity value. This is the case even when considering the second order effect on stock value through lower borrowing costs.

This perspective is helpful for building intuition about the role of the Merton auxiliary model and the identification of the bailout parameters more broadly. It is also not the complete story. Moral hazard can complicate the relationship between implicit guarantee probabilities and CDS rates. The estimation manages this problem by targeting moments related to multiple time horizons in the term structure of credit spreads. In general, short-term CDS spreads offer a cleaner glimpse at the bailout put option per se, whereas longer term spreads carry more information about moral hazard. Changes in risk-taking decisions can take greater than a year to respond fully to a protracted increase in bailout probabilities. This is due largely to risk adjustment and underwriting costs. It is also due to the state dependence of moral hazard, which may present not as an immediate course of action but a subsequent *inaction* in the face of deteriorating fundamentals. Lengthier time horizons thus provide insurers more breadth for moral hazard. This is compounded by the mean-reverting nature of bailout probabilities, as costly risk adjustment implies that the effects of moral hazard can outlive the protection that motivated it in the first place.[11] These features of the model pivotal for identification derive a measure of external validity from the financial crisis evidence presented in Chapter 1. Both the relative sensitivity of debt compared to equity and the downward-sloping term structure of credit risk reductions are consistent with the main implications of the April 8 event studies.

Two moments help to pin down μ. The first moment is the OLS coefficient from a regression of 1-year CDS spreads on a dummy variable that is one if ATM put option IV is two standard deviations above its mean and zero otherwise. The second moment for μ is the OLS coefficient from a regression of 6-month CDS spreads on the interaction of general account liabilities and OTM put IV. This second regression also includes an intercept and controls for excess stock returns. Neither of these moments most informative for μ use the Merton auxiliary model. They do, however, built on the intuition undergirding it. Both OLS coefficients are decreasing in μ. This reflects a declining price sensitivity to equity volatility for short-term CDS spreads as the central tendency of bailout probabilities rise. The relationship is stable even for large values of μ, as the resulting benefits of steady government protection outweigh the risks of moral hazard on average. The cleanest signals come from shorter CDS time horizons.

Three moments are most informative for bailout parameter ν. The first is the 95th percentile of 1-year CDS spreads, which is decreasing in ν. The greater dispersion of bailout probabilities from a rising ν extends their right tail.[12] One way in which this manifests most clearly in the data is through a mitigation of typical worst-case scenarios; hence the shrinking right tail of short-term CDS rates. The second moment for ν is the variance of the total 5-year Merton model gap, which I define as the sum of the 5-year Merton gap computed using OTM put IV and the 5-year Merton gap computed using ATM put IV. This moment is increasing and convex in ν. The informativeness of this variance is derived in part by the Merton gap's direct proxy relation with implicit guarantee probabilities. This variance also captures investor expectations about moral hazard, which can make the gap more volatile as well. The third moment for ν is the OLS coefficient from a regression of the 6-month OTM Merton model gap (i.e., the Merton gap for the 6-month time horizon computed using OTM put option IV) on excess stock returns. This regression beta captures convexity in the proxy relationship between the Merton gap and bailout probabilities that follow from greater dispersion in these probabilities.

[11] Note that this does not mean that long-term CDS rates grow arbitrarily insensitive to bailout probabilities for the states in which the moral hazard effect dominates. Rather, their reaction to a shock in bailout probability can be sizable in magnitude, but of the opposite sign than a decline in default risk would imply.

[12] I assume that nontrivial right-tail censoring of $\widetilde{\phi}$ is economically implausible as this would entail periods in which short-term CDS spreads are zero—a feature that is not observed in real world data.

I use two moments to identify ρ_ϕ, which governs the persistence of shocks to bailout probabilities. The first moment is the OLS coefficient from a regression of the 3-year OTM Merton gap on the 6-month OTM Merton gap. The second moment is similar: it is the OLS beta from a regression of the 5-year OTM Merton gap on the 6-month OTM Merton gap. Both moments are increasing in ρ_ϕ; what differentiates the two is the degree of convexity in their relationship with the parameter of interest. The informativeness of these moment restrictions is intuitive. More persistent shocks to bailout probabilities increase the co-movement between short-term and long-term auxiliary model-based proxies for bailout probabilities.

Before continuing on to the estimation results, some of the additional moment restrictions are worth highlighting. Targets for skewness of key variables and the average market-to-book (MTB) equity ratio are motivated in part by institutional details of the life insurance industry. These moments also have some added relevance for the present exercise. In addition to moments related to the put options data, skewness provides a glimpse at the heavily asymmetric payoffs of the insurance business. Due to the central importance of credit risk for the parameters of interest, a reasonable characterization of the higher moment risks faced by life insurers is also worth emphasizing. Note, however, that these moments receive a relatively light weight in the SMD objective function as skewness requires a generous amount of data to measure precisely. As pointed out in Section 2.2.3, market-to-book ratios tend to trade around one in the public life insurance sector. Because of the model's net worth default condition as well as the general significance of recoverable assets available to creditors, I require that the estimation attempts to match this institutional feature.

2.4 Estimation Results

2.4.1 Moments and Parameter Estimates

To assess the estimation results, I begin by examining how well the model fits the data. Table 2.4 reports simulated moments juxtaposed with their data counterparts both the large and small sample estimations. The accuracy with which the model describes the data is highly consistent between the two samples. The model performs reasonably well in most cases for moments related to key accounting variables such as the core spread, general account liabilities, and insurance liabilities. Particularly notable in this regard are the mean of general account liabilities and, especially in the case of large insurers, the variance of insurance liabilities. The model tends to understate the level of the core spread. This buys the estimation a better match for the mean of operating profit, however, which is not targeted. This is unsurprising as the level of underwriting expenditure tends to be higher in the data. Interestingly, the core spread exhibits considerably less serial correlation in the simulated data. This appears not to be driven in the model by extensive risk adjustment so much as the low estimates for ρ_ϵ. The high persistence in the data may reflect the smoothing influence of accounting.

On balance, the model provides a reasonable description of capital market prices. The model tends to perform best with moments related to long-term CDS spreads. Simulated data exhibit the implied volatility "smirk" observed in real world data, though in both samples the unconditional means for put option IV are roughly 10 percentage points lower in the model than in the data. The elevated put option IV in the data hints at market segmentation. The model also overstates the magnitude of covariation between capital market prices and aggregate productivity, tracing out the limits of a single-factor pricing kernel in the present application. For most of the moments particularly informative for the bailout parameters, which draw heavily from the capital market side, the model tends to fit the data quite well. This is particularly noticeable for moments that are chiefly informative for μ.

The simulated minimum distance parameter estimates for the large and small life insurer samples are presented in Table 2.5. The economic interpretation of parameters governing fundamentals benefits from some perspective on the similarities and differences between the two sample estimates. The 1.046 point estimate for cash flow asymmetry parameter η in the large carrier sample stands in contrast to the lower 0.735 value in the small carrier sample. These estimates imply that large life insurers face fatter downside tail risk. Insofar as fragility concerns tail risks in particular, this supports the notion that fragility in the life insurance sector is more severe among large carriers. It is also in line with anecdotal evidence from the financial crisis where the most extraordinary losses were concentrated among a few of the largest life insurance companies. The lower estimates for risk carrying cost parameter, ζ, and liability regulatory cost, κ, among large insurers is consistent with extant evidence on their propensity for regulatory arbitrage. The lower regulatory cost of liabilities faced by large carriers in particular likely reflects the value of their captive reinsurance activities. In contrast to the case with insurance liabilities, the point estimate for risk exposure adjustment cost parameter α_σ is slightly lower for large life insurers. This may hint at the liberating nature of enhanced regulatory arbitrage tools or the more general notion that the risks entailed by the non-traditional activities in which large insurers heavily engage are less costly to adjust over time. However, it is clear from the standard errors that the parameter is measured too imprecisely to weigh these interpretations with any statistical confidence.

Among large life insurance carriers, the estimates for implicit guarantee parameters μ and ν are 5.7% and 14.5%, respectively. Recall that in the model the unconditional distribution for the probability of a bailout conditional on default is censored normal. These parameter estimates therefore imply that the unconditional mean and standard deviation of bailout probabilities are 9.1% and 10.3%, respectively. The corresponding standard errors for these first two unconditional moments are 2.1% and 1.7%, respectively. The point estimate for ρ_ϕ is 0.183 for large insurers, indicating that the shocks to bailout probabilities unrelated to productivity are relatively transient. In the model the unconditional probability of a bailout is 0.36% on a per annum basis. The expected government equity injection conditional on a bailout is 1.9% of general account assets. For perspective, MetLife's general account assets at fiscal year-end 2020 was $595.2 billion; were the large life insurer to receive an emergency government rescue this conditional expected bailout in dollar terms is $11.3 billion. In addition, the standard deviation of a net worth shortfall conditional on a bailout is 2.5% in simulated data. To continue with the MetLife example, a two-standard deviation bailout as of Q4 2020 general account assets translates to $41.1 billion in dollar terms.

These model-implied investor expectations of support are broadly consistent with the extent of government aid provided to large life insurers during the crisis. There are also two ways in which the agreement between the structural estimates and the event study results of Chapter 1 are notable. First, the relative size of dispersion parameter ν with respect to central tendency parameter μ reveals that for implicit backstops dynamics are critical. This, in addition to the low persistence of shocks, is consistent with the term structure of government support implied by CDS reactions to the April 8 announcement. Second, lower bounds for government bailout probabilities recovered from these event studies fall within a reasonable subinterval of the model's support. Nearly all of the Chapter 1 estimates are within three standard deviations of the unconditional distribution in the model, with the majority of these estimates falling within two. In comparing the two sets of analyses, it is worth bearing in mind that bailout probabilities from the TARP evidence in Chapter 1 are risk neutral whereas structural estimates reported here are for physical probabilities.

The small life insurer sample estimates provide an additional source of external validation. If the above estimates indeed reflect a "too big to fail" premium, it stands to reason that the methodology should imply more modest levels of support for small insurers. The SMD estimates imply that the unconditional mean and standard deviation of the censored normal bailout probabilities among small life insurers are 4.1% and 2.3%, respectively. These unconditional moments are measured very precisely, with corresponding standard errors of 0.5% and 1.0%, respectively. Interestingly, the 0.631 estimate for ρ_ϕ implies that shocks to bailout probabilities per se are more persistent for small insurers. The unconditional probability of a bailout for small life insurance companies according to the model is 0.21% per annum—less than 60% of the value for large life insurers. Not only are small carriers nearly half as likely to receive bailouts, simulations reveal that the extent of support conditional on an emergency rescue is also more limited. Conditional on a bailout, the mean and standard deviation of net worth shortfalls are 1.2% and 1.0%, respectively, of general account assets. This is in contrast to the 1.9% conditional mean and 2.5% conditional standard deviation among large life insurers.

A comparison of simulated data from the two estimations also reveals interesting differences in the nature and extent of bankruptcy risk. Notwithstanding the greater tail risk and moral hazard among large carriers, these institutions are less risky than their smaller peers according to some metrics. In particular, small life insurers are 1.6 times more likely to enter a state of financial distress. Overall differences in risk between the two samples take on nuance, however, in view of the more severe losses conditional on default among large insurers. These nuances help to explain why a cursory examination of the data moments fails to yield clear hints of a "too big to fail" premium.

2.4.2 No Bailout Counterfactual Analysis

How might the U.S. life insurance industry look in the absence of implicit government guarantees? To address this question, I solve and simulate the model under the same parameterization as the SMD estimates above but with the probability of bailouts fixed at zero. The counterfactual analysis quantifies the extent of moral hazard generated by the implicit support.

Table 2.6 displays simulated sample statistics for large life insurers under the no bailout counterfactual alongside their values according to the baseline parameterization. These sample statistics include variables that are observable in real world data, such as CDS spreads and option IV, as well as variables that are unobservable, including risk exposure level, σ, and the unconditional probability of financial distress. What is evident from Table 2.6 is that even in full view of unobservable quantities most of the differences between the counterfactual and the baseline scenario are not striking. Consistent with moral hazard, general account liabilities on average are slightly higher with implicit guarantees at 87.91% as opposed to 87.86% under the counterfactual. But this almost negligible difference between these two parameterizations is strictly smaller in magnitude than the distance between either scenario and the corresponding moment in real world data. Average levels of put option IV and CDS spreads at each tenor are elevated in the presence of bailouts, indicating that there is indeed a modest increase in downside risk due to moral hazard in risk-taking practices. Here again, these differences appear economically insignificant. For example, the mean 5-year CDS under the baseline and the counterfactual scenarios are 1.06% and 1.04%, respectively. Average OTM put option IV is 29.07% under the baseline compared to 29.01% under the counterfactual.

A direct comparison of market prices understates the moral hazard because in addition to reflecting differences in fundamental risk they also reflect differences in implicit guarantees. However, even the mean of risk exposure variable σ remains virtually unchanged in the counterfactual. One of the reasons why moral hazard is not readily apparent from most unconditional statistics is that the moral hazard itself is not unconditional. A clear message from the parameter estimates is that the fixed aspect of bailout probabilities is dwarfed by the time-varying, state-dependent component. The counterfactual analysis reveals that this has two implications for the associated excess in risk-taking. First, the moral hazard practices themselves inherit much of the dynamic properties of the subsidies that motivate them. The excess in risk therefore manifests less in a steady fashion than in a state contingent one. Second, the right tail of bailout probabilities exerts a disproportionately heavy influence on life insurers' risk decision-making.

These features add context to the idea that a comparison of the likelihood and severity of distress scenarios provides one of the more informative perspectives in this regard. Table 2.6 indicates the per annum probability of distress declines from 2.3% under the baseline scenario to 2.1% in the counterfactual. The standard error for this probability is slightly less than 0.1%, placing the counterfactual value just outside of the 95% confidence interval of the baseline estimate. The conditional average shortfall in the event of distress is 2.0% under the baseline parameters compared to 1.9% according to the no bailout counterfactual, though this difference is less notable from a statistical perspective. While these distress statistics reveal that the influence of moral hazard is less negligible than it appears on first glance, it is clear that the overall extent of excess risk-taking is nonetheless limited.

2.5 Conclusion

Building on companion research documented in Chapter 1, the structural results presented here provide evidence of market expectations of implicit backstops that are at times economically substantial, concentrated among large life insurers, and characterized by extensive time variation. My analysis also provides supporting evidence that the regulatory costs of leverage and risk-taking are lower for large life insurers despite greater tail risk fragility. A counterfactual analysis identifies limited but non-negligible moral hazard in risk-taking practices among large insurers, which reflect the dynamic nature of the implicit support that motivates them. Overall, the results contribute a new dimension of evidence that investors hold beliefs about large life insurance companies that are consistent with a "too big to fail" subsidy. The findings suggest further that because this subsidy exhibits extensive time-variation, its influence in tranquil periods may create a misleading impression of its full range.

Table 2.1: External parameters

Parameter description	Notation	Value: Large sample	Value: Small sample
Parameters estimated separately			
Risk-free rate (annualized)	r_f	0.0134	0.0134
Debt liability share	$\bar{d}$	0.0752	0.0547
Weight on idiosyncratic disturbances	ι	0.5758	0.7303
Calibrated parameters			
General account assets in financial distress	ξ	1	1
Equity issuance cost	χ	0.0446	0.0446
Persistence of TFP	ρ_z	0.979	0.979
Conditional std. dev. of TFP	s_z^2	0.0072	0.0072
Constant price of risk parameter	m_0	3.22	3.22
Time-varying price of risk parameter	m_1	-14.75	-14.75

This table reports external model parameter values used in the structural estimations of the large and small life insurer samples.

Table 2.2: SMD moment restrictions, table 1 of 2

	Moment	Type	Notes
1	Mean ATM put IV	–	
2	Mean OTM put IV	–	
3	Mean 6m CDS spread	–	
4	Mean 1y CDS spread	–	
5	Mean 5y CDS spread	–	
6	Variance of 6m CDS spread	–	
7	Variance of 5y CDS spread	–	
8	Variance of excess stock returns	–	
9	Variance of adj. core spread	–	Adjusted core spread is the core spread divided by policy liabilities
10	Correlation of 6m CDS and core spread	–	
11	Skewness of ATM put IV	–	
12	Skewness of excess stock returns	–	
13	Skewness of core spread	–	
14	Mean of equity MTB ratio	–	
15	75th percentile of 6m CDS spread	–	
16	75th percentile of 5y CDS spread	–	
17	Mean of core spread	θ	
18	95th percentile of ATM put IV	η	
19	Quasi-conditional variance of OTM call IV	η	Variance of OTM call IV interacted with indicator that is one if core spread is below its mean
20	OLS coef., CP IV gap regression	η	Regression of call-put OTM IV gap on z-score of adj. core spread squared and interacted with its sign, intercept included
21	Persistence parameter for core spread	ρ_ϵ	OLS Persistence parameter from AR(1) autoregression of core spread
22	Correlation of core spread and ΔTFP	ω	ΔTFP is first differences in standardized (i.e., z-score) TFP

Table 2.2 is the first of two tables that, together, provide a comprehensive list of moment restrictions used for the simulated minimum distance estimations. The notes column provides important additional information for instances in which a moment is not self-explanatory. The type column indicates for which parameter the corresponding moment is most informative, where an uncategorized moment need not be disproportionately informative for any individual parameter in particular.

Table 2.3: SMD moment restrictions, table 2 of 2

	Moment	Type	Notes
23	Correlation of stock returns and ΔTFP	ω	
24	Correlation of 5y CDS rate and TFP	ω	TFP expressed in levels
25	OLS coef., Merton gap 3y term struct.	ρ_ϕ	Regression of 3y OTM Merton gap on 6m OTM Merton gap
26	OLS coef., Merton gap 5y term struct.	ρ_ϕ	Regression of 5y OTM Merton gap on 6m OTM Merton gap
27	OLS coef., 1y CDS and ATM IV	μ	Regression of 1y CDS rate on indicator that is one if ATM put IV is two std. deviations above mean
28	OLS coef., 6m CDS and OTM IV	μ	Regression of 6m CDS rate on interaction of liabilities and OTM put IV, with intercept and controlling for excess stock returns
29	95th percentile of 1y CDS spread	ν	
30	Variance of 5y total Merton spread gap	ν	Total Merton gap is the sum of ATM Merton gap and DOTM Merton gap
31	OLS coef., Merton gap and stock returns	ν	Regression of 6m OTM Merton gap on excess stock returns
32	95th percentile of core spread	ζ	
33	95th percentile of excess stock returns	ζ	
34	Variance of ATM put IV	α_σ	
35	Quasi-conditional variance of core spread	α_σ	Variance of core spread interacted with an indicator that is one if core spread is within interquartile range
36	Correlation of 6m and 5y CDS rates	α_σ	
37	Variance of insurance liabilities	α_λ	
38	Mean of general account liabilities	κ	

Table 2.3 is the second of two tables that, together, provide a comprehensive list of moment restrictions used for the simulated minimum distance estimations. The notes column provides important additional information for instances in which a moment is not self-explanatory. The type column indicates for which parameter the corresponding moment is most informative, where an uncategorized moment need not be disproportionately informative for any individual parameter in particular.

