Essays on Financial Markets and Monetary Policy

Knox, Benjamin

Document Version
Final published version

Publication date:
2021

License
Unspecified

Link to publication in CBS Research Portal

ESSAYS ON FINANCIAL MARKETS AND MONETARY POLICY

PhD Series 15.2021

Benjamin Knox

ESSAYS ON FINANCIAL MARKETS AND MONETARY POLICY

CBS PhD School

PhD Series 15.2021

CBS COPENHAGEN BUSINESS SCHOOL
HANDELSHØJSKOLEN

Essays on Financial Markets and Monetary Policy

Benjamin Knox

A thesis presented for the degree of

Doctor of Philosophy

Supervisor: Lasse Heje Pedersen

Ph.D. School in Economics and Management

Copenhagen Business School

Benjamin Knox
Essays on Financial Markets
and Monetary Policy

The CBS PhD School is an active and international research environment at Copenhagen Business School for PhD students working on theoretical and empirical research projects, including interdisciplinary ones, related to economics and the organisation and management of private businesses, as well as public and voluntary institutions, at business, industry and country level.

Abstract

This thesis concerns financial markets and monetary policy. It consists of three chapters on the topics of insurance pricing, asset pricing and the economic effects of monetary policy respectively. The chapters can be read independently.

The first chapter considers the investment strategies of insurance companies and their impact on the pricing of insurance contracts. Our paper proposes and tests a new theory of insurance pricing, which we call "asset-driven insurance pricing". Consistent with the theory, we show empirically that (1) insurers with more stable insurance funding take more investment risk and, therefore, earn higher average investment returns; (2) insurance premiums are lower when expected investment returns are higher, both in the cross section of insurance companies and in the time series. Our findings indicate that the assets and liabilities of insurance companies are more connected than previously thought.

The second chapter presents a new decomposition approach for stock returns that is based on the sensitivity of the stock price with respect to expected returns and dividends at various horizons. Our method splits unexpected stock returns into news about cashflows and news about discount rates using observables. This decomposition, which is computed from the prices of traded financial products, avoids many of the model-implied assumptions associated with standard decomposition approaches. We apply our new decomposition in 2020, shedding light on the evolution of the return on US stocks during the COVID crisis.

The third chapter considers the effects of monetary policy on the economy. I document rich heterogeneity in business cycles across U.S. states. As a result, state-level Taylor rules imply very different optimal monetary policies across states. To exploit the cross-sectional variation, I present a granular approach to monetary policy identification. The intuition behind the approach is that shocks to economic activity in one state can lead to changes in monetary policy, which are exogenous monetary policy shocks from the perspective of other states. I implement this approach in the United States and find large effects of

monetary policy changes on future unemployment rates.

Acknowledgements

I would like to take the opportunity to say thank you to some of those who have helped and supported me through the process of writing this thesis. I have incurred some debts that I'm afraid will never be paid. The least I can do is acknowledge them here.

First, I would like to thank Lasse Heje Pedersen and Annette Vissing-Jorgensen, who have each been so influential on my thesis. Together they have taught me how to do research. To have Lasse as my advisor has been a genuine privilege. The time, thought and commitment he has put into my development as an economist I will always be deeply thankful for. It has also been fun along the way.

Annette initially sponsored me to visit Berkeley Haas in 2019. What started as a semester visit has become a great friendship and co-authorship. I feel very fortunate to have worked together, and must thank you for your belief, patience and support. All three have been invaluable.

Many others at Copenhagen Business School deserve a mention, including David Lando, Paul Whelan and Peter Feldhutter. David, in particular, should take great credit (and pride) for having built the FRIC research centre. Thanks also to the many PhDs (past and present) for the good times, and especially to Jakob Ahm Sørensen, who has been a great classmate, officemate, co-author, friend and councillor.

My wife Roberta deserves the most special thanks. There is no doubting its a tough life being the partner of a PhD candidate. I wouldn't wish it on anyone, and have great guilt having inflicted it upon you. On the plus side you've become the world's most knowledgeable architect on the topic of insurance funding. Not everyone can say that.

Last but not least, thanks to my family. I would like to dedicate this thesis to them.

Summaries in English

Asset-Driven Insurance Pricing

Insurance companies receive premiums from consumers at the start of insurance contracts, and, in exchange, promise to pay claims on the contracts at future dates. There are two important features of these contracts. The first is that the claims are uncertain, which creates risk for the insurance company. To compensate themselves for this risk, insurers can charge contract prices that are greater than the expected claims, and thus generate insurance underwriting profits. The second feature is the timing of the cashflows. Insurance companies receive premiums at the start of contracts and paying claims later. In effect, they are borrowing money from their consumers. Insurance companies can thus generate investment profits on insurance contracts by investing premiums before claims come due.

This chapter of the thesis considers the impact of insurance company investment strategies on their pricing of insurance contracts. A growing body of evidence in the literature has shown that there is significant risk in the asset portfolios of insurers (Ellul et al. (2011), Becker and Ivashina (2015), Becker et al. (2020), Ge and Weisbach (2020), Ellul et al. (2020)). However, the implications of these risky asset portfolios for insurance pricing is under-explored. To help bridge this gap, we propose and test a new theory of insurance pricing, which we call "asset-driven insurance pricing". This pricing behaviour shows that insurance companies set insurance premiums lower when their expected investment returns on risky assets are higher.

The previous assumption in the literature was that all the profits from risky investment strategies should go to the owners of the insurance company (i.e. the shareholders). However, we show that some of these risky investment profits instead go to the consumers of the insurance company (i.e. the policyholders). To explain this pricing behaviour, we argue that insurance companies have a competitive advantage investing into illiquid asset

markets due to the stable funding that insurance underwriting provides. Consistent with this interpretation, we find that the insurance companies with the most stable insurance funding (a) invest a larger faction of their assets into illiquid assets, and (b) set lower premiums relative to their competitors when returns to illiquid assets are higher.

In summary, we contribute to the literature by uncovering a new stylized fact and presenting theory that explains this fact: insurance premiums are asset driven.

A Stock Return Decomposition Using Observables

What news drives fluctuations in the price of the stock market? This chapter of the thesis contributes to this central question in the asset pricing literature.

In theory, the price of the stock market is the present value of all its future dividends. A fall in the price can therefore be due to investors expecting future dividends to be lower (cashflow news), or the value of the dividends to investors decreasing (discount rate news). Standard approaches to decomposing these two sources of variation are based around methods first proposed by Campbell and Shiller (1988). However, as shown by Chen and Zhao (2009), there are several problems with these approaches that arise due to the sensitivity of results to the inputs used in the models.

To overcome the issues, we show that one can get a long way towards a decomposition of unexpected stock market returns into cash flow news and discount rate news using the prices of traded financial products. Our contribution is therefore to provide a model-free method of decomposition that does not require some of the assumptions usually made in the literature. Our approach has been made possible by the introduction of many new financial products to financial markets in recent decades, including index-linked government bonds, dividend futures and equity options. We show how to convert the prices of these products into an overall stock market return decomposition.

We use our approach to understand the evolution of the stock market over the COVID crisis in 2020. Our decomposition reveals three key facts. First, risk premium increased sharply as the crisis intensified in March, contributing 14 percent of the 26 percent market decline up to March 18. Second, the market recovery was heavily influenced by declining real rates, which contributed a positive 18 percent to the realized stock return for the year. Third, news about dividends out to 10 years had a modest effect with a larger role for a decline and subsequent recovery of expectations for more distant dividends.

A Granular Identification of Monetary Policy

Monetary policy is action that a country's central bank takes to influence how much money is in the economy and how much it costs to borrow. Conventional wisdom is that monetary policy decisions matters qualitatively for the economy. However, the quantitative impact of monetary policy changes is still poorly understood.

In this chapter of the thesis, I contribute to the literature by presenting a new approach for identifying the quantitative impact of monetary policy on the economy. Any identification approach needs random variation in monetary policy, and the idea behind mine is to use cross-sectional variation in economy activity within a currency union. A shock to the economic growth of one region in a currency union can lead to a change in monetary policy, with this change in monetary policy in effect random from the perspective of the other regions.

There are many monetary policy identification methods that have been applied previously in the literature,[1] which all rely on truley random variations in monetary policy. However, Ramey (2016) highlights that monetary policy shocks are small and rare due to the predictable nature of monetary policy decisions, which makes identification this way challenging. An appealing contribution of my approach is therefore that it works even in the extreme case where monetary policy is completely predictable.

I implement the approach in the U.S. and find a strong effect of monetary policy on unemployment rates. When the central bank increases the interest rate by 1 percentage point, I find that unemployment increases by 1.8 percentage points over the next 15 months. The effect is larger than existing estimates in the literature, indicating a more important role of monetary policy than previously estimated.

---

[1]including narrative methods (Friedman and Schwartz (1963)), VAR model's (Christiano et al. (1999)), deviations from Taylor rules (Romer and Romer (2004) Coibion (2012)), and high-frequency methods (Gertler and Karadi (2015) Nakamura and Steinsson (2018))

Summaries in Danish

Aktiv-drevne Forsikringspræmier

Ved indgåelsen af en forsikringskontrakt modtager forsikringsselskabet en forsikringspræmie fra kunden mod et løfte om en fremtidig udbetaling. En sådan forsikringskontrakt har to vigtige kendetegn. For det første er størrelsen på den fremtidige udbetaling usikker, hvilket udgør en risiko for forsikringsselskabet. Som kompensation for denne risiko kan forsikringsselskabet kræve en forsikringspræmie som overstiger det forventede tab på kontrakten, og således generere en profit. Det andet kendetegn er timingen af betalingerne. Forsikringsselskabet modtager forsikringspræmien umiddelbart efter kontrakten er indgået, og udbetaler først senere. Det betyder at forsikringsselskaber i praksis låner penge af deres kunder. Forsikringsselskaber kan således generere en profit ved at investere forsikringspræmien inden en eventuel udbetaling indtræffer.

Dette kapitel i afhandlingen omhandler forsikringsselskabers investeringsstrategier og hvordan disse påvirker prisen på forsikringer. Et stadigt stigende antal akademiske artikler har påvist at forsikringsselskaber påtager sig betydelige risici i forbindelse med deres investeringer ((Ellul et al. (2011), Becker and Ivashina (2015), Becker et al. (2020), Ge and Weisbach (2020), Ellul et al. (2020)). Hvordan dette påvirker prisen på forsikringer er dog et uudforsket spørgsmål. Vi besvarer spørgsmålet i denne artikel ved både at fremlægge en ny teori for prissætning af forsikringskontrakter, og ved at teste denne teori empirisk. Vi kalder denne nye teori for "aktiv-drevne forsikringspræmier". Vi viser både teoretisk og empirisk at forsikringsselskaber sætter lavere priser på deres forsikringer når de har højere forventede afkast på deres aktiver.

Den akademiske litteratur på området har tidligere antaget at al profit fra risikable investeringer ville tilfalde forsikringsselskabets ejere (dvs. aktionærerne). I modsætning til denne antagelse påviser vi at en del af den profit der generes gennem forsikringsselskabets investeringer, tilfalder selskabets kunder gennem lavere forsikringspræmier. Vi

argumenterer i artiklen for at denne prissætning kan forklares ved at forsikringsselskaber har en konkurrencemæssig fordel når det kommer til at investere i illikvide aktiver på grund af den stabile finansieringskilde som forsikringspræmierne udgør. Vi finder, i overensstemmelse med denne fortolkning, at forsikringsselskaber med stabil finansiering fra forsikringspræmier (a) investerer en større andel af deres portefølje i illikvide aktiver, og (b) sætter lavere præmier end deres konkurrenter når de forventede afkast på illikvide aktiver stiger. For at opsummere, vi bidrager til litteraturen på området ved at fremlægge et nyt stiliseret faktum, og en teori som forklarer dette faktum: forsikringspræmier er aktivdrevne.

En dekomposition af aktieafkast ved brug af handlede aktiver

Hvilke nyheder driver ændringer af prisen på aktiemarkedet? Dette kapitel i afhandlingen besvarer til dette centrale spørgsmål i litteraturen om prissætning af finansielle aktiver.

I teorien er prisen på det samlede aktiemarked den nutidige værdi af alle fremtidige dividender. Et fald i aktiemarkedet kan derfor skyldes et fald i investorernes forventninger til de fremtidige dividender (dividendenyheder), eller et fald i nutidsværdien af de fremtidige dividender (diskonteringsnyheder). Den klassiske tilgang til adskillelse af disse to kilder til variation i aktiepriser er baseret på metoder som blev introduceret første gang af Campbell and Shiller (1988). Der er dog, som vist i Chen and Zhao (2009), adskillige problemer med den klassiske tilgang grundet resultaternes følsomhed over for modelinput.

Vi viser i dette papir at man i høj grad kan overkomme problemerne med den klassiske tilgang ved at bruge priser på handlede finansielle aktiver til at adskille dividendenyheder fra diskonteringsnyheder. Vores bidrag er at udvikle en modelfri dekompositionsmetode, som ikke beror på de antagelser som litteraturen traditionelt set har gjort brug af. Vores metode er muliggjort af en række finansielle aktiver som er blevet introduceret i løbet af de sidste årtier såsom indekserede statsobligationer, dividendefutures og aktieoptioner. Vi viser hvordan man kan anvende priserne på disse produkter til en dekomposition af det samlede aktiemarked.

Vi bruger vores nye metode til at forstå bevægelserne i aktiemarkedet under Covidkrisen i 2020. Vores dekomposition afslører tre centrale observationer. Den første observation vi gør er at risikopræmien steg dramatisk i marts da krisen intensiveredes. Konkret udgjorde stigningen i risikopræmien 14 af de 26 procent som markedet faldt frem mod den 18 marts. Den anden observation er at markedets genopretning efter den 18 marts

var stærkt påvirket af en faldende realrente som alene bidrog med 18 procent til årets realiserede aktieafkast. Den tredje og sidste observation vi gør er at nyheder om dividender til udbetaling inden for de næste ti år havde en moderat effekt på markedets fald og dets efterfølgende genopretning, mens nyheder om dividender med længere tidshorisonter spillede en større rolle.

En granulær identifikation af pengepolitik

Pengepolitik er handlinger som et lands centralbank foretager for at påvirke pengemængden i økonomien, og de overordnede låneomkostninger. Den konventionelle visdom har hidtil været at pengepolitik har en kvalitativ betydning for økonomien, mens den kvantitative betydning af pengepolitik endnu ikke har været tilstrækkeligt forstået.

I dette kapitel af afhandlingen bidrager jeg til den akademiske litteratur ved fremlægge en ny tilgang til at identificere pengepolitikkens kvantitative indflydelse på økonomien. Enhver identifikation beror på variation i pengepolitikken, og idéen bag min tilgang er at benytte variation i økonomisk aktivitet på tværs af medlemsstater i en valutaunion. En ændring i den økonomiske vækst i en medlemsstat kan føre til en ændring i pengepolitikken for hele valutaunionen. Denne ændring i pengepolitikken kan betragtes som fuldstændig tilfældig af de andre medlemmer i valutaunionen.

Den akademiske litteratur har tidligere forsøgt at identificere effekten af pengepolitik. ,[2] som alle gør brug af ren eksogen variation i pengepolitikken. Ramey (2016) pointerer dog at denne slags ren eksogen variation i pengepolitik er sjælden og lille i størrelse. Dette skyldes at pengepolitik i høj grad er forudsigelig, hvilket besværliggør identifikationen. Min metode har den tiltalende egenskab at den fungerer selv i det ekstreme scenarie hvor pengepolitikken er fuldt ud forudsigelig.

Jeg implementerer min tilgang på data fra USA og finder at pengepolitikken har en stærk effekt på arbejdsløsheden. Når den amerikanske centralbank øger renten med et procentpoint, stiger arbejdsløsheden med 1.8 procentpoint over de følgende 15 måneder. Denne effekt er større end hvad man tidligere har kunnet påvise i den akademiske litteratur, hvilket indikerer at pengepolitik er vigtigere end hidtil antaget.

---

[2]Disse inkluderer *narrativ metoder* (Friedman and Schwartz (1963)), VAR modeller (Christiano et al. (1999)), afvigelser fra Taylor reglen (Romer and Romer (2004) Coibion (2012)), og metoder der gør brug af højfrekvens data (Gertler and Karadi (2015) Nakamura and Steinsson (2018))

Contents

Chapter 1

Asset-Driven Insurance Pricing

Benjamin Knox and Jakob Ahm Sørensen[1]

We develop a theory that connects insurance premiums, insurance companies' investment behavior, and equilibrium asset prices. Consistent with the model's key predictions, we show empirically that (1) insurers with more stable insurance funding take more investment risk and, therefore, earn higher average investment returns; (2) insurance premiums are lower when expected investment returns are higher, both in the cross section of insurance companies and in the time series. We show our results hold for both life insurance companies and, using a novel data set, for property and casualty insurance companies. Consistent findings across different regulatory frameworks helps identify asset-driven insurance pricing while controlling for alternative explanations.

1.1 Introduction

This paper proposes and tests a new theory of insurance pricing, which shows that insurance premiums are lower when insurance companies have higher expected investment returns. We call this way of setting premiums "asset-driven insurance pricing". Our theory and evidence connects two important functions of the insurance industry, namely the pricing of insurance products and the allocation of its assets. Insurance products facilitate risk-sharing for 95% of all US households, and the premiums fund large asset portfolios, with US insurers holding marketable asset worth $11.2 trillion as of Q4 2019.[2] Hence, insurance companies are both economically important asset allocators and facilitators of risk sharing, and we show that these two functions are more connected than previously thought.

The traditional view of insurers is that their main business – and therefore their main source of risk and return - is insurance underwriting. Such a view has little consideration for insurer's asset allocation decisions in the context of insurance premium pricing. However, recent evidence shows that there is significant risk in the asset portfolios of insurers (Ellul et al. (2011), Becker and Ivashina (2015), Becker et al. (2020), Ge and Weisbach (2020), Ellul et al. (2020)). Indeed, contrary to the traditional view, risk-free assets make up only 10% of investment portfolios, with insurers instead investing heavily in illiquid credit markets. This behaviour in their investment portfolios motivates our two main research questions: (1) Why do insurers have such high exposure to credit and liquidity risk in their asset portfolios? (2) Do the expected investment returns on these portfolios affect how they set premiums?

We address these questions by considering a model of insurance premiums and illiquid asset prices and by presenting consistent empirical evidence. We show asset-driven insurance pricing holds in both the time series and the cross section of insurance companies, in good and bad times, and for both life insurance companies and the property and casualty (P&C) industry. The P&C results use novel data, which, due to the industry's distinct regulatory framework relative to the Life Insurance industry, helps us to identify asset-driven insurance pricing from alternative mechanisms of insurance pricing. We also present evidence of asset-driven insurance pricing following changes to investment returns

---

[2]For a sense of the order of magnitude, note that the total value of insurer marketable assets is in excess of 40% of the US Treasury and corporate bond markets combined. Data sources: Insurance Information Institute, Financial Accounts of the United States (Fed Reserve), SIFMA Fact Book.

due to mergers.

Our model features two types of agents, investors and insurance companies. There are also two assets, one liquid and one illiquid. All agents face an exogenous cost from selling the illiquid asset before maturity, and, in the spirit of Diamond and Dybvig (1983), investors are *ex-ante* uncertain whether they are early or late consumers. These assumptions combine to generate an endogenous liquidity risk premium. The key insight of the model is that insurers enjoy relatively more certainty on the timing of cash flows due to the diversification benefit of underwriting many homogeneous insurance policies. This diversification creates stable insurance funding, which is an advantage when investing in illiquid assets.

Insurance companies with more stable insurance funding are able to extract more value from illiquid assets and therefore allocate a greater fraction of assets to illiquid investments (Proposition 1). In the time series, when the excess return on the illiquid asset is higher, the marginal cost of supplying insurance is lower, insurers compete for funding, and insurance premiums are set lower in the aggregate (Proposition 2). In the cross section, insurance companies that take more investment risk and have higher expected returns are able to set lower premiums relative to competitors (Proposition 3). The model's predictions rest on a violation of the Modigliani and Miller (1958) capital irrelevance theorem. We argue that an investor's funding structure matters when a illiquidity return premium is available in asset markets, and insurers' funding choices determine their ability to earn the illiquidity return premium.

To test Proposition 1, we calculate rolling 5-year estimates of the standard deviation of insurer's underwriting profitability. Using data from 2001-2018, we find that insurers with more stable underwriting profitability have lower allocations to cash assets and higher allocations to credit assets (and take more credit risk within their credit portfolios). Our results extend on Ge and Weisbach (2020), who show that large insurers take more investment risk. Assuming large insurers have more diversification benefits in their underwriting businesses, this initial result is consistent with our model prediction. However, our findings take this a step further, showing that, even when comparing firms of equal size, the insurer with less volatile underwriting performance takes more investment risk. The finding provides evidence that insurer's asset allocation decision depends on firm-level characteristics, and specifically on the stability of cash flows in their underwriting business. According to our model, the explanation is that insurers use the stability of the

insurance funding to earn liquidity premium on their assets.

To test Proposition 2 and the time series of premiums, we use credit spreads as a proxy for industry-wide expected investment returns. Figure 1.1 presents an illustrative example in the life insurance industry, plotting the industry average insurance premium against credit spreads (on an inverse axis scale). The figure shows that insurance premiums are lower when insurance companies have higher expected investment returns. Our main dependent variable in the Life Insurance industry are annuity markups as calculated in Koijen and Yogo (2015). Across products, we find a 100bp increase in credit spreads leads to a 50bp decrease in an annualised annuity markup on average, with a t-statistic of 4.03 controlling for other effects. The average markup is 1%, and hence the 50bps decrease mean insurers drop their markups by half when they can earn 100bp more buying corporate bonds. This sensitivity is an economically significant effect. In the P&C industry, we use insurers' reported underwriting profitability as the main dependent variable. This measure is the ratio of their insurance underwriting profit to their insurance underwriting liabilities. We interpret lower underwriting profit as evidence of lower premiums. We find that the industry average underwriting profitability ratio falls by 1.31 standard deviations (t-statistic of 4.68 with full controls) when lagged credit spreads increase by one standard deviation.

To test Proposition 3, we use insurer's reported accounting investment returns to measure cross sectional variation in investment opportunities. The analysis utilizes a rich heterogeneity in investment portfolios across insurers. At any point in time, we show that the level of credit risk in credit portfolios explains the majority of variation in accounting returns, and that this variation predicts future returns, consistent with our interpretation that accounting returns captures insurers' expected investment returns.[3] We consistently find that the insurers with higher expected investment returns set lower insurance prices. In the life insurance industry, an insurer with an expected investment return one standard deviation higher than competitors reduces their relative markup by 0.05 standard deviations (t-statistic 2.77). In the P&C industry, we find an insurer with a one standard deviation higher expected investment return has an underwriting profitability ratio 0.03 standard deviations lower than competitors (t-statistic 5.45). The magnitudes are not as large as in the time series, showing that investment returns have more affect on industry

---

[3]Anecdotal evidence from market participants also tells us that insurers consider accounting returns to reflect future expected investment returns.

average premium, rather than relative pricing in the cross section of premiums.

We provide further evidence of asset-driven insurance pricing with three extensions to our analysis. First, in the cross section of P&C insurers, we implement an instrumental variable estimation, using underwriting funding volatility and firm size (from the test of Proposition 1) as instruments for insurer's investment returns. We show that when instrumented investment returns are 100bps higher, insurance premiums are 0.3 percentage points lower. Second, in the cross section of life insurers, we use a series of shocks to investment return due to mergers. When insurer companies are purchased by other insurers, their investment returns change as their portfolios adapt to the investment strategy of their acquiring insurance company. Using a difference in difference analysis, we show how insurance premiums fall (rise) in response to increases (decreases) in investment returns that are driven specifically by merger events. Third, in the time series, we show that the sensitivity to credit spreads is driven by expected excess return on bonds, as proxied by the Gilchrist and Zakrajšek (2012) excess bond premium, rather than the component of credit spreads that reflects expected default risk.

To understand our contribution, it useful to think of insurance premiums as the product of:

$$\text{Premium} = \underbrace{\frac{E\left[Claim\right]}{1 + R^F}}_{\substack{\text{Actuarial price:} \\ \text{(Hill, 1979)} \\ \text{(Kraus and Ross, 1982)}}} \times \underbrace{\text{Markup}}_{\substack{\text{Imperfect competition} \\ \text{(Mitchell et al., 1999)}}} \times \underbrace{\text{Shadow Cost}}_{\substack{\text{Regulatory capital constraints} \\ \text{(Froot and O'Connell, 1999)} \\ \text{(Koijen and Yogo, 2015)} \\ \text{(Ge, 2020)}}} \times \underbrace{\frac{1 + R^F}{1 + R^I}}_{\substack{\text{Asset-driven insurance pricing} \\ \text{(this paper)}}}$$

The first term is the expected claim discounted at the risk free rate. It is typically considered to be the insurers' marginal cost of underwriting a policy. The basic intuition is that an insurer can invest premiums received in a portfolio of Treasury bonds that replicate the expected liabilities. Due to the time value of money, the marginal cost is therefore lower than the expected claim. The second term results from imperfect competition, and the third term rests on theories of financial frictions. When insurers are capital constrained and their access to external finance is costly, they deviate from their optimal unconstrained premium price in order to improve their regulatory capital position. The contribution of this paper is to return to the fundamental question of what insurance companies consider to be their time value of money. We challenge whether it is the risk-free rate, as the actuarial price suggests, instead arguing that insurers' also use the liquidity premium in their expected investment return, R^I, such that the discount rate is higher than the risk-free rate. The rational is based on there being a liquidity friction in asset markets, with

insurance companies able to take advantage of this due to their unique funding source.

We consider the other channels of insurance pricing in our analysis, with particular focus on capital constraints (Froot and O'Connell (1999), Koijen and Yogo (2015), Ge (2020)), which has previously been shown to drive insurance prices. To guide the empirical analysis, we first extend the model with a statutory capital constraint that, in the spirit of Koijen and Yogo (2015), shows how insurance premiums can change when the constraint is binding. To rule out this mechanism capital constraints as the driver of our empirical results, we show that asset driven insurance pricing is present in the P&C markets industry, where binding capital constraints should result in higher premiums, thus alleviating the confounding variable problem. We further show that our results hold in periods where insurance companies are unlikely to have been capital constrained. We therefore argue that while capital constraints play an important role in insurance pricing, they are not the only factor. Instead, insurance companies also account for expected returns when setting prices, and this mechanism is especially important when insurance companies are unconstrained by regulatory capital requirements.

Two other alternative mechanisms we consider empirically are differences in the demand for insurance and also reinsurance activity. A possible explanation of our cross sectional results is that the insurance companies which take more investment risk are more likely to default themselves. Lower insurance premiums could thus be driven by relatively lower demand for insurance relative to their competitors. To rule out this alternative mechanism, we use AM Best capital strength ratings, showing that our results hold for the subset of highly rated firms in the life industry. The results also hold after controlling for measures of balance sheet strength in the full sample of P&C insurers. Regarding reinsurance activity, a potential alternative hypothesis is that insurance companies that are better able to reinsure their liabilities are therefore able to set lower premiums.[4] We show our results are robust to controlling for the fraction of an insurer's underwriting premiums that are reinsured.

Our paper is also related to Stein (2012), Hanson, Shleifer, Stein, and Vishny (2015), and Chodorow-Reich, Ghent, and Haddad (2020) who also study the comparative advantage of intermediaries investing in illiquid assets. As in our paper, these theories rest on a violation of the Modigliani and Miller (1958) capital irrelevance theorem, with an asset's value dependent on the funding structure of the investor. In particular, intermediaries are

---

[4]We thank Stefano Rossi for this observation.

able to earn excess returns relative to other investors. However, in the referenced papers, the value generated flows to the equity holder of the intermediary by assumption. The key contribution of our paper is to document that the value from stable funding can flow to the insurer's policy holders, rather than just the equity holders. Our finding has potential welfare implications, with insurers offering cheaper insurance to households when financial markets are distressed.

Novy-Marx and Rauh (2011) and Rauh (2016) document how US pension funds increase the discount rate on their existing liabilities to reduce the present value of their reported liabilities. We instead study how insurance companies set the price on new liabilities, highlighting the interconnectedness of an insurer's assets and liabilities. In this sense, our paper relates to Kashyap, Rajan, and Stein (2002), who show study the synergies of banks assets and liabilities. While their paper focuses on how banks provide immediate liquidity on both liabilities and assets (i.e. credit lines), we argue insurer's stable liabilities mean they can take liquidity risk on their assets.

More broadly, our results relates to the intermediary asset pricing literature. Constraints on the liability side of intermediary's balance sheets affect their asset preferences (Brunnermeier and Pedersen (2009) He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014)) which ultimately ends up changing asset prices (Ellul et al. (2011), Adrian et al. (2014), He et al. (2017) and Greenwood and Vissing-Jorgensen (2018)) due to intermediary's position as marginal investors in segmented markets. We not only study how intermediaries affect asset prices, but also consider how asset markets affect intermediary liability prices. The findings of our paper therefore sheds further light on the interdependencies of intermediaries and asset markets that has been widely discussed post financial crisis.

In summary, we contribute to the literature by uncovering a new stylised fact and presenting theory that explains this fact: insurance premiums are asset driven.

1.2 Model of Insurance Premiums and Illiquid Asset Prices

The economy has three periods, $t = 0$, 1 and 2, two types of agents, investors and insurance companies, and two asset markets.

Assets. There is a liquid asset with exogenous return R^F, and an illiquid asset with fixed

supply S. The illiquid asset pays one unit of wealth at maturity $t = 2$, and the price at $t = 0$ is determined endogenously. The defining characteristic is that the illiquid asset incurs a cost if sold before before maturity (i.e. sold at $t = 1$). The seller of the asset receives their initial investment less a cost of $\frac{1}{2}\lambda x^2$ dollar for every x dollar sold. The parameter λ therefore captures liquidity conditions in the secondary market of the illiquid asset.

Investors. A continuum of risk-neutral investors, each endowed e, are identical at $t = 0$. In the spirit of Diamond and Dybvig (1983), they learn at $t = 1$ if they are early or late consumers. Early consumers only care about consumption at $t = 1$, while late consumers only care about consumption at $t = 2$. Each investor knows at $t = 0$ the probability ω of being an early consumer.

If the investor chooses to a buy dollar amount θ of the illiquid asset their consumption is

$$c = \begin{cases} e\left(1 + R^F\right) - \frac{1}{2}\lambda\theta^2 & \text{with probability } \omega \qquad \text{(early consumer)} \\ e\left(1 + R^F\right) + \theta R & \text{with probability } 1 - \omega \quad \text{(late consumer)} \end{cases} \tag{1.1}$$

where

$$R = \frac{1}{\text{Asset Price}} - \left(1 + R^F\right) \tag{1.2}$$

is the equilibrium excess return on the illiquid asset.

In the first case of equation (1.1), the investor learns they are an early consumer and sells all assets at time 1, paying the associated transaction costs on their illiquid asset holdings. In the second case, the investor learns they are a late consumer and holds all assets to maturity, earning the excess return on their illiquid asset holdings.

The problem facing the investor is to choose θ to maximise expected consumption

$$\max_{\theta} \quad \mathbb{E}\left[c\right] = e\left(1 + R^F\right) + \left(1 - \omega\right)\theta R - \frac{1}{2}\omega\lambda\theta^2. \tag{1.3}$$

Insurance Companies. The economy's other agent is a representative insurance company. The risk-neutral insurer receives premiums on insurance policies at $t = 0$ and pays the policy claims at either $t = 1$ or 2. The premium P is set by the insurance company,

and the number of policies sold is determined by the exogenously given downward sloping demand curve

$$Q\left(P\right) = kP^{-\epsilon} \tag{1.4}$$

where $\epsilon > 1$ is the elasticity of demand.

The insurer is endowed with equity capital E at $t = 0$ such that their total liabilities

$$L = E + QP \tag{1.5}$$

are the sum of equity and the funding generated from the insurance underwriting business. The total future claims underwritten are defined

$$C = Q\bar{C}. \tag{1.6}$$

where $\bar{C}$ is the policy claim on each individual contract.

We assume that the insurance business is sufficiently diversified that we can think of total claims, C, as being a known constant. Insurance companies are thus not worried about the size of the claims to be paid, but instead face liquidity risk as claims can arrive at either $t = 1$ or $t = 2$. We define the fraction of total claims arriving time 1 as $\tau \in \{\bar{\tau} - \sigma, \bar{\tau} + \sigma\}$ and assume that each state occurs with equal probability. The remaining fraction of claims, $(1 - \tau)$, arrive at time 2. Claims are on insurance products such as car or household insurance, which are not related to the investment liquidity risk, λ, and are held by households outside of the model.

The insurer buys dollar amount $\Theta \geq 0$ in the illiquid asset and puts remaining wealth $L - \Theta \geq 0$ in the liquid asset. We assume both allocations are greater than or equal to zero, so the insurer's only source of balance sheet leverage is the funds generated from insurance underwriting.

The insurer's final wealth depends on the dollar amount τC of claims to be paid at $t = 1$ relative to the dollar amount $L - \Theta$ invested in the liquid asset. If the insurer holds more liquid assets than early claims, there is no sale of illiquid assets at $t = 1$. However, if early claims exceed liquid asset holdings, the insurer is forced to sell a fraction of illiquid assets before maturity. The final wealth is thus expressed with two cases

$$W = \begin{cases} L\left(1 + R^F\right) - C + \Theta R & \text{if } \tau C \leq L - \Theta \\ L\left(1 + R^F\right) - C + \left(L - \tau C\right) R - \frac{1}{2}\lambda\left(\tau C - \left(L - \Theta\right)\right)^2 & \text{if } \tau C > L - \Theta. \end{cases} \tag{1.7}$$

The first case shows the simple outcome in which the insurer holds enough liquid assets to cover early claims and all illiquid asset holdings therefore earn the liquidity premium R.

In the second case, the insurer sells all their liquid assets plus a portion of their illiquid asset portfolio to cover remaining $t = 1$ claims. Dollar amount $\tau C - (L - \Theta)$ of illiquid assets are sold before maturity and incur the associated sale cost, which we assume is paid at $t = 2$. The dollar amount of unsold illiquid assets is the initial holdings minus the sold holdings: $\Theta - (\tau C - (L - \Theta)) = L - \tau C$. These illiquid assets still earn the liquidity premium.

The insurer's objective function is to choose P and Θ to maximise their expected final wealth

$$\max_{P,\,\Theta} \ \mathbb{E}\left[W\right] \tag{1.8}$$

where wealth W is defined in equation (1.7).[5]

Equilibrium. We conclude this section by defining the equilibrium in the economy. The competitive equilibrium in the illiquid asset market is given by the market clearing condition

$$\theta^* + \Theta^* = S \tag{1.9}$$

where investor demand θ^* and insurer demand Θ^* are given by the optimisation problems (1.3) and (1.8) respectively. Supply S of the illiquid asset is exogenously given. Equilibrium in the insurance market is also where demand equals supply, with supply given by the insurers profit maximisation (1.8) and demand exogenously given from demand curve (1.4).

1.3   Theoretical Results

We begin by considering the asset allocation decision of the two agents in the model. All proofs are in Appendix 1.11.

Proposition 1 (illiquid asset allocations).

1. The investor's equilibrium dollar investment in the illiquid asset is

$$\theta^* = \frac{(1 - \omega)}{\omega} \frac{R}{\lambda}. \tag{1.10}$$

---

[5]We could also have insurance equity bought by investors,and insurance companies maximising the present value of final wealth. As along as the discount rate is a fixed required return (for example, the liquid return R^F or the illiquid return $R^F + R$), it is therefore independent of the insurance company's asset allocation, and the qualitative results of the model are unchanged. A fixed required return results from the fact that agents are risk-neutral.

2. The insurer's equilibrium dollar investment in the illiquid asset is

$$\Theta^* = L - (\bar{\tau} + \sigma)\,C + \frac{R}{\lambda}. \tag{1.11}$$

The investor and insurer both increase their illiquid asset allocation in the illiquid asset excess return, R, and reduce their illiquid asset allocation in the cost λ of selling the illiquid asset in secondary markets. The investor and insurer also decrease their illiquid allocation in the probability of early consumption ω and the expected fraction of claims $\bar{\tau}$ to be paid early. These parameters increase the chance of costly $t = 1$ sales of the illiquid asset. For the insurer, the variance σ of claims arriving early also matters for the illiquid investment allocation. The more volatile an insurer's funding (i.e. higher σ), the less illiquid assets they hold.

We next consider the insurer's pricing decision on insurance policies. We assume that the insurer treats the excess return on the illiquid asset R as a fixed constant — that is, they do not internalize the incremental impact of their choices on the magnitude of the excess return. First-order conditions of equation (1.8) with respect to P therefore yields the following proposition.

Theorem 1 (asset-driven insurance pricing). *The equilibrium insurance premium P of a policy with claim $\bar{C}$ is*

$$P = \frac{\bar{C}}{1 + R^F}\left(\frac{\varepsilon}{\varepsilon - 1}\right)\left(\frac{1 + R^F}{1 + R^I}\right) \tag{1.12}$$

where R^I is the insurer's expected investment return on their asset holdings that are funded by premiums

$$R^I = \frac{1 + R^F + R}{1 + (\bar{\tau} + \sigma)\,R} - 1 > 0. \tag{1.13}$$

We can see that the insurance premium is the product of three components. The first term, the actuarial price, is the claim discounted by the risk-free rate. The second term, $\frac{\varepsilon}{\varepsilon - 1} > 1$, is the markup the insurer can charge due to imperfect competition.[6] The final term, $\frac{1 + R^F}{1 + R^I} < 1$, is related to the insurer's expected excess return on their illiquid asset holdings. Given that the fraction of claims $\tau \in \{\bar{\tau} - \sigma, \bar{\tau} + \sigma\}$ arriving at $t = 1$ can not exceed one, we know that $R^I > 0$. This means that insurers set lower premiums when illiquid investment returns are higher. We call this *asset-driven insurance pricing.*

---

[6]As the elasticity of demand for insurance tends to infinity, the insurer has no market power and the markup tends to one.

Asset-driven insurance pricing means that the premium depends on the illiquid asset excess return R, and the funding characteristics ($\bar{\tau}$ and σ) of the insurer. The insurer's borrowing costs through insurance underwriting are now dependent on their asset allocation and funding decisions. This Modigliani and Miller (1958) violation occurs because insurance companies can earn a risk-free liquidity premium on illiquid investments due to their stable funding.

To understand the mechanism, note that the maximum amount of claims to be paid by the insurer at $t = 1$ is $(\bar{\tau} + \sigma)\,C$. This observation leads to the lower bound $\underline{\Theta}$ on the insurer's illiquid asset holdings

$$\underline{\Theta} = L - (\bar{\tau} + \sigma)\,C. \tag{1.14}$$

Investing less than this in illiquid assets would mean forgoing liquidity premium that is available to the insurer risk-free, so $\Theta^* \geq \underline{\Theta}$. Other investors in the economy, on the other hand, face the risk of selling all assets at $t = 1$. The $\underline{\Theta}$ component of the illiquid allocation is therefore the insurer's source of competitive advantage relative to other investors in the illiquid asset market. Indeed, as $\underline{\Theta}$ investments earn insurers R with zero risk, these investments lower the insurer's marginal cost of underwriting . Insurers therefore compete for funding and insurance premiums are set lower when R is higher.

The special case where $\bar{\tau} + \sigma = 1$ illuminates the point. In this case, the insurer faces the risk that all claims arrive at $t = 1$ and they thus have no competitive advantage. The expected investment return on the asset holdings funded by premiums is $R^I = R^F$, and our result nests Modigliani and Miller (1958). The insurance premium is priced by discounting the claim by the exogenously given liquid risk-free rate, and is no longer dependent on the insurer's illiquid asset allocation Θ or the equilibrium liquidity premium R.

The model's next prediction follows directly from the partial derivative of insurance premium with respect to illiquid asset returns. While insurance companies take the illiquid asset return as a fixed constant in their pricing decision, we also show how the illiquid asset return moves in equilibrium with respect to exogenous shocks to liquidity.

Proposition 2 (time series of insurance premiums and illiquid asset returns).
Insurance companies set lower premiums when the expected excess returns on illiquid asset are higher

$$\frac{\partial P}{\partial R} < 0, \tag{1.15}$$

with increases in equilibrium illiquid asset returns resulting from

1. *an exogenous increase in transaction costs for the illiquid asset $\frac{\partial R}{\partial \lambda} > 0$; or*

2. *an exogenous increase in demand for liquidity from other investors $\frac{\partial R}{\partial \omega} > 0$.*

Proposition 2 allows us to make predictions for the average insurance premium price, which we expect to fluctuate over time in response to expected illiquid asset returns. When illiquid asset returns increase, either due to exogenous shocks to liquidity or exogenous shocks to liquidity demand from investors, insurers reduce premiums and increase funding. Note that this behaviour makes the insurer a counter-cyclical liquidity investor. When liquidity conditions deteriorate, insurers increase their balance sheet and illiquid asset holdings, dampening the impact of negative liquidity shocks on equilibrium returns.

We now consider the cross section of insurance premiums. We introduce a small insurer to the model, which we will denote with subscript i. We assume that they have mass zero, such that they do not affect equilibrium, and that the small insurer has less stable funding relative to competitors (i.e. $\sigma_i > \sigma$). We can see from equation (1.13) that this means $R_i^I < R^I$. The next proposition follows from this observation.

Proposition 3 (**cross section of insurance premiums and illiquid asset returns**). *For insurer i, with an expected investment return on illiquid investments lower than that of the industry average ($R_i^I < R^I$), the insurance premium will be set higher relative to competitors ($P_i > P$).*

Proposition 3 allows us to make predictions for the cross section of insurance premiums, which we expect to vary in relation to individual insurer expected investment returns relative to their competitors.

Numerical Example. We conclude the model by illustrating how insurers' stable funding, σ, and exogenous shocks to asset market liquidity, λ, affect insurance premiums by way of a numerical example. We choose parameters as follows: asset supply is $S = 1$, investors have $\omega = 0.2$ probability of being early consumers, insurance claims arrive at $t = 1$ with probability $\bar{\tau} = 0.5$, elasticity of insurance demand is $\epsilon = 15$, the fixed parameter in the demand function is $k = 1$, claims are $\bar{C} = 1$, and the insurer is endowed with equity capital $E = 0.25$.

In Figure 1.2, Panels A, we investigate how the expected return on the illiquid asset, R, depends on the transaction costs of selling the illiquid asset, λ. We show the solution for

three choices of funding stability of the insurer: $\sigma = 0.1$, $\sigma = 0.3$ and $\sigma = 0.5$. A lower σ means the insurer has more stable insurance funding. We see that the illiquid asset return increases as transaction costs increase in the secondary market. However, the sensitivity is less steep when insurer's funding is more stable and σ is lower.

In Panel B, we see that insurer's illiquid asset allocation also increases in λ, as the higher expected return encourages them to increase their exposure to the asset. The effect is stronger the more stable the insurer's funding is. The insurer's stable funding therefore makes them a counter-cyclical investor, increasing allocations when expected returns are higher. This feedback affects the equilibrium return, explaining why the return on the illiquid asset is less sensitive to λ when the insurer has more stable funding. The insurer absorbs more of the illiquid asset when liquidity conditions deteriorate, dampening the effect of liquidity on the equilibrium illiquid asset return.

Panel C shows that the insurance premium markup falls as λ increases. The insurer is able to extract more illiquid investment returns on their assets, and thus the marginal cost of underwriting the claim $\tilde{C}$ falls. In the case $\sigma = 0.5$, the insurer has no funding advantage, with $\bar{\tau} + \sigma = 1$ meaning they face the risk that all claims arrive at $t = 1$. The premium markup and insurer asset allocation are no longer dependent on λ, with our model nesting Modigliani and Miller (1958). The equilibrium return R is also now a linear function of λ, with no dampening impact of a counter-cyclical insurer allocation to the asset.

1.4   Data and Methodology

1.4.1   Measuring Insurance Prices

Life Insurance. To measure the price of life and term annuities we use the markups, which are defined as the percent deviation of the quoted price to the actuarial price. The actuarial price is defined as the expected claims discounted at the risk-free rate:

$$\text{Actuarial Price}_t = \sum_{k=1}^{T} \frac{E_t\left[C_{t+k}\right]}{\left(1 + R_{t+k}^f\right)^k} \qquad (1.16)$$

where C_{t+k} is the policy's claim k periods from its inception t, and R_{t+k}^f is the k-period risk-free rate at time t.

In addition to absolute markups, we also use annualised markups in our study. These

are the markup divided by the duration of the expected cash flows of the product. Following Koijen and Yogo (2015), we calculate expected cash flows and present values based on appropriate mortality table from the American Society of Actuaries and the zero-coupon Treasury curve Gürkaynak, Sack, and Wright (2007).