Table 2.4: Simulated minimum distance estimation moments

	Moment	Type	Large insurers Data	Large insurers Model	Small insurers Data	Small insurers Model
1	Mean ATM put IV	–	0.3123	0.1969	0.3409	0.2190
2	Mean OTM put IV	–	0.3942	0.2907	0.4401	0.3000
3	Mean 6m CDS spread	–	0.0085	0.0034	0.0096	0.0034
4	Mean 1y CDS spread	–	0.0086	0.0047	0.0103	0.0048
5	Mean 5y CDS spread	–	0.0124	0.0106	0.0153	0.0161
6	Variance of 6m CDS spread	–	9.0e-04	7.5e-04	0.0012	7.2e-04
7	Variance of 5y CDS spread	–	4.5e-04	6.0e-04	5.5e-04	4.4e-04
8	Variance of excess stock returns	–	0.0399	0.0321	0.0379	0.1790
9	Variance of adj. core spread	–	1.5e-04	1.2e-04	2.9e-04	1.6e-04
10	Correlation of 6m CDS and core spread	–	-0.1975	-0.0015	-0.1347	0.1899
11	Skewness of ATM put IV	–	3.8575	3.3787	3.3372	5.6809
12	Skewness of excess stock returns	–	1.2603	3.1949	2.2778	63.675
13	Skewness of core spread	–	0.4638	-2.8824	1.0046	-0.8142
14	Mean of equity MTB ratio	–	1.1236	1.8651	1.1005	1.6685
15	75th percentile of 6m CDS spread	–	0.0046	4.1e-05	0.0097	9.1e-08
16	75th percentile of 5y CDS spread	–	0.0131	0.0146	0.0178	0.0145
17	Mean of core spread	θ	0.0140	0.0051	0.0176	0.0056
18	95th percentile of ATM put IV	η	0.6824	0.3071	0.7361	0.3710
19	Quasi-conditional variance of OTM call IV	η	0.0419	0.0123	0.0796	0.0140
20	OLS coef., CP IV gap regression	η	-0.0038	9.9e-04	0.0029	4.8e-04
21	Persistence parameter for core spread	ρ_ϵ	0.8229	0.3612	0.9325	0.4756
22	Correlation of core spread and ΔTFP	ω	0.0859	-0.0613	0.0234	-0.0922
23	Correlation of stock returns and ΔTFP	ω	0.1771	0.8195	0.2089	0.3492
24	Correlation of 5y CDS rate and TFP	ω	-0.1781	-0.5325	-0.1320	-0.7173
25	OLS coef., Merton gap 3y term struct.	ρ_ϕ	0.7184	0.8561	0.7630	0.5177
26	OLS coef., Merton gap 5y term struct.	ρ_ϕ	0.6247	0.7980	0.6876	0.4736
27	OLS coef., 1y CDS and ATM IV	μ	0.0896	0.1030	0.0930	0.1059
28	OLS coef., 6m CDS and OTM IV	μ	0.1001	0.1042	0.1142	0.1196
29	95th percentile of 1y CDS spread	ν	0.0302	0.0180	0.0282	0.0096
30	Variance of 5y total Merton spread gap	ν	0.0015	0.0023	0.0029	0.0016
31	OLS coef., Merton gap and stock returns	ν	7.3e-04	-0.0074	0.0044	0.1538
32	95th percentile of core spread	ζ	0.0313	0.0122	0.0438	0.0168
33	95th percentile of excess stock returns	ζ	0.2638	0.3838	0.2825	0.3799
34	Variance of ATM put IV	α_σ	0.0392	0.0068	0.0499	0.0116
35	Quasi-conditional variance of core spread	α_σ	4.2e-05	1.2e-05	6.8e-05	1.2e-05
36	Correlation of 6m and 5y CDS rates	α_σ	0.9583	0.9389	0.9650	0.7097
37	Variance of insurance liabilities	α_λ	0.0042	0.0022	0.0075	0.0015
38	Mean of general account liabilities	κ	0.8703	0.8791	0.8464	0.8564

Table 2.4 reports values of moments targeted in the simulated minimum distance estimations of the large and small life insurer samples. For each estimation sample, the data column provides the value of the moment in the real world data. The model column reports the corresponding moment value in simulated data. For added reference, the type column indicates for which parameter the corresponding moment is most informative, where an uncategorized moment need not be disproportionately informative for any individual parameter per se.

Table 2.5: Simulated minimum distance parameter estimates

Parameter description	Notation	Value: Large sample	Value: Small sample
Constant core spread parameter	θ	1.7295	1.8288
		(0.0052)	(0.0079)
Insurance cash flow asymmetry	η	1.0455	0.7352
		(0.0120)	(0.0237)
Persistence of insurance cash flow shock	ρ_ϵ	0.0695	0.3646
		(0.0137)	(0.0004)
Risk carrying cost parameter	ζ	67.989	74.019
		(0.2042)	(0.5551)
Risk adjustment cost	α_σ	1095.5	1298.4
		(53.400)	(94.941)
Insurance liability adjustment cost	α_λ	1.2323	1.1954
		(0.0139)	(0.0105)
Liability regulatory cost	κ	0.0002	0.0003
		(0.0000)	(0.0000)
Complement of weight on TFP shock	ω	0.7490	0.7156
		(0.0038)	(0.0001)
Bailout probability central tendency parameter	μ	0.0574	0.0403
		(0.0270)	(0.0064)
Bailout probability dispersion parameter	ν	0.1448	0.0238
		(0.0228)	(0.0109)
Persistence of bailout shocks	ρ_ϕ	0.1825	0.6309
		(0.0605)	(0.1398)

Table 2.5 reports parameter estimates recovered by simulated minimum distance from the large and small life insurer samples. Corresponding standard errors are provided in parentheses.

Table 2.6: No bailout counterfactual analysis, large life insurers

Variable and moment		Model: Baseline	Model: Counterfactual
Distress and bailouts			
Probability of distress (annual)		2.26%	2.05%
Probability of bailout (annual)		0.36%	0.00%
Shortfall conditional on distress	Mean	-1.96%	-1.91%
	Std. dev.	2.43%	2.39%
Accounting and state variables			
General account liabilities ($\lambda + d$)	Mean	87.91%	87.86%
	Std. dev.	5.10%	5.09%
Risk exposure level (σ)	Mean	1.24%	1.24%
	Std. dev.	0.08%	0.09%
Core spread (ψ)	Std. dev.	0.90%	0.88%
Stock and option pricing			
Excess stock returns (annual)	Std. dev.	35.86%	34.29%
ATM put IV, 91-day	Mean	19.69%	19.64%
OTM put IV, 91-day	Mean	29.07%	29.01%
OTM call IV, 91-day	Mean	19.63%	19.63%
Credit spreads			
CDS spread, 6-month	Mean	0.34%	0.30%
CDS spread, 1-year	Mean	0.47%	0.41%
CDS spread, 2-year	Mean	0.68%	0.63%
CDS spread, 3-year	Mean	0.84%	0.80%
CDS spread, 5-year	Mean	1.06%	1.04%

This table reports the results of a no bailout counterfactual analysis for large life insurance carriers. Statistics in the baseline column are computed on data simulated according to the estimated structural parameters. The counterfactual column reports the corresponding statistics from simulation data in which the model is solved and simulated under a modified parameterization in which bailout central tendency parameter, μ, and dispersion parameter, ν, are each set to zero.

Chapter 3

The Changing Pre-FOMC Announcement Drift: Policy Impact and Anticipation

3.1 Introduction

Seminal work by Lucca and Moench (2015), hereinafter abbreviated as LM, documents that about 80% of S&P 500 index excess returns between 1994 and 2011 were earned during the 24-hour run up to the Federal Open Market Committee's (FOMC) bi-quarterly[1] meeting announcements. Though considerable progress has been made since this discovery (e.g., Ai and Bansal, 2018; Boguth, Grégoire, and Martineau, 2019; Liu, Tang, and Zhou, 2022), the phenomenon evades comprehensive explanation. Subsequent literature also points out that this pre-FOMC announcement drift has all but vanished in more recent periods (Kurov, Wolfe, and Gilbert, 2021). The current balance of evidence invites three questions. Is the pre-FOMC drift no more? What changed after 2011? Finally, what do these changes reveal about the forces behind this stunning phenomenon?

Answers to the above questions warrant concession to the obvious: not all monetary policy decisions are created equally. This chapter embraces this point with analysis informed by recent advances in the event study literature. Using intraday stock price data extending into the COVID-19 pandemic, I first show that the pre-announcement drift has rebounded partially in more recent periods. The mean simple holding period return during 24-hour pre-FOMC windows rose from 4 basis points (bps) between April 2011 and December 2017 to 37 bps during the 2018 to 2020 period. This is compared to a 52 bps average between September 1994 and March 2011. Next, using options-based event studies, I quantify the impact of individual FOMC announcements on stock prices and the degree to which these decisions were anticipated by investors. Time series analysis of these estimates indicate a prominent role for risk premia, which constitute between 32% and 77% of the mean close-to-announcement

[1]Rare exceptions to the FOMC meeting schedule arise in extraordinary periods, as was the case toward the beginning of the COVID-19 pandemic.

ex-dividend drift[2] from 1996 to 2011. Returns from informed trading, by contrast, represent only 1% to 12% of this average. I find that as much as 30% to 55% of this sample mean was due to an anomaly that—consistent with an arbitrage explanation—disappeared after 2011. Changes in risk premia and the disappearance of this anomaly appear to be the main drivers of recent fluctuation in the pre-FOMC drift.

In the forward guidance policy era FOMC announcements are likely well-anticipated by investors. This presents an empirical challenge. Borochin, Celik, Tian, and Whited (2022) show that traditional methods to quantify the asset price impact of an event are inadequate when the event in question is highly anticipated. My methodology therefore follows in the spirit of their approach. The analysis models an FOMC decision as a binary outcome in which one of two stock price and volatility pairs, (S, σ), prevail upon realization of the announcement. In this framework, an asset price is the risk neutral probability weighted average of its present value under the two alternatives. Because options embed investor beliefs about the likelihood and impact of future events, state-contingent parameters and ex ante probabilities over states can be recovered from the distribution of options prices.

The FOMC event studies provide a time series of parameter estimates indicating which state prevailed on each announcement day and how its realization compared to ex ante investor expectations. The results explain the component of each close-to-close S&P 500 announcement day return that is attributable to the FOMC decision. To further distinguish returns driven by the arrival of policy information from risk premia, I extract physical probabilities over FOMC outcomes from the event study estimates with a power utility-based pricing kernel. With these structural components, I further decompose ex-dividend FOMC returns into close-to-announcement and announcement-to-close portions using restricted time series regressions. The decomposition analysis provides useful perspective on how much information about FOMC decisions, on average, were reflected in prices pre-announcement compared to their post-announcement counterpart. It further assesses the extent to which risk premia can account for pre-announcement returns. I repeat the empirical exercise for different calibrations of the pricing kernel.

The time series decomposition analysis yields a number of interesting findings. Between January 1996 and March 2011, 15% to 23% of close-to-close FOMC returns attributable to increases in price informativeness were earned *prior* to official release times; i.e., during the

[2] Though much of the existing literature focuses on the 24-hour pre-announcement window, my decomposition of pre-FOMC returns is limited to the close-to-announcement period. This is necessitated by data constraints as the historical options prices used in the event analyses are available only at daily frequency. As Table 3.1 demonstrates, a sizable share of the 24-hour pre-FOMC drift is earned during the close-to-announcement period. Note that for parsimony my decomposition analysis is restricted to ex-dividend returns. Dividends constitute a negligible share of pre-FOMC returns, as also shown in Table 3.1.

corresponding close-to-announcement sub-windows. In this period structural factors—which comprise information about the FOMC statement, its impact, and risk premia—account for less than 20% of the variation in pre-announcement returns. In contrast, event information factors explain about 70% of the variation in post-announcement returns during this time. This is somewhat surprising. The results run counter to the suspicion that most of the information about FOMC decisions is impounded into prices pre-announcement, as visual inspection of returns (see Figure 3.1) might suggest. Potential first impressions aside, the predominantly ex post materialization of information is consistent with theoretical implications of existing volatility and trade volume evidence (Lucca and Moench, 2015).

What cannot be explained by the analysis is informative in itself. An anomalous component of the pre-FOMC drift may be captured indirectly by intercepts for pre-announcement returns in the decomposition regressions described above. These intercept estimates for pre-FOMC returns range from 11 to 20 bps during the 1996 to 2011 period. It warrants emphasis that more reasonable pricing kernel calibrations for this sample correspond to the high end of this range.[3] For perspective the average close-to-announcement ex-dividend return during this time was 37 bps. These intercepts vanish after March 2011. The result suggests a McLean and Pontiff (2016) interpretation in which academic research exposed a pattern of mispricing that has since been claimed by arbitrage. Adding further support for this view, structural factors explain roughly twice as much variation in pre-announcement returns in the more recent sample.

The disappearance of the anomaly was a significant factor in the dramatic retreat of the pre-FOMC drift between April 2011 and December 2017. In fact, if attention is restricted to the most economically sensible pricing kernel calibrations for each sample period, it explains about 70% of the decline in close-to-announcement returns. Interpretation of this assessment comes with a caveat. The inherently nonconstructive measurement of the anomaly leaves open the possibility that it is less a mispricing phenomenon than it is a manifestation of unmodeled investor risk preferences or frictions. Data limitations notwithstanding, the balance of evidence supports the tentative conclusion that a surge in the price of FOMC risk was most responsible for the return of the pre-announcement drift after 2017.

Adding to the intrigue surrounding the decline in the pre-FOMC drift is that it occurred amid significant changes to the Federal Reserve's communication policy. Beginning in April 2011, the Federal Reserve Chair began holding press conferences after some FOMC decisions in the service of "additional transparency and accountability" (Bernanke, 2011; see also

[3]Because the pricing kernel calibrations I examine were selected to facilitate comparison across sample periods, the central tendency of economically sound estimates for any particular period may differ from what the full spectrum of reported results suggests.

Boguth, Grégoire, and Martineau, 2019). The Federal Reserve also departed from its usual practice of releasing meeting statements at 2:15 pm and began providing these announcements at other times including 12:30 pm and 2:00 pm. This raises questions about whether these changes played a role in the pre-FOMC drift's concurrent diminution. Were announcements better anticipated by investors after 2011? Did a crackdown on leaks reduce returns due to informed trading? In neither case do my results support an answer in the affirmative. I find no evidence for better market anticipation on business days prior to announcements. There is also no evidence for a reduction in informed trading post-2011.

The 2011 policy changes invite additional questions. Did the FOMC's new communication policy, as Kurov, Wolfe, and Gilbert (2021) suggest, lower a pre-announcement premium by way of reduced uncertainty? Between the prominent role of the now-vanished anomaly and a preponderance of evidence that the cost of (FOMC) risk grew after 2011, my results suggest otherwise. One might also ask what the analysis presented in this paper reveals about this potential relationship between pre-FOMC returns and the interest rate environment. Figure 3.4 illustrates, for example, that periods during which pre-announcement returns are most elevated tend to be characterized by declining interest rates. Suggestive evidence on changes in investor risk aversion imply that this apparent relationship is unlikely to be mediated through risk premia that vary inversely with interest rates. More definitive evidence runs contrary to the explanation that the relationship follows from informed trading in the run-up to surprise interest rate cuts.

Previous research by Bernile, Hu, and Tang (2016) and Kurov, Sancetta, Strasser, and Wolfe (2019) finds evidence of informed trading activity during information embargo periods before FOMC announcements. Although my results may appear to contravene these findings, there is less distance than prima facie comparison might suggest. By itself the modest contribution that informed trading makes to the pre-FOMC drift understates the evidence for informed trading activity that I detect. Mean returns due to informed trading need not be low strictly because of limited information. Another reason is that FOMC announcements need not be good news for stockholders. The event study estimates reveal that only 60% of FOMC decisions between 1996 and 2011 were unambiguously beneficial for stock prices. In addition, despite the larger presence of risk premia in the sample average, information-based factors explain more of the variation in pre-announcement returns than risk premia. It should be noted that none of the evidence in this chapter on informed trading can distinguish the effect of leaks from more deliberate informal guidance. This is worth bearing in mind as Vissing-Jorgensen (2020a,b) argues that central banks may be unable to avoid including informal communication in their policy repertoire.

A number of potential explanations for the pre-announcement drift have been proposed.

LM briefly consider risk-based interpretations, for which their empirical evidence is mixed. In more recent literature, Laarits (2019), Hu, Pan, Wang, and Zhu (2022), and Liu, Tang, and Zhou (2022) advance the risk premium view. Though my findings on risk premia lend some support for these explanations, my analysis is largely silent on potential sources of systematic risk. This is also the case for risk reallocation stories in the vein of Duffie (2010). Kurov, Wolfe, and Gilbert (2021) consider but dismiss the view that arbitrage following academic discovery of the phenomenon is what led to its recent near-disappearance. My results suggest that this hypothesis should be revisited.[4]

The good news explanation for the pre-FOMC drift posits that it reflects investor reaction to unexpectedly accommodative monetary policy. This could be reconciled with the large pre-announcement fraction of returns, LM point out, if the drift reflects information that was already publicly available but not yet reflected in prices due to rational inattention (Sims, 2003; Kacperczyk et al., 2009). More recently Morse and Vissing-Jorgensen (2020) adopt a good news view, arguing that the pre-announcement drift is part of bi-weekly FOMC cycle driven not by rational inattention but informal central bank communication. My evidence refutes the good news explanation.[5] Not only do my event study estimates imply that good news was more limited than meets the eye, I find that FOMC news enters prices primarily after the official announcements.

The remainder of this chapter is organized as follows. In Section 3.2 I extend high frequency stock price data for FOMC announcement windows through 2020 and provide an overview of recent changes in the pre-announcement drift. Section 3.3 describes the methodology and summarizes direct results for the structural event analysis of FOMC announcements. In Section 3.4 I examine event study implications for the pre-announcement drift. I conclude in Section 3.5.

3.2 The Changing Pre-announcement Drift

In this section I examine intraday S&P 500 index data that extends from 1994 through 2020. This updated sample is particularly useful for evaluating recent changes in the pre-FOMC drift because it includes the first recession observed since the conclusion of the original LM sample period. Before discussing results, I provide a brief description of my data sources.

My high frequency price data come from the Trade and Quote (TAQ) database which

[4] A small irony is that the Kurov, Wolfe, and Gilbert (2021) advocated for this arbitrage explanation in an earlier draft of their paper.