P&C Insurance. For most types of P&C contracts neither actual nor actuarially fair prices are readily available, making it impossible to calculate a markup. However, P&C insurers do track their pricing and underwriting performance through a measure called *combined ratio*, which is reported quarterly to the market. It is defined as:

$$\text{Combined Ratio} = \frac{\text{Losses} + \text{Expenses}}{\text{Premium Earned}} \tag{1.17}$$

where *losses* are the claims paid out on policies in the quarter (plus any significant revisions to future expected claims), *expenses* are the operating expenses of running the underwriting business and *premium earned* are the premium received on policies spread evenly over the life of the contracts. For example, if an insurer receives premium $P_{t,n}$ at time t on a policy that has a life of n quarterly reporting periods, then the reported *premium earned* on this contract in future reporting periods t' will be

$$\text{Premium Earned}_{t'} = \begin{cases} \dfrac{P_{t,n}}{n}, & \text{if } t < t' \leq t + n. \\ 0, & \text{otherwise.} \end{cases} \tag{1.18}$$

Premium earned is used in the *combined ratio* to ensure that realised claims are offset against the premiums that were received to cover their payment, and prevents the measure from being biased by changes in an insurers' underwriting volume. If an insurer doubles the size of their underwriting business, premiums received, $P_{t,n}$, double immediately while realised claims, at that time, are unaffected. Calculating the combined ratio with premiums received would therefore suggest a sudden improvement in underwriting (high inflows to outflows) even though the profitability of the underwriting business is unchanged. Premium earned, on other hand, increase in future periods, at the same time that claims are increasing due to the increased volume of business.

In our empirical analysis, we define underwriting profitability as:

$$\text{Underwriting Profitability}_t = \frac{\text{Premium Earned}_t - \text{Losses}_t - \text{Expenses}_t}{\text{Insurance Liabilities}_{t-1}} \tag{1.19}$$

which is the profit from underwriting divided by the size of the underwriting business. Insurance liabilities are reported by insurance companies and are the sum of "management's best estimate" of future losses and reinsurance payables (Odomirok et al., 2014).

An increase in an insurer's underwriting profit can either be created by higher premiums relative to expected claims, or realised claims that are lower than insurer expectations. The latter generates some noise in our measure of insurance premiums, but we assume the noise from claim risk is uncorrelated with investment returns for our empirical analysis.

Our theory states that the predictive variables for premiums should reflect expected investment returns at the time the policies are written, not when the earning from these policies are reported. In our regression analysis, we therefore use annual averages over the preceding 12 months, since the Property and Casualty insurance is usually short maturity contracts. For example, auto-mobile insurance policies (42% of the total P&C market) are typically standardised to have one year duration. We therefore only need expected returns over the previous four quarters for our regression analysis.

1.4.2 Data

Life Annuity Pricing. Koijen and Yogo (2015) collate data on annuity products prices from WebAnnuities Insurance Agency over the period 1989 to 2011. There is pricing for 3 types of annuities: term annuities (i.e. products that provide guaranteed income for a fixed term), life annuities (i.e. products that provide guaranteed income for an unfixed term that is dependent on survival) and guarantee annuities (i.e. products that provide guaranteed income for fixed term and then for future dates dependent on survival). The maturity of term annuities range from 5 to 30 years, whilst guarantees are of term 10 or 20 years. Further, for life and guarantee annuities, pricing is distinguished for males and females, and for ages 50 to 85 (with every five years in between). The time series consists of roughly semi-annual observations, except for the life annuities (with and without guarantees) which is also semi-annual, but with monthly observations during the years around the financial crisis, 2007-2009, which is the focus of Koijen and Yogo (2015). To summarize we have 96 insurers quoting prices on 1, or more, of 54 different annuity products at 73 different dates, which gives us 1380 company-date observations.

P&C Insurer Financial Statements. Insurance entities are required to report financial statements to regulatory authorities on a quarterly basis. S&P Global: Market Intelligence collates and provides this data. Our sample period is 2001 to 2018 for both Life Insurance and P&C Insurance companies.

In total, there are 3,951 individual P&C insurance entities in our sample. Large insur-

ance groups often have many separately regulated insurance entities under their overall company umbrella. We aggregate the entities up to their P&C insurance groups. For example, the two largest P&C insurance groups in our sample, State Farm and Berkshire Hathway, have been aggregated from 10 and 68 individual insurance entities respectively. To aggregate dollar financial variables we sum across entities. To aggregate percentages and ratios (such as investment yield) we use the asset-value weighted average.

Our final P&C sample consists of 1,070 insurance groups running P&C businesses over 68 quarters from March 2001 through to December 2017. In total we have 44,780 firm-quarter observations, with a minimum of 184 insurance groups available in any given quarter and a maximum of 735. To get to this final sample we have excluded insurance companies with less than 4 years of data, companies who never exceed $10 million in net total assets, company-year observations where the company has less than $1 million in earned premium over the year, and observations with non-positive net total assets and net premium earned. We do this to ensure that the companies we are looking at are relatively large and active. All financial statement variables are winsorized at the 5th and 95th percentiles in each quarterly reporting period.

The financial statements provides balance sheet and net income variables. For cross sectional analysis, our main variable is the accounting investment returns as described in Section 1.5. We also use their average credit portfolio rating[7], asset allocations and various measures of balance sheet strength: Size (log of total assets), Asset Growth (annual change in total assets), Leverage Ratio, Risk-Based Capital, Amount of Deferred Annuities (Life insurers only)[8], Unearned Premium to Earned Premium ratio[9] and reinsurance activity (net premiums reinsured / net premiums received). The last two are for P&C insurers only.

For cross sectional analysis on life insurance companies, we merge S&P Global financial statement data with the annuity markup data provided in Koijen and Yogo (2015). In the period 2000 to 2011, the intersection of our two datasets, we are able to merge both data with investment yields and annuity markups for 16 companies. Consistent with the P&C data construction, we have excluded insurance companies with less than 4 years of data.

---

[7]The insurance regulator assigns bonds into six broad categories (categories 1 through 6) based on their credit ratings, with higher categories reflecting higher credit risk. Level 1 is credit AAA-A, level 2 is BBB, level 3 BB, level 4 is B, level 5 CCC and level 6 is all other credit.

[8]these unprofitable products caused constraints in the financial crisis

[9]this gives an indication of the remaining unpaid liabilities relative to current volume of business

Financial Market and Macroeconomic Variables. To measure the credit spread we use Moody's Seasoned Baa corporate bond yield relative to 10-Year Treasury and retrieved from St. Louis Fed's website (fred.stlouisfed.org). We also use the excess bond risk premium portion of credit spreads as provided in (Gilchrist and Zakrajšek (2012)). Other right-hand side variables include the 6-Month to 10-Year Treasury Constant Maturity Rates and TED spread (downloaded from St. Louis Fed's website), to proxy for funding costs and the shadow cost of funding respectively. The TED spread is the difference between the three-month Treasury bill and the three-month LIBOR based in US dollars. The CAPE ratio, which is real earnings per share over a 10-year period, is retrieved from the Robert Shiller website.

Mergers and Acquisitions. We have hand collected data on mergers and acquisitions across our sample of life insurers with annuity pricing. The insurer net yields on invested assets around these assets are taken from our S&P Global: Market Intelligence dataset (where available) or directly from insurer financial reports on line. The list of events that we use in our analysis is shown in table (1.12).

1.4.3 Summary Statistics

Table 1.1 presents summary statistics for the key variables in our empirical analysis. The average annuity markup on an absolute basis is 6.75%, 5.31% and 4.24% for fixed term, life and guarantee annuities respectively. On an annualised basis, these markups are 1.03%, 1.12% and 0.50% respectively. Our main dependent variable in P&C markets is underwriting profitability, which across this sample has a mean of 0.31% and standard deviation of 3.24%. The average 5-year rolling standard deviations of underwriting profitability at an insurer-level is 2.35%. In our cross sectional analysis, the main independent variable is insurance companies investment return. This averages 2.75% in the P&C industry and 5.97% in our sub-sample of life insurers.

1.5 Preliminary Evidence

Before testing the model propositions in section 1.6, in this section we provide preliminary evidence that shows the importance of investment returns to the insurance business model.

Table 1.2 presents the aggregated industry balance sheets for the Life Insurance industry and P&C Insurance industry. There are two key takeaways that are relevant for our analysis. First, we see that the large asset portfolios are predominately funded by insurance underwriting. The Life Insurance industry has an average equity ratio of 9% and the P&C industry has an equity ratio of 38%, with the dominating source of leverage in both cases being insurance liabilities. Second, we see that insurance companies take lots of investment risk in their asset portfolios. Risk-free asset allocations (cash and Treasuries) are only 8% for the Life Insurance industry and 14% for the P&C industry. Instead, insurers invest in risky and often illiquid assets. Corporate bonds, mortgage loans and other credit (such as MBS, RMBS and municipal bonds) make up 75% and 42% of the balance sheets for the Life and P&C industries respectively.

Figure 1.3 next presents the P&C industry's aggregated net income. The total net income is split between the earnings reported from the asset portfolio investments, the earnings reported on the insurance underwriting business and (the residual) other income. The striking feature of Figure 1.3 is that the industry often loses money through insurance underwriting, and is only profitable once investment income is included. It should be noted that the underwriting losses shown in Panel A do not take time value of money into account. The industry standard for reporting on their underwriting is to ignore this. In Panel B, we adjust for this, increasing (decreasing) underwriting (investment) income by the value of insurance liabilities multiplied by the risk-free rate. Even after this adjustment, we see that returns on investment portfolios are of first order importance to the insurance business model.[10]

Figure 1.4 presents boxplots of insurers' investment returns in each reporting quarter of our sample, highlighting both the time series trends in insurer investment returns, and the rich heterogeneity in investment returns in the cross section of insurers. In any given quarter in our sample, the range between the 25th and 75th percentiles of investment returns is in excess of 150 bps. These investment returns are insurer's accounting investment returns, which are reported on a quarterly basis. For fixed income assets, the accounting treatment of investment returns is to report the yield at purchase amortised smoothly over the life of the bond. If the bond defaults or the insurer sells with a gain/loss, this

---

[10]Life industry insurance companies don't report underwriting profits is the same way as P&C insurers, so the equivalent analysis is not possible in this industry. Refer to Appendix 1.10.1 for a discussion of profitability in the Life Insurance Industry.

is also included in their investment return. However, so long as the insurer does not sell or the issuer does not default on the bond, the investment return methodology protects the insurer from mark-to-market volatility on their credit assets.[11] This treatment reflects insurers' long-term buy and hold approach to investing,[12] and is consistent with Chodorow-Reich, Ghent, and Haddad (2020) view of insurers as "asset insulators" that can ride out transitory dislocations in market prices. It is also consistent with our model of insurers being able to earn liquidity premium on illiquid investments.

Table 1.3 Panel A shows how variation in insurers' asset allocations explain cross sectional variation in insurer investment returns. We regress insurer investment returns (in bps) on asset allocations (in percent) with controls for time fixed effects. We see that insurers with large credit allocations have higher investment returns, while large allocations to treasuries and cash mean lower investment returns. For example, column 1 shows that a 1 percentage point increase in credit and cash allocations result in a 1.25 bps increase and 1.50 bps decrease in investment returns respectively. In column 2 of Table 1.3 we interact credit allocations with the credit portfolios value-weighted average credit rating.[13] We can see that the effect of credit allocations on investment returns is largely driven by the level of credit risk in these portfolios. Finally, in column 3 of Table 1.3, we interact credit rating interacted with credit allocation with the previous quarter's credit spread. The effect of credit portfolios on investment returns is larger when credit spreads are higher.

Table 1.3 Panel B explains the time series variation in individual insurance company's investment returns. Columns 1-2 show that there is a high degree of persistence in insurer investment returns, with an insurer's current quarter investment return explaining 37% of their next quarter investment return. Given insurer accounting returns predict next periods investment returns, we interpret cross sectional variation in this measure as cross sectional variation in insurer's *expected* investment return. The auto correlation of investment returns at an insurer level is not surprising given the accounting treatment of investment returns on fixed income assets.

Columns 3-4 of Table 1.3 Panel B show the macro-level time series drivers of invest-

---

[11]Refer to 1.10.2 for a more detailed description of how accounting investment returns are calculated by insurers.

[12]Schultz (2001) and John Y. Campbell (2003) estimate that insurers hold between 30% and 40% of corporate bonds and yet account for only about 12% of trading volume

[13]The insurance regulator, NAIC, assigns credit into six broad categories (level 1 through 6) based on their credit ratings, with higher categories reflecting higher credit risk. Level 1 bonds are rated AAA-A, level 2 is BBB, level 3 BB, level 4 is B, level 5 CCC and level 6 is all other credit

ment returns. We see that the large fixed income allocations in insurer portfolios make the risk-free rate, the slope of the yield curve and the credit spread on corporate bonds all very significant drivers of investment returns. On the other hand, the CAPE ratio (capturing expected equity returns) and the TED spread (capturing financial market distress) are unimportant. Our finding that credit spreads predict insurer investment returns is consistent with previous work that show that corporate bonds deliver excess returns to treasuries over the long-term (Krishnamurthy and Vissing-Jorgensen (2012), Gilchrist and Zakrajšek (2012)). In the long-term, the insurers accounting return on investments must equal their economic return. If credit spread only reflected default losses, then credit spreads would have no predictability for insurer investment returns on average.

1.6 Empirical Results

1.6.1 Stable Insurance Funding and Illiquid Asset Allocations

We first test Proposition 1's prediction for insurance companies asset allocation decision: insurers with more stable insurance funding hold more illiquid assets. We take this prediction to the data using P&C insurers' historical volatility on insurance underwriting as a proxy for stable funding. For each insurer, we calculate rolling 5-year volatility estimates of insurance underwriting profitability (as defined in equation (1.19)). We then use volatility lagged one quarter as the independent variable. Our two variables for capturing insurer investment risk is their cash allocation and their credit allocation multiplied by the average credit rating of this portfolio.[14] We report the results in columns 1-6 of Table 1.4 Panel C.

We see that stable funding predicts low cash allocations and large allocations to risky credit. For example, an insurer with underwriting profitability volatility 1 standard deviation higher than competitors has a 0.22 standard deviations (or by 3 percentage points) higher cash allocation compared to competitors. Following Ge and Weisbach (2020), we include firm size and other variables that capture insurers balance sheet strength as controls. Consistent with their work, we find strong evidence that the size of an insurer is a determinant in the amount of risk in an insurer's investment portfolios. Assuming large insurers have more diversified and stable underwriting businesses, this result is consistent with our model prediction. However, our results take this a step further, showing that

---

[14]We use a numeric measure of average credit rating, as assigned by the insurance regulator.

even when comparing firms of equal size, the insurer with less volatile underwriting performance takes more investment risk in their credit portfolio. This finding also holds after controlling for a vector of balance sheet strength variables.

Columns 7-9 of Table 1.4 show that large insurers and those with stable underwriting cash flows also realise higher investment returns. In other words, the increased investment risk translates into higher investment returns. To give a sense of the order of magnitude, an insurer with funding volatility one standard deviation lower than competitors has an investment return that is 21bps higher than it's competitors.

In summary, in this subsection we have documented a relationship between the stability of the funding generated by insurance underwriting and the asset allocation decisions of insurance companies. Insurance companies that are large and have more stable funding take more investment risk and earn higher investment returns. According to our model, the explanation is that insurers use the stability of the insurance funding to earn liquidity premium on their assets.

1.6.2   Investment Returns Drive the Time Series of Premiums

We next test Proposition 2's prediction for insurance prices and illiquid investment returns in the time series: high expected asset returns mean lower insurance premiums. We take this prediction to the data using credit spreads as a proxy for illiquid investment expected returns.

Figure 1.1 illustrates our central time series finding using our longest available sample. The figure presents the industry average markup on a 10 year fixed term annuity against the 10 year BAA credit spread from 1989 to 2011. Markups are defined as the quoted price relative to their actuarially fair price. The negative correlation between the markup (left hand axis) and credit spreads (right side axis, inverse) is obvious. In fact, the R-squared from the single variable regression of markups on credit spreads is as high as 77%.

We now show the relationship between annuity markups and credit spreads is present across different life products and sample periods, and robust to controls for other market returns and macroeconomic variables. Motivated by our theory, we focus on the impact of expected investment returns. We control for the global financial crisis using a dummy variable, as it was a period where financially constrained life insurers charged very low markups Koijen and Yogo (2015), which may confuscate our results.[15] We also control for

---

[15]Section 1.7 considers the impact of capital constraints within the context of asset-driven insurance

unemployment rate to proxy for shifts in the demand for insurance.

Table 1.5 reports the parameter estimates from the following regression:

$$m_{ikt} = \beta_c \cdot CS_t + \beta_{GFC} \cdot \mathbb{1}_{GFC} + \beta_{cGFC} \cdot CS_t \times \mathbb{1}_{GFC} + B' \cdot X_t + FE_i + FE_k + \epsilon_{ikt}$$

where m_{ikt} is the annualised markup set by insurer i at time t for an annuity which is in subproduct category k. Subproducts vary depending on age, sex and maturity of the annuities. CS_t is Moody's credit spread of BAA corporate bonds, and $\mathbb{1}_{GFC}$ is an indicator variable set to one over the global financial crisis (November 2008 through February 2010). We include a vector of time series controls, X_t, which includes the risk-free rate, the slope of the yield curve, the TED spread, the CAPE ratio (to capture other drivers of expected investment returns) and US unemployment rate (to capture time variation in the demand for insurance). We also include lagged markups in the control vector to control for potential autocorrelation in the dependent variable. Columns 1-3 report the parameter estimates from time series regressions where for the dependent variable, $\overline{m}_t$, we have averaged across insurers and subproduct categories in each time period. Columns 4-5 report full panel specifications. Panel A, B and C show the results for markups on life, guarantee and fixed-term annuity products respectively.

Across specifications, we see that a 100bps increase in credit spreads lowers annualised markups by 52bps (t-statistic of 5.34). Given that annualised markups are 1% on average, this means that markups fall by 50% when insurers can earn more on their credit portfolios.[16] The explanatory power is also very large. Taking life annuities as an example, the credit spread alone explains 80% of the variation in levels (see the adjusted r-squared in column 1 of Panel A). The main result of this section is also robust to including the vector, X_t, of time series controls. We report estimates for all variables in vector X_t in Appendix Table 1.13. Note that the risk-free rate is not significant as the effect of risk-free rates on premiums is captured in the actuarial price (equation 1.16), which is used in our dependent variable.[17]

---

pricing in detail.

[16]We use annualised markups (rather than absolute markups) so that it is easier to interpret coefficients across products with different durations. However, all results are qualitatively consistent to specifications with absolute markups.

[17]Table 1.15 in the appendix presents results from identical specifications as table 1.5, but with markups and investment returns in changes rather than levels. Our results are robust to this specification, with estimated sensitivities of similar magnitudes. We proceed with analysis in levels throughout the rest of the empirical results.

Koijen and Yogo (2015) highlight that the financial crisis saw a dramatic fall in markups from November 2008 through to February 2010. Figure 1.1 shows the annualised 10yr annuity markup fell from 1.25% to -0.75% across the dates. In Columns 3 and 5 we interact credit spreads with $\mathbb{1}_{GFC}$, which is an indicator variable set to one over the same period. The estimated coefficient on the interaction is positive, and generally we find it to be statistically significant. The positive interaction coefficient shows that the baseline coefficient is less negative in the financial crisis. Said differently, the negative relationship between premiums and credit spreads is stronger outside of the global financial crisis period. Nevertheless, our results suggests that credit spreads were still important in this period, with roughly 40% of this drop in markups due to sensitivity of markups to credit spreads. The remaining 60% was due other factors such as capital constraints.[18] We therefore argue that while capital constraints play an important role in insurance pricing, they are not the only factor. Instead, insurance companies also account for expected returns when setting prices, and this mechanism is especially important when insurance companies are unconstrained by regulatory capital requirements.

Table 1.6 shows how insurance premiums in the P&C industry vary with credit spreads. The table has the same five column specifications as the previously discussed Table 1.5. In the P&C industry we do not observe prices directly but instead use underwriting profitability (1.19) as the main dependent variable. This measure is the ratio of their underwriting profit relative to their insurance liabilities. We interpret lower underwriting profitability as lower prices. Given that underwriting profitability reflects insurance premium pricing over the previous year, we use lagged credit spreads on the right hand side of the regression. We find a statistically significant impact of credit spreads, with a 100bps increase in credit spreads lowering underwriting profitability by one percentage point. For a one standard deviation increase in credit spreads, the industry's underwriting profitability decreases by 1.3 standard deviations. Table 1.14 presents full specification results, including the control vector coefficients.

In summary, in this subsection we find an economically and statistically significant negative relationship between the time series of insurance premiums and the investment returns insurance companies expect to earn on their investment portfolios.

---

[18]Credit spreads and markups changed by 320bps and -200bps respectively. The credit spread coefficient, adjusting for the interaction coefficient, is $-0.59 + 0.36 = -0.23$ in the global financial crisis, and thus we see credit spreads account for $0.23 * 320 = 74$bps of the markup change.

1.6.3 Investment Returns Drive the Cross Section of Premiums

We next test Proposition 3's prediction for insurance prices and investment returns in the cross section of insurers: insurers with higher expected investment returns set relatively lower prices. As with the time series results, we begin with an illustration of our core finding. We use the P&C industry because it is our richest cross section, grouping the 1,240 insurers into 20 portfolios ranked on their investment return. For each portfolio, we then calculate equal weighted underwriting profitability and investment returns. Figure 1.5 presents a binned scatter graph of the portfolio averages with underwriting performance on the vertical axis and investment returns on the horizontal axis. There is a clear negative correlation with insurers with higher investment yields also reporting lower underwriting profitability.