[5] Though not a good news story, my findings on risk premia may support a prominent role for informal communication if the mere possibility of leaks during pre-FOMC windows presents a source of systematic risk for which stock investors demand compensation.

contains stock-level records of S&P 500 constituents. Center for Research in Security Prices (CRSP) data provide information on shareholder distributions, among other items. Note that cum-dividend returns reported in this chapter include not only ordinary cash dividends but any distributions, cash or noncash, used in the calculation CRSP's total holding period returns. The risk-free rate comes from Ken French's daily Fama and French factors via Wharton Research Data Services (WRDS).

FOMC announcement times for the original LM sample period are drawn from the authors' Internet Appendix. I update these data through 2021 with three sources. For announcements between April 2011 and December 2016 I use Dow Jones and Thomson Reuters newswires are archived in Factiva and ProQuest. After 2016 I use official FOMC statement release times published on the Federal Reserve System website. For ease of comparison I exclude any FOMC meeting dates excluded by LM. My analysis of more recent periods discards observations for FOMC statements released outside of market trading hours. This excludes announcements on March 15, 2020, March 23, 2020, and August 27, 2020.

The effective federal funds rate is drawn from the Federal Reserve Economic Data (FRED) online database. I also use official federal funds rate targets and NBER U.S. business cycle dates published by FRED. In analysis below I sometimes represent interest rate cuts and hikes in succinct notation with indicator functions $\mathbf{1}_{cut}$ and $\mathbf{1}_{hike}$, respectively. Likewise, $\mathbf{1}_{no\ change}$ indicates no change in the federal funds rate target. Empirical analysis in subsequent sections of this chapter employ data from a few additional sources. Historical options data come from OptionMetrics, which are available from January 1996 through December 2021. These data are necessary for the FOMC event studies I describe in the next section. The Chicago Board Options Exchange's VIX index is provided by Yahoo Finance.

A brief note on sample nomenclature is in order. Due to slight variation in historical coverage among data sources, sample labels in this chapter are context-dependent. The original LM data extended from September 1994 to March 2011. For analysis involving options data, my examination of this earlier period is limited to FOMC announcements after January 1996 (inclusive). References to the LM sample in this chapter may refer to the full period beginning in 1994 or the limited period beginning in 1996. Similarly, the post-LM sample refers to the April 2011 to December 2020 period for analysis involving intraday TAQ data, and the extended April 2011 to December 2021 period otherwise. For instances where sample labels are not clear from the context I specify the relevant dates.

My investigation begins with visual assessment of average cumulative excess returns for the S&P 500 index over three-business-day windows centered on FOMC announcement dates. To highlight how the pre-FOMC drift has evolved over time, I group the analysis into three separate periods. Figure 3.1 displays cumulative FOMC returns from September 1994 to

March 2011, the original LM sample period. This chart is analogous to Figure 1 in LM. Consistent with previous research S&P 500 returns exhibit a large positive drift in the 24-hour run up to FOMC announcements. Though the drift has grown familiar since its first depiction in the literature, it is an arresting reminder of why LM's discovery remains compelling. Figure 3.1 serves an additional purpose. Because LM use intraday S&P 500 data from alternate sources (Thomson Reuters TickHistory and Tickdata.com), it provides visual assurance that the TAQ-based methodology employed in this chapter accurately replicates their findings.

Corresponding charts for the April 2011 to December 2017 period and the January 2018 to December 2020 period are presented in Figures 3.2 and 3.3, respectively. In addition to highlighting recent changes in the pre-FOMC drift, this post-LM period split in 2018 aids comparison with previous literature. This earlier sub-period overlaps closely with Boguth, Grégoire, and Martineau (2019), whose extended sample runs from April 2011 to September 2017. It is clear in Figure 3.2 that between 2011 and 2017 the pre-FOMC drift is not what it once was. Though much less dramatic than what is observed in the earlier sample, it does not appear to have vanished entirely. Three other features are apparent. First, compared to the LM period a more significant fraction of pre-statement returns is earned early in the business day prior to announcement dates. Boguth, Grégoire, and Martineau (2019) document a similar finding. The second feature is that S&P 500 returns appear to jump close to the market open on statement dates. Third, FOMC returns are less persistent during this period as ex post business days exhibit some evidence of reversals. Confidence bands also expand noticeably on these ex post days.

Figure 3.3 shows that although the pre-announcement drift remains subdued in comparison to the LM period it has rebounded partially since 2017. Magnitudes aside, cumulative FOMC returns during this time inherent many of the features that characterize the interim 2011–2017 period. The apparent jump at market open on FOMC announcement dates remains, as does visual evidence of reversals—though reversals appear to commence earlier in this sample. My analysis in this chapter does not directly address these jumps or reversals. The reversals do nonetheless corroborate the limited role of informed trading in the pre-statement drift I document in Section 3.4, at least for any case in which FOMC statements have a lasting impact. Investigation of these features in future research may also prove useful for understanding the nature of the pre-announcement premium quantified below.

The full sample time series for pre-FOMC announcement returns is plotted in Figure 3.4. This chart also includes a 12-month moving average (MA) of these returns. As in LM, pre-announcement returns shown are computed over a 24-hour window ending fifteen minutes prior to the FOMC statement release time. A minor point of departure with earlier literature is that these returns are simple holding period returns rather than log returns. For added context the chart also includes the effective federal funds rate. Recession periods defined by NBER business cycle designations are shaded. While it may be tempting to associate muted pre-announcement returns with protracted low interest rate environments Figure 3.4 reveals that this relationship is not without exception. It does appear, however, that periods of elevated returns often cluster near recessions. The juxtaposition of 12-month MA returns with raw values also indicates that the pre-FOMC drift's rebound is driven to significant degree by one unusually large return early in the COVID-19 pandemic. This is worth bearing in mind when considering the recent increase in the drift.

To assess the pre-FOMC drift throughout the updated sample in more precise terms I tabulate sample statistics for each time period. These statistics, which include sample averages, standard deviations, and an array of percentiles, are presented in Table 3.1. As with Figure 3.4, returns are computed as simple holding period returns. For 24-hour cum-dividend pre-announcement returns, the LM period average is 52 bps. Notwithstanding the minor methodological differences, this is in line with values reported in previous literature. Conversely the mean pre-FOMC return for the post-LM period, April 2011 to December 2020, is 14 bps. A comparison of sub-periods confirms that much of this decline is concentrated in the earlier portion of the post-LM sample. Whereas the 2018–2020 average according to this return definition is 37 bps, it is only 4 bps in the April 2011 to December 2017 period. Sample averages for alternative return definitions exhibit a broadly similar pattern across time periods, although the gulf between the LM drift and the post-LM drift is narrower for close-to-announcement returns.

Another salient feature of Table 3.1 is that 24-hour pre-announcement returns appear considerably more volatile during the LM period, at least in absolute terms. For example, the standard deviation of cum-dividend returns of this kind is 132 bps in the LM sample compared to 67 bps in the post-LM sample. This difference is only remarkable for 24-hour pre-announcement returns. The standard deviation for close-to-announcement cum-dividend returns is only 69 bps for the LM period. The corresponding value is 62 bps for the subsequent period. The much lower volatility for close-to-announcement returns in the LM sample is notable because it is not obvious from Figure 3.1 that this would be the case. On the other contrary, the post-LM period's muted sensitivity to holding period in this regard is less surprising. In this case it is likely due to the prominent influence of the "market open jump" described above—a phenomenon that is included in all of the return definitions I examine. It should be cautioned that close-to-announcement returns may be more difficult to compare since variation in statement release times after March 2011 implies that they can reflect meaningfully different holding periods.

The evidence presented in this section confirms that the pre-announcement drift has declined considerably since the conclusion of the original LM sample period. Examination of the most recent data reveals that the drift has recovered to a limited but nontrivial degree. I also show that this partial rebound is due in large part to returns earned during the COVID-19 pandemic. To better understand what drives variation in pre-FOMC returns I perform event study analysis of FOMC announcements in the next section.

3.3 Impact and Anticipation of FOMC News

This section analyzes FOMC announcements from an event study perspective. The purpose is to estimate both the impact of individual FOMC statements on stock prices and the extent to which the outcomes were anticipated by investors. After describing the methodology I provide a summary of event study results. In the next section I examine these results further, with an emphasis on their implications for the pre-announcement drift.

3.3.1 Methodology

My approach in this section is similar to existing methodology in the event study literature. The theoretical framework is virtually identical to the one employed in Borochin et al. (2022). Minor departures in terminology are shaped by implementation; whereas Borochin et al. (2022) investigate the heterogeneous impact of an aggregate event at the individual stock-level, all event analysis in this chapter is performed at the aggregate index level. Here I provide a brief overview of the model before describing how I identify and estimate the parameters in my empirical setting.

Consider a stock index facing an uncertain future event—an FOMC announcement in the present context. Upon realization of the announcement, assume that the price behavior of the index under the risk-neutral probability measure converges to one of two geometric Brownian motions with corresponding index price and volatility parameters $(S_\theta, \sigma_\theta)$, $\theta \in \{u, d\}$. To fix nomenclature, let $S_u \geq S_d$ so that states u and d may be referred to as the upside and downside outcomes, respectively. This is without loss of generality.[6] Though somewhat arbitrary, these labels help to simplify discussion of the results that follow. The price process for the stock index may be expressed piecewise:

[6]The inequality is not to be confused with identifying assumptions which are described later in this section. Because this paper focuses on the aggregate response of the S&P 500 with analysis restricted to the index level, I anchor terminology to the question of whether or not the index value is increased by the FOMC decision. This means that the key identifying challenge in the estimation here is not determining if $S_u \geq S_d$ or $S_u < S_d$ as is sometimes the case in stock-level event research (e.g., Borochin et al., 2022). Rather, the challenge is determining whether the S&P 500 converged to state u or d. For the present application the distinction is semantic.

$$dS = \begin{cases} r_f S_u dt + \sigma_u S_u dW, & \text{with probability } q \\ r_f S_d dt + \sigma_d S_d dW, & \text{otherwise.} \end{cases} \tag{3.1}$$

Above, r_f denotes the instantaneous risk-free rate.

At any time prior to the FOMC announcement, the price of the index and any derivative thereof is a function of investor expectations of model parameters $\{q, S_u, S_d, \sigma_u, \sigma_d\}$. In particular the ex ante price of any such asset is the risk-neutral probability weighted average of its price under the two alternative scenarios. The index value at time t, for example, is $S_t = \mathbb{E}_t[q]\mathbb{E}_t[S_u] + (1 - \mathbb{E}_t[q])\mathbb{E}_t[S_d]$. Price equations for derivative securities referencing the stock index are analogous. This highlights an attractive feature of the theoretical framework. For any set of state-contingent parameters and corresponding event probability, ex ante stock and European option prices admit closed-form solutions. Note that it need not be the case that expectations of individual parameters are constant throughout the ex ante period. In fact, all ex ante price fluctuations in this framework necessarily follow from changes in expectations about q or about state-contingent prices or volatilities.

Discrete time series of ex ante stock or option prices may therefore be rationalized by the model with a corresponding time series of parameters $\{\mathbb{E}_t[q], \mathbb{E}_t[S_u], \mathbb{E}_t[S_d], \mathbb{E}_t[\sigma_u], \mathbb{E}_t[\sigma_d]\}$ at the same frequency or lower. Post-announcement stock and options prices are likewise characterized more parsimoniously with state-contingent parameters $\{S_u, \sigma_u\}$ if the price process converges to upside outcome u, and $\{S_d, \sigma_d\}$ otherwise. As I explain shortly, my identification strategy makes use of ex post price information to pin down the outcomes associated with individual FOMC announcements.

Given a consistent notion of which state ensues after a particular FOMC decision, the above parameters may be recovered from the ex ante prices for an array of index options expiring after the announcement. The high-level intuition is straightforward. An option price embeds investor expectations about possible future events that may arise over the course of its maturity horizon. Because features of the underlying asset payoff distribution influence options at contrasting levels of moneyness differently, investor beliefs about the likelihood and impact of future events are reflected in the distribution of option prices. An out-of-the money put option, for example, is more sensitive to a decline in the downside state-contingent price S_d than an in-the-money put. It is also easy to see that the prices of at-the-money calls are generally more informative about upside state-contingent parameters than downside counterparts, while the reverse is the case for at-the-money put options.

The main identification challenge may be summarized by two questions which are, in the present context, equivalent. For a particular FOMC announcement, did the price behavior converge to the upside state or the downside state? How does one know that the parameter estimate for $\mathbb{E}[S_u]$ that rationalizes ex ante prices accurately concerns investor expectations about upside state price S_u rather than the downside state price, S_d? This risk of label switching is a known problem of mixture models (Stephens, 2000; Jasra, Holmes, and Stephens, 2005), including the present binary event framework (Borochin and Golec, 2016; Borochin et al., 2022). The issue in general is more insidious than meets the eye as Stephens (2000) illustrates that even "obvious" solutions need not resolve the problem.

In this paper, I address the label switching problem with a rational expectations approach. My method imposes two conditions on the estimation. First, I require that the converged model explain ex post prices. If the state converges to outcome d, for example, state-contingent parameters $\{S_d, \sigma_d\}$ alone must adequately fit post-announcement stock and options prices. Second, ex ante expectations of state-contingent parameters must be within a reasonable neighborhood of their ex post values. In other words, one of the two pairs of state-contingent parameters that characterize the cross-section of pre-announcement options prices must describe ex post prices better than the alternative pair. For each FOMC event I therefore estimate the model subject to these conditions under both potential convergence scenarios. My approach then infers what outcome carried the day from which fitted scenario better rationalizes both ex ante and ex post data. The identifying assumption is that the ex ante expectations of the parameters to which the state converged must be significantly nearer to their true (ex post) realizations than pre-announcement expectations of the alternative state-contingent parameters.[7] Hence, a form of rational expectations is required of investors. By anchoring terminal states to post-announcement data this approach also mitigates the risk of contamination from investor beliefs about unrelated future events.

For each individual FOMC announcement in my sample I estimate the above parameters for the S&P 500 index using daily historical data. The ultimate goal of the exercise is to shed light on pre-announcement FOMC returns with an emphasis on time series heterogeneity. I therefore perform the estimation at the aggregate index level only using index options.[8] My

[7] For some intuition on how the data are informative about the outcome realization, consider the second identifying condition above in maximally restrictive form: i.e., assume investors have perfect foresight about state-contingent parameters. In this case, conditional on convergence to outcome θ, the estimation requires that $\mathbb{E}_t[S_\theta] = S_\theta$ and $\mathbb{E}_t[\sigma_\theta] = \sigma_\theta$ for all t prior to the announcement. This means that the *only* way for the model-implied return on the index to be positive on the FOMC date is if the state converges to upside outcome u. A positive announcement day index return in real world data accordingly influences the estimates in this direction. Likewise if population parameters satisfy $\sigma_d > \sigma_u$, ex post options data characterized by elevated option-implied volatility would be more consistent with convergence to downside scenario d. Insofar as investor expectations are sufficiently accurate, this intuition extends to cases in which the assumption of perfect foresight about state-contingent parameters is relaxed.

[8] This approach is less computationally taxing than one that makes direct use of data for index constituents. It also capitalizes on the enhanced liquidity of index options.

research design also focuses on one-business-day ex ante parameters, since expectations as of earlier dates are unnecessary for analyzing close-to-announcement FOMC returns. To simplify the discussion that follows let ψ represent the model parameter vector, which includes all ex ante expected parameters and the pair of ex post state-contingent parameters.

The estimation targets price and return data for a balanced set of N call options and N put options in addition to the index itself. Return targets are limited to announcement day returns. I also target prices for one ex ante business day and $\tau + 1 \geq 1$ ex post business days, inclusive of the announcement date. Option prices are targeted for the ex post period only. That is, ex ante information about options enters the estimation exclusively via returns. Subscript $t \in \{T-1, T, \ldots, T+\tau\}$ indexes dates in the event window for FOMC announcement date T. Let $p_{i,t}(\psi)$ and $r_{i,t}(\psi)$ denote the model-implied price and return, respectively, for asset i at time t. The corresponding price and return observed in the data are denoted by $\hat{p}_{i,t}$ and $\hat{r}_{i,t}$. Finally, $w_{i,t}^{j}$ is the estimation weight for target type $j = p, r$ corresponding to asset i at time t. As I detail shortly, weights are set to establish parity across targets in the objective function. Conditional on the final convergence scenario, parameters for the FOMC event are estimated by minimizing the following objective:

$$\min_{\psi} \left\{ \sum_{i=1}^{2N+1} \left(w_{i,T}^{r} \left| r_{i,T}(\psi) - \hat{r}_{i,T} \right| + \sum_{t=T-1}^{T+\tau} w_{i,t}^{p} \mathbf{1}_{i,t}^{p} \left| p_{i,t}(\psi)/\hat{p}_{i,t} - 1 \right| \right) \right\}, \tag{3.2}$$

where $\mathbf{1}_{i,t}^{p}$ is an indicator function that is one if the price target i at time t is included in the objective function and zero otherwise. Because ex ante options prices per se are not targeted, $\mathbf{1}_{i,t}^{p}$ is zero when $t = T-1$ and i is an option. Parameters are estimated under both convergence scenarios subject to the rational expectations conditions described above. The FOMC outcome type is then determined by which convergence scenario provides the best objective value.

In words, parameters are estimated by minimizing the weighted absolute percent errors for targeted prices and the weighted absolute errors for targeted returns in the event window. As follows from the discussion above most of the identifying information for ex ante expected parameters comes from the option return targets. The converged state-contingent parameters are pinned down primarily by ex post price targets, which are also important for ascertaining the outcome type. Notwithstanding differences in the label identification strategy, the loss function itself is a natural extension of the stock-level absolute error objectives in existing event study literature (Borochin and Golec, 2016; Borochin et al., 2022). My inclusion of return targets and the allowance for targets to be weighted by type comprise the main source of distinction in this regard.

For each business day in the estimation window $[T-1, T+\tau]$, I use $N = 5$ index options of each exercise type. I restrict the analysis to short-term options with time to expiration ranging from 20 days (inclusive) to 90 days (exclusive). Options of each type are selected according to their liquidity and their contribution to overall dispersion in moneyness. To allow the estimation to target returns, options in the sample are organized into overlapping two-day selection windows. My estimation targets date T log returns for the S&P 500 index and all ten index options that meet the sample inclusion criteria for the $[T-1, T]$ return period. One ex ante price target is used for the level of the index. For each day in the $(\tau+1)$-day ex post window, I also target the level of the index and prices of the ten options with two-day selection windows ending on that day. I perform the estimation using ex post windows of $\tau + 1 = 3$ business days. The main results presented in the next section are not sensitive to either the choice of N or τ.

The rational expectations condition described above requires that the converged state-contingent parameters are within a neighborhood of their ex ante expected values. My implementation constrains the log return on S_θ implied by its deviation from $\mathbb{E}[S_\theta]$ to be within the 1st and 99th daily percentiles predicted by the ex ante average implied volatility (IV) of at-the-money (ATM) options. Note that I suppress $T-1$ subscripts for the expectation operator in the interest of succinct notation. Likewise, I require that σ_θ and $\mathbb{E}[\sigma_\theta]$ deviate no more than the 1st and 99th daily percentiles according to the lagged 10-day moving standard deviation of ATM option IV. The disciplining device therefore establishes a time-consistent notion of outcome scenarios while accommodating imperfect investor foresight.

Depending on its moneyness the volatility of an option price can differ significantly with the underlying asset and with other options of similar maturity and exercise type. This means that in practice an unweighted implementation of the error loss function above would *implicitly* weight high-volatility options more heavily than other assets. To mitigate this issue, my objective function weighs each error by historical return volatility for the corresponding security type. For the full sample of short-term S&P 500 index options in OptionMetrics, I categorize options according to their exercise type and level of moneyness. Next, I compute the average 10-day moving standard deviation of log returns for each option category. Weights $w_{i,t}^j$ for options used in the estimation are then calculated as the inverse of this volatility measure adjusted to daily frequency and lagged by one business day.[9] Similarly, weights for underlying index errors are computed as the inverse of lagged daily volatility implied by the 10-day moving standard deviation of S&P 500 log returns.

In summary, for an FOMC announcement released on date T the estimation provides

[9] It is necessary to compute option weights according to their category as the most liquid options in the event window need not in general have sufficient price history to estimate their security-specific volatility.

$T-1$ ex ante investor expectations for the parameters $\{q, S_u, S_d, \sigma_u, \sigma_d\}$. It further yields date T estimates for either $\{S_u, \sigma_u\}$ or $\{S_d, \sigma_d\}$, depending on the state to which the outcome converged.

3.3.2 Event Study Results Overview

Before delving into more precise implications for the pre-announcement drift, I first provide an overview of the event study results. This surface level look addresses two questions. First, what are the time series characteristics of FOMC announcement impact and investor anticipation thereof? Second, how do the event study estimates relate to observable features of the economy, stock market, and monetary policy?

As a first pass method to quantify investor anticipation I examine the extent to which risk-neutral probability $q_{T-1} \equiv \mathbb{E}_{T-1}[q]$ accurately predicts the outcome. Define $\Delta q_T = \mathbf{1}_u - q_{T-1}$, where $\mathbf{1}_u$ is an indicator function that is one of the state converged to the upside outcome and zero otherwise. Abstracting from risk aversion, larger magnitudes of $\Delta q_T \in [-1, 1]$ indicate that investors were more surprised by the announcement. Figure 3.5 provides a time series of Δq_T estimates for FOMC announcements between 1996 and 2021. Note that Δq_T is negative only in the event of bad news (i.e., θ converges to state d) and positive only in the event of good news (i.e., θ converges to state u). The most salient feature of this time series is that investor anticipation of FOMC decisions appears quite volatile. The chart suggests that some announcements were very well anticipated whereas others were near total shocks, with few clear patterns. To get a sense for how broader trends in anticipation evolved over time, the figure also includes a 12-month moving average of $|\Delta q_T|$. The overall level of surprise appears elevated in the post-LM period, but not to a statistically significant extent.

The dearth of visual evidence for a structural break around 2011 may be informative in itself. As noted above in Section 3.1, the Federal Reserve instituted changes to its FOMC communications policy in 2011 (see, Boguth, Grégoire, and Martineau, 2019). My event study results are at variance with the idea that this communication policy change coincided with a significant shift in the market anticipation of FOMC announcements. Figure 3.5 also does not support the view that sophisticated investors have leveraged technological advances to improve greatly their prediction of FOMC decisions over time. These observations notwithstanding, the time series of Δq_T provides few clues about the evolution of the pre-announcement drift.

One way to measure the FOMC's scope for impacting stock prices is to consider the distance between $\mathbb{E}[S_u]$ and $\mathbb{E}[S_d]$. Figure 3.6 presents a time series of this state-contingent price difference scaled by the model-implied ex ante index level. The chart additionally includes a 12-month moving average of this measure, which I refer to as the impact gap. As with the previous figure, this metric varies dramatically from one FOMC announcement to the next. Suggestive of nontrivial crash risk, there are a number of FOMC announcements with very large impact gaps upwards of 20%. This is particularly the case between 1996 and early 2001, which invites speculation about how the rise and fall of the dot-com bubble might relate to these findings. Over the full course of the sample the impact gap exhibits a broadly declining trend. Unsurprisingly, the most notable exceptions to this trend materialize toward the end of the 2008–2009 recession and in the midst of the COVID-19 pandemic period.

Table 3.2 provides basic sample statistics for the event study estimates grouped by sample period. For additional context, this table also includes statistics for close-to-announcement returns, r^e_{pre}, and announcement-to-close returns, r^e_{post}, as well as other variables related to monetary policy and the economy. Note that while my intraday data extends from 1994 to 2020, my options data—and hence event study estimates—range from 1996 to 2021. Because Table 3.2 reports the union of these samples, variables from separate sources should be compared with caution. In general these statistics confirm the patterns, or lack thereof, described in my analysis above. There are slightly fewer instances of good news announcements for stockholders, as represented by indicator variable $\mathbf{1}_u$, in the post-LM period (61%) compared to the LM period (64%). The statistics also reveal that the downward trend in the impact gap is driven in roughly equal measure by changes in S_u and S_d. Aside from serving as a reminder that the LM sample contained more recessionary periods and more instances of rate cuts than the post-LM sample, the statistics in this table are largely unremarkable.

To see how the event study estimates relate to observable measures of monetary policy, Table 3.3 provides a cross-tabulation of federal funds rate target decisions and model-implied announcement outcomes. Of the 25 interest rate cuts observed in my sample from 1996 to 2021, the event studies classify 15 of these announcements as upside scenarios for stockholders. Insofar as low interest rates are typically considered accommodating to stocks, this is broadly in line with what one might expect to find. Somewhat more surprising is what the event studies imply about rate increases. Of the 33 interest rate hikes between 1996 and 2021, only 13 represent downside stock market scenarios according to the model. While this may raise concerns about the methodology's capacity to distinguish good news from bad, it is worth bearing in mind that the magnitude of rate changes throughout this period are highly asymmetric as Figure 3.4 points out. The results suggest that for many rate hike decisions, stock investors held ex ante concerns about a more aggressive tightening regime.

The rate hike findings are also a reminder that the information contained in FOMC statements may extend well beyond contemporaneous interest rate targets. This is underscored by the instances in which no rate changes were announced. For the full sample, and the LM sub-period in particular, almost twice as many no change announcements were categorized by the model as good news for stocks. Even in the most narrow interpretation of how central banks affect stock markets, the result supports the idea that the projected future path of interest rates often dominates.

For a sharper perspective on how the event studies interpret FOMC news, I now turn brief attention to model-implied surprises. I define an FOMC announcement as a surprise if Δq_T exceeds one standard deviation. Table 3.4 provides sample statistics conditional on surprise announcements. The statistics are partitioned by the type of surprise: good news and bad news. In this case the relationship between model-implied news and contemporaneous federal funds targets remains nuanced but appears less counterintuitive. Of the 18 bad news shocks between 1996 and 2021, 11% coincided with interest rate cuts and 22% coincided with interest rate hikes. Among the 35 good news shocks during this period, 17% coincided with interest rate cuts and 14% coincided with interest rate hikes. It is also worth pointing out that only 6% of bad news surprises occurred during recessions compared to 14% for good news surprises. This is consistent with the view that the Federal Reserve may be particularly careful not to alarm financial markets during periods of economic distress.

Table 3.4 also provides some clues about the relationship between the pre-FOMC drift and information-based returns. On the one hand, at 40 bps the average pre-announcement returns are almost twice as large for good news surprises compared to bad news (22 bps). This may be a sign that some informed trading occurred in the run up to these announcements. However this is difficult to judge without assessing the scope for risk premia, as the greater overlap of recessionary periods and good news shocks underscores. On the other hand, the statistics reveal that for both types of surprises the post-announcement fraction of returns dominates. The average post-announcement return is -90 bps for bad news shocks and 85 bps for good news shocks. This hints at an important finding that I document in the next section: most information about FOMC announcement outcomes enters stock prices ex post.

To assess briefly how model-implied news and investor anticipation relate to FOMC returns, I report correlation coefficients for close-to-announcement returns, announcement-to-close returns, and estimates for q_{T-1}, and the outcome. These correlations, which are partitioned by sample period, are presented in Table 3.5. As expected, ex post returns are positively correlated with upside outcomes in both sample periods. This is also the case for ex ante returns, which is again suggestive of pre-announcement informed trading. For a more quantitative perspective, linear probability model estimates (not otherwise reported) reveal that a one standard deviation shock in ex ante returns is associated with a 9.6% change in the probability of an upside FOMC announcement during the LM sample period. In the post-LM sample period, a one standard deviation shock in ex ante returns is associated with a 12.8% change in the probability of an upside outcome. Interestingly, there is no statistically significant correlation between ex ante returns and ex post returns in either period. Simple correlations provide no evidence, therefore, that the reversal patterns described above begin prior to the close of announcement days.

3.4 Implications for the Pre-FOMC Drift

I now focus attention on a more precise assessment of what the FOMC event studies imply for the pre-announcement drift. My analysis uses vector time series regressions of pre-announcement returns and post-announcement returns on model-implied FOMC return components. The analysis asks whether, and to what extent, the event study estimates can explain the ex ante and ex post fraction of announcement day returns.

3.4.1 Return Decomposition

The preceding event study results provide only a risk neutral perspective on the relationship between FOMC announcements and stock prices. If there exists an FOMC risk premium, as other literature argues (e.g., Laarits, 2019; Hu et al., 2022; Liu, Tang, and Zhou, 2022), these components alone will provide an incomplete framework for examining pre-announcement returns. The analysis in this section therefore allows for the presence of risk premia with the added treatment of investor risk preferences. Suppose that assets are priced by a representative investor who has isoelastic utility over wealth:[10]

$$u(w) = \frac{w^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0, \tag{3.3}$$

where w is the value of the wealth portfolio. Under the additional assumption that the S&P 500 index is a reasonable stand-in for the wealth portfolio, the physical probability, p, of the upside FOMC outcome can be extracted from the event study estimates by

$$p = \frac{q}{(1-q)\,(S_d/S_u)^{\gamma} + q}. \tag{3.4}$$

I conduct all analysis in this section according to three different risk aversion parameter calibrations: $\gamma = 3, 5, 8$. As I elaborate below, these these values are based on how aspects of the fitted model compare to theoretical predictions.

Now consider the composition of holding period returns on FOMC announcement days.

[10]This is equivalent to power utility over consumption in a two-period ($t = 1, 2$) representative agent model, as in Liu, Tang, and Zhou (2022). Under this interpretation, it follows from market clearing that the value of $t = 2$ consumption is equal to that of the wealth portfolio.

To simplify the analysis I focus on ex dividend returns. For convenience, denote the gross risk-free rate by $R_f \equiv 1 + r_f$. Let $\mathbf{1}_\theta$ represent the indicator function that state $\theta \in \{u, d\}$ ensues upon realization of the announcement on date T. Suppose that the component of the return that is unrelated to the FOMC event may be represented as an additive error term, e_T, that satisfies $\mathbb{E}_{T-1}[e_T] = 0$. This mean-zero error assumption implicitly requires that risk premia unrelated to FOMC statements are negligible during the one business day windows containing these announcements. Under the binary event model in Section 3.3, the realized excess announcement day return, $r^e \equiv r - r_f$, may then be decomposed as

$$\begin{aligned} r^e &= \left(\frac{\mathbf{1}_u S_u + \mathbf{1}_d S_d}{S_{T-1}} + e_T \right) - R_f \\ &= \frac{\mathbf{1}_u S_u + \mathbf{1}_d S_d - \mathbb{E}_{T-1}[S_T]}{S_{T-1}} + \left(\frac{\mathbb{E}_{T-1}[S_T]}{S_{T-1}} - R_f \right) + e_T. \end{aligned} \tag{3.5}$$

The first term in Equation 3.5 represents the component of the return due to the resolution of uncertainty about the state. The second term in this equation is the FOMC risk premium, $\mathbb{E}_{T-1}\left[r - r_f\right]$.

Realized returns can be further broken down to distinguish the expectational error about the terminal state-contingent price from the collapse in probabilities over states:

$$r^e = \varepsilon_p + \varepsilon_\theta + \mathbb{E}_{T-1}\left[r - r_f\right] + e_T, \tag{3.6}$$

where

$$\begin{aligned} \varepsilon_p &\equiv \frac{(\mathbf{1}_u \mathbb{E}_{T-1}[S_u] + \mathbf{1}_d \mathbb{E}_{T-1}[S_d]) - \mathbb{E}_{T-1}[S_T]}{S_{T-1}}, \\ \varepsilon_\theta &\equiv \frac{\mathbf{1}_u \left(S_u - \mathbb{E}_{T-1}[S_u]\right) + \mathbf{1}_d \left(S_d - \mathbb{E}_{T-1}[S_d]\right)}{S_{T-1}}. \end{aligned}$$

In the results that follow, the distinction between ε_p and ε_θ proves less than illuminating. However, separation of these factors can accommodate more nuanced interpretations in principle. For example, it is possible that whereas ε_p reflects largely nonpublic knowledge pre-announcement, ε_θ is driven in part by publicly available information about the economy's sensitivity to the FOMC event. In other words, a prominent role for ε_θ in ex ante returns may constitute less conclusive evidence for leaks. It might further be hypothesized that asynchronous counteraction between ε_p and ε_θ could shed light on the post-announcement reversals described in Section 3.2. To summarize, close-to-close announcement day returns consist of an error component unrelated to the FOMC event, e_T, a risk premium, $\mathbb{E}_{T-1}[r - r_f]$, and two information factors, ε_p and ε_θ, which capture the distance between the ex post outcome and investors' pre-announcement expectations over states.

Abstracting from second order compounding terms, r^e is the summation of the pre-announcement return, r^e_{pre}, and the post-announcement return, r^e_{post}. A natural question is to ask how the close-to-close structural factors of Equation 3.6 are allocated on average across these ex ante and ex post portions. This question may be answered at the aggregate level with time series projections of these intraday returns on estimates for ε_p, ε_p, and $\mathbb{E}_{T-1}\left[r - r_f\right]$. Suppressing unnecessary time subscripts, I examine vector projections of the form

$$\begin{pmatrix} r_{pre} \\ r_{post} \end{pmatrix} = \begin{pmatrix} \beta_0^{pre} \\ \beta_0^{post} \end{pmatrix} + \begin{pmatrix} \beta_p^{pre} \\ \beta_p^{post} \end{pmatrix} \hat{\varepsilon}_p + \begin{pmatrix} \beta_\theta^{pre} \\ \beta_\theta^{post} \end{pmatrix} \hat{\varepsilon}_\theta + \begin{pmatrix} \beta_{\mathbb{E}r}^{pre} \\ \beta_{\mathbb{E}r}^{post} \end{pmatrix} \times \hat{\mathbb{E}}_{T-1}\left[r - r_f\right] + \begin{pmatrix} e^{pre} \\ e^{post} \end{pmatrix}, \tag{3.7}$$

subject to the condition that both pre- and post- variants of the coefficients β_p, β_θ, and $\beta_{\mathbb{E}r}$ are nonnegative. I define the ex ante return window as the period between the previous day's market close and fifteen minutes prior to the FOMC announcement time.[11] The ex post return window extends from fifteen minutes before the announcement to the market close. I estimate this regression with the following additional restrictions:

$$\beta_p^{pre} + \beta_p^{post} = 1, \tag{3.8}$$

$$\beta_\theta^{pre} + \beta_\theta^{post} = 1, \tag{3.9}$$

$$\beta_{\mathbb{E}r}^{pre} = 1, \;\; \beta_{\mathbb{E}r}^{post} = 0. \tag{3.10}$$

The first two restrictions above require that there is no "overcounting" of information-based components in the total daily return. The corresponding coefficients therefore have the interpretation of ex ante and ex post weights. The third restriction follows from theory. Because uncertainty about the FOMC announcement is resolved at the beginning of the ex post window, any risk premium tied to the event should be earned entirely in the ex ante potion. That is to say, $\beta_{\mathbb{E}r}^{post} = 0$. Further, if $\beta_{\mathbb{E}r}^{pre}$ were greater (less) than one, investors must have been more (less) risk averse than the pricing kernel calibration implies.

If the event model estimates and the pricing kernel calibrations are appropriate for the analysis at hand, the regression coefficients should lend themselves to the above restrictions without difficulty. For some perspective in this regard, I estimate the model in both restricted

[11] This definition's exclusion of the fifteen minute period immediately prior to the FOMC statement release time is in line with previous literature (Lucca and Moench, 2015).

and unrestricted form. The constraints also guide the range of pricing kernel calibrations. Specifically, in neither subsample do the data reject the restriction in Equation 3.10 at the 5% significance level for any of the γ values examined.

The analysis hinges critically on the soundness of the event study estimates in Section 3.3. If the estimates are not reasonable proxies for their true corresponding values, for example, the scope for measurement error bias in the regression above may be as unmanageable as it is pervasive. While restriction fit may provide some reassurance, this concern reaches beyond the purview of constraints. Sensitivity analysis provides one way to gauge validity. In robustness tests available upon request, I find that the results in this section are not sensitive to two important parameters of the event estimation: the number of ex post days used for price targets and the number of index options used. The range of different γ calibrations provide another dimension along which sensitivity may be evaluated. Although some results are more sensitive in this regard, unrestricted estimates of $\beta_{\mathbb{E}r}^{pre}$ and $\beta_{\mathbb{E}r}^{post}$ are useful for identifying the most economically sound specifications for each sample period.

3.4.2 Results

Before delving into the results, I note that Table 3.6 reports sample statistics for the dependent and independent variables. To allow for direct comparison with the regression sample, the statistics are computed not on the union of all available data but the intersection. I therefore use statistics from this table for the purpose of interpreting results. One aspect of Table 3.6 worth noting is that statistics are reported for ε_p and $\mathbb{E}_{T-1}[r - r_f]$ for each calibration of the pricing kernel. This provides additional quantitative perspective on how different calibrations for risk aversion parameter γ translate to stock return composition for each sample period.

Tables 3.7 and 3.8 present the regression results for the LM and post-LM sample periods, respectively. In each table, Panel A provides estimates for the restricted model whereas Panel B reports estimates for the unrestricted model. Columns are grouped by pricing kernel calibration and sub-categorized by dependent variable (i.e., ex ante and ex post returns). Panel A also provides p-values for linear restrictions in Equations 3.8 and 3.9 and the first of the two equality restrictions in Equation 3.10 (i.e., $\beta_{\mathbb{E}r}^{pre} = 1$). Also included in Panel A are partial R^2 values for each of the main explanatory variables. These partial R^2 estimates quantify how much of the variation in the dependent variable that is unexplained by all other regressors is explained by the one in focus. Newey-West standard errors are reported in parentheses. It should be cautioned, however, that these errors are not corrected for the sampling variability of the event estimations. In my discussion that follows results refer to the restricted model estimates unless otherwise noted.