We now formally test the relationship between insurance prices and the investment returns for both the Life Insurance industry and P&C industry, beginning with the Life Insurance. Table 1.7 reports the parameter estimate from the following panel regression using the cross section of life insurers:

$$m_{ikt} = \beta_y \cdot y_{it} + \beta_{yFC} \cdot y_{it} \times \mathbb{1}_{GFC} + B' \cdot X_{it-1} + FE_i + FE_k + FE_t + \epsilon_{ikt}$$

where m_{ikt} is the annualised markup set by insurer i at time t for an annuity which is in sub-product category k, y_{it} is the insurer's expected investment return, and X_{it} is a vector of lagged variables that have been shown to capture balance sheet strength (Koijen and Yogo (2015)). The control vector includes variables squared to capture any non-linear effects of capital constraints. We additionally control for date fixed effects, product fixed effects and firm fixed effects, and report within group R-squared. Panel A, B and C show the results for markups on fixed-term, guarantee and life annuity products respectively. Columns 4-5 interact investment return with an indicator variable $\mathbb{1}_{GFC}$ set equal to one during the financial crisis. Across specifications and products, we see that an insurer with a investment return 100bps higher than competitors sets annualised markups 3bps lower. In the majority of specifications the relationship is statistically significant.

Table 1.8 tests the cross sectional relationship between insurance pricing and insurers' expected investment returns in the P&C industry. The table follows the same structure as Table 1.7, but with insurer underwriting profitability replacing markups as the dependent variable. We also include a variable that controls for the level of reinsurance activity by insurance companies. In the P&C industry, we find that an insurance company with a

100bps higher expected investment return compared to competitors reports underwriting profitability that is 10bps lower. To compare the cross sectional results in both the life industry and P&C industry, we have also calculated standardized coefficients. In the life insurance industry, a one standard deviation higher insurer investment return reduces an insurer's relative markup by 0.05 standard deviations. In the P&C industry, we find an insurer with a one standard deviation higher insurer investment return has an underwriting profitability ratio 0.03 standard deviations lower than competitors.

Table 1.9 extends the P&C cross sectional analysis with an instrumental variable estimation. We use the analysis from Section 1.6.1 as the first stage, instrumenting the expected investment returns of insurance companies by their underwriting volatility (Table 1.9 Column 1) and both underwriting volatility and firm size (Table 1.9 Column 2). As previously discussed, the stability of insurance underwriting allows insurers to take more investment risk and earn higher investment returns. Columns 3 and 4 report the parameter estimate from regressing instrumented investment returns on insurance premiums. We see that, with full controls, the insurance premiums fall by 0.27 percentage points when investment returns are 100bps higher. The coefficients in Table 1.9 are more negative than those estimated in Table 1.8. This is likely because the instrumented investment returns provide a cleaner estimate of the impact of investment strategy on insurance premiums.

Table 1.9 reports the Cragg-Donald Wald F-statistic, and in the case where we have two instrumental variables (Column 4), we report the p-value from the Sargan's χ^2 test of overidentifying restrictions. The large F-statistics indicate that our instruments do not have weak instrument concerns. Running the specification with controls but only volatility as the instrument results in weak instrument concerns Stock and Yogo (2005). We therefore add firm size as an additional instrument. As discussed in Section 1.6.1, we view this variable as another proxy for the stability of insurance underwriting fund that insurers enjoy. The large p-value in the Sargan test indicates that the instruments are also uncorrelated with the structural error term.

In summary, in this subsection we have shown that the negative time series relationship between insurance premiums and expected investment returns is also present in the cross section of insurance companies. Insurance companies that expect to earn higher returns on their investment portfolios set lower premiums relative to their competitors.

1.6.4 Evidence from Mergers and Acquisitions

In this section, we present evidence on how changes in investment returns due to merger and acquisition affect insurance premiums. We argue that these exogenous shocks to the insurance companies allow us to extract a cleaner estimates of the cross sectional relationship between premiums and investment returns. Figure (1.6) presents a representative example from our sample. American Heritage was acquired by AllState Insurance in October 1999. In the 12 months proceeding the acquisition, American Heritage earned a return of 7.22% on their investment portfolio and AllState Insurance earned 5.80%. The figure shows that American Heritage's investment returns fell post acquisition, reflecting the more defensive strategy of their acquirer. Critically, the figure also shows an adjustment in pricing on 10yr fixed term annuities. American Heritage were consistently selling annuities at a discount to the industry pre-acquisition. However, following the acquisition, their markup pricing increased significantly.

We next show evidence consistent with the case study but with multiple merger events in Table 1.10. We have five merger events in our sample, and study the premium impact on three products: 20yr fixed term annuity, life annuity for males aged 50, and 10 year guarantee life annuity for a male aged 50. In a difference-in-differences approach, we use life insurance companies involved in a merger and acquisition event as our treatment group, and other insurance companies as the control group. The treatment period is the two years after the merger event, and the control period is the two years before the merger event. Table 1.10 reports the parameter estimate from the following regression:

$$m_{ikt} = \beta_D \cdot D_{it} + FE_i + FE_k + FE_t + \epsilon_{ijt}$$

where m_{ikt} is the markup set by insurer i at time t on product k. Our explanatory variable, D_{it}, is the investment return differential between the treatment group insurance company and the other insurance company involved in the transaction. It is set equal to this value for the treatment insurer and treatment time period (i.e. in the two years following the merger for the treatment insurer), and set to zero in all other cases (i.e. two years pre merger event for the treatment group, and in all observations for control group insurers). The interpretation of a positive investment return differential is that insurer i is being acquired by an insurer with a more risky investment strategy, and thus going forward their own investment returns are expected to be higher. In each of our observations, we confirm that investment return differential do indeed lead to a change in the insurers

investment returns post transaction in-line with this interpretation. This is illustrated in Figure 1.6 with the American Heritage example.

By controlling for time and firm fixed effects in the regression, the coefficient β_D captures the impact of the merger induced change in investment return on the treatment insurer's relative markup pricing as compared to the industry average. We see from Table 1.10 that a 1% increase in investment return following merger activity results in a 0.22% fall in an insurers markup relative to the industry. The t-statistic is 3.44. The coefficient is larger than in Table 1.8, suggesting the merger sample is better able to identify the relationship between investment returns and insurance premiums. We also note that the sample includes examples of where the investment return differential is both negative and positive. This helps rule out competing interpretations of the results. For example, one could imagine insurers discount products ahead of a merger to increase the value of the merger, which would lead to an increase in markups post merger. However, this can't explain the observations in the sample with an increase in investment returns and fall in markups.

In summary, in this subsection we extend our analysis to show the negative relationship between insurance premiums and expected investment returns in the cross section of insurance companies holds following exogenous shocks to returns and premiums that are due to merger activity.

1.6.5   Evidence from Excess Bond Returns

Credit spreads can be split into spread that compensates investors for expected default losses and a premium in excess of this. It is the latter component, the excess return, that our model predicts is driving the correlation between credit spreads insurance premiums. Insurance companies use their stable insurance funding to extract liquidity premium on their asset portfolios, with some of the excess return passed on to policyholders through lower premiums.

We test this interpretation in Table 1.11 by re-estimating the regression specifications in Table 1.5, but splitting credit spreads between excess bond premium and the fair credit spread given the underlying default risk (Gilchrist and Zakrajšek (2012)). As per our previous analysis, we run specifications with time series averages and the full panel of insurers, as well as specifications with / without an interaction with the financial crisis period. We see that negative correlation between premiums and credit spreads is driven

entirely by the excess bond return component of credit spreads, with 100bps increase in excess bond returns reducing the markup by 50bps depending on specification and product. The default risk component of credit is statistically significant only in the panel C (fixed-term annuities). The coefficient on the excess bond return suggests that insurance companies pass back 50bps of excess returns on their credit portfolios to policyholders, and maintain 50bps for equity holders.

In summary, in this subsection we extend our analysis to show the time series correlation between credit spreads and insurance premiums are driven by the excess bond return component of credit spreads. This finding is strong evidence in support of asset-driven insurance pricing.

1.7 Introducing Insurer Capital Constraints

1.7.1 Theoretical Background

Capital constraints also affect insurance premiums (Gron (1994), Froot and O'Connell (1999), Koijen and Yogo (2015) andGe (2020). We embed this additional premium pricing mechanism into our existing framework by subjecting the insurer to a statutory capital constraint. The statutory value of each insurance policy is

$$\bar{V} = \frac{\bar{C}}{1 + R^S} \tag{1.20}$$

where R^S is the statutory discount rate for claims. The total statutory value of all Q claims is therefore $V = Q\bar{V}$. In the spirit of Koijen and Yogo (2015), the insurance company faces a capital constraint

$$\frac{V}{L} \leq \phi \tag{1.21}$$

where $\phi \leq 1$ is the maximum statutory leverage ratio and L is their total liabilities (equation 1.5). The likelihood of this constraint binding is decreasing in the statutory discount rate R^S. A higher discount rate reduces the statutory value of each policy and therefore reduces statutory leverage.

The first-order condition of equation (1.8) with respect to P when the insurer is subject to (1.21) yields the following proposition.

Proposition 4 (insurance premium with capital constraints). *In equilibrium, a*

policy with claim $\bar{C}$ will be underwritten with premium

$$\hat{P} = \frac{\bar{C}}{1+R^I}\left(\frac{\varepsilon}{\varepsilon-1}\right)\left(\frac{1+(\bar{\tau}+\sigma)R+\frac{\eta}{\phi(1+R^S)}}{1+(\bar{\tau}+\sigma)R+\frac{\eta}{(1+R^I)}}\right). \tag{1.22}$$

where $\eta \geq 0$ be the Lagrange multiplier on the capital constraint (1.21).

Note that when the capital constraint is not binding, then $\eta = 0$ and therefore $\hat{P} = P$ as defined in (1.12). However, our interest in this section is for the case $\eta > 0$, which we explore in detail below.

Proposition 5 (**capital constraints vs. no capital constraints**). *When an insurer is capital constrained so eq. (1.21) holds with equality, the optimal price $\hat{P}$ relative to optimal price in the unconstrained case P depends on the relationship between the insurers time value of money R^I, the statutory discounting of claims R^S, and the maximum statutory leverage ratio ϕ. In particular:*

(i) When $(1+R^I) < \phi(1+R^S)$ then $\hat{P} < P$

(ii) When $(1+R^I) > \phi(1+R^S)$ then $\hat{P} > P$

(iii) When $(1+R^I) = \phi(1+R^S)$ then $\hat{P} = P$

This three case proposition extends the main theoretical result of Koijen and Yogo (2015), showing that the impact of insurer investment returns R^I is also important when the regulatory constraint binds. We describe the economic mechanisms below.

Case 1: $(1+R^I) < \phi(1+R^S)$**.** In this case the discount rate applied to statutory liabilities is higher than the expected return on assets multiplied by a factor of $\phi^{-1} > 1$. A new policy increases liabilities by $\bar{V}\phi^{-1}$ and increases assets by the premium received P. A higher R^S reduces $\bar{V}$ through the statutory discounting, and if R^S is sufficiently high it can mean new policies create an instantaneous improvement in an insurers statutory capital position. The result is that constrained insurers write policies at cheaper prices than an unconstrained competitor. Although writing policies cheaper reduces final wealth, insurer do it due to the temporary statutory capital relief it creates. Koijen and Yogo (2015) provide a detailed description of the calculation of R^S for different products in the life industry, showing that it was particularly high in the financial crisis. Consistent with their model prediction, they find constrained life insurers reduced annuity and guarantee markups significantly during the financial crisis.

Case 2: $\left(1 + R^I\right) > \phi\left(1 + R^S\right).$ In this case capital constraints lead to an increase in insurance prices. If the insurer sets the unconstrained premium price P, a new policy creates more statutory liabilities than assets, as $\tilde{V}\phi^{-1} > P$. Constrained insurers are therefore forced to increase prices to a level such that the premium received offsets the increase in liabilities. Froot and O'Connell (1999) provide an example of such a case by documenting how supply of catastrophe insurance fell following a negative shock to insurers' capital.

Case 3: $\left(1 + R^I\right) = \phi\left(1 + R^S\right).$ This is a special case where the mechanisms underlying case 1 and 2 offset each other. It means that a binding capital constraint has no impact on an insurer's optimal premium.

Our main time series empirical implementation uses credit spreads, which are likely to be positively correlated with capital constraints, to proxy R^I. Proposition 2 predicts lower premiums when credit spreads (expected returns) increase. At the same time, proposition 5 case 1 predicts lower premiums with higher credit spreads (assuming higher credit spreads mean more financial constraints and lower insurance capital). The predicted impact of capital constraints on premiums is therefore the same as asset-driven insurance pricing, which makes it hard to empirically separate the two channels. However, in case 2 of proposition 5, the sign of the effect of capital constraints is reversed. This means that the asset-driven insurance pricing and capital constraint effects move in opposite directions.

1.7.2 Controlling for Capital Constraints Empirically

The financial crisis was a period of particularly high capital constraints in the Life Insurance industry (Koijen and Yogo (2015)). We have therefore been careful to separate out the financial crisis in all of our previously discussed results. We show that our findings are robust across periods and apply in *normal* times only. In fact, in most specifications, we find the negative relation between insurance prices and investment returns is less strong in the financial crisis. Said differently, the asset-driven insurance pricing effect holds stronger in normal times where capital constraints are less prevalent. To see this, note the coefficient on credit spreads interacted with the financial crisis indicator is positive and statistically significant. For example, in Table 1.5 Panel A, we find a coefficient of 0.31 (t-statistic 2.85).

Proposition 5 highlights that the impact of capital constraints on the insurance pre-

miums depends on the level of statutory discount rates relative to expected investment returns. In the second case of the proposition 5, capital constraints predict higher premiums in times of stress, while asset-driven pricing predicts premiums are lower when credit spreads are higher. Empirical settings where insurers are in case two therefore makes it easier to disentangling capital constraints and asset-driven pricing empirically. For P&C markets, liabilities are not discounted ($R^S = 0$) for typical products such as car insurance, with the regulator making no adjustment for the premium's time-value of money (NAIC (2018)). This regulatory feature of the industry means case two always applies in this market. Our time series empirical results in the P&C industry, as documented in Table 1.6, therefore help to identify asset-driven insurance pricing while controlling for the potential impact of capital constraints.

In the cross sectional analysis, the result that insurer-specific asset portfolios affects relative insurance pricing across insurers is evidence that insurer investment portfolios matter for insurance pricing. However, it is possible that insurers with higher investment returns are also financially constrained and *gambling on resurrection*. To control for this potentially confounding factor, we include standard controls for insurer capital constraints (i.e. leverage, risk-cased statutory capital, asset growth). The results are once again robust.

1.8 Alternative Mechanisms

1.8.1 Insurer Default Risk

An alternative interpretation of our cross sectional results is that the insurers taking increased investment risk have higher probability of default, and thus face less demand for the insurance contracts they offer. In respect to this possible channel of insurance pricing, it is important to first note that the insurance industry is tightly regulated from a capital standpoint, with the key purpose of minimising the risk of insurer defaults on policyholders. Insurers are regulated on a risk based capital measure, and have to hold more capital when taking increased risk (including in their investment portfolios). In fact, the measure of investment portfolio credit risk that we use on the right hand side in Table 1.3 Panel A is the variable used by regulators when assessing how much capital insurers must hold for their credit portfolio investment risk. This means an increase in investment returns is also associated with an increase in the regulatory capital buffer an insurer must

hold. All else equal, this should reduce the probability of default.

Finally, A.M. Best provide all insurers with a financial strength rating that ranges from A++ to C-. A lower rating would signify a higher probability of default. The life insurance data we have, taken from from Koijen and Yogo (2015), is for the subset of insurers with an A rating. The fact that we see a sensitivity between investment returns and insurance premiums *within* this group is further evidence of asset-driven insurance pricing at play, controlling for default risk. Further, our cross sectional specifications control for financial variables that demonstrate balance sheet strength. These should also absorb the impact of insurer default risk.

1.8.2 Reinsurance

Insurance companies use reinsurance markets to hedge or remove some of the underlying risk on the contracts they write. The level of reinsurance activity could therefore be expected to affect profitability of insurance underwriting. To rule out this alternative hypothesis as a driver of our results in Table 1.8, we include the fraction of underwriting premiums which are reinsured as a control variable. We find that while premiums are significantly lower when an insurer's reinsurance activity is higher (Table 1.16), our main result that insurance premiums are lower when investment return are higher is still robust to the inclusion of this variable. The negative effect of reinsurance on premiums suggest that insurance companies that hedge more of the risks on their liabilities through reinsurance are able to charge lower premiums.[19]

1.9 Conclusion

Asset-driven insurance pricing is a new channel of insurance pricing, which shows that insurance premiums are lower when insurance companies have higher expected investment returns. In a violation of the Modigliani and Miller (1958) capital irrelevance theorem, the pricing of insurer liabilities depends on the expected returns on their asset portfolios. Specifically, insurance companies use the stable nature of insurance funding to take advantage of liquidity premium in illiquid asset markets. When expected returns are higher, insurers compete for funding, and insurance premiums fall.

A recent directive in Solvency II insurance regulation[20] means life insurers can now ap-

---

[19]We thank Stefano Rossi for this observation.

[20]see Solvency II, art. 77b and 77c

ply for a *matching adjustment* on some products, which allows them to apply to discount liabilities with the expected return on assets:

"The matching adjustment is an adjustment made to the risk-free interest rate when the insurer sets aside a portfolio of assets to back a predictable portion of their liabilities. It is based on the yield spread over the risk-free rate credit spread of the assigned portfolio of matching assets, minus a fundamental spread that accounts for expected default and downgrade risk. It is designed to reflect the fact that long-term, buy-and-hold investors only bear downgrade and default risks as they seek to hold assets to maturity, and allows them to capture other aspects of the spread such as the liquidity premium" – The Actuary [21]

The matching adjustment directive shows that insurers also think about their funding and investing in a similar manner to the arguments put forth in this paper.

We conclude by noting that asset-driven insurance pricing has two potential welfare implications. Firstly, insurers act as pro-cyclical investors, increasing asset allocations to illiquid investments when liquidity premium are higher, dampening asset market volatility. Second, insurers provide households with cheaper access to insurance when financial markets are distressed. These interesting macroeconomic implications of our findings offer interesting avenues for future research.

---

[21]https://www.theactuary.com/features/2016/06/2016/05/23/matching-adjustment-fit?fbclid=IwAR1GqbTH3ZrG5zaxWJz34YJNMhoip054u-IBRxsHFBda5EwePlmvfNm69tc

Figure 1.1: Expected investment returns drive the time series of insurance premiums. This figure shows the relation between insurance premiums and insurer expected investment returns as proxied by credit spreads. Panel A plots the two time series in levels. Panel B plots a scatter plot of the two time series in changes. Insurance premiums are measured as the percent deviation of the quoted price from actuarially fair value. We use the industry average 10 year fixed term annuity markup of Koijen and Yogo (2015). The credit spread variable is Moody's BAA 10-year corporate bonds yield over 10-year treasury yield (fred.stlouisfed.org).

(a) Time-Series Graph (Levels)

(b) Scatter Plot (Changes)

51

Figure 1.2: Model predictions. This figure presents numerical solutions of the model with the parameters: asset supply $S = 1$, investors have $\omega = 0.2$ probability of being early consumers, insurance claims arrive at $t = 1$ with probability $\bar{\tau} = 0.5$, elasticity of insurance demand is $\epsilon = 15$, the fixed parameter in the demand function is $k = 1$, claims are $\tilde{C} = 1$, and the insurer is endowed with equity capital $E = 0.25$. Panel A, B and C plot the expected return on the illiquid asset, R, the insurance company's share in illiquid asset, Θ/S, and the premium markup relative to the expected claim, $P/\bar{C}-1$, respectively. In each panel the variable is plotted as a function of of the asset market illiquidity, λ, with three choices of funding stability, σ.

(a) Return on the Illiquid Asset

(b) Insurer Illiquid Asset Allocation

(c) Premium Relative to Actuarial Fair Price

52

Figure 1.3: Investment income drives total net income. This figure plots the P&C industry's aggregate net income split between the main contributing sources. The three components are earnings generated from i) insurance underwriting, ii) investment portfolios, iii) other. Together they constitute the total net income of the industry. In Panel A, the profits on insurance underwriting are the premiums earned minus losses and expenses. As per the industry reporting standard, it does not include any adjustment for the time-value of money of underwriting. In Panel B, we increase (decrease) underwriting (investment) income by the value of insurance liabilities multiplied by the risk-free rate. The data comes from US insurance company statutory filings and is provided by SNL Global. Individual company data has been aggregated to show the industry-wide net income.

(a) Net Income as reported by insurance companies

(b) Adjusting for the time-value of money of underwriting funding

Figure 1.4: Variation in the expected investment returns of insurance companies. This figure illustrates variation in the expected investment returns of insurance companies in both the time series and cross section. In each reporting quarter of our sample, the figure presents a boxplot of expected investment returns. Our sample includes firm-level data for 1,104 P&C insurers in total. Expected investment returns are measured as the net yield on invested assets, as reported in insurance company financial accounts. The data comes from US insurance company statutory filings and is provided by SNL Global.

Figure 1.5: Expected investment returns drive the cross section of insurance premiums. This figure presents a binned scatter plot of insurer's insurance premiums against their expected investment returns. Insurance companies have been grouped into 20 equal sized portfolios based on the ranking of their investment portfolio returns. The figure plots each portfolio's average premium against its average investment return. Insurance premiums are measured as the ratio of an insurer company's insurance underwriting profit to their insurance liabilities. The sample includes firm-level data for 1,104 Property & Casualty (P&C) insurers over the period Q1 2001 to Q4 2017, with a total of 44,780 observations. The data is reported in US insurance company statutory filings and is provided by SNL Global.

Figure 1.6: Mergers & acquisitions evidence - american heritage acquisition case study. This figure plots American Heritage's excess markup on a 10yr annuity and their investment portfolio return. The sample period is 1995/2001. On October 1999 American Heritage was acquired by AllState Insurance. The acquisition is denoted by vertical line in the figure. A markup m_{ikt} for insurer i at time t on product k is the percentage deviation of the insurer's quoted price relative to the actuarial fair price. The excess markup $m_{ikt}^{ex} = m_{ikt} - \bar{m}_{kt}$ is the insurer's markup minus the industry average markup at time t on product k. The investment return is the investment portfolio income over the total value of invested assets. Markup data is provided by Koijen and Yogo (2015) and investment returns are collected this from insurer financial statements.

Table 1.1: Summary statistics

This table presents summary statistics of the variables used in the empirical analysis. The markups on life insurance are available biannually from 1989 through 2011 (Koijen and Yogo (2015)). Financial variables (for both P&C and Life insurance) are available quarterly from March 2001 through December 2017. The financial market and macroeconomic variables are available at monthly frequencies and have been collected from various sources.

	Count	Mean	SD	p05	p25	p50	p75	p95
Annuity Markups								
Life	19,923	6.75	7.07	-24.49	2.45	7.12	11.37	32.34
Life (ann.)	19,923	1.03	0.98	-1.92	0.38	0.96	1.62	4.36
Term	2,927	5.31	5.00	-17.32	2.65	5.79	8.41	32.64
Term (ann.)	2,927	1.12	1.06	-1.73	0.37	0.99	1.81	5.55
Guarantee	10,221	4.24	6.43	-24.70	0.41	4.94	8.34	32.35
Guarantee (ann.)	10,221	0.50	0.68	-2.00	0.05	0.52	0.94	2.93
Property& Casualty Financial Variables								
Underwriting Profitability	44,780	0.31	3.24	-5.09	-1.27	0.14	1.70	6.23
Underwriting Profits Volatility	27,787	2.35	1.34	0.58	1.25	2.17	3.23	4.85
Investment Return	44,780	3.08	1.29	0.95	2.13	3.08	3.97	5.22
Credit Allocation	44,780	54.09	22.40	13.17	37.68	57.98	72.58	84.68
Credit Risk	44,780	1.72	0.97	1.04	1.19	1.38	1.81	3.77
Cash Allocation	44,780	13.59	13.38	1.26	4.29	8.65	17.78	46.58
Treasuries Allocation	44,780	15.98	15.02	0.22	4.33	11.25	23.70	48.59
Stocks Allocation	44,780	11.57	11.38	0.00	1.34	8.72	17.98	36.31
Other Allocation	44,780	3.77	4.90	0.00	0.00	1.76	5.79	14.78
Size (t-1)	41,589	4.92	1.87	2.40	3.33	4.63	6.19	8.53
Asset Growth (t-1)	37,044	6.32	20.61	-11.78	0.00	5.63	11.78	27.29
Leverage (t-1)	41,589	42.54	14.44	21.17	31.62	40.51	52.24	70.58
Risk Based Capital (t-1)	41,589	4.74	2.95	1.32	2.56	3.96	6.03	11.75
Unearned Premia (t-1)	41,589	1.94	0.84	0.36	1.50	1.97	2.31	3.56
Reinsurance Activity (t-1)	41,589	0.13	0.40	-0.73	0.00	0.13	0.33	0.76
Life Financial Variables								
Investment Return	258	5.97	1.68	4.15	5.19	5.62	6.42	8.49
Size	258	16.36	1.12	14.69	15.36	16.38	17.36	18.10
Asset Growth	258	8.30	12.86	-7.99	0.11	7.34	12.91	30.98
Leverage	258	90.86	4.22	83.00	88.19	91.35	93.99	96.97
Risk Based Capital	258	14.60	45.86	-39.00	-24.00	2.00	50.00	102.00
Deferred Annuities	258	11.03	14.24	0.49	1.77	5.87	14.34	45.41
Financial Market and Macroeconomic Variables								
Credit Spread (BAA)	403	2.33	0.72	1.29	1.77	2.20	2.76	6.01
Risk Free (1yr)	469	4.65	3.73	0.10	1.30	4.63	6.64	16.72
Risk Free (5yr)	469	5.54	3.52	0.62	2.54	5.09	7.71	15.93
Slope (5yr - 1yr)	469	0.89	0.74	-1.63	0.38	0.87	1.46	2.50
TED Spread	403	0.57	0.42	0.12	0.26	0.46	0.73	3.35
Excess Bond Risk Premia	434	0.06	0.55	-1.14	-0.31	-0.04	0.28	3.00
US Unemployment Rate	469	6.22	1.68	3.60	5.00	5.70	7.30	10.80
CAPE ratio	469	22.35	8.43	6.64	16.43	22.42	26.79	44.20

Table 1.2: Insurance funding is invested in illiquid credit assets

This table shows the aggregated balance sheets of the Life Insurance industry and the P&C Insurance Industry as of December 2017. The assets are split by the largest investment allocations, and the liabilities are split into insurance liabilities and other liabilities. The shaded regions highlight two important observations: a) there is a significant amount of credit and liquidity risk taken in insurer asset portfolios, and b) the asset portfolios are predominantly funded by insurance liabilities. The data comes from US insurance company statutory filings and is provided by SNL Global. Individual company data has been aggregated to show the industry-wide balance sheet.

	Life Insurance ($bn)	Property and Casualty ($bn)	Life Insurance (%)	Property and Casualty (%)
Total Assets	**4301**	**1998**	**100%**	**100%**
Cash & Short Term Investments	105	116	2%	6%
Bonds - US Government	235	162	5%	8%
Bonds – Corporate	2199	414	51%	21%
Bonds – Other Credit	539	404	13%	20%
Mortgage Loans	477	17	11%	1%
Stocks	105	415	2%	21%
Other Investments	414	163	9%	9%
Total Cash & Investments	4075	1691	95%	85%
None-Financial Assets	227	306	5%	15%
Total Liabilities	**4301**	**1998**	**100%**	**100%**
Insurance Liabilities	3294	1021	77%	51%
Other Liabilities	615	211	14%	11%
Capital And Surplus (Equity)	393	765	9%	38%

Table 1.3: Understanding the investment returns of insurance companies

This table explains variation in the investment returns of insurance companies. Panel A reports the parameter estimate from the following panel regression:

$$y_{it} = \mathrm{B}'W_{it} + \beta_r \cdot risk_{it} + \beta_{w^c r} \cdot w_{it}^{credit} \times risk_{it} + \beta_{w^c rCS} \cdot w_{it}^{credit} \times risk_{it} \times CS_{t-1} + FE_t + \epsilon_{it}$$

where y_{it} is insurer i's investment return at time t and W_{it} is a vector of asset allocations including the allocation to credit, w_{it}^{credit}. We also include a numeric measure of the credit risk in the insurer's credit portfolio, $risk_{it}$, and the previous period credit spread, CS_{t-1}. All specifications in Panel A include time fixed effects FE_t. Investment returns are measured in bps, asset allocations are in percent, and the measure of credit risk range from 1-6 (and are as assigned by the insurance regulator).

Panel B reports the parameter estimate from the following panel regression:

$$y_{it} = \beta_y \cdot y_{i,t-k} + \mathrm{B}' \cdot X_t + FE_i + \epsilon_{it}$$

where $y_{i,t-k}$ is lagged insurer returns, X_t is a vector of time series variables that capture insurer investment opportunities or macroeconomic conditions, and FE_i captures firm fixed effects. All variables in panel B are measured in percent. The sample consists of quarterly observations from March 2001 through March 2018. t-statistics are reported in the brackets and are calculated using standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

Panel A: Investment Returns: Asset Allocation and Credit Portfolio Risk

	Investment Return (bps)		
	(1)	(2)	(3)
Credit Allocation	1.25***	0.54*	0.50*
	(11.24)	(1.90)	(1.81)
Cash Allocation	-1.50***	-1.42***	-1.41***
	(-7.89)	(-6.02)	(-5.96)
Credit Risk		14.62***	14.12***
		(4.99)	(4.90)
Credit Allocation × Credit Risk		0.93***	-0.11
		(5.61)	(-0.42)
Treasuries Allocation		-0.99***	-1.00***
		(-3.06)	(-3.16)
Stocks Allocation		-0.19	-0.20
		(-0.61)	(-0.66)
Other Allocation		-0.81*	-0.74*
		(-1.94)	(-1.81)
Credit Allocation × Credit Risk × Credit Spread (t-1)			0.40***
			(4.50)
Date FE	yes	yes	yes
Adj R-sq (Within)	0.168	0.202	0.207
Observations	44,780	44,780	44,780

Panel B: Investment Returns: Persistence and Time Series Variation

	Investment Return (it)			
	(1)	(2)	(3)	(4)
Investment Return (i,t-1)	0.61***	0.47***		
	(17.46)	(16.63)		
Investment Return (i,t-5)		0.19***		
		(9.21)		
Credit Spread (t-1)			0.39***	0.25*
			(7.14)	(1.85)
Risk-free Rate (t-1)			0.52***	0.51***
			(17.36)	(13.68)
Slope (t-1)			0.45***	0.48***
			(6.19)	(6.53)
TED (t-1)				0.10
				(0.61)
CAPE (t-1)				-0.03
				(-1.31)
Firm FE	yes	yes	yes	yes
Adj R-sq (Within)	0.371	0.395	0.341	0.346
Observations	37,044	37,044	37,044	37,044

Table 1.4: Insurers with stable funding take more investment risk

This table shows the relation between insurer's investment allocation and their insurance funding. The table reports the standardized parameter estimates from the following panel regression:

$$y_{it} = \beta_{vol} \cdot Volatility_{i,t-1} + \beta_{Size} \cdot Size_{i,t-1} + B' \cdot X_{i,t-1} + FE_t + \epsilon_{it}$$

where y_{it} is either insurer i's cash allocation at time t (columns 1-3), insurer i's credit asset allocation at time t multiplied by a numeric measure of the credit risk in these portfolios at time t (columns 4-6), or insurer i's investment return at time t (columns 7-9). Independent variables include, the historical 5-year volatility of insurer i's underwriting profitability up to time and including time $t-1$, $Volatility_{i,t-1}$, the insurers size (log assets), and a vector of other balance sheet measures, X_{it}, that capture balance sheet strength. All specifications include time fixed effects FE_t. Asset allocations and funding volatility are measured in percentage and investment returns are measured in bps. Credit risk is insurer i's credit portfolio value-weighted average credit rating, with bonds assigned a number from 1-6 dependent on their credit risk (as assigned by the insurance regulator, NAIC). The sample consists of quarterly observations from March 2001 through December 2017. t-statistics are reported in the brackets and are calculated using standard errors clustered by firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	Cash Allocation (perc.)			Credit Assets $\times$ Risk			Investment Return (bps)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Underwriting Volatility (i,t-1)	0.22***		0.06*	-0.26***		-0.09***	-0.16***		-0.06***
	(7.56)		(1.79)	(-8.91)		(-2.97)	(-7.18)		(-2.66)
Size (t-1)		-0.34***	-0.30***		0.36***	0.31***		0.21***	0.17***
		(-12.57)	(-9.40)		(11.61)	(8.95)		(9.80)	(7.22)
Reinsurance Activity (t-1)			0.08***			-0.00			-0.04**
			(2.75)			(-0.13)			(-2.25)
Risk Based Capital (t-1)			-0.14***			0.03			0.05**
			(-5.28)			(1.07)			(2.43)
Asset Growth (t-1)			0.05***			-0.03***			-0.02**
			(4.63)			(-2.71)			(-2.29)
Unearned Premia (t-1)			-0.05*			0.01			-0.00
			(-1.76)			(0.47)			(-0.11)
Date FE	yes	yes	yes	yes	yes	yes	yes	yes	yes
Adj R-sq (Within)	0.048	0.114	0.146	0.067	0.128	0.135	0.036	0.065	0.075
Observations	25,091	25,091	25,091	25,091	25,091	25,091	25,091	25,091	25,091

Table 1.5: Investment returns drive the time series of premiums: life insurance

This table shows the time series relation between insurance premiums, as measured by the markups on annuities issued by life insurers, and credit spreads. It reports the parameter estimates from the following regression:

$$m_{ikt} = \beta_{CS} \cdot CS_t + \beta_{GFC} \cdot \mathbb{1}_{GFC} + \beta_{csGFC} \cdot CS_t \times \mathbb{1}_{GFC} + B' \cdot X_t + FE_i + FE_k + \epsilon_{ikt}$$

where m_{ikt} is the annualised markup set by insurer i at time t for an annuity which is in sub-product k. Sub-products vary depending on age, sex and maturity of the annuities. CS_t is Moody's credit spread of 10 year BAA corporate bonds yields over treasuries, and $\mathbb{1}_{GFC}$ is an indicator variable set to one over the global financial crisis (November 2008 through February 2010). We include a vector of time series controls X_t which includes the risk-free rate, the slope of the yield curve, the TED spread, the CAPE ratio and US unemployment rate. We also include lagged markups in the control vector. Columns 1-3 report the parameter estimates from time series regressions where $\overline{m}_t$ is the average markup across insurers and sub-product categories in each time period. Columns 4-5 report full panel specifications. Panel A, B and C show the results for markups on life, guarantee and fixed-term annuity products respectively. The sample consists of biannual observations from January 1989 through July 2011. The t-statistics in the time series regressions are calculated using Newey and West (1987) standard errors with automatic bandwith selection. The panel regression also includes firm and fixed effects and standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

Panel A: Life Annuity Markups and Credit Spreads

	$\overline{m}_t$			m_{ikt}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.44***	-0.38***	-0.50***	-0.29***	-0.44***
	(-11.49)	(-5.66)	(-5.47)	(-4.58)	(-4.03)
$\mathbb{1}_{GFC}$			-1.01*		-0.66
			(-1.93)		(-1.51)
Credit Spread $\times$ $\mathbb{1}_{GFC}$			0.23*		0.21*
			(1.83)		(1.80)
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Product FE				yes	yes
Adj R-sq (Within)	0.800	0.871	0.876	0.596	0.603
Observations	72	72	72	12,460	12,460

Panel B: Guarantee Annuity Markups and Credit Spreads

	$\overline{m}_t$			m_{ikt}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.46***	-0.32***	-0.43***	-0.26***	-0.41***
	(-12.97)	(-5.43)	(-4.21)	(-4.93)	(-3.99)
$\mathbb{1}_{GFC}$			-1.06***		-0.66
			(-3.18)		(-1.60)
Credit Spread $\times$ $\mathbb{1}_{GFC}$			0.24**		0.20*
			(2.43)		(1.82)
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Product FE				yes	yes
Adj R-sq (Within)	0.799	0.875	0.883	0.655	0.664
Observations	53	53	53	14,529	14,529

Panel C: Fixed-Term Annuity Markups and Credit Spreads

	$\overline{m}_t$			m_{ikt}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.54***	-0.40**	-0.62***	-0.31***	-0.57***
	(-9.20)	(-2.62)	(-4.50)	(-2.89)	(-4.83)
$\mathbb{1}_{GFC}$			-0.87		-1.13***
			(-1.56)		(-2.72)
Credit Spread $\times$ $\mathbb{1}_{GFC}$			0.37**		0.44***
			(2.47)		(3.58)
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Product FE				yes	yes
Adj R-sq (Within)	0.861	0.857	0.873	0.432	0.458
Observations	45	45	45	2,557	2,557

Table 1.6: Investment returns drive the time series of premiums: P&C Insurance

This table shows the time series relation between insurance premiums, as measured by P&C insurer's underwriting profitability, and credit spreads. It reports the parameter estimates from the following time series regression:

$$u_{it} = \beta_{cs} \cdot CS_t + \beta_{GFC} \cdot \mathbb{1}_{GFC} + \beta_{csGFC} \cdot CS_t \times \mathbb{1}_{GFC} + B' \cdot X_t + FE_i + \epsilon_{it}$$

where u_{it}, is the underwriting profitability for insurer i in quarter t. Underwriting profitability is defined as underwriting profits (premiums earned minus losses and expenses) divided by the premiums earned. c_t is the 1-year rolling average of Moody's credit spread of BAA corporate bonds, $\mathbb{1}_{GFC}$ is an indicator variable set to one over the financial crisis (November 2008 through February 2010), and X_t is a vector of time series controls including 1-year rolling averages of investment returns and macroeconomic variables. Columns 1-3 report parameter estimates from the time series regression where the dependent variable, $\overline{u}_t$, is the average underwriting profitability in quarter t across all insurers. Columns 4-5 report parameter estimates from panel regressions with insurer fixed effects. The sample consists of quarterly observations from March 2001 through December 2017. t-statistics are reported in the brackets and are calculated using Newey and West (1987) standard errors in the time-series specifications, and standard errors clustered by date and firm in the panel specifications. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	$\overline{u}_t$			u_{it}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.44***	-0.83***	-1.08***	-0.74***	-1.06***
	(-2.71)	(-3.32)	(-4.85)	(-2.94)	(-4.73)
FC			-1.80		-1.64
			(-1.28)		(-0.80)
Credit Spread × FC			0.85**		0.87*
			(2.57)		(1.72)
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Adj R-sq (Within)	0.119	0.222	0.293	0.031	0.039
Observations	67	67	67	41,589	41,589

Table 1.7: Investment returns drive the cross section of premiums: Life Insurance

This table shows the cross section relation between insurance premiums, as measured by the markups on annuities issued by life insurers, and firm-specific expected investment returns. It reports the parameter estimate from the following panel regression:

$$m_{ikt} = \beta_y \cdot y_{it} + \beta_{yGFC} \cdot y_{it} \times \mathbb{1}_{GFC} + B' \cdot X_{it-1} + FE_i + FE_k + FE_t + \epsilon_{ikt}$$

where m_{ikt} is the annualised markup set by insurer i at time t for an annuity which is in sub-product category k, y_{it} is the insurer's investment return, $\mathbb{1}_{GFC}$ is an indicator variable set to one over the global financial crisis (November 2008 through February 2010), and X_{it} is a vector of lagged variables that capture balance sheet strength (leverage, risk-based capital, asset growth and deferred annuities). The control vector includes squared variables to capture non-linear effects of capital constraints. We additionally control for date fixed effects, product fixed effects and firm fixed effects, and report within group r-squared. Panel A, B and C show the results for markups on fixed-term, guarantee and life annuity products respectively. The sample consists of quarterly observations from March 2001 through March 2018. t-statistics are reported in bracket and calculated using standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

Panel A: Fixed Term Annuities

	(1)	(2)	(3)	(4)	(5)
Investment Return	-0.03***	-0.03***	-0.01	-0.03***	-0.01
	(-2.63)	(-2.86)	(-1.31)	(-2.77)	(-1.20)
Investment Return $\times \mathbb{1}_{Fin.Crisis}$				-0.06	-0.10*
				(-0.99)	(-1.74)
Firm Controls Vector		yes		yes	
Firm FE			yes		yes
Date FE	yes	yes	yes	yes	yes
Product FE	yes	yes	yes	yes	yes
Adj R-sq (Within)	0.010	0.078	0.007	0.078	0.009
Observations	955	955	955	955	955

Panel B: Guarantee Annuities

	(1)	(2)	(3)	(4)	(5)
Investment Return	-0.01***	-0.01***	-0.01***	-0.01***	-0.01***
	(-3.86)	(-4.35)	(-3.21)	(-4.34)	(-2.72)
Investment Return $\times \mathbb{1}_{Fin.Crisis}$				0.00	-0.05***
				(0.20)	(-4.70)
Firm Controls Vector		yes		yes	
Firm FE			yes		yes
Date FE	yes	yes	yes	yes	yes
Product FE	yes	yes	yes	yes	yes
Adj R-sq (Within)	0.121	0.229	0.165	0.229	0.168
Observations	5,989	5,989	5,989	5,989	5,989

[table continued on next page...]

Panel C: Life Annuities

	(1)	(2)	(3)	(4)	(5)
Investment Return	0.00	-0.02***	-0.02***	-0.02***	-0.02***
	(0.37)	(-2.97)	(-3.15)	(-3.48)	(-3.31)
Investment Return $\times$ $\mathbb{1}_{Fin.Crisis}$				0.08***	0.03
				(3.55)	(1.54)
Firm Controls Vector		yes		yes	
Firm FE			yes		yes
Date FE	yes	yes	yes	yes	yes
Product FE	yes	yes	yes	yes	yes
Adj R-sq (Within)	0.001	0.069	0.004	0.072	0.005
Observations	3,410	3,410	3,410	3,410	3,410

Table 1.8: Investment returns drive the cross section of premiums: P&C Insurance

This table shows the cross section relation between insurance premiums, as measured by P&C insurer's underwriting profitability, and firm-specific expected investment returns. It reports the parameter estimate from the following panel regression:

$$u_{it} = \beta_y \cdot y_{it} + \beta_{yGFC} \cdot y_{it} \times \mathbb{1}_{GFC} + B' \cdot X_{it-1} + FE_i + FE_t + \epsilon_{it}$$

where u_{it} is the underwriting profitability for insurer i at time t, and y_{it} is the insurer's investment return. We additionally control for date fixed effects, firm fixed effects and X_{it}, which is a vector of lagged variables that capture balance sheet strength (leverage, risk-based capital, asset growth and unearned premiums). This includes variables squared to control for non-linear effects of capital constraints. We also include a control for the level of reinsurance activity insurance company i engages in at time t. The samples consist of quarterly observations from Q1 2001 through Q4 2017. In columns 4-5 we interact investment return with an indicator variable $\mathbb{1}_{GFC}$ set equal to one during the global financial crisis (Q4 2008 through Q1 2010). t-statistics are reported in bracket and calculated using standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
Investment Return	-0.10**	-0.12***	-0.11***	-0.13***	-0.12***
	(-2.37)	(-3.09)	(-5.19)	(-3.37)	(-5.70)
Investment Return × FC				0.10	0.12**
				(1.52)	(2.56)
Firm Controls Vector		yes		yes	
Firm FE			yes		yes
Time FE	yes	yes	yes	yes	yes
Adj R-sq (Within)	0.001	0.071	0.001	0.071	0.001
Observations	37,044	37,044	37,044	37,044	37,044

Table 1.9: P&C Insurance cross section: instrumented variable estimation

This table shows the cross section relation between insurance premiums, as measured by P&C insurer's underwriting profitability, and the instrumented expected investment returns of individual insurance companies. Columns (3) and (4) report the parameter estimate from the following instrumental variable panel regression:

$$u_{it} = \beta_y \cdot y_{it} + B' \cdot X_{it-1} + FE_t + \epsilon_{it}$$

where u_{it} is the underwriting profitability for insurer i at time t, and y_{it} is the instrumented investment return of insurer i at time t. Columns (1) and (2) report the first-stage results from the regression

$$y_{it} = \beta_{vol} \cdot Volatility_{i,t-1} + \beta_{Size} \cdot Size_{i,t-1} + B' \cdot X_{it-1} + FE_t + \epsilon_{it}$$

where the instruments are the historical 5-year volatility of insurer i's underwriting profitability up to and including time $t-1$, $Volatility_{i,t-1}$, and the insurers size (log assets) at $t-1$. First stage results in Columns (1) and (2) correspond to the second-stage results in Columns (3) and (4) respectively . We control for date fixed effect in all specifications, and in (2) and (4) we include an untabulated vector, X_{it-1}, of lagged variables that capture balance sheet strength (leverage, risk-based capital, asset growth and unearned premiums), and the level of reinsurance activity insurance company i engages in at time t. The samples consist of quarterly observations from Q1 2001 through Q4 2017. For the second stage, we report the Cragg-Donald Wald F-statistic, and in the case where we have two instrumental variables (Column 4), we report the p-value from the Sargan's χ^2 test of overidentifying restrictions. t-statistics are reported in bracket and calculated using standard errors clustered by firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	First Stage:		Second Stage:	
	(1)	(2)	(3)	(4)
Underwriting Volatility (t-1)	-0.16***	-0.07***		
	(-7.47)	(-2.63)		
Size (t-1)		0.17***		
		(6.69)		
Investment Return			-0.35***	-0.27**
			(-3.41)	(-2.36)
Control Vector		yes		yes
Date FE	yes	yes	yes	yes
Adj R-sq (Within)	0.040	0.075		
Cragg-Donald F-stat			101.576	2042.461
Sargan test p-value				0.478
Observations	25,091	25,091	25,091	25,091