I begin with a description of the results for the LM sample period. Table 3.7 reveals that overall the model is well-behaved. This is particularly the case for lower calibrations for risk aversion parameter γ. Unrestricted estimates for β_p and β_θ are close to their restricted counterparts in all cases. With a point estimate of 1.09, the unrestricted value for $\beta_{\mathbb{E}r}^{pre}$ is nearest to its restricted value (one) under the $\gamma = 3$ parameterization. None of the three main restrictions are rejected by the data for any pricing kernel calibration that I consider. It should be noted, however, that contrary to theory the unrestricted value for $\beta_{\mathbb{E}r}^{post}$ is statistically different from zero for the $\gamma = 7$ model. Total R^2 values for the restricted model are declining in γ, which lends added support for the view that the $\gamma = 3$ calibration is the most suited of the three for this sample period.

Three features of Table 3.7 are immediately salient. The first is the large and significant positive intercept for ex ante returns. This intercept, which indicates the presence of an anomaly, ranges from 11 bps to 20 bps across different pricing kernel calibrations. Recall that the average close-to-announcement return during this period was 37 bps. The second feature is that structural factors explain much less of the variation in ex ante returns than ex post returns. In the $\gamma = 3$ model, for example, the R^2 is 19.6% for pre-announcement returns compared to 70.3% for post-announcement returns. The third feature is that information-based factors are reflected predominantly in the ex post fraction of returns. On average only 16.8% to 23.1% of ε_p enters returns pre-announcement. Likewise, only 14.7% to 22.6% of ε_θ is reflected in ex ante returns. It bears repeating that neither of these latter two results might have appeared likely from Figure 3.1.

One way to ask how the results explain the pre-announcement drift is to compare it with average fitted values component-wise. In this perspective, the anomaly accounts for 29.8% to 54.7% of the 37 bps sample average. The analysis implies an FOMC risk premium that ranges from 12 bps to 28 bps, or 32.0% to 77.0% of the sample mean. Information-based returns constitute a much less prominent share of the pre-FOMC drift during this period. If both ε_p and ε_θ are counted among information-based returns, these represent between 0.4 and 3.2 bps on average—that is, 1.1% to 8.7% of the close-to-announcement drift. Alternatively, if one adopts the view that ε_θ may be driven by largely public information about how the stock market is exposed to the announcement, ε_p may be considered in isolation. In that case, average ex ante returns due to informed trading range from 2.2 to 4.4 bps, or 6.1% to 12.0% of the sample mean.

Notwithstanding their limited contribution to the drift per se, information-based returns explain more of the variation in pre-announcement returns than do risk premia. This can be seen from direct comparison of partial R^2 values for ε_p and $\mathbb{E}_{T-1}[r - r_f]$. The difference is even more noticeable if one puts marginal effects into perspective. For example, according to the $\gamma = 3$ model, a one standard deviation shock to ε_p is associated with a 21 bps change in ex ante returns. By contrast, a one standard deviation perturbation to $\mathbb{E}_{T-1}[r - r_f]$ translates to an 8 bps movement in ex ante returns. The results thus provide evidence for pre-announcement informed trading. The wrinkle is that neither the extent of informed trading nor the prevalence of good news from FOMC announcements are sufficient to account for a sizable fraction of the drift.

I now turn attention to results for the post-LM period, which are presented in Table 3.8. As with the earlier sample, the model is generally well-behaved. One caveat, however, is that the data reject the first two parameter restrictions. This is not obvious from the unrestricted results, as most of these estimates appear to be within a reasonable distance of their restricted counterparts. The most notable exception is $\beta^{pre}_{\mathbb{E}r}$, which has unreasonably high point estimates for lower values of γ. Though imperfect, the $\gamma = 7$ pricing kernel appears to offer the most fitting risk premia for this period. As I argue below, this apparent rise in the price of risk is due in part to the COVID-19 pandemic sub-period. Restrictions aside, total R^2 values for pre-announcement returns are improved considerably for this sample. Indeed if we compare the calibrations best suited for each period, the restricted model R^2 nearly doubles from 19.6% in the LM period to 39.0% in the LM period.

The most striking difference from the LM sample results is that the pre-announcement return intercept is now gone. This is the case for each of the three pricing kernel calibrations and for both restricted and unrestricted model estimates. Aside from the aforementioned evidence for greater investor risk aversion and the larger ex ante return R^2 values, additional differences with the LM sample results are few. Point estimates tend to be larger for β^{pre}_p and β^{pre}_θ, but not to a statistically significant extent. The parameters are not measured precisely enough from the limited data to establish whether pre-announcement informed trading increased since 2011.

For some visual perspective on these results, I provide two additional charts. Figure 3.7 displays the time series of pre-announcement and post-announcement returns. Figure 3.8 plots results for a full-sample estimation of the restricted model with the risk aversion parameter set to $\gamma = 5$. Included in the upper plot area are moving averages of model-implied ex ante return components. A moving average of the residual is displayed in the lower plot area. Note that care should be taken in comparing magnitudes between the two plot areas as the separate axes are of different length.

Unsurprisingly, Figure 3.8 illustrates that the risk premium soars dramatically during the 2008–2009 financial crisis. Such episodes aside, it is also clear that for any fixed level of investor risk aversion the risk premium exhibits a downward trend mirroring that of the policy impact gap displayed in Figure 3.6. The chart also provides visual confirmation of the comparatively modest role of information-based returns. Not only are they frequently dwarfed by the risk premium, Figure 3.8 highlights the often countervailing influences of ε_p and ε_θ. Good news for stockholders is frequently tempered by a more restrained appraisal of how good it is.

It is interesting that the risk premium remains flat throughout the 2001 recession. In stark contrast, the residual surged above 60 bps at the onset of this recession before retreating in subsequent months. Another interesting feature about this episode is that, more than any other period, ex post returns appear to offset much of the returns earned ex ante. This can be seen easily in Figure 3.7. The episode also motivates a broader point about the interpretation of the anomaly documented in this chapter. The time series behavior of the residual—which is also remarkably elevated during the financial crisis and the 2020 pandemic period—suggests that a share of the anomaly may represent risk premia not captured by my generally static treatment of investor risk preferences.[12] If it is nonetheless the case that mispricing is the prime animator of the anomalous returns, it also raises the possibility that a source of this mispricing may be intricately related to changes in risk premia.

3.4.3 What Changed?

The results highlight a prominent role for anomalous returns and risk premia in the pre-announcement drift. This leaves open the question of what changed in recent years. What drove the large decline in the pre-FOMC drift after 2011? What of its subsequent return after 2017? To better analyze changes in the post-LM sample, Table 3.9, provides sample statistics for returns and model-implied return components for the separate sub-periods of April 2011 to December 2017 and January 2018 to December 2020. This table is the direct analogue of Table 3.6.

What is clear from the above findings is that the role of informed trading in these recent shifts is limited. Evidence on the relative contribution of risk premia and the anomaly in the evolving pre-FOMC drift is more speculative. As the data can support a wide range of values for risk premia, it is difficult to disentangle these components with much precision. Here I argue that the disappearance of the anomaly played the most significant role in the post-2011 retreat in the pre-FOMC drift. Though this structural break is suggestive of a mispricing corrected by arbitrage, it should be reiterated that my analysis cannot rule out alternatives involving more subtle changes in investor risk preferences. In my appraisal of the results that follows, I advance the tentative conclusion that the more recent growth in the pre-FOMC drift was driven by a pandemic-related increase in the cost of risk.

[12]Added emphasis for this view comes from Liu, Tang, and Zhou (2022) who argue that the FOMC risk premium is heavily time-varying.

One way to assess what changed is to focus on the γ calibrations best suited for each period. This means taking account of trends in investor risk aversion by comparing the $\gamma = 3$ results for the LM period with the $\gamma = 7$ results for the post-LM period. In this case, the average risk premium barely budges from 11.8 bps in the LM period to 11.1 bps between April 2011 and December 2017. The net average of information-based returns declines from 3.2 bps to 0.3 bps for these periods. This change accounts for only 10% of the 29.2 bps fall in the close-to-announcement pre-FOMC drift in the 2011–2017 sample. This leaves only the anomaly returns which, according to this model comparison, can explain about 70% of this change.

The risk premium's seemingly negligible role in the pre-FOMC drift's 2011–2017 retreat is counterintuitive given its nontrivial contribution to the drift more generally. It must be cautioned, however, that this assessment is sensitive to which pricing kernel calibrations are compared. For example, if one limits attention to the $\gamma = 5$ case for both of the main sample periods, the results imply that the average risk premium falls by 12.2 bps in the 2011–2017 sub-period. Contrasting heavily with the previous comparison, this represents about 42% of the close-to-announcement drift's decline during this time. In this case, the disappearance of the anomaly accounts for a still-prominent but more limited 53% of the drift's 2011–2017 retreat.

Data limitations make the post-2017 growth in the pre-FOMC drift more challenging to assess. One clear message from Table 3.9 is that unless investors grew *less* risk averse after 2017—an unlikely scenario, as I argue below—average information-based returns run counter to the pre-announcement drift's re-emergence. This follows from a basic comparison of average ε_p and ε_θ estimates across sub-periods and pricing kernel calibrations.[13] Beyond this finding, the results offer fewer direct clues about this most recent change. Table 3.9 reveals that for the $\gamma = 7$ model, for example, the average risk premium grew by 2.1 bps. But even with headwinds from information-based returns aside, this still leaves 90% of the pre-announcement drift's 21 bps increase unexplained.

Reserved though the direct findings for this sub-period may be, the results are nonetheless more consistent with a price of risk explanation than plausible alternatives. The rationale follows from circumstantial evidence viewed through the lens of established asset pricing theory. Countercyclical risk aversion predicts that if there is a systematic pre-FOMC risk, the price of this risk should be elevated during the extraordinary turmoil of the COVID-19 pandemic (Cohn et al., 2015). Not only do the above findings provide evidence for this systematic risk

[13]The implicit assumption here about the stability of coefficients for these variables throughout the post-LM period is weaker than it may appear. Note that even phenomenal changes in these parameters would struggle to account for a meaningful share of the pre-FOMC drift's post-2017 increase. This is especially the case for greater levels of investor risk aversion.

in general, residual time series behavior is highly consistent with this prediction. Recall that pre-announcement return residuals, which follow from a model in which investors have fixed levels of constant relative risk aversion, climb precipitously in 2020 as illustrated by Figure 3.8. Additional unrestricted model analysis of risk premium loading $\beta_{\mathbb{E}r}^{pre}$ (not reported) also lends support for a 2020 surge in investor risk aversion. Finally, for reasons noted above, my results imply that an informed trading explanation is inconsistent with countercyclical risk aversion. Overall, the preponderance of evidence suggests that a pandemic-related spike in the price of risk was most responsible for the recent increase in the pre-FOMC drift.

3.5 Conclusion

The remarkable prominence of the pre-FOMC drift in stock market returns represents a puzzle at the intersection of asset pricing and macrofinancial policy—one that can be neither overlooked nor readily explained. Recent changes in this drift raise additional questions about the phenomenon while offering researchers novel opportunities to learn of its underlying nature. Guided by these observations, this chapter's event-based methodology leverages the joint dynamics of pre-FOMC returns, option-implied investor expectations, and monetary policy decisions to shed new light on this mystery.

My findings identify a sizable pre-announcement return anomaly that disappeared after the conclusion of the original LM sample period. The results indicate a prominent and enduring role for an FOMC risk premium in the pre-announcement drift while revealing a much more limited influence of informed trading in advance of good news. The analysis implies that recent changes in pre-announcement returns are due in large part to dynamics in risk premia and the disappearance of anomalous returns. Whether the bygone anomaly in pre-announcement returns was driven by mispricing or alternative explanations is left to future research.

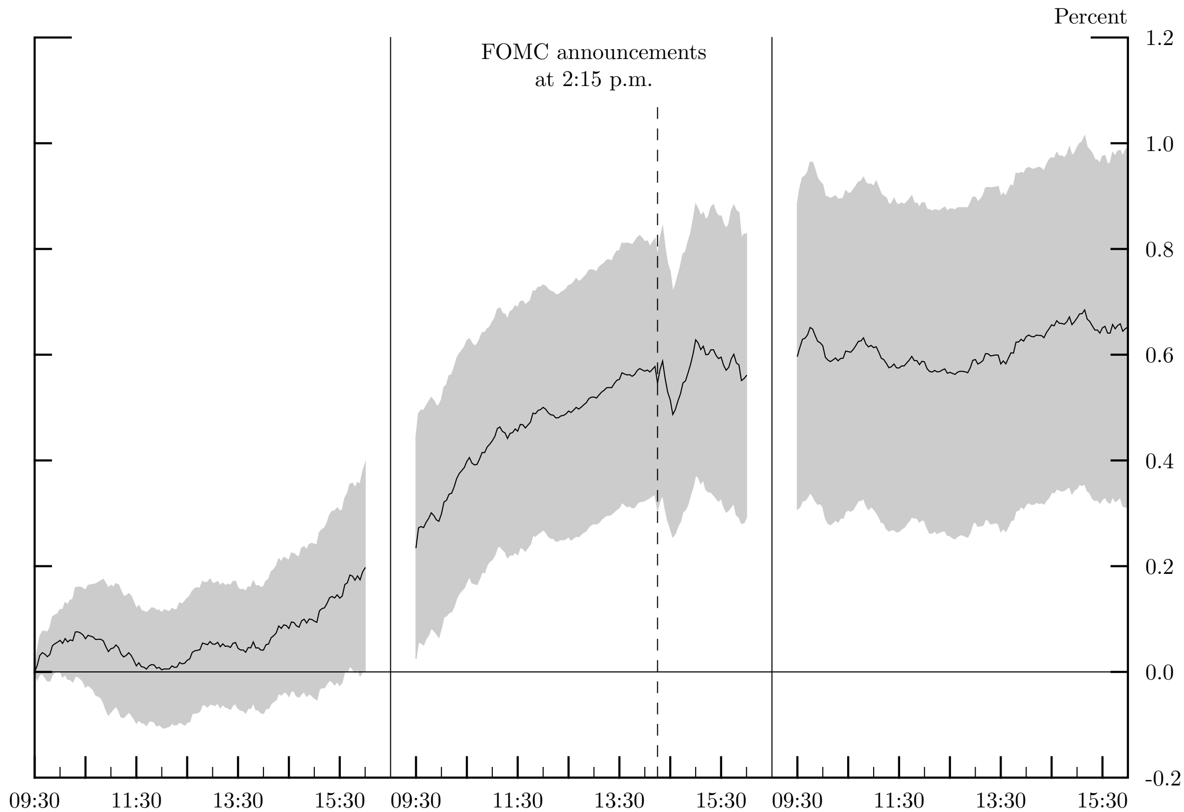

Figure 3.1: S&P 500 Index cumulative returns, Sept. 1994 – Mar. 2011. This chart displays average cumulative excess returns on the S&P 500 index at intraday frequency over three-business-day windows centered on FOMC announcement dates. Index returns in this figure are cum-dividend. The shaded areas represents 95% confidence bands based as determined by pointwise standard errors. The dashed vertical line indicates the official 2:15 p.m. FOMC announcement time for this sample period.

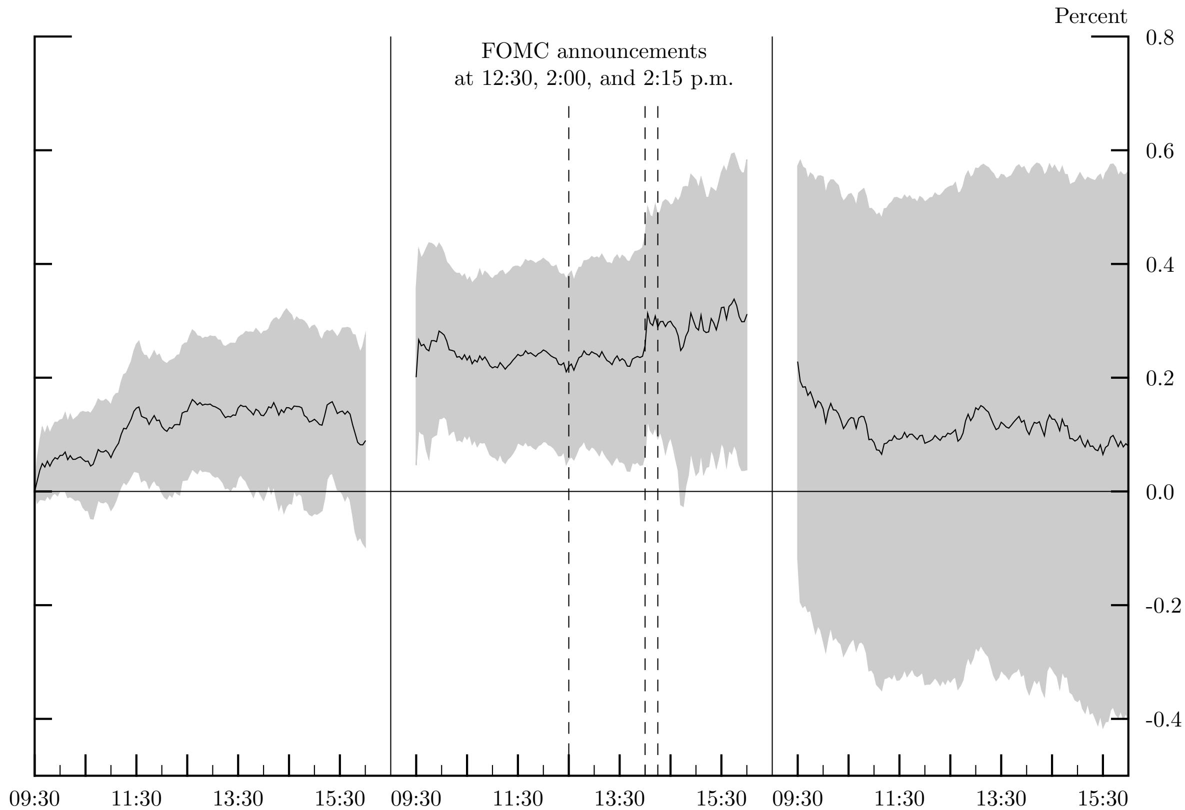

Figure 3.2: S&P 500 Index cumulative returns, Apr. 2011 – Dec. 2017. This chart displays average cumulative excess returns on the S&P 500 index at intraday frequency over three business day windows centered on FOMC announcement dates. Index returns in this figure are cum-dividend. The shaded areas represents 95% confidence bands based as determined by pointwise standard errors. Official FOMC announcement times for this sample period are indicated by dashed vertical lines.

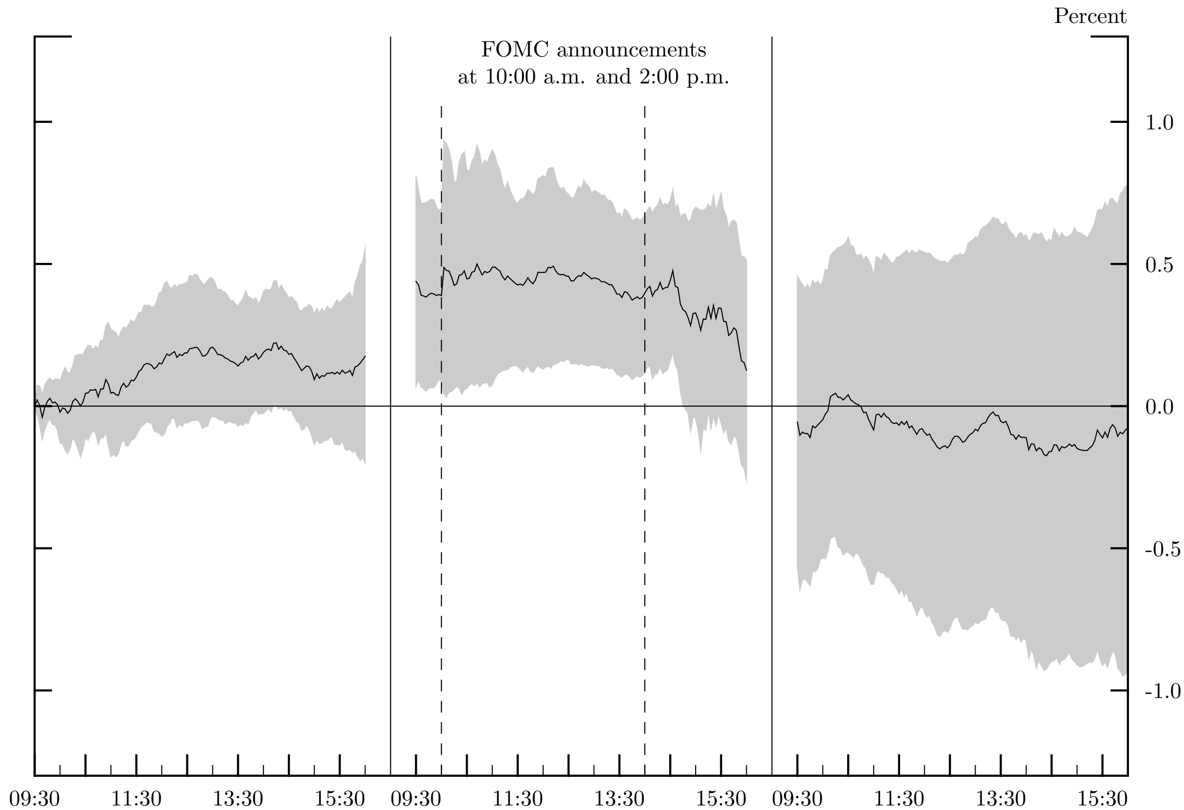

Figure 3.3: S&P 500 Index cumulative returns, Jan. 2018 – Dec. 2020. This chart displays average cumulative excess returns on the S&P 500 index at intraday frequency over three business day windows centered on FOMC announcement dates. Index returns in this figure are cum-dividend. The shaded areas represents 95% confidence bands based as determined by pointwise standard errors. Official FOMC announcement times for this sample period are indicated by dashed vertical lines.

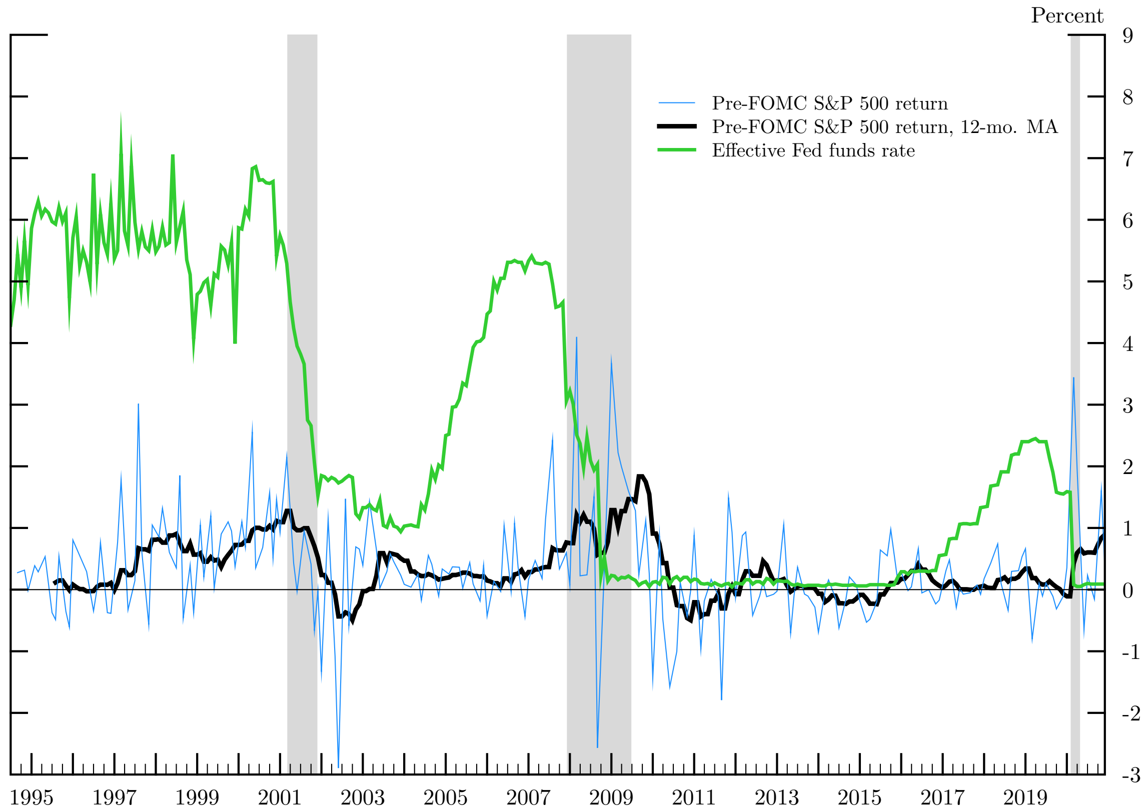

Figure 3.4: Pre-FOMC announcement returns over time. This figure illustrates the time series characteristics of pre-announcement returns on the S&P 500 index between Sept. 1994 and Dec. 2020. The narrow blue line represents 24-hour excess pre-announcement holding period returns. S&P 500 index returns in this figure are cum-dividend. A 12-month moving average of these returns is indicated by the thick black line. The effective federal funds rate is represented by the green line. Shaded regions demarcate U.S. recessions according to NBER business cycle definitions. An outlier return for the FOMC announcement on October 29, 2008 is excluded from the plot, consistent with previous literature (Lucca and Moench, 2015, Figure 6).

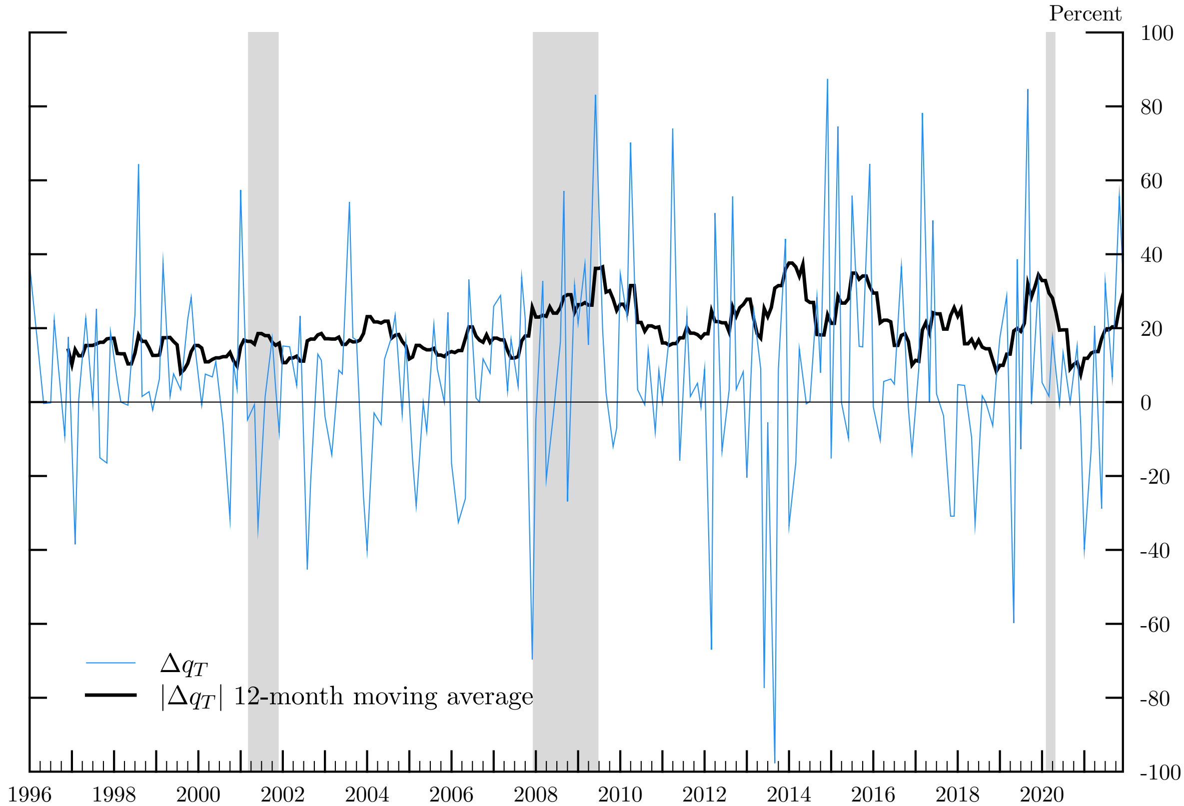

Figure 3.5: Changes in risk neutral probabilities of upside FOMC decision outcomes. This chart displays event study estimates for the distance between the one business day ex ante risk neutral probability of upside outcome, q_{T-1}, and the realized outcome implied by the model. The time series estimates of $\Delta q_T \equiv \mathbf{1}_u - q_{T-1}$ are plotted above as a narrow blue line. A 12-month moving average of $|\Delta q_T|$ is represented by the thick black line. Shaded regions demarcate U.S. recessions according to NBER business cycle definitions.

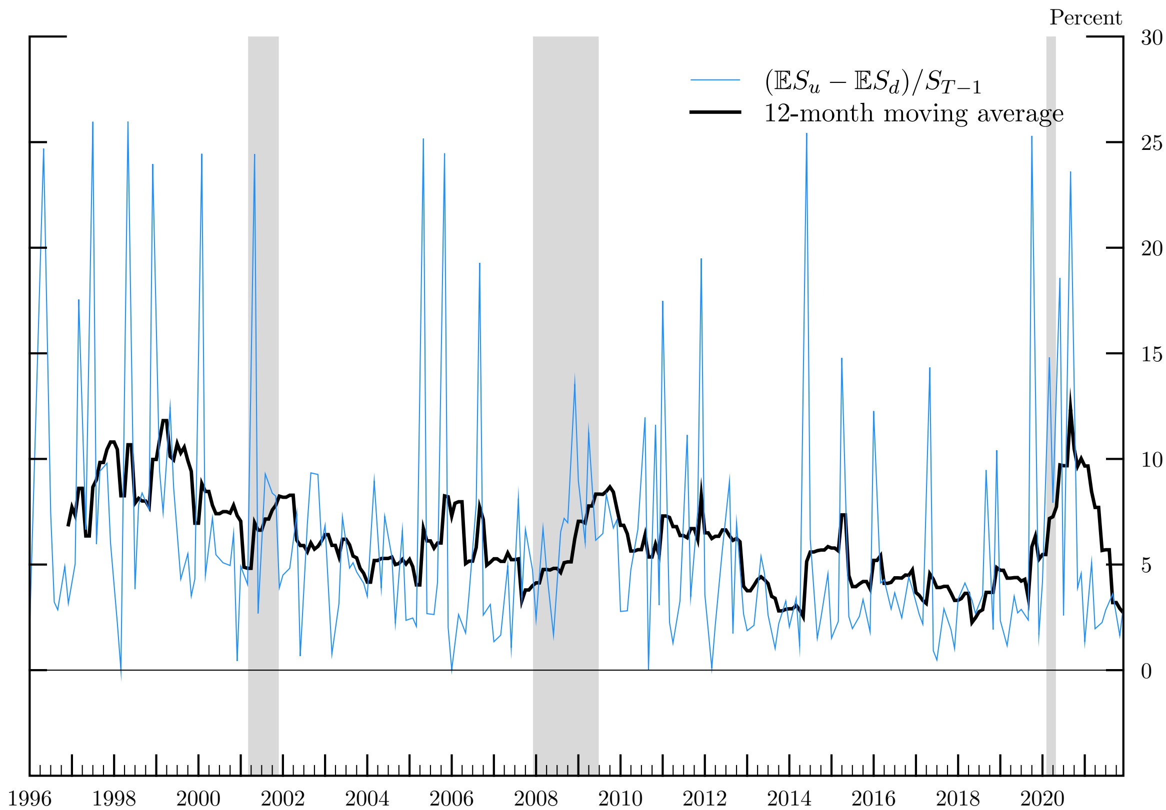

Figure 3.6: Time series estimates of the FOMC policy impact gap. The policy impact gap denotes the relative distance between model-implied ex ante expectations of upside and downside scenarios, i.e., $(\mathbb{E}S_u - \mathbb{E}S_d)/S_{T-1}$. The impact gap is represented over time by a narrow blue line. The thick black line represents a 12-month moving average. Shaded regions demarcate U.S. recessions according to NBER business cycle definitions.

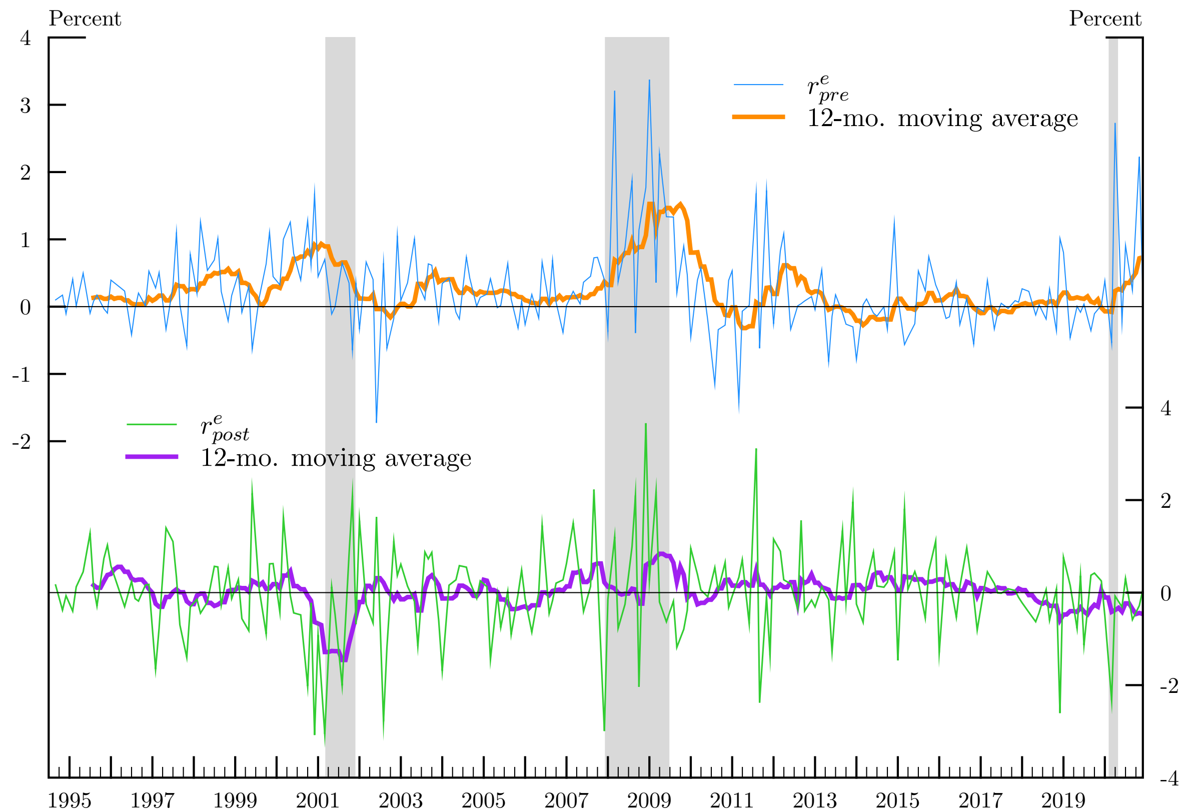

Figure 3.7: Time series of pre-announcement and post-announcement FOMC returns. Excess close-to-announcement ex ante returns are plotted as a narrow blue line in the upper plot area. Announcement-to-close ex post excess returns are represented by a narrow green line in the lower plot area. Index returns in this chart are ex-dividend holding period returns. The thick orange and purple lines represent 12-month moving averages of ex ante and ex post returns, respectively. Shaded regions demarcate U.S. recessions according to NBER business cycle definitions.