Table 1.10: Life Insurance Cross Section: evidence from mergers and acquisitions

This table shows the relation between the annuity markups and investment returns using a difference-in-differences approach around merger events. The treatment group is the life insurance companies involved in a merger and acquisition event over our sample, and the control group is all other life insurance companies. The control time period is the two years pre-mergers, and the treatment is the two years following merger. The table reports the parameter estimate from the following regression:

$$m_{ikt} = \beta_D \cdot D_{it} + FE_i + FE_k + FE_t + \epsilon_{ijt}$$

where m_{ikt} is the markup set by insurer i at time t on product k, and D_{it} is a variable set equal to zero for all observations except for treatment group insurance companies in the treatment period (the two years following their merger or acquisition event). For these observations, the variable is set equal to the treatment group insurance company's investment return minus the investment return of the other insurance company involved in the transaction (i.e. it is the investment return differential). For each individual mergers, we select the two years either side of the event for our sample, with our total sample made up of the union of the individual merger samples. This leads to 941 observations across 20 quarterly dates, with 5 treatment group entities and 48 control group entities. We use one annuity product type for each of our three broad categories of annuity - 20yr fixed term annuity, life annuity for males aged 50, and 10 year guarantee life annuity for a male aged 50. We control for time, company and product fixed effects. Standard errors are clustered by insurance company and date. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	markup (ikt)
Δ Investment Return(it)	-0.22***
	(-3.44)
Firm FE	yes
Date FE	yes
Product FE	yes
Adj R-sq (Within)	0.007
Observations	2318

Table 1.11: Evidence from excess bond risk premium

This table shows the relation between the markups on annuities issued by life insurers and the expected return component of credit spreads. It reports the parameter estimate from the

$$m_{jt} = \beta_e \cdot EBP_t + \beta_{df} \cdot DF_t + \beta_{eGFC} \cdot EBP_t \times \mathbb{1}_{GFC} + \beta_{dGFC} \cdot DF_t \times \mathbb{1}_{GFC} + FE_i + FE_k + \epsilon_{jt}$$

where $j = (i, k)$ and m_{jt} is the annualised markup set by insurer i at time t for an annuity which is in sub-product category k. Sub-products vary depending on age, sex and maturity of the annuities. EBP_t is the Gilchrist and Zakrajšek (2012) credit spread attributed to excess bond risk premium, DF_t is the credit spread attributed to default losses, and $\mathbb{1}_{GFC}$ is an indicator variable set to one over the global financial crisis (November 2008 through February 2010). We include a vector of time series controls X_t which includes the risk-free rate, the slope of the yield curve, the TED spread, the CAPE ratio and US unemployment rate. We also include lagged markups in the control vector. Columns 1-2 report the parameter estimates where markups, $\overline{m}_t$, are averaged across insurers and sub-products in each time period. Columns 3-4 report full panel specifications. Panel A, B and C show the results for markups on life, guarantee and fixed-term annuity products respectively. The sample consists of biannual observations from January 1989 through July 2011. t-statistics in the time series regressions are calculated using Newey and West (1987) standard errors with automatic bandwith selection. The panel regression also includes firm and product fixed effects and standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

Panel A: Life Annuities

	(1)	(2)	(3)	(4)
Excess Bond Risk Premia	-0.36***	-0.61***	-0.31***	-0.46***
	(-4.35)	(-5.39)	(-4.20)	(-5.07)
Default Risk	-0.10	0.27	-0.03	0.18
	(-0.73)	(1.45)	(-0.39)	(1.22)
$\mathbb{1}_{Fin.Crisis}$		0.71		1.21**
		(1.39)		(2.31)
Excess Bond Risk Premia $\times$ $\mathbb{1}_{Fin.Crisis}$		0.48***		0.42***
		(4.11)		(4.35)
Default Risk $\times$ $\mathbb{1}_{Fin.Crisis}$		-0.43**		-0.48***
		(-2.51)		(-2.75)
Entity FE			yes	yes
Product FE			yes	yes
Time Series Controls	yes	yes	yes	yes
Adj R-sq (Within)	0.871	0.895	0.600	0.618
Observations	72	72	12460	12460

Panel B: Guarantee Annuities

	(1)	(2)	(3)	(4)
Excess Bond Risk Premia	-0.27***	-0.45***	-0.27***	-0.33***
	(-4.21)	(-10.08)	(-4.75)	(-5.08)
Default Risk	-0.17	-0.06	-0.10	-0.21
	(-1.56)	(-0.29)	(-1.22)	(-1.52)
$\mathbb{1}_{Fin.Crisis}$		0.12		0.18
		(0.20)		(0.34)
Excess Bond Risk Premia $\times$ $\mathbb{1}_{Fin.Crisis}$		0.37***		0.23***
		(4.25)		(2.92)
Default Risk $\times$ $\mathbb{1}_{Fin.Crisis}$		-0.15		-0.05
		(-0.72)		(-0.27)
Entity FE			yes	yes
Product FE			yes	yes
Time Series Controls	yes	yes	yes	yes
Adj R-sq (Within)	0.884	0.914	0.670	0.685
Observations	53	53	14529	14529

Panel C: Fixed Term Annuities

	(1)	(2)	(3)	(4)
Excess Bond Risk Premia	-0.63***	-0.68***	-0.52***	-0.56***
	(-4.97)	(-5.01)	(-6.60)	(-6.59)
Default Risk	0.35*	0.69*	0.27***	0.37**
	(2.00)	(1.78)	(3.32)	(2.09)
$\mathbb{1}_{Fin.Crisis}$		0.60		1.34**
		(0.70)		(2.41)
Excess Bond Risk Premia $\times$ $\mathbb{1}_{Fin.Crisis}$		0.26*		0.58***
		(1.85)		(5.25)
Default Risk $\times$ $\mathbb{1}_{Fin.Crisis}$		-0.41		-0.55***
		(-1.22)		(-3.08)
Entity FE			yes	yes
Product FE			yes	yes
Time Series Controls	yes	yes	yes	yes
Adj R-sq (Within)	0.910	0.911	0.472	0.475
Observations	45	45	2557	2557

1.10 Appendix A: Institutional Background

1.10.1 Underwriting Profit in Life Insurance

Mitchell, Poterba, Warshawsky, and Brown (1999) and Koijen and Yogo (2015) have documented markups an average of 6 to 10 percent on specific life insurance products, which is after adjusting for a time value of money (assumed to be the risk-free rate). While these markups make life insurance underwriting look profitable at first glance, it is important to note that they are gross of operating expenses and commissions. Expenses on the specific products of their studies are not available to make a direct net of expenses assessment. However, on an aggregated basis, the life insurance industry reported commission and expense costs that were 20% of premiums in 2018 (SNL Statutory Files). It is therefore not unreasonable to assume that life insurance, like P&C, is dependent on asset returns for overall profitability.

Indeed, for comparision, in figure 1.7 we plot P&C underwriting income between its three main components - claims and expenses (outflows) and earned premium (inflows). It shows that expesnes are significant fraction of premiums, ranging from 25%-30% across the sample. P&C underwriting performance gross of expenses looks extremely profitable. In other words, expenses are critical for an overall understanding of underwriting performance.

1.10.2 Accounting Treatment of the Investment Returns of Insurance Companies

For cross sectional comparisons of insurer expected investment return, we use their self-reported *Net Yield on Invested Assets*. This is their accounting return on assets, and is defined as dollar net income from investments over the dollar book value of invested assets. Anecdotally, we know from market participants that it is the key metric from which insurance companies assess their expected investment portfolio performance.

For fixed income assets, which are the average insurers' main asset allocation, net yield on any asset is simply the amortisation of the purchase yield. Such treatment of assets reflects that insurers are buy and hold investors and can weather mark to market fluctuations. If the insurer does sell a bond before maturity, in the reporting period of sale the realised mark to market gain/loss is also included in the net yield measure. Further, if there are significant revisions to the prospects for a bond (i.e. default appears likely), adjustments may also be made in reported investment income. For equity investments,

the net yield is the dividend rate, with mark to market fluctuations once again realised at the point of sale.

To capture insurers' expected returns at an industry-level we use the credit spread on corporate bonds. This is the average insurers' main source of investment risk and thus is our best proxy for industry wide investment opportunities. We also use the excess bond risk premium portion of credit spreads as provided in (Gilchrist and Zakrajšek (2012)), which is a way to strip out expected default loss from the credit spread.

1.11 Appendix B: Proofs

Proof of Proposition 1

i) Investor illiquid allocation.

The first-order condition for the investor's illiquid asset allocation in equation (1.3) is

$$0 = (1 - \omega) R - \omega \lambda \theta \tag{1.23}$$

from which the optimal allocation (1.10) follows. ∎

ii) Insurer illiquid allocation.

We have already defined in equation (1.14) the lower bound on the insurer's optimal asset allocation. By a similar logic we can also define an upper bound. To see this, note that $\tau = \bar{\tau} - \sigma$ is the minimum fraction of claims that will arrive early. The insurer therefore knows they will be forced to sell assets of at least $(\bar{\tau} - \sigma) C$ at time 1. Optimally they hold at least this amount in liquid assets, which leads to the following definition

$$\overline{\Theta} = \begin{cases} L - (\bar{\tau} - \sigma) C & \text{if } L - (\bar{\tau} - \sigma) C < S \\ S & \text{otherwise.} \end{cases} \tag{1.24}$$

In the first case, investing $\Theta > \overline{\Theta}$ would mean paying sales costs on illiquid assets of amount $\Theta - \overline{\Theta}$ with no expectation of earning the liquidity premia R. In the second case, the insurer knows that if they invest more than total size of the illiquid asset market, it requires other investors to go short the asset. This would result in a negative R, which makes the asset more unattractive to the insurer. They therefore cap their investment at the total size S of the illiquid asset market.

The key implication of the upper bound $\overline{\Theta}$ is that the insurer does not sell illiquid assets when $\tau = \bar{\tau} - \sigma$ realizes. We can therefore restate wealth (1.7) in two cases that

depend on the fraction τ of claims arriving early

$$W = \begin{cases} L\left(1 + R^F\right) - C + \Theta R & \text{if } \tau = \bar{\tau} - \sigma \\ L\left(1 + R^F\right) - C + \underline{\Theta} R - \frac{1}{2}\lambda\left(\Theta - \underline{\Theta}\right)^2 & \text{if } \tau = \bar{\tau} + \sigma \end{cases} \tag{1.25}$$

with both cases occurring with equal probability. The first case shows the simple outcome where the insurer holds enough liquid assets to cover early claims. In the second case, the insurer sells all their liquid assets plus a portion of their illiquid asset portfolio to cover remaining $t = 1$ claims. Dollar amount $\tau C - (L - \Theta) = (\bar{\tau} + \sigma) C - (L - \Theta)$ of illiquid assets are sold early. Substituting in equation (1.14) this can be restated $\Theta - \underline{\Theta}$. The residual $\underline{\Theta}$ illiquid assets are held to maturity, earning the liquidity premia R.

The insurers objective function (1.8) is therefore be restated

$$\max_{P,\,\Theta} \quad L\left(1 + R^F\right) - C + \frac{1}{2}\left(\Theta + \underline{\Theta}\right) R - \frac{1}{4}\lambda\left(\Theta - \underline{\Theta}\right)^2. \tag{1.26}$$

The first-order condition for the illiquid asset dollar investment is

$$0 = \frac{1}{2}R - \frac{1}{2}\lambda\left(\Theta - \underline{\Theta}\right) \tag{1.27}$$

and thus the optimal solution Θ^* in equation (1.11) follows. ∎

Note the solution holds for any required return on insurer equity providing that the required return is independent of the insurer's asset allocation decision. We have this in this model due to risk neutral investors. However, it would hold in any model with a flat security market line.

Proof of Theorem 1

The proof is shown with the insurer facing a generalised convex cost function of selling illiquid assets. We now assume that the insurer pays $\lambda f(x)$ dollar for every x dollar sold of the illiquid asset, where $f'(x) > 0$ and $f''(x) > 0$. The generalised version of the insurer's objective function (1.26) is thus

$$\max_{P,\,\Theta} \quad L\left(1 + R^F\right) - C + \frac{1}{2}\left(\Theta + \underline{\Theta}\right) R - \frac{1}{2}\lambda f(x) \tag{1.28}$$

where $x = \Theta - \underline{\Theta}$ is the dollar amount of illiquid assets sold.

The first-order condition with respect the illiquid asset allocation Θ is

$$0 = \frac{1}{2}R - \frac{1}{2}\lambda f'(x) \frac{\partial x}{\partial \Theta}$$

where we have used the chain rule and assumed the insurer takes the illiquid asset return R as fixed. Given $\frac{\partial x}{\partial \Theta} = 1$, the first-order condition solves to

$$R = \lambda f'(x). \tag{1.29}$$

From this condition we can see the marginal benefit R of an extra dollar of illiquid investment is equal to the marginal cost $\lambda f'(x)$ of an extra dollar of illiquid investment. The insurer optimally increases their illiquid investment allocation until this holds for any convex cost function.

Meanwhile, for fixed illiquid asset allocation, the first-order condition on (1.28) for the insurance price is

$$0 = \frac{\partial L}{\partial P}\left(1 + R^F\right) - \frac{\partial C}{\partial P} + \frac{1}{2}\frac{\partial \Theta}{\partial P}R - \frac{1}{2}\lambda f'(x)\frac{\partial x}{\partial P} \tag{1.30}$$

where $\frac{\partial x}{\partial P} = -\frac{\partial \underline{\Theta}}{\partial P}$. Using the envelope theorem, we now substitute in condition 1.29 from the optimal illiquid asset decision to simplify to

$$0 = \frac{\partial L}{\partial P}\left(1 + R^F\right) - \frac{\partial C}{\partial P} + \frac{\partial \underline{\Theta}}{\partial P}R. \tag{1.31}$$

Note that the only impact of the excess return R on the optimal insurance price comes via the lower bound of illiquid investment $\underline{\Theta}$. This is the portion of the assets that the insurer knows it will not be forced to sell at $t = 1$. Substituting in the lower bound of the illiquid asset allocation (1.14) we have

$$0 = \frac{\partial L}{\partial P}\left(1 + R^F + R\right) - \frac{\partial C}{\partial P}\left(1 + (\bar{\tau} + \theta)R\right). \tag{1.32}$$

and using equations 1.5 and 1.6 and the product rule, the first order condition is thus

$$0 = \left(Q + \frac{\partial Q}{\partial P}P\right)\left(1 + R^F + R\right) - \frac{\partial Q}{\partial P}\bar{C}\left(1 + (\bar{\tau} + \theta)R\right) \tag{1.33}$$

$$= P(1 - \epsilon)\left(1 + R^F + R\right) + \epsilon\bar{C}\left(1 + (\bar{\tau} + \theta)R\right) \tag{1.34}$$

where the second line has been multiplied through by $\frac{P}{Q}$ and uses

$$\epsilon = -\frac{\partial \log Q}{\partial \log P} > 1. \tag{1.35}$$

Equation 1.34 is rearranged to give the final solution 1.12. ∎

Proof of Proposition 2

Equation 1.15 proof

By the chain rule we have

$$\frac{\partial P}{\partial R} = \frac{\partial P}{\partial R^I} \frac{\partial R^I}{\partial R}. \tag{1.36}$$

From equation 1.12 we can see

$$\frac{\partial P}{\partial R^I} = -\frac{P}{1 + R^I} < 0 \tag{1.37}$$

and from 1.13 we can see

$$\frac{\partial R^I}{\partial R} = \frac{1 - \bar{\tau} - \sigma}{\left(1 + (\bar{\tau} + \sigma) R\right)^2} > 0 \tag{1.38}$$

given $\bar{\tau} + \sigma < 1$. ∎

Exogenous shocks to equilibrium asset returns

The asset market clearing condition (1.9) states

$$S = \frac{(1 - \omega)}{\omega} \frac{R}{\lambda} + L - (\bar{\tau} + \sigma) C + \frac{R}{\lambda} \tag{1.39}$$

$$= \frac{1}{\omega \lambda} R + \underline{\Theta}. \tag{1.40}$$

In the first line we have used the equilibrium asset demands (1.10) and (1.11). In the second line we have substituted in $\underline{\Theta}$ from equation (1.14) and rearranged.

We therefore have the equilibrium condition

$$R = \omega \lambda \left(S - \underline{\Theta} \right) \tag{1.41}$$

with recognition that $\underline{\Theta}(R)$ is endogenous. The derivative with respect λ[22] is therefore

$$\frac{\partial R}{\partial \lambda} = \omega \left(S - \underline{\Theta} - \lambda \frac{\partial \underline{\Theta}}{\partial R} \frac{\partial R}{\partial \lambda} \right) \tag{1.42}$$

where we have used the product rule, and chain rule with respect the endogenous variable. The derivative rearranges to

$$\frac{\partial R}{\partial \lambda} = \frac{\omega \left(S - \underline{\Theta} \right)}{1 + \lambda \frac{\partial \underline{\Theta}}{\partial R}} \tag{1.43}$$

and we can see that to show $\frac{\partial R}{\partial \lambda} > 0$, we require to show both

1. $S > \underline{\Theta}$

2. $\frac{\partial \underline{\Theta}}{\partial R} > -\frac{1}{\lambda}$

---

[22] or derivative wrt ω. The proof for each variable from here is identical. We proceed by showing with λ.

Part 1. holds by definition 1.24. The insurer will not hold more than the total illiquid asset market. The rest of the proof focuses on part 2.

We will, in fact, show that $\frac{\partial \underline{\Theta}}{\partial R} > 0$. The result is intuitive. If R increases then insurer's set cheaper insurance (see 1.15), which increases the number of contracts they underwrite. Stable funding is constant fraction of claims. An increase in claims is therefore an increase in stable funding, which allows the insurer to invest more in illiquid assets (i.e. $\underline{\Theta}$ increases).

To show the following result

$$\frac{\partial \left[E + QP - (\bar{\tau} + \sigma) Q\tilde{C} \right]}{\partial R} > 0 \tag{1.44}$$

we can see that we must show

$$\frac{\partial QP}{\partial R} - (\bar{\tau} + \sigma) \tilde{C} \frac{\partial Q}{\partial R} > 0. \tag{1.45}$$

To proceed from here, we use $Q = kP^{-\epsilon}$ from equation (1.4) and the chain rule to show

$$\frac{\partial Q}{\partial R} = \frac{\partial Q}{\partial P} \frac{\partial P}{\partial R}$$
$$= -\varepsilon \frac{Q}{P} \frac{\partial P}{\partial R}.$$

Using this result and the product rule we also show

$$\frac{\partial QP}{\partial R} = \frac{\partial Q}{\partial R} P + \frac{\partial P}{\partial R} Q$$
$$= -\varepsilon \frac{\partial P}{\partial R} Q + \frac{\partial P}{\partial R} Q$$
$$= (1 - \varepsilon) \frac{\partial P}{\partial R} Q.$$

Substituting these two derivatives into inequality (1.45), we thus have:

$$(1 - \varepsilon) \frac{\partial P}{\partial R} Q + (\bar{\tau} + \sigma) \tilde{C} \varepsilon \frac{Q}{P} \frac{\partial P}{\partial R} > 0$$

and dividing through by (the negative) $\frac{Q}{P} \frac{\partial P}{\partial R}$ we have

$$P (1 - \varepsilon) + \varepsilon (\bar{\tau} + \sigma) \tilde{C} < 0$$

and dividing through by (the negative) $(1 - \varepsilon)$ we have

$$P - M (\bar{\tau} + \sigma) \tilde{C} > 0$$

where we have used $M = \frac{\varepsilon}{\varepsilon - 1}$. Finally, we substitute the equilibrium premium price 1.12 and simplify

$$M\frac{1 + R(\bar{\tau} + \sigma)}{1 + R}\tilde{C} - M(\bar{\tau} + \sigma)\tilde{C} > 0$$

$$1 + R(\bar{\tau} + \sigma) - (1 + R)(\bar{\tau} + \sigma) > 0$$

$$1 - (\bar{\tau} + \sigma) > 0$$

which we know holds. The fraction $\tau \in \{\bar{\tau} - \sigma, \bar{\tau} + \sigma\}$ of insurer claims arriving at time 1 can not exceed one. ∎

Proof of Proposition 4

The Lagrangian for the insurer's optimisation problem (1.8) when subject to (1.21) is

$$\mathcal{L}(P, \Theta, \eta) = W + \eta\left(L - \frac{C}{1 + R^S}\phi^{-1}\right). \tag{1.46}$$

Following the proof of proposition **??**, the corresponding first order condition for the insurance premium can be stated

$$0 = \frac{\partial L}{\partial P}(1 + R) - \frac{\partial C}{\partial P}(1 + (\bar{\tau} + \theta)R) + \eta\left(\frac{\partial L}{\partial P} - \frac{\frac{\partial C}{\partial P}}{1 + R^S}\phi^{-1}\right). \tag{1.47}$$

Using equations 1.5, 1.6 and 1.4, and the product rule, the first order condition is rearranged to

$$P(1 - \epsilon)(1 + R^I)\left(1 + \frac{\eta}{1 + R^I}\frac{1}{1 + (\bar{\tau} + \sigma)R}\right) = -\epsilon\tilde{C}\left(1 + \frac{\eta}{\phi(1 + R^S)}\frac{1}{1 + (\bar{\tau} + \sigma)R}\right) \tag{1.48}$$

and we rearrange this formula to solve the equilibrium price (1.22).

Proof of Proposition 5

This result follows straight from the equilibrium price 1.22, with the cases depending on whether

$$\frac{1 + (\bar{\tau} + \sigma)R + \frac{\eta}{\phi(1 + R^S)}}{1 + (\bar{\tau} + \sigma)R + \frac{\eta}{(1 + R^I)}}$$

is greater or less than 1.

1.12   Appendix C: Further Figures and Tables

Figure 1.7: P&C Insurance - Industry Wide Insurance Underwriting Cashflows. This figures plots the industry-wide insurance underwriting cashflows in Property & Casualty markets. Total income from insurance underwriting are the premiums received minus the claims paid and the operating expenses associated with the running of an insurance underwriting business (pricing, reserving, marketing, operations etc.). The data comes from quarterly US insurance company statutory filings 2001:2018 and is provided by SNL Global. Individual company data has been aggregated to show the industry-wide net income.

Figure 1.8: Entities in the Cross-Section

This figures plots the number of entities observed in the cross-section for each time-period. Panel A plots the number of life insurance companies in annuity cross-sectional regressions. Panel B plots the number of Property & Casualty entities.

(a) Life Insurers (annuities)

(b) P&C Insurers

Table 1.12: Mergers and Acquisitions Sample

This table shows the sample of mergers and acquistions that exist for our life insurance company dataset. The insurance companies underlined are those for which we have markup data for both pre and post the event.

Company A	Company B	Deal Type	Date of Completion
American Heritage	AllState Insurance	Acquisition	October 1999
General Electric Capital Assurance	Genworth Financial	Acquisition	January 2003
John Hancock	ManuLife	Acquisition	April 2004
Jefferson-Pilot	Lincoln National	Merger	April 2006

Table 1.13: Life Insurance Time Series - Full Specification Estimates

This table shows the relation between the markups on annuities issued by life insurers and credit spreads. It reports the parameter estimates from the following regression:

$$m_{ikt} = \beta_c \cdot c_t + \beta_{GFC} \cdot \mathbb{1}_{GFC} + \beta_{cGFC} \cdot c_t \times \mathbb{1}_{GFC} + B' \cdot X_t + FE_i + FE_k + \epsilon_{ikt}$$

where m_{ikt} is the annualised markup set by insurer i at time t for an annuity which is in sub-product category k. Sub-products vary depending on age, sex and maturity of the annuities. c_t is Moody's credit spread of BAA corporate bonds, and $\mathbb{1}_{GFC}$ is an indicator variable set to one over the global financial crisis (November 2008 through February 2010). We include a vector of time series controls X_t which includes the risk-free rate, the slope of the yield curve, the TED spread and US unemployment rate. Columns 1-3 report the parameter estimates from time series regressions where for the dependent variable, $\overline{m}_t$, we have averaged across insurers and sub-product categories in each time period. Columns 4-5 are full panel specifications. Panel A, B and C show the results for markups on fixed-term, guarantee and life annuity products respectively. The sample consists of biannual observations from January 1989 through July 2011. The t-statistics in the time series regressions are calculated using Newey and West (1987) standard errors with automatic bandwith selection. The panel regression also includes firm and fixed effects and standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

Panel A: Life Term Annuities

	$\overline{m}_t$			m_{ikt}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.44***	-0.38***	-0.50***	-0.29***	-0.44***
	(-11.49)	(-5.66)	(-5.47)	(-4.58)	(-4.03)
$\mathbb{1}_{GFC}$			-1.01*		-0.66
			(-1.93)		(-1.51)
Credit Spread $\times$ $\mathbb{1}_{GFC}$			0.23*		0.21*
			(1.83)		(1.80)
markup (j,t-1)		0.23***	0.18**	0.47***	0.46***
		(2.74)	(2.06)	(7.63)	(7.57)
Risk Free (5yr)		0.11	0.11	0.12***	0.09***
		(1.38)	(1.59)	(4.01)	(2.92)
Slope (5yr - 1yr)		0.18	0.22**	0.16***	0.23***
		(1.48)	(2.06)	(2.79)	(3.97)
Ted Spread		-0.08	-0.03	0.00	0.07
		(-0.94)	(-0.32)	(0.04)	(0.73)
CAPE ratio		0.01**	0.02***	0.01	0.01*
		(2.42)	(3.22)	(1.25)	(1.81)
Unemployment Rate		0.10**	0.14***	0.12***	0.11***
		(2.39)	(2.93)	(3.04)	(2.72)
Duration (j,t)	-0.34***	0.12	0.16		
	(-3.08)	(0.48)	(0.78)		
Constant	5.10***	-0.88	-1.30		
	(5.57)	(-0.36)	(-0.61)		
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Product FE				yes	yes
Adj R-sq (Within)	0.800	0.871	0.876	0.596	0.603
Observations	72	72	72	12,460	12,460

[table continued on next page...]

Panel B: Guarantee Annuities

	$\overline{m}_t$			m_{ikt}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.46***	-0.32***	-0.43***	-0.26***	-0.41***
	(-12.97)	(-5.43)	(-4.21)	(-4.93)	(-3.99)
$\mathbb{1}_{GFC}$			-1.06***		-0.66
			(-3.18)		(-1.60)
Credit Spread $\times \mathbb{1}_{GFC}$			0.24**		0.20*
			(2.43)		(1.82)
markup (j,t-1)		0.23*	0.23*	0.45***	0.43***
		(1.77)	(1.75)	(6.73)	(6.44)
Risk Free (5yr)		0.12*	0.14	0.14***	0.12*
		(1.70)	(1.47)	(3.14)	(1.91)
Slope (5yr - 1yr)		0.06	0.13***	0.14**	0.20***
		(0.85)	(2.75)	(2.43)	(4.08)
Ted Spread		-0.08	-0.02	-0.06	-0.01
		(-0.94)	(-0.26)	(-0.79)	(-0.12)
CAPE ratio		0.01	0.01	0.00	0.01
		(0.99)	(0.84)	(0.40)	(0.69)
Unemployment Rate		0.10**	0.15***	0.12***	0.12**
		(2.38)	(3.44)	(3.08)	(2.46)
Duration (j,t)	-0.04	-0.13*	-0.19**		
	(-0.42)	(-1.75)	(-2.69)		
Constant	2.06**	1.26	1.56*		
	(2.24)	(1.58)	(1.73)		
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Product FE				yes	yes
Adj R-sq (Within)	0.799	0.875	0.883	0.655	0.664
Observations	53	53	53	14,529	14,529

Panel C: Term Annuities

	$\overline{m}_t$			m_{ikt}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.54***	-0.40**	-0.62***	-0.31***	-0.57***
	(-9.20)	(-2.62)	(-4.50)	(-2.89)	(-4.83)
$\mathbb{1}_{GFC}$			-0.87		-1.13***
			(-1.56)		(-2.72)
Credit Spread $\times \mathbb{1}_{GFC}$			0.37**		0.44***
			(2.47)		(3.58)
markup (j,t-1)		0.13	0.15	0.39***	0.39***
		(1.02)	(1.13)	(4.72)	(4.79)
Risk Free (5yr)		0.06	0.06	0.17***	0.13***
		(1.14)	(1.60)	(4.03)	(4.08)
Slope (5yr - 1yr)		0.05	0.11	0.14*	0.20***
		(0.49)	(0.98)	(1.76)	(3.07)
Ted Spread		-0.17	-0.28**	-0.13	-0.23*
		(-0.83)	(-2.25)	(-0.77)	(-1.98)
CAPE ratio		0.01	0.01	0.02**	0.02***
		(1.34)	(1.65)	(2.19)	(3.52)
Unemployment Rate		0.01	0.00	0.10**	0.09**
		(0.30)	(0.04)	(2.60)	(2.63)
Duration (j,t)	-0.28***	-0.23***	-0.19***		
	(-6.28)	(-3.48)	(-2.84)		
Constant	4.22***	2.83***	2.99***		
	(17.09)	(3.02)	(3.54)		
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Product FE				yes	yes
Adj R-sq (Within)	0.861	0.857	0.873	0.432	0.458
Observations	45	45	45	2,557	2,557

Table 1.14: P&C Time Series - Underwriting Profitability and Credit Spreads

This table shows the relation between quarterly P&C insurance underwriting profitability and credit spreads. Columns 1-3 report the parameter estimate from the following time series regression:

$$\overline{u}_t = \alpha + \beta_c \cdot c_t + \beta_{cFC} \cdot c_t \times \mathbb{1}_{FC} + B' \cdot X_t + \epsilon_t$$

where $\overline{u}_t$ is the average underwriting profitability in quarter t across all insurers. Underwriting profitability is defined as underwriting profits (premiums earned minus losses and expenses) divided by the premiums eared. c_t is the 1-year rolling average of Moody's credit spread of BAA corporate bonds, $\mathbb{1}_{FC}$ is an indicator variable set to one over the financial crisis (November 2008 through February 2010), and X_t is a vector of time series controls with 1-year rolling averages of investment returns and macroeconomic variables. we also run the regression in the full panel of insurance companies by estimating the model:

$$u_{it} = \beta_c \cdot c_t + \beta_{cFC} \cdot c_t \times \mathbb{1}_{FC} + B' \cdot X_t + FE_i + \epsilon_{it}$$

where u_{it}, is the underwriting profitability for insurer i in quarter t. Reported adjusted r-squared are within groups for panel specifications. The sample consists of quarterly observations from 2001Q1 through to 2018Q3. T-statistics are reported in the brackets and are calculated using Newey and West (1987) standard errors in the time-series specifications when possible, and standard errors clustered by date and firm in the panel specifications. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	$\overline{u}_t$			u_{it}	
	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.44***	-0.83***	-1.08***	-0.74***	-1.06***
	(-2.71)	(-3.32)	(-4.85)	(-2.94)	(-4.73)
Risk Free (1yr)		-0.35**	-0.34***	-0.17*	-0.18*
		(-2.63)	(-2.81)	(-1.70)	(-1.97)
Ted Spread		1.10***	0.05	0.76*	-0.36
		(2.71)	(0.10)	(1.67)	(-0.73)
Slope (5yr - 1yr)		-0.26	-0.35**	0.12	-0.02
		(-1.24)	(-2.08)	(0.65)	(-0.09)
Unemployment Rate		-0.05	-0.07	-0.11	-0.12
		(-0.56)	(-0.75)	(-1.36)	(-1.47)
Reinsurance Activity (t-1)		0.28	0.63	-0.10*	-0.10
		(0.16)	(0.39)	(-1.68)	(-1.65)
Risk Based Capital (t-1)		-0.51*	-0.42	0.22***	0.22***
		(-1.69)	(-1.56)	(12.33)	(12.60)
FC			-1.80		-1.64
			(-1.28)		(-0.80)
Credit Spread × FC			0.85**		0.87*
			(2.57)		(1.72)
Constant	1.52***	5.69***	6.35***		
	(3.54)	(3.30)	(3.85)		
Time Series Controls Vector		yes	yes	yes	yes
Entity FE				yes	yes
Adj R-sq (Within)	0.119	0.222	0.293	0.031	0.039
Observations	67	67	67	41,589	41,589

Table 1.15: Life Insurance Time Series - Estimates in Changes

This table shows the relation between the markups on annuities issued by life insurers and credit spreads. It reports the parameter estimates from the following regression:

$$\Delta m_{jt} = \beta_c \cdot \Delta c_t + \beta_{FC} \cdot \mathbb{1}_{FC} + \beta_{cFC} \cdot \Delta c_t \times \mathbb{1}_{FC} + B' \cdot \Delta X_t + FE_i + FE_k + \epsilon_{ikt}$$

where $j = (i, k)$ and Δm_{jt} is the change in the annualised markup set by insurer i at time t for an annuity which is in sub-product category k. Sub-products vary depending on age, sex and maturity of the annuities. Δc_t is the change in the Moody's credit spread of BAA corporate bonds, and $\mathbb{1}_{FC}$ is an indicator variable set to one over the financial crisis (November 2008 through February 2010). We include a vector of time series controls ΔX_t in changes, which includes the risk-free rate, the slope of the yield curve, the TED spread and US unemployment rate. Columns 1-3 report the parameter estimates from time series regressions where for the dependent variable, $\overline{m}_t$, we have averaged across insurers and sub-product categories in each time period. Columns 4-5 are full panel specifications. Panel A, B and C show the results for markups on fixed-term, guarantee and life annuity products respectively. The sample consists of biannual observations from January 1989 through July 2011. The t-statistics in the time series regressions are calculated using Newey and West (1987) standard errors with automatic bandwith selection. The panel regression also includes firm and fixed effects and standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

Panel A: Life Annuities

	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.51***	-0.22**	-0.32***	-0.32***	-0.41***
	(-3.73)	(-2.22)	(-2.82)	(-5.37)	(-4.66)
$\mathbb{1}_{Fin.Crisis}$			0.16**		0.12
			(2.47)		(1.43)
Credit Spread $\times \mathbb{1}_{Fin.Crisis}$			0.16		0.11
			(1.11)		(1.16)
Entity FE				yes	yes
Product FE				yes	yes
Time Series Controls		yes	yes	yes	yes
Adj R-sq (Within)	0.239	0.521	0.527	0.420	0.426
Observations	72	71	71	11388	11388

Panel B: Guarantee Annuities

	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.41***	-0.31***	-0.49***	-0.32***	-0.47***
	(-4.23)	(-3.37)	(-3.82)	(-4.64)	(-5.46)
$\mathbb{1}_{Fin.Crisis}$			0.09		0.14*
			(1.25)		(1.76)
Credit Spread $\times$ $\mathbb{1}_{Fin.Crisis}$			0.21*		0.16**
			(1.78)		(2.03)
Entity FE				yes	yes
Product FE				yes	yes
Time Series Controls		yes	yes	yes	yes
Adj R-sq (Within)	0.302	0.404	0.387	0.397	0.415
Observations	53	52	52	12927	12927

Panel C: Fixed-Term Annuities

	(1)	(2)	(3)	(4)	(5)
Credit Spread	-0.49***	-0.34***	-0.39**	-0.36***	-0.45***
	(-4.76)	(-3.46)	(-2.32)	(-4.24)	(-4.75)
$\mathbb{1}_{Fin.Crisis}$			0.38**		0.33***
			(2.35)		(3.14)
Credit Spread $\times$ $\mathbb{1}_{Fin.Crisis}$			0.05		0.12
			(0.41)		(1.18)
Entity FE				yes	yes
Product FE				yes	yes
Time Series Controls		yes	yes	yes	yes
Adj R-sq (Within)	0.343	0.657	0.662	0.373	0.383
Observations	45	44	44	2247	2247

Table 1.16: Investment returns drive the cross section of premiums: P&C Insurance

This table shows the relation between quarterly returns to P&C insurance underwriting and firm-specific expected investment returns. It reports the parameter estimate from the following panel regression:

$$u_{it} = \beta_y \cdot y_{it} + \beta_{yFC} \cdot y_{it} \times \mathbb{1}_{FC} + B' \cdot X_{it-1} + FE_i + FE_t + \epsilon_{it}$$

where u_{it} is the underwriting profitability for insurer i at time t, and y_{it} is the insurer's investment return. We additionally control for date fixed effects, firm fixed effects and X_{it}, which is a vector of lagged variables that capture balance sheet strength (leverage, risk-based capital, asset growth and unearned premiums). This includes variables squared to control for non-linear effects of capital constraints. The samples consist of quarterly observations from March 2001 through March 2018. In columns 4-5 we interact investment return with an indicator variable $\mathbb{1}_{FC}$ set equal to one during the financial crisis (Q4 2008 through Q1 2010). t-statistics are reported in bracket and calculated using standard errors clustered by date and firm. *, **, and *** indicate statistical significance at the 10%, 5% and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
Investment Return	-0.10**	-0.12***	-0.11***	-0.13***	-0.12***
	(-2.37)	(-3.09)	(-5.19)	(-3.37)	(-5.70)
Size (t-1)		-0.07***		-0.07***	
		(-2.94)		(-2.93)	
Reinsurance Activity (t-1)		-0.22*		-0.22*	
		(-1.95)		(-1.95)	
Reinsurance Activity (t-1)		0.00			
		(.)			
Risk Based Capital (t-1)		0.41***		0.41***	
		(6.38)		(6.37)	
Asset Growth (t-1)		0.01***		0.01***	
		(4.16)		(4.15)	
Unearned Premia (t-1)		-0.01		-0.01	
		(-0.15)		(-0.15)	
(Risk Based Capital)^2		-0.01*		-0.01*	
		(-1.69)		(-1.68)	
(Asset Growth)^2		0.00**		0.00**	
		(2.26)		(2.27)	
(Leverage)^2		-0.00		-0.00	
		(-1.00)		(-0.99)	
Investment Return × FC				0.10	0.12**
				(1.52)	(2.56)
Reinsurance Activity (t-1)				0.00	
				(.)	
Firm Controls Vector		yes		yes	
Entity FE			yes		yes
Time FE	yes	yes	yes	yes	yes
Adj R-sq (Within)	0.001	0.071	0.001	0.071	0.001
Observations	37,044	37,044	37,044	37,044	37,044

Chapter 2

A Stock Return Decomposition Using Observables

Benjamin Knox[1] and Annette Vissing-Jorgensen[2]

We propose a new decomposition approach for stock returns that is based on the sensitivity of the stock price with respect to expected returns and dividends at various horizons. The decomposition does not rely on log-linearization or VAR estimation, and can be implemented at a daily frequency using observable data on the term structure of real rates, the Martin (2017) lower bound of equity risk premia, and dividend futures. We apply our approach to shed light on the evolution of the return on US stocks during the COVID crisis in 2020. The equity risk premium increased sharply in the near term as the crisis intensified in March, contributing 14 percent of the 26 percent market decline up to March 18. The market recovery was heavily influenced by declining real rates even at long maturities, with lower real rates contributing a positive 18 percent to the realized stock return for the year. News about dividends out to 10 years had a modest effect with a larger role for a decline and subsequent recovery of expectations for more distant dividends.

2.1 Introduction

A central theme in asset pricing is what news drives fluctuations in asset prices. The standard approach to assessing this is to exploit the Campbell and Shiller (1988) decomposition of unexpected returns into cash flow news and discount rate news. This decomposition is commonly implemented by estimating a vector-autoregressive (VAR) model that includes realized equity returns and predictors of equity returns. A problem with this approach is that results tend to be sensitive to which predictors are included, as shown by Chen and Zhao (2009). Any misspecification of the process for expected returns results in imprecise estimates of not only discount rate news but also cash flow news since the latter is calculated as a residual.

To overcome these issues, we argue that one can get a long way towards a decomposition of unexpected returns into cash flow news and discount rate news without making assumptions about return predictors and without estimating a VAR. The stock price is the present value of expected dividends discounted using the expected return on stocks which in turn equals the real riskless rate plus the equity risk premium. Therefore, in order to decompose unexpected returns into riskless rate news, risk premium news and cash flow news, one needs data on the evolution of the term structures of the real riskless rate, the equity risk premium, and expected dividends.

A lot of information is available about each of these inputs. The term structure of the real riskless rate can be measured out to around 30 years from data on either nominal Treasuries and inflation swaps, or data on inflation-indexed Treasuries (TIPS). The term structure of the equity risk premium is not directly observable but Martin (2017) provides a lower bound on the equity risk premium based on S&P500 index options. He argues that this lower bound is approximately tight and thus is close to the actual equity risk premium.[3] While Martin studies the equity risk premium out to 1 year, this can be extended out to around 2 years in recent years, based on available S&P500 options. If fluctuations in the equity risk premium are concentrated at the short end of the term structure, we can

---

[3]We supplement Martin's analysis with theoretical analysis of how the *change* in the Martin lower bound relates to the true change in the equity risk premium. In particular, we show that for the CRRA log-normal case, the same parameters that ensure that the lower bound is in fact a lower bound (Martin's negative correlation condition) also ensure that the change in the lower bound is smaller than the change in the true risk premium. This suggests that our approach will understate the role of risk premium changes for realized returns to the extent that Martin's lower bound is not tight.

estimate most of the risk premium news based on S&P500 options (with less transitory fluctuations, one can combine this with assumptions about the speed of mean-reversion in risk-premia past year 2). Finally, some information about expected dividends is available from dividend futures, available out to 10 years. The residual unexpected return not explained by any of the measured components will then capture news about dividends past year 10, as well as any news about the real riskless rate or equity risk premium past the horizons stated.

To implement this idea, we derive a new decomposition of returns that maps more directly to available data than the Campbell-Shiller decomposition (and avoids log-linearization). Result 1 shows that the effect of an instantaneous change to the expected return for year $t + k$ on today's stock price can be expressed as a function of one minus the fraction of the stock price paid for dividends out to year k.[4] Result 2 shows that effect of an instantaneous change to expected dividend for year $t + k$ on today's stock price. Result 3 combines the above to decompose realized returns into its expected component and the three unexpected components: real risk-free rate news, risk premia news and dividend news.

We use our approach to understand the evolution of the stock market over the COVID crisis in 2020. We provide a decomposition of daily returns and document the cumulative series for each of the return components over the year. The evolution of the US stock market during the COVID crisis in 2020 has been dramatic. Figure 2.1 graphs the cumulative return on the S&P500 index over the year 2020. The market fell 31 percent from January 1 to March 23, before rebounding sharply. It had full recovered by June 8 and ended the year with a 16 percent annual realized return. Figure 2.1 also graphs the cumulative return of the contributors of stock market return as set out in Result 3. While the financial press has covered an apparent disconnect between the recovery of the stock market against the continued struggles of the real economy in 2020, our decomposition can go a long way to explaining the realised stock market returns.

The decomposition reveals three key facts. First, the equity risk premium increases sharply until March 18 and had a substantial role in the market crash. We estimate that from the start of the year up to March 18, the equity risk premium for the one-year

---

[4]We map the fraction of the stock price paid for dividends out to year k to dividend futures. Past the 10-year point, we assume a Gordon growth model and constant expected growth of dividends to estimate the fraction of the stock price paid for dividends out to year $10 + k$.

horizon increased from 2.6% to 15.6%, with further increases in the year-2 risk premium. Together, the increase in the risk premium for the first two years contributed a minus 14.3 percent effect on the stock return up to March 18. A downward sloping equity risk premium During the recovery period, these risk premia decline quickly. An "A-shaped" pattern for the equity risk premium thus helps explain the V-shaped pattern of the stock price. Equity premia remain somewhat higher at the end of 2020 than at the start of the year.

Second, with the exception of an upward spike in long rates from March 9-18, real riskless rates drop dramatically across all maturities and do not recover by the end of the year. The 10-year real riskless rate declines over 100 bps over the year and real forward rates fall even out to the 30-year horizon. The forward real rate from year 21 to 30 drops over 50 bps. For the year 2020, the decline in the term structure of real rates out to year 30 contributes a 18.3 percent increase in the stock market. Evidence from 50-year UK inflation-linked bonds suggests that real rates fell even beyond year 30.