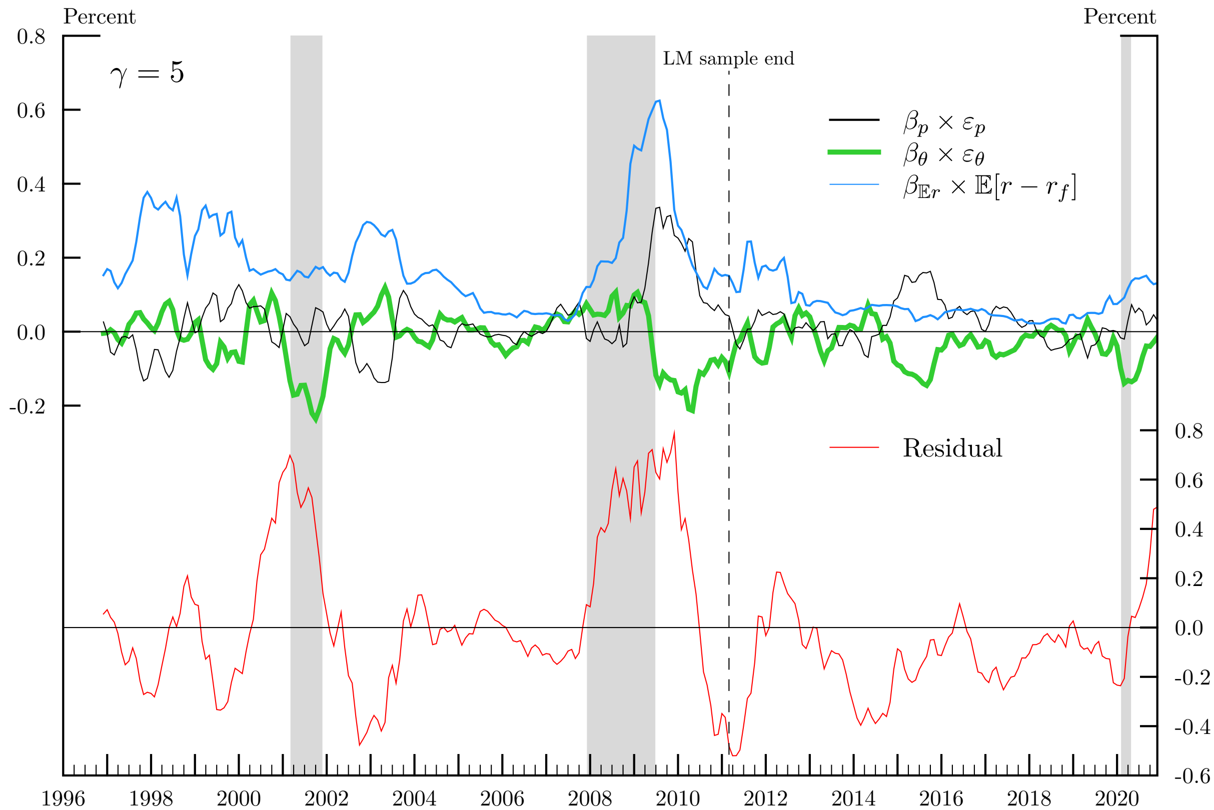

Figure 3.8: Pre-announcement return components from a projection of FOMC returns on structural factors. The coefficients and residual displayed in this figure correspond to a full sample estimation of the restricted model with risk aversion parameter calibrated to $\gamma = 5$. In the upper plot area are model-implied pre-announcement return components. The blue line is the FOMC risk premium. The narrow black line and thick green lines represent the component of ex ante returns due to informational factors ε_p and ε_θ, respectively. The pre-announcement return residual is indicated by a thin red line in the lower plot area. Point estimates for information return loadings in this figure are $\beta_p^{pre} = 0.201$ and $\beta_\theta^{pre} = 0.201$. The coefficient on the risk premium, $\beta_{\mathbb{E}r}^{pre}$, is restricted to one. Shaded regions demarcate U.S. recessions according to NBER business cycle definitions.

Table 3.1: Pre-FOMC announcement drift, Sept. 1994 – Dec. 2020

A. 24-hour pre-announcement excess return

		Sample period			
		LM 1994 – 2011	2011 – 17	2018 – 20	Post-LM 2011 – 20
	N obs.	132	54	24	78
Cum-dividend	Mean	0.520	0.044	0.368	0.144
	Std. dev.	1.316	0.519	0.896	0.669
	p10	-0.484	-0.481	-0.336	-0.481
	p25	-0.015	-0.172	-0.108	-0.155
	p75	0.913	0.206	0.481	0.342
	p90	1.574	0.876	1.609	0.931
Ex-dividend	Mean	0.513	0.035	0.362	0.136
	Std. dev.	1.316	0.517	0.895	0.668
	p10	-0.492	-0.488	-0.336	-0.488
	p25	-0.019	-0.179	-0.109	-0.159
	p75	0.908	0.195	0.471	0.342
	p90	1.572	0.839	1.597	0.913

B. Close-to-announcement excess return

		Sample period			
		LM 1994 – 2011	2011 – 17	2018 – 20	Post-LM 2011 – 20
	N obs.	132	54	24	78
Cum-dividend	Mean	0.353	0.086	0.292	0.149
	Std. dev.	0.688	0.518	0.787	0.615
	p10	-0.328	-0.448	-0.357	-0.448
	p25	0.003	-0.251	-0.155	-0.185
	p75	0.640	0.210	0.327	0.264
	p90	1.104	0.753	1.022	0.910
Ex-dividend	Mean	0.347	0.077	0.286	0.141
	Std. dev.	0.689	0.517	0.787	0.615
	p10	-0.337	-0.455	-0.361	-0.455
	p25	-0.009	-0.253	-0.155	-0.185
	p75	0.638	0.196	0.324	0.264
	p90	1.102	0.744	1.013	0.910

This table reports pre-FOMC announcement return statistics for different sample periods. Panel A provides these statistics for 24-hour ex ante excess returns. Panel B reports statistics for close-to-announcement ex ante excess returns. Sample periods are grouped by column. Within each panel, results are sub-categorized according to whether returns are cum-dividend or ex-dividend. Values are reported in percentage points.

Table 3.2: FOMC announcement sample statistics

A. LM sample period (Sept. 1994 – Mar. 2011)

	N obs.	Mean	Std. dev.	p10	p90
Model estimates					
q_{T-1}	121	0.564	0.361	0.029	0.962
$\mathbf{1}_u$	121	0.636	0.483	0	1
Δq_T	121	0.072	0.228	-0.209	0.328
$\mathbb{E}[S_u]/S_{T-1}$	121	1.038	0.064	1.002	1.087
$\mathbb{E}[S_d]/S_{T-1}$	121	0.969	0.031	0.923	0.999
$\mathbb{E}\sigma_u$	121	0.204	0.116	0.099	0.327
$\mathbb{E}\sigma_d$	121	0.170	0.135	0.088	0.262
FOMC returns, other items					
r^e_{pre}	132	0.347%	0.689%	-0.337%	1.102%
r^e_{post}	132	0.009%	1.039%	-1.153%	1.200%
Effective fed funds rate	132	3.509%	2.168%	0.190%	5.810%
VIX	132	21.23%	8.60%	12.19%	30.43%
$\mathbf{1}_{recession}$	132	0.129	0.336	0	1
$\mathbf{1}_{cut}$	132	0.174	0.381	0	1
$\mathbf{1}_{hike}$	132	0.197	0.399	0	1
$\mathbf{1}_{no\ change}$	132	0.629	0.485	0	1

B. Post-LM sample period (Apr. 2011 – Dec. 2021)

	N obs.	Mean	Std. dev.	p10	p90
Model estimates					
q_{T-1}	86	0.532	0.365	0.018	0.965
$\mathbf{1}_u$	86	0.605	0.492	0	1
Δq_T	86	0.072	0.329	-0.309	0.557
$\mathbb{E}[S_u]/S_{T-1}$	86	1.029	0.056	1.002	1.049
$\mathbb{E}[S_d]/S_{T-1}$	86	0.979	0.027	0.938	0.999
$\mathbb{E}\sigma_u$	86	0.157	0.198	0.082	0.235
$\mathbb{E}\sigma_d$	86	0.122	0.081	0.073	0.199
FOMC returns, other items					
r^e_{pre}	78	0.141%	0.615%	-0.455%	0.910%
r^e_{post}	78	-0.007%	0.897%	-1.134%	0.966%
Effective fed funds rate	86	0.599%	0.769%	0.080%	1.930%
VIX	86	17.40%	6.32%	11.81%	26.04%
$\mathbf{1}_{recession}$	86	0.023	0.152	0	0
$\mathbf{1}_{cut}$	86	0.047	0.212	0	0
$\mathbf{1}_{hike}$	86	0.105	0.308	0	1
$\mathbf{1}_{no\ change}$	86	0.849	0.360	0	1

This table presents sample statistics for the FOMC event analysis. Each FOMC announcement included in my analysis corresponds to one unit of observation in the sample. Panels A and B report sample statistics for the LM (Sept. 1994 – Mar. 2011) and post-LM sample periods (Apr. 2011 – Dec. 2021), respectively. Note that $\mathbf{1}_x$ denotes an indicator variable that is one conditional on x and zero otherwise.

Table 3.3: Fed Fund Rate Changes and Model-implied Stock Market Outcomes

A. Full sample period (Jan. 1996 – Dec. 2021)		
	$\mathbf{1}_d$	$\mathbf{1}_u$
Rate cut	10 4.83%	15 7.25%
Rate hike	13 6.28%	20 9.66%
No rate change	55 26.57%	94 45.41%
B. LM sample period (Jan. 1996 – Mar. 2011)		
	$\mathbf{1}_d$	$\mathbf{1}_u$
Rate cut	8 6.61%	13 10.74%
Rate hike	10 8.26%	14 11.57%
No rate change	26 21.49%	50 41.32%
C. Post-LM sample period (Apr. 2011 – Dec. 2021)		
	$\mathbf{1}_d$	$\mathbf{1}_u$
Rate cut	2 2.33%	2 2.33%
Rate hike	3 3.49%	6 6.98%
No rate change	29 33.72%	44 51.16%

Presented in this table is cross-tabular analysis of event study implied FOMC announcement outcomes (columns) and official Federal Reserve policy regarding federal funds rate targets (rows). Panel A provides results for the full sample. Panels B and C present results for the LM sample period (Jan. 1996 – Mar. 2011) and post-LM sample period (Apr. 2011 – Dec. 2021), respectively.

Table 3.4: Statistics Conditional on Model-implied FOMC Announcement Surprises

A. Bad news surprises ($\Delta q_T < -1s$)

	N obs.	Mean	Std. dev.	p10	p90
Model estimates					
q_{T-1}	18	0.455	0.201	0.288	0.773
Δq_T	18	-0.455	0.201	-0.773	-0.288
$\mathbb{E}[S_u]/S_{T-1}$	18	1.016	0.010	1.000	1.034
$\mathbb{E}[S_d]/S_{T-1}$	18	0.987	0.011	0.967	0.997
$\mathbb{E}\sigma_u$	18	0.123	0.060	0.063	0.194
$\mathbb{E}\sigma_d$	18	0.139	0.055	0.072	0.219
FOMC returns, other items					
r^e_{pre}	16	0.221%	0.460%	-0.142%	0.830%
r^e_{post}	16	-0.903%	1.195%	-2.782%	0.894%
Effective fed funds rate	18	2.041%	2.037%	0.070%	5.170%
VIX	18	18.43%	7.64%	10.20%	35.82%
$\mathbf{1}_{recession}$	18	0.056	0.236	0	0
$\mathbf{1}_{cut}$	18	0.111	0.323	0	1
$\mathbf{1}_{hike}$	18	0.222	0.428	0	1
$\mathbf{1}_{no\ change}$	18	0.667	0.485	0	1

B. Good news surprises ($\Delta q_T > 1s$)

	N obs.	Mean	Std. dev.	p10	p90
Model estimates					
q_{T-1}	35	0.506	0.185	0.218	0.712
Δq_T	35	0.494	0.185	0.288	0.782
$\mathbb{E}[S_u]/S_{T-1}$	35	1.018	0.013	1.005	1.041
$\mathbb{E}[S_d]/S_{T-1}$	35	0.981	0.017	0.962	0.995
$\mathbb{E}\sigma_u$	35	0.156	0.086	0.087	0.227
$\mathbb{E}\sigma_d$	35	0.131	0.069	0.070	0.200
FOMC returns, other items					
r^e_{pre}	32	0.400%	0.773%	-0.391%	1.234%
r^e_{post}	32	0.845%	0.957%	-0.006%	2.100%
Effective fed funds rate	35	1.724%	2.114%	0.090%	5.400%
VIX	35	19.25%	8.63%	12.50%	29.05%
$\mathbf{1}_{recession}$	35	0.143	0.355	0	1
$\mathbf{1}_{cut}$	35	0.171	0.382	0	1
$\mathbf{1}_{hike}$	35	0.143	0.355	0	1
$\mathbf{1}_{no\ change}$	35	0.686	0.471	0	1

This table provides sample statistics for FOMC events conditional on announcement surprises. An announcement is defined as a surprise if the distance between risk neutral probability q_{T-1} and the model-implied realized outcome $\mathbf{1}_u$ exceeds one standard deviation. Panel A reports statistics conditional on a bad news surprise (i.e., $\Delta q_T < -1s$). Panel B presents statistics conditional on a good news surprise (i.e., $\Delta q_T > -1s$). Note that $\mathbf{1}_x$ denotes an indicator variable that is one conditional on x and zero otherwise.

Table 3.5: FOMC-record Sample Correlations

A. LM sample period (Sept. 1994 – Mar. 2011)

	q_{T-1}	$\mathbf{1}_u$	r^e_{pre}	r^e_{post}
q_{T-1}	1.000 – 121			
$\mathbf{1}_u$	0.893 $p < .0001$ 121	1.000 – 121		
r^e_{pre}	0.159 $p = 0.082$ 121	0.198 $p = 0.029$ 121	1.000 – 132	
r^e_{post}	0.176 $p = 0.054$ 121	0.388 $p < .0001$ 121	-0.116 $p = 0.186$ 132	1.000 – 132

B. Post-LM sample period (Apr. 2011 – Dec. 2021)

	q_{T-1}	$\mathbf{1}_u$	r^e_{pre}	r^e_{post}
q_{T-1}	1.000 – 86			
$\mathbf{1}_u$	0.743 $p < .0001$ 86	1.000 – 86		
r^e_{pre}	0.208 $p = 0.067$ 78	0.260 $p = 0.022$ 78	1.000 – 78	
r^e_{post}	-0.0427 $p = 0.710$ 78	0.198 $p = 0.083$ 78	0.099 $p = 0.390$ 78	1.000 – 78

This table presents sample correlations for model-implied risk neutral probabilities of FOMC outcomes and pre-announcement and post-announcement returns. Note that q_{T-1} represents the one business day ex ante probability of an upside outcome whereas $\mathbf{1}_u$ denotes the terminal realization of the state. Panel A reports correlations for the LM sample period (Sept. 1994 – Mar. 2011). Panel B provides correlations for the post-LM sample period (Apr. 2011 – Dec. 2021).

Table 3.6: FOMC Returns and Model-implied Components

A. LM sample period (Jan. 1996 – Mar. 2011)							
		Mean	Std. dev.	p10	p25	p75	p90
r^e_{pre}		0.369	0.713	-0.341	0.006	0.648	1.105
r^e_{post}		-0.013	1.071	-1.183	-0.471	0.550	1.200
ε_p	$\gamma = 3$	0.263	1.256	-0.978	-0.229	0.687	1.286
	$\gamma = 5$	0.182	1.250	-1.199	-0.284	0.628	1.151
	$\gamma = 7$	0.097	1.263	-1.286	-0.345	0.588	1.028
ε_θ		-0.081	1.034	-1.465	-0.372	0.575	0.915
$\mathbb{E}[r - r_f]$	$\gamma = 3$	0.118	0.149	0.008	0.025	0.154	0.296
	$\gamma = 5$	0.199	0.247	0.014	0.042	0.265	0.506
	$\gamma = 7$	0.284	0.353	0.020	0.059	0.389	0.794

B. Post-LM sample period (Apr. 2011 – Dec. 2020)							
		Mean	Std. dev.	p10	p25	p75	p90
r^e_{pre}		0.141	0.615	-0.455	-0.185	0.264	0.910
r^e_{post}		-0.007	0.897	-1.134	-0.424	0.361	0.966
ε_p	$\gamma = 3$	0.162	0.913	-0.606	-0.242	0.446	1.129
	$\gamma = 5$	0.129	0.910	-0.641	-0.310	0.420	1.050
	$\gamma = 7$	0.093	0.913	-0.676	-0.358	0.398	1.031
ε_θ		-0.135	0.935	-0.924	-0.492	0.397	0.940
$\mathbb{E}[r - r_f]$	$\gamma = 3$	0.048	0.080	0.006	0.011	0.052	0.107
	$\gamma = 5$	0.082	0.130	0.010	0.019	0.086	0.177
	$\gamma = 7$	0.117	0.182	0.013	0.029	0.126	0.292

Reported above are sample statistics for pre-announcement (r^e_{pre}) and post-announcement (r^e_{post}) FOMC returns and model-implied total FOMC return components. Panel A provides statistics for the LM sample period from Jan. 1996 to Mar. 2011. Panel B presents statistics for the post-LM period between Apr. 2011 and Dec. 2020. See Section 3.4.1 for model-implied return component definitions. Statistics for risk premia and information factor ε_p are reported for each pricing kernel calibration, i.e. $\gamma = 3, 5, 7$. Values are reported in percentage points.

Table 3.7: FOMC Return Decomposition, LM Sample Period (Jan. 1996 – Mar. 2011)

A. Restricted Model

	$\gamma = 3$		$\gamma = 5$		$\gamma = 7$	
	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}
β_0	0.202***	-0.160***	0.155***	-0.110**	0.110**	-0.063
	(0.0542)	(0.0503)	(0.0542)	(0.0504)	(0.0548)	(0.0510)
β_p	0.168***	0.832***	0.194***	0.806***	0.231***	0.769***
	(0.0604)	(0.0604)	(0.0636)	(0.0636)	(0.0694)	(0.0694)
	$R^2 = 0.126$	$R^2 = 0.776$	$R^2 = 0.139$	$R^2 = 0.758$	$R^2 = 0.175$	$R^2 = 0.724$
β_θ	0.147*	0.853***	0.186**	0.814***	0.226***	0.774***
	(0.0751)	(0.0751)	(0.0784)	(0.0784)	(0.0837)	(0.0837)
	$R^2 = 0.088$	$R^2 = 0.674$	$R^2 = 0.108$	$R^2 = 0.656$	$R^2 = 0.126$	$R^2 = 0.622$
$\beta_{\mathbb{E}r}$	1.000	–	1.000	–	1.000	–
	$R^2 = 0.104$	–	$R^2 = 0.134$	–	$R^2 = 0.137$	–
$\beta_p^{pre} + \beta_p^{post} = 1$	$p = 0.2487$		$p = 0.3327$		$p = 0.3582$	
$\beta_\theta^{pre} + \beta_\theta^{post} = 1$	$p = 0.8162$		$p = 0.9481$		$p = 0.9931$	
$\beta_{\mathbb{E}r}^{pre} = 1$	$p = 0.4298$		$p = 0.8147$		$p = 0.7971$	
N obs.	121		121		121	
R^2	0.196	0.703	0.190	0.693	0.141	0.666

B. Unrestricted Model

	$\gamma = 3$		$\gamma = 5$		$\gamma = 7$	
	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}
β_0	0.199***	-0.173***	0.207***	-0.178***	0.217***	-0.186***
	(0.0674)	(0.0634)	(0.0672)	(0.0632)	(0.0672)	(0.0631)
β_p	0.229***	0.823***	0.238***	0.816***	0.247***	0.808***
	(0.0770)	(0.0648)	(0.0771)	(0.0650)	(0.0773)	(0.0652)
β_θ	0.241**	0.803***	0.245**	0.800***	0.250**	0.796***
	(0.1005)	(0.0842)	(0.1006)	(0.0845)	(0.1008)	(0.0847)
$\beta_{\mathbb{E}r}$	1.088**	0.080	0.697***	0.414	0.522***	0.562***
	(0.4536)	(0.4325)	(0.2624)	(0.2521)	(0.1810)	(0.1748)
N obs.	121		121		121	
R^2	0.211	0.705	0.206	0.703	0.201	0.702

This table presents time series regression results for the LM sample period of Jan. 1996 to Mar. 2011. Panel A provides results for the restricted model. Panel B reports estimates for the unrestricted model. See Section 3.4.1 for model and variable definitions. Columns are grouped by risk aversion parameter γ calibration and sub-categorized by dependent variable (i.e., ex ante and ex post FOMC returns). Newey-West standard errors are reported in parentheses. Partial R^2 values are reported for explanatory variables in the restricted model. These indicate the fraction of variability in the dependent variable that is unexplained by other variables but explained by the particular variable in question. Also provided in Panel A are p-values for each of the three main parameter restrictions.

Table 3.8: FOMC Return Decomposition, Post-LM Period (Apr. 2011 – Dec. 2020)

A. Restricted Model

	$\gamma = 3$		$\gamma = 5$		$\gamma = 7$	
	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}
β_0	0.012	-0.001	-0.005	0.015	-0.020	0.029
	(0.0474)	(0.0466)	(0.0481)	(0.0472)	(0.0488)	(0.0480)
β_p	0.266***	0.734***	0.270***	0.730***	0.289***	0.711***
	(0.0806)	(0.0806)	(0.0783)	(0.0783)	(0.0760)	(0.0760)
	$R^2 = 0.305$	$R^2 = 0.685$	$R^2 = 0.299$	$R^2 = 0.683$	$R^2 = 0.308$	$R^2 = 0.671$
β_θ	0.308***	0.692***	0.303***	0.697***	0.307***	0.693***
	(0.1029)	(0.1029)	(0.1003)	(0.1003)	(0.0978)	(0.0978)
	$R^2 = 0.203$	$R^2 = 0.680$	$R^2 = 0.209$	$R^2 = 0.679$	$R^2 = 0.220$	$R^2 = 0.669$
$\beta_{\mathbb{E}r}$	1.000	–	1.000	–	1.000	–
	$R^2 = 0.135$	–	$R^2 = 0.199$	–	$R^2 = 0.243$	–
$\beta_p^{pre} + \beta_p^{post} = 1$	$p = 0.0029$		$p = 0.0024$		$p = 0.0016$	
$\beta_\theta^{pre} + \beta_\theta^{post} = 1$	$p = 0.0071$		$p = 0.0054$		$p = 0.0042$	
$\beta_{\mathbb{E}r}^{pre} = 1$	$p = 0.3331$		$p = 0.2200$		$p = 0.0834$	
N obs.	78		78		78	
R^2	0.353	0.712	0.377	0.710	0.390	0.700

B. Unrestricted Model

	$\gamma = 3$		$\gamma = 5$		$\gamma = 7$	
	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}	r^e_{pre}	r^e_{post}
β_0	0.004	-0.027	0.003	-0.027	0.008	-0.030
	(0.0628)	(0.0693)	(0.0610)	(0.0684)	(0.0577)	(0.0653)
β_p	0.328***	0.817***	0.336***	0.813***	0.348***	0.808***
	(0.0801)	(0.0817)	(0.0804)	(0.0827)	(0.0811)	(0.0844)
β_θ	0.294***	0.779***	0.297***	0.778***	0.303***	0.775***
	(0.0896)	(0.1002)	(0.0908)	(0.1008)	(0.0923)	(0.1016)
$\beta_{\mathbb{E}r}$	2.580**	-0.137	1.655**	0.260	1.206***	0.453
	(1.1398)	(1.4423)	(0.6355)	(0.8380)	(0.3979)	(0.5470)
N obs.	78		78		78	
R^2	0.438	0.721	0.421	0.720	0.411	0.719

This table presents time series regression results for the post-LM sample period of Apr. 2011 to Dec. 2020. Panel A provides results for the restricted model. Panel B reports estimates for the unrestricted model. See Section 3.4.1 for model and variable definitions. Columns are grouped by risk aversion parameter γ calibration and sub-categorized by dependent variable (i.e., ex ante and ex post FOMC returns). Newey-West standard errors are reported in parentheses. Partial R^2 values are reported for explanatory variables in the restricted model. These indicate the fraction of variability in the dependent variable that is unexplained by other variables but explained by the particular variable in question. Also provided in Panel A are p-values for each of the three main parameter restrictions.

Table 3.9: FOMC Returns and Model-implied Components, Post-LM Sub-periods

A. Apr. 2011 – Dec. 2017

		Mean	Std. dev.	p10	p25	p75	p90
r^e_{pre}		0.077	0.517	-0.455	-0.253	0.196	0.744
r^e_{post}		0.140	0.907	-0.811	-0.248	0.464	1.155
ε_p	$\gamma = 3$	0.187	0.978	-0.606	-0.242	0.446	1.142
	$\gamma = 5$	0.155	0.973	-0.641	-0.286	0.431	1.127
	$\gamma = 7$	0.122	0.975	-0.676	-0.332	0.406	1.112
ε_θ		-0.105	0.884	-0.863	-0.510	0.273	0.760
$\mathbb{E}[r - r_f]$	$\gamma = 3$	0.046	0.088	0.005	0.011	0.049	0.079
	$\gamma = 5$	0.077	0.142	0.008	0.019	0.083	0.130
	$\gamma = 7$	0.111	0.197	0.011	0.027	0.119	0.197

B. Jan. 2018 – Dec. 2020

		Mean	Std. dev.	p10	p25	p75	p90
r^e_{pre}		0.286	0.787	-0.361	-0.155	0.324	1.013
r^e_{post}		-0.336	0.795	-1.134	-0.591	0.205	0.376
ε_p	$\gamma = 3$	0.107	0.763	-0.395	-0.239	0.452	0.996
	$\gamma = 5$	0.070	0.764	-0.414	-0.334	0.396	0.962
	$\gamma = 7$	0.029	0.773	-0.625	-0.388	0.349	0.849
ε_θ		-0.205	1.058	-2.015	-0.456	0.519	1.061
$\mathbb{E}[r - r_f]$	$\gamma = 3$	0.054	0.061	0.007	0.020	0.072	0.126
	$\gamma = 5$	0.091	0.101	0.011	0.033	0.110	0.270
	$\gamma = 7$	0.132	0.147	0.016	0.045	0.150	0.364

Reported above are sample statistics for pre-announcement (r^e_{pre}) and post-announcement (r^e_{post}) FOMC returns and model-implied total FOMC return components. Panel A provides statistics for the Apr. 2011 to Dec. 2017 post-LM sub-period. Panel B presents statistics for the Jan. 2018 to Dec. 2020 sub-period. See Section 3.4.1 for model-implied return component definitions. Statistics for risk premia and information factor ε_p are reported for each pricing kernel calibration, i.e. $\gamma = 3, 5, 7$. Values are reported in percentage points.

Bibliography

Acemoglu, Daron, Simon Johnson, Amir Kermani, James Kwak, and Todd Mitton. 2016. "The value of connections in turbulent times: Evidence from the United States." *Journal of Financial Economics* 121 (2):368–391.

Acharya, Viral V, Deniz Anginer, and A Joseph Warburton. 2016. "The end of market discipline? Investor expectations of implicit government guarantees." Manuscript, New York University.

Agrawal, Anup, Jeffrey F Jaffe, and Gershon N Mandelker. 1992. "The post-merger performance of acquiring firms: a re-examination of an anomaly." *The Journal of finance* 47 (4):1605–1621.

Ahern, Kenneth R. 2009. "Sample selection and event study estimation." *Journal of Empirical Finance* 16 (3):466–482.

Ai, Hengjie and Ravi Bansal. 2018. "Risk preferences and the macroeconomic announcement premium." *Econometrica* 86 (4):1383–1430.

Atkeson, Andrew G, Adrien d'Avernas, Andrea L Eisfeldt, and Pierre-Olivier Weill. 2019. "Government guarantees and the valuation of american banks." *NBER Macroeconomics Annual* 33 (1):81–145.

Bernanke, Ben S. 2011. "Transcript of Chairman Bernanke's press conference." April 27, 2011. <https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20110427.pdf>. Retrieved June 5, 2022.

Berndt, Antje, Darrell Duffie, and Yichao Zhu. 2019. "The decline of too big to fail." Manuscript, Stanford University.

Bernile, Gennaro, Jianfeng Hu, and Yuehua Tang. 2016. "Can information be locked up? Informed trading ahead of macro-news announcements." *Journal of Financial Economics* 121 (3):496–520.

Bloomberg BusinessWeek. 2009. "Stocks Finish Higher." *Business Source Complete* April 9.

Boguth, Oliver, Vincent Grégoire, and Charles Martineau. 2019. "Shaping expectations and coordinating attention: The unintended consequences of FOMC press conferences." *Journal of Financial and Quantitative Analysis* 54 (6):2327–2353.

Borochin, Paul, Murat Alp Celik, Xu Tian, and Toni M Whited. 2022. "Identifying the heterogeneous impact of highly anticipated events: Evidence from the Tax Cuts and Jobs Act." SSRN Working Paper No. 3806560.

Borochin, Paul and Joseph Golec. 2016. "Using options to measure the full value-effect of an event: Application to Obamacare." *Journal of Financial Economics* 120 (1):169–193.

Business Wire. 2009. "Fitch: TARP for U.S. Life Insurers Positive for Financial Strength." *Gale General OneFile* April 9.

Cetina, Jill, Arthur Fliegelman, Jonathan Glicoes, and Ruth Leung. 2016. "Mind the Gaps: What Do New Disclosures Tell Us About Life Insurers' Use of Off-Balance-Sheet Captives?" Tech. rep., Office of Financial Research Brief Series, US Department of the Treasury.

Christie, Rebecca and Andrew Frye. 2009. "Life Insurers Eligible for Aid: U.S. Treasury." *The Financial Post* April 8.

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel André Maréchal. 2015. "Evidence for countercyclical risk aversion: An experiment with financial professionals." *American Economic Review* 105 (2):860–85.

Copeland, Thomas E and David Mayers. 1982. "The value line enigma (1965–1978): A case study of performance evaluation issues." *Journal of Financial Economics* 10 (3):289–321.

Cummins, J David and Mary A Weiss. 2013. "Systemic Risk and the insurance Industry." In *Handbook of insurance*. Springer, 745–793.

Duffie, Darrell. 2010. "Presidential address: Asset price dynamics with slow-moving capital." *The Journal of Finance* 65 (4):1237–1267.

Ellul, Andrew, Chotibhak Jotikasthira, and Christian T Lundblad. 2011. "Regulatory pressure and fire sales in the corporate bond market." *Journal of Financial Economics* 101 (3):596–620.

Financial Stability Board. 2019. "Vulnerabilities associated with leveraged loans and collateralised loan obligations." December 19.

Financial Stability Oversight Council. 2014. "Basis for the Financial Stability Oversight Council's Final Determination Regarding Metlife, Inc." December 18.

Foley-Fisher, Nathan, Nathan Heinrich, and Stéphane Verani. 2020. "Capturing the Illiquidity Premium." Manuscript, Board of Governors of the Federal Reserve System.

Foley-Fisher, Nathan, Borghan Narajabad, and Stéphane Verani. 2020. "Self-fulfilling runs: Evidence from the US life insurance industry." *Journal of Political Economy* 128 (9).

Froot, Kenneth A. 2007. "Risk management, capital budgeting, and capital structure policy for insurers and reinsurers." *Journal of Risk and Insurance* 74 (2):273–299.

Giglio, Stefano. 2016. "Credit default swap spreads and systemic financial risk." European Systemic Risk Board Working Paper Series, No. 15.

Hartley, Daniel, Anna Paulson, and Richard J Rosen. 2017. "Measuring Interest Rate Risk in the Life Insurance Sector: The U.S. and the U.K." In *The Economics, Regulation, and Systemic Risk of Inusrance Markets*, edited by Felix Hudfeld, Ralph SJ Koijen, and Christian Thiman. Oxford University Press, 124–150. Chapter 6.

Houweling, Patrick and Ton Vorst. 2005. "Pricing default swaps: Empirical evidence." *Journal of International Money and Finance* 24 (8):1200–1225.

Hu, Grace Xing, Jun Pan, Jiang Wang, and Haoxiang Zhu. 2022. "Premium for heightened uncertainty: Explaining pre-announcement market returns." NBER Working Paper No. 25817.

Jasra, Ajay, Chris C Holmes, and David A Stephens. 2005. "Markov chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling." *Statistical Science* 20 (1):50–67.

Jones, Christopher S and Selale Tuzel. 2013. "Inventory investment and the cost of capital." *Journal of Financial Economics* 107 (3):557–579.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp. 2009. "Rational attention allocation over the business cycle." NBER Working Paper No. 15450.

Kardos, Donna. 2009. "Genworth Shares Fall On Ineligibility For Treasury Capital." *Dow Jones Institutional News* April 14.

Kim, Yunjeen. 2016. "Bank bailouts and moral hazard? Evidence from banks' investment and financing decisions." Manuscript, Boston University.

King, Robert G and Sergio T Rebelo. 1999. "Resuscitating real business cycles." *Handbook of macroeconomics* 1:927–1007.

Koijen, Ralph SJ and Stijn Van Nieuwerburgh. 2020. "Combining life and health insurance." *The Quarterly Journal of Economics* 135 (2):913–958.

Koijen, Ralph SJ and Motohiro Yogo. 2016. "Shadow insurance." *Econometrica* 84 (3):1265–1287.

———. 2017. "Risk of life insurers: Recent Trends and Transmission Mechanisms." In *The Economics, Regulation, and Systemic Risk of Inusrance Markets*, edited by Felix Hudfeld, Ralph SJ Koijen, and Christian Thiman. Oxford University Press, 79–99. Chapter 4.

———. 2021. "The fragility of market risk insurance." Manuscript, Becker Friedman Institute for Research in Economics Working Paper No. 2018-9.

Kurov, Alexander, Alessio Sancetta, Georg Strasser, and Marketa Halova Wolfe. 2019. "Price drift before US macroeconomic news: Private information about public announcements?" *Journal of Financial and Quantitative Analysis* 54 (1):449–479.

Kurov, Alexander, Marketa Halova Wolfe, and Thomas Gilbert. 2021. "The disappearing pre-FOMC announcement drift." *Finance Research Letters* 40:101781.

Laarits, Toomas. 2019. "Pre-announcement risk." Working paper, NYU Stern School of Business.

Lawder, David. 2009. "Treasury says some insurers qualify for TARP." *Reuters* April 8.

Liu, Hong, Xiaoxiao Tang, and Guofu Zhou. 2022. "Recovering the FOMC risk premium." *Journal of Financial Economics* 145 (1):45–68.

Lucca, David O and Emanuel Moench. 2015. "The pre-FOMC announcement drift." *The Journal of Finance* 70 (1):329–371.

MacKinlay, A Craig. 1997. "Event studies in economics and finance." *Journal of Economic Literature* 35 (1):13–39.

McDonald, Robert and Anna Paulson. 2015. "AIG in Hindsight." *Journal of Economic Perspectives* 29 (2):81–106.

McLean, R David and Jeffrey Pontiff. 2016. "Does academic research destroy stock return predictability?" *The Journal of Finance* 71 (1):5–32.

Merrill, Craig B, Taylor D Nadauld, René M Stulz, and Shane Sherlund. 2012. "Did capital requirements and fair value accounting spark fire sales in distressed mortgage-backed securities?" Manuscript, NBER.

Michaelides, Alexander and Serena Ng. 2000. "Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators." *Journal of Econometrics* 96 (2):231–266.

Michaels, Ryan, T Beau Page, and Toni M Whited. 2019. "Labor and capital dynamics under financing frictions." *Review of Finance* 23 (2):279–323.

Morse, Adair and Annette Vissing-Jorgensen. 2020. "Information Transmission from the Federal Reserve to the Stock Market: Evidence from Governors' Calendars." Working paper, University of California at Berkeley.

Moyer, Liz. 2009. "Treasury May Extend TARP To Life Insurers." *Forbes* April 7.

Nagel, Stefan. 2012. "Evaporating liquidity." *The Review of Financial Studies* 25 (7):2005–2039.

Naubert, Christopher and Linda L Tesar. 2019. "The Value of Systemic Unimportance: The Case of MetLife." *Review of Finance* 23 (6):1069–1078.

Nazareth, Rita. 2009. "U.S. Stocks Gain as Shares of Life Insurers, Centex Lead Rally." *Bloomberg* April 8.

Neuberg, Richard, Paul Glasserman, Benjamin Kay, Sriram Rajan et al. 2016. "The market-implied probability of european government intervention in distressed banks." Manuscript, Office of Financial Research.

Patterson, Scott. 2009. "Questions Dog Insurers on TARP." *The Wall Street Journal* April 9.

Patterson, Scott and Leslie Scism. 2009. "The Next Big Bailout Decision: Insurers." *The Wall Street Journal* March 12.

Patterson, Scott, Deborah Solomon, and Leslie Scism. 2009. "U.S. to Offer Aid to Life Insurers." *The Wall Street Journal* April 8.

Paulson, Anna, Thanases Plestis, Richard J Rosen, Robert McMenamin, and Zain Mohey-Deen. 2014. "Assessing the Vulnerability of the U.S. Life Insurance Industry." In *Modernizing Insurance Regulation*, edited by John H Biggs and Matthew P Richardson. John Wiley & Sons, Inc Hoboken, NJ, 61–83. Chapter 6.

Peirce, Hester. 2014. "Securities Lending and the Untold Story in the Collapse of AIG." Manuscript, Mercatus.

Phillips, Richard D, J David Cummins, and Franklin Allen. 1998. "Financial Pricing of Insurance in Multiple Line Insurance Companies." *Journal of Risk and Insurance* 65:597–636.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3):665–690.

Stephens, Matthew. 2000. "Dealing with label switching in mixture models." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62 (4):795–809.

Tauchen, George. 1986. "Finite state markov-chain approximations to univariate and vector autoregressions." *Economics letters* 20 (2):177–181.

Tibken, Shara. 2009. "Insurers Jump On Possible OK To Apply For TARP Funds." *Dow Jones Institutional News* February 9.

Tsesmelidakis, Zoe and Robert C Merton. 2012. "The value of implicit guarantees." Manuscript, Massachusetts Institute of Technology.

Vissing-Jorgensen, Annette. 2020a. "Central Banking with Many Voices: The Communications Arms Race." Paper presented at XXIII Annual Conference of the Central Bank of Chile, Santiago, Chile.

———. 2020b. "Informal central bank communication." NBER Working Paper No. 28276.

Williams Walsh, Mary. 2009. "Months In, Insurers Ask Whether Aid Is on Its Way." *The New York Times* February 26.

Wutkowski, Karey and Patrick Rucker. 2009. "Treasury to approve insurers for TARP funds: sources." *Reuters* February 7.