Third, changes to expected dividends out to year 10 have a modest effect on the market, contributing minus 2.5 percentage to the stock return over the year and never more than minus 4.5 percent during the year. This is unsurprising given that the first decade's dividends generally contribute only about 1/5 of the value of the stock market. More interestingly, we can get a sense of how important changes to expected dividends past year 10 were as these will drive the residual component in our return decomposition after accounting for the expected return, the riskfree rate news component, the equity risk premium news component and the effects of news about dividends out to year 10. We estimate that the more distant dividends contributed about 20 percentage point of the stock market crash but that this effect fully reverted by the end of the year. About 7 percentage points of the reversion occurred in early November following the presidential election and the news about the BioNTech/Pfizer vaccine.

Aside form its link to the long literature on stock-return decomposition,[5] our paper is related to an evolving literature on the stock market during the COVID crash and recovery. Several papers have constructed measures of the cash flow impact and argued that it is difficult to explain the sharp decline in the market in March. Landier and Thesmar (2020) analyze analyst earnings forecasts (up to May 2020). They document that downward revisions occurred smoothly and affected mainly earnings estimates for

---

[5]for seminal work see Campbell (1991) Campbell and Ammer (1993) Vuolteenaho and Campbell (2004)

2020-2022, with longer-term forecasts remaining stable. Gradual and modest reductions of earnings expectations are inconsistent with the sharp market decline and recovery. Cox et al. (2020) studies the COVID crisis using the estimated structural asset pricing model of Greenwald et al. (2019). They conclude that it is difficult to explain the V-shaped trajectory of the stock market over the COVID crisis with plausible fluctuations in economic activity, corporate profit shares, or short-term interest rates. A central input to their estimation is that, based on data from the Survey of Professional Forecasters, the COVID shock was expected to be quite transitory based on GDP growth forecasts for 2020:Q2 and 2020:Q3. Gormsen and Koijen (2020) study dividend futures. They show that to explain the decline in the stock market from February 12 to March 12, the value of dividends past year 10 must have declined substantially. Furthermore, during the recovery period up to July 20, they show that the value of near-term (up to year 10) dividends do not recover, implying that price recovery must be due to recovery in the value of distant (past year 10) dividends. Our contributions compared to this literature is to provide a simple return decomposition framework that allows for quantification of each of the components of realized returns using observable data. Compared to Landier and Thesmar (2020) we take the complementary approach of focusing on measuring discount rate news rather than cash flow news. Relative to Cox et al. (2020) we avoid the need for a structural model by arguing that many of the inputs to a return decomposition can be estimated directly from available data. Our focus on measuring discount rate news supplements Gormsen and Koijen (2020) in that discount rate news drives the changes in the value of dividends they document. Consistent evidence is also found in Gormsen et al. (2021).

In recent years, survey data on the subjective expectations of investors have been used to revisit stock-return decomposition questions. Contrary to the previous consensus Cochrane (2011) that discount rates movements primarily stock market volatility, both Bordalo et al. (2020) and De La O and Myers (2021) find evidence that variation in cashflow news is instead the principle driver of stock movements. Our results highlight an important role of discount rates during 2020, which supports the more traditional view of stock decompositions. Dahlquist and Ibert (2021) also find consistent results using subjective survey expectations. Using the long-term return expectations of asset management firms, they show expected equity premium adjusted upward by 2.4 percentage points in March, before quickly reversing as equity markets recovered. Our option-implied estimates of

93

equity risk premium in the COVID crisis are qualitatively similar.

2.2 A new stock return decomposition

We derive a new result for how changes in expected stock returns affect the stock price. This result enables a simple decomposition of returns into riskfree rate news, risk premium news and cash flow news components. We compare the new decomposition to the Campbill-Shiller decomposition and argue that the former maps directly to available data and may be more accurate since it does rely on log-linearization around the historical average of the log dividend-price ratio.

2.2.1 The effect of expected return and expected dividend changes on the stock price

Start from the present value formula, valuing the stock market as the present value of expected dividends

$$P_t = \sum_{n=1}^{\infty} \frac{E_t\left[D_{t+n}\right]}{1 + R_t^{(n)}}$$

where

$$1 + R_t^{(n)} = E_t \prod_{k=1}^{n} \left(1 + R_{t+k}\right)$$

and time is in years. R_{t+k} is the 1-year discount rate for year $t+k$, i.e., from date $t+k-1$ to date $t+k$. Unless otherwise noted, all dividends and returns are in real terms. We then characterize the effect of expected return changes on stock returns as follows.

Result 1 (expected return news and stock returns).
If $E_t \prod_{k=1}^{n} \left[1 + R_{t+k}\right] = \prod_{k=1}^{n} E_t \left[1 + R_{t+k}\right]$, then the effect of an instantaneous change to the expected return for year $t+k$ on the stock return can be expressed as:

$$\frac{\partial P_t / P_t}{\partial E_t R_{t+k}} \simeq -\frac{1}{E_t\left[1 + R_{t+k}\right]} \left(1 - \sum_{n=0}^{k-1} w_{n,t}\right) \tag{2.1}$$

where

$$w_{n,t} = \frac{E_t\left[D_{t+n}\right] / \left(1 + R_t^{(n)}\right)}{P_t} = \frac{F_{n,t} / \left(1 + y_{n,t}^{nom}\right)^n}{P_t} \quad \text{for all } n \tag{2.2}$$

with $F_{n,t}$ denoting the date t price of a dividend future paying the nominal dividends for year $t+n$ at $t+n$ and $y_{n,t}^{nom}$ is the (annualized) riskless nominal yield at date t for a n-year investment.

94

Proof: See appendix.

Result 1 is related to the standard bond pricing formula that relates bond price changes to duration and yield changes. However, in Result 1, expected returns are allowed to differ across future years and we derive the effect of a change to the expected return for one future year. To see the intuition, consider a change in the expected return for year $t + k$, $E_t R_{t+k}$, of one percentage point. With the higher discount rate for year $t+k$, all dividends to be paid at $t + k$ or later will now be discounted by one percentage point more when we discount back from $t + k$ to $t + k - 1$. Therefore, if there were no dividends before date $t + k$, then $\frac{\partial P_t / P_t}{\partial E_t R_{t+k}}$ would simply be -1 (ignoring the term $\frac{1}{E_t[1+R_{t+k}]}$). However, if there are dividends before date $t + k$, their present value is unaffected by the change in the expected return for year $t+k$, resulting in a smaller effect of $E_t R_{t+k}$ on P_t. The factor $(1 - w_{1,t} - ... - w_{k-1,t})$ captures the fraction of today's price P_t that is due to dividends at date $t + k$ and later.

The numerator $E_t D_{t+n} / \left(1 + R_t^{(n)}\right)$ in the weight $w_{n,t}$ is the price at t of a dividend strip paying D_{t+n} at $t + n$. It is well known that dividend strips (which are not traded) can be valued from dividend futures (e.g. van Binsbergen et al. (2013)). Since dividend futures pay off at maturity $(t + n)$, dividend strips and dividend futures prices are related by $E_t [D_{t+n}] / \left(1 + R_t^{(n)}\right) = F_{n,t} / \left(1 + y_{t,n}^{\text{nom}}\right)^n$. In this expression, $F_{n,t}$ is nominal (since actual dividend futures contracts pay the nominal dividend) and therefore discounted using the nominal yield.

The assumption used to derive Result 1 , $E_t \prod_{k=1}^n [1 + R_{t+k}] = \prod_{k=1}^n E_t [1 + R_{t+k}]$, states that realized returns are independent, conditional on information known at date t. Importantly, this does not rule out time-variation in expected returns and expected returns for different maturities can update in a correlated fashion. What needs to hold is that *realized* returns in one year are not informative for realized returns in another year, conditional on what is known at t. For example, $E_t [(1 + R_{t+1}) (1 + R_{t+2})] = E_t [1 + R_{t+1}] E_t [1 + R_{t+2}] + cov_t (R_{t+1}, R_{t+2})$.[6] Thus, the assumption holds for horizon

---

[6]

$$
\begin{aligned}
E_t [(1 + R_{t+1}) (1 + R_{t+2})] &= 1 + E_t [R_{t+1}] + E_t [R_{t+2}] + E_t [R_{t+1} R_{t+2}] \\
&= 1 + E_t [R_{t+1}] + E_t [R_{t+2}] + E_t [R_{t+1}] E_t [R_{t+2}] + cov_t [R_{t+1}, R_{t+2}] \\
&= E_t [1 + R_{t+1}] E_t [1 + R_{t+2}] + cov_t [R_{t+1}, R_{t+2}]
\end{aligned}
$$

n=2 if $cov_t(R_{t+1}, R_{t+2}) = 0$, i.e., if the distance of R_{t+1} from its conditional mean is uninformative for the distance of R_{t+2} from its conditional mean.

We next characterize the effect of expected dividend changes on stock returns as follows.

Result 2 (dividend news and stock returns).

The effect of an instantaneous change to the expected dividend for year $t + k$ on the stock return can be expressed as:

$$\frac{\partial P_t / P_t}{\partial E_t D_{t+k}} = \frac{1/P_t}{1 + R_t^{(k)}} \tag{2.3}$$

A unit change in expected dividend $t + k$ has less impact on the aggregate stock market return as the maturity k of the dividend increases. This is because later cash flows are discounted more, and thus contribute less weight to the overall value of the stock market.

2.2.2 The stock return decomposition

The expected stock return for year $t + k$ can be expressed as

$$E_t R_{t+k} = f_{t+k} + ep_{t+k} \tag{2.4}$$

where f_{t+k} denotes the forward rate for a riskless 1-year investment in year $t + k$ and ep_{t+k} denotes the equity risk premium for year $t + k$. Result 1 holds whether changes to $E_t R_{t+k}$ are due to changes in f_{t+k} or ep_{t+k}. Therefore, we can use Results 1 and 2 to decompose the capital gain over a one-day period (assumed short enough that the first-order approximation is accurate) as follows

$$\text{Realized capital gain}_d - \text{Expected capital gain}_d \tag{2.5}$$

$$= \underbrace{\Sigma_{k=1}^{\infty} \frac{\partial P_t / P_t}{\partial E_t R_{t+k}} \partial f_{t+k}}_{\text{Riskfree rate news}_d} + \underbrace{\Sigma_{k=1}^{\infty} \frac{\partial P_t / P_t}{\partial E_t R_{t+k}} \partial ep_{t+k}}_{\text{Risk premium news}_d} + \underbrace{\Sigma_{k=1}^{\infty} \frac{\partial P_t / P_t}{\partial E_t D_{t+k}} \partial E_t D_{t+k}}_{\text{Cash flow news}_d}$$

The realized return on day d is then

$$\text{Realized return}_d = \text{Realized capital gain}_d + \text{Realized dividend yield}_d$$

We then have the following decomposition.

Result 3 (Realised return decomposition).

$$\text{Realized return}_d \tag{2.6}$$

$$= \text{Expected return}_d + \underbrace{\Sigma_{k=1}^{\infty}\frac{\partial P_t/P_t}{\partial E_t R_{t+k}}\partial f_{t+k}}_{\text{Riskfree rate news}_d} + \underbrace{\Sigma_{k=1}^{\infty}\frac{\partial P_t/P_t}{\partial E_t R_{t+k}}\partial ep_{t+k}}_{\text{Risk premium news}_d} + \underbrace{\Sigma_{k=1}^{\infty}\frac{\partial P_t/P_t}{\partial E_t D_{t+k}}\partial E_t D_{t+k}}_{\text{Cash flow news}_d}$$

In our application to 2020, we implement the return decomposition for each day of 2020 and then aggregate each component across days.

2.2.3 Comparison to the Campbell-Shiller decomposition

Let $r_t = \ln(1 + R_t)$, $p_t = \ln P_t$, $d_t = \ln D_t$

$$
\begin{aligned}
r_{t+1} &= \ln(P_{t+1} + D_{t+1}) - \ln P_t \\
&= \ln\left(P_{t+1}\left[1 + \frac{D_{t+1}}{P_{t+1}}\right]\right) - \ln P_t \\
&= p_{t+1} - p_t + \ln(1 + \exp(d_{t+1} - p_{t+1}))
\end{aligned}
$$

The Campbell-Shiller decomposition is based on a first-order Taylor approximation of $\ln(1 + \exp(d_{t+1} - p_{t+1}))$ around the historical average value of $d-p$ (denote this by $\overline{d-p}$)

$$
\begin{aligned}
&\ln(1 + \exp(d_{t+1} - p_{t+1})) \\
&= \ln\left(1 + \exp\left(\overline{d-p}\right)\right) + \frac{\exp\left(\overline{d-p}\right)}{1 + \exp\left(\overline{d-p}\right)}\left[d_{t+1} - p_{t+1} - \left(\overline{d-p}\right)\right] \\
&= k + (1 - \rho)\left[d_{t+1} - p_{t+1}\right]
\end{aligned}
$$

with

$$
\begin{aligned}
\rho &= \frac{1}{1 + \exp\left(\overline{d-p}\right)} \\
k &= -\ln\rho - (1 - \rho)\ln\left(\frac{1}{\rho} - 1\right)
\end{aligned}
$$

This implies,

$$r_{t+1} = p_{t+1} - p_t + k + (1 - \rho)\left[d_{t+1} - p_{t+1}\right] \implies p_t = \rho p_{t+1} + k + (1 - \rho)d_{t+1} - r_{t+1}$$

Iterating forward

$$p_t = \frac{k}{1 - \rho} + \Sigma_{j\geq 0}\rho^j(1 - \rho)d_{t+1+j} - \Sigma_{j\geq 0}\rho^j r_{t+1+j}$$

Therefore,

$$\frac{\partial p_t}{\partial E_t r_{t+k}} = -\rho^{k-1} \tag{2.7}$$

(2.7) is close to our Result 1 since

$$\frac{\partial P_t / P_t}{\partial E_t R_{t+k}} = \frac{\partial P_t / P_t}{\partial \ln\left(1 + E_t R_{t+k}\right)} \frac{\partial \ln\left(1 + E_t R_{t+k}\right)}{\partial E_t R_{t+k}} = \frac{\partial P_t / P_t}{\partial \ln\left(1 + E_t R_{t+k}\right)} \frac{1}{\left(1 + E_t R_{t+k}\right)}.$$

Result 1 therefore implies $\frac{\partial P_t / P_t}{\partial \ln(1 + E_t R_{t+k})} \simeq -\left(1 - \sum_{n=0}^{k-1} w_{n,t}\right)$. Like the w_t weights in Result 1, the $\rho < 1$ in the Campbell and Shiller (1988) (and Campbell (1991)) approach captures the fact that the price effect of changes in expected returns in a later period are smaller the more dividends are received before that period. Our Result 1 makes this more transparent than the Campbell-Shiller approach. Furthermore, the w_t weights map directly to dividend futures at t as we have laid out, whereas ρ in the Campbell-Shiller approach is a historical average. Gao and Martin (2020) argue that the Campbell-Shiller log-linearization can be inaccurate when the log price-dividend ratio is far from its historical average.[7] This issue may be particularly relevant for the year 2020 and the years to come given the COVID recession. One could consider a version of the Campbell-Shiller approach in which the $\ln\left(1 + \exp\left(d_{t+k} - p_{t+k}\right)\right)$ was log-linearized around $E_t\left(d_{t+k} - p_{t+k}\right)$. Then

$$\frac{dp_t}{dE_t r_{t+k}} = -\rho_{t+1}\rho_{t+2}...\rho_{t+k-1}$$

with

$$\rho_{t+1} = \frac{1}{1 + \exp\left(E_t\left(d_{t+1} - p_{t+1}\right)\right)}, \quad \rho_{t+2} = \frac{1}{1 + \exp\left(E_t\left(d_{t+2} - p_{t+2}\right)\right)} \text{ etc.}$$

This would be more accurate than the standard Campbell-Shiller approach but because $d_{t+1} - p_{t+1}$ in ρ_{t+1} is in logs, d_{t+1} does not map directly to dividend futures. Furthermore, p_{t+1} is a future price so implementing $E_t\left(d_{t+1} - p_{t+1}\right)$ would require assumptions about price expectations (similarly for ρ_{t+2} etc.).[8]

2.3   Implementation

2.3.1   Data

We implement the stock return decomposition in Result 3 for the S&P500 for year 2020. We use forward rates f_{t+k} calculated from zero-coupon Treasury yields obtained from the

---

[7]Gao and Martin (2019) propose a different log-approach and use it to understand the expected growth rate investors must have to be happy hold the market at a given point in time.

[8]One would also need a lower bound for expected log returns, as opposed to expected returns, but that is possible based on Gao and Martin (2020).

Federal Reserve and inflation swaps from Bloomberg.[9] For ep_{t+k} we use the methodology of Martin (2017) who calculates a lower bound on the risk premium using prices of stock market index options and argues that this lower bound is approximately equal to the true risk premium. Martin's data covers the period 1996-2012. We extend his series to 2020 using data from OptionMetrics for 2013-2019 and from the CBOE for 2020. We are able to almost exactly replicate Martin's series over his sample period. Appendix B details our data construction.

For the cash flow news component of stock returns, expected dividends are extracted from dividend futures

$$E_t D_{t+k} = \frac{1 + R_t^{(k)}}{\left(1 + y_{n,t}^{\text{nom}}\right)^n} F_{t+k} \tag{2.8}$$

where discount rates $R_t^{(k)}$ are implied from the risk-free and risk premium data described above. Dividend futures are obtained from Bloomberg and exist out to the 10 year maturity. The effect of changes in expectations of long-dated dividends are therefore not captured by observables and the cash flow news past year 10 is instead calculated as the residual in the stock return decomposition.

In our baseline estimation we assume that riskless forward rates do not change past year 30 and that the equity risk premium does not change past year 5. This is motivated by data availability (more on this below), but we will argue empirically that these horizons will allow us to capture the majority or discount rate news given actual mean-reversion of riskless rates and equity risk premia. Any changes to real rates past year 30 or risk premia past year 5 will also enter the residual in the decomposition.

2.3.2 Dividend Weights

To calculate the effect of changing discount rates using Result 1, we need daily estimates of $E_t [1 + R_{t+k}]$, and dividend weights, $w_{k,t}$, out to $k = 30$ years. For the discount rates, we use the data for riskless rates and the equity premium as stated, assuming that ep_{t+k} has mean-reverted to equal its pre-crisis average past $k = 5$. For the weights, we can calculate $w_{1,t},...w_{10,t}$ from (2.2) using available data on dividend futures, zero-coupon Treasury yields and the stock price. To estimate dividend weights past 10 years, we assume a Gordon growth model for dividends past 10 years. The value of long-term dividends is

---

[9]https://www.federalreserve.gov/data/nominal-yield-curve.htm

then

$$L_t = \sum_{k=11}^{\infty} \frac{E_t D_{t+k}}{1 + R_t^{(k)}}$$

$$= \frac{E_t D_{t+10}}{1 + R_t^{(10)}} \left(\frac{1 + g_{11}}{1 + R_{t+11}} + \frac{(1 + g_{11})(1 + g_{12})}{(1 + R_{t+11})(1 + R_{t+12})} + \ldots \right)$$

$$= \frac{E_t D_{t+10}}{1 + R_t^{(10)}} \left(\frac{1 + g}{R - g} \right)$$

where g is a constant dividend growth rate and R is a constant forward discount rate. Rearranging this equation, the growth rate of dividends past year 10 is

$$g = \frac{R - x}{1 + x}$$

where

$$x = \frac{F_{n,t} / \left(1 + y_{n,t}^{\mathrm{nom}}\right)^n}{L_t}$$

is the ratio of the 10 year dividend strip value to the sum of the long-term dividend values. To estimate g, we set the constant forward discount rate equal to the observed year 11 (real) forward discount rate, $R = R_{t+11}$, and compute the value of long-term dividends as the difference between aggregated stock price and the sum of dividend prices up to 10 years

$$L_t = P_t - \sum_{k=1}^{10} \frac{F_{t+k}}{1 + y_{k,t}^{\mathrm{nom}}}.$$

Given the above estimated growth rate in dividends past year 10, dividend weights are

$$w_k = w_{10} \left(\frac{1 + g}{1 + R} \right)^{k-10} \qquad \text{for } k > 10.$$

Figure 2.2 plots the cumulative weight of dividends (averaged over the year) at each dividend maturity. Dividends up to 10 (30) years have averaged a combined weight of 17% (44%) of the total stock market value. These weights are very similar to those extracted by van Binsbergen (2020). We update the weight schedule daily when implementing the stock return decomposition.

2.3.3   Estimating changes to equity risk premia

Martin (2017) starts from the fact that the time t price of a claim to a cash flow X_T at time T can either be expressed using the stochastic discount factor M_T as

$$\mathrm{Price}_t = E_t \left(M_T X_T \right)$$

100

or using risk-neutral notation as

$$\text{Price}_t = \frac{1}{R_{f,t}} E_t^* (X_T)$$

where the expectation E_t^* is defined by

$$E_t^* (X_T) = E_t (R_{f,t} M_T X_T).$$

The return on an investment can similarly be written in terms of the SDF or using risk-neutral notation

$$
\begin{aligned}
1 &= E_t (M_T R_T) \\
&= \frac{1}{R_{f,t}} E_t (R_{f,t} M_T R_T) \\
&= \frac{1}{R_{f,t}} E_t^* (R_T).
\end{aligned}
$$

The conditional risk-neutral variance can be expressed as

$$\text{var}_t^* R_T = E_t^* R_T^2 - (E_t^* R_T)^2 = R_{f,t} E_t (M_T R_T^2) - R_{f,t}^2$$

The risk premium expressed as a function of the risk-neutral variance is

$$
\begin{aligned}
E_t R_T - R_{f,t} &= \left[E_t (M_T R_T^2) - R_{f,t} \right] - \left[E_t (M_T R_T^2) - E_t R_T \right] \\
&= \frac{1}{R_{f,t}} \text{var}_t^* R_T - cov_t (M_T R_T, R_T) \\
&\geq \frac{1}{R_{f,t}} \text{var}_t^* R_T \text{ if } cov_t (M_T R_T, R_T) \leq 0
\end{aligned}
$$

Thus $\frac{1}{R_{f,t}} \text{var}_t^* R_T$ provides a lower bound on $E_t R_T - R_{f,t}$ if $cov_t (M_T R_T, R_T) \leq 0$, denoted the "negative correlation condition" (NCC).

Martin (2017) shows that the lower bound $\frac{1}{R_{f,t}} \text{var}_t^* R_T$ can be calculated from put and call option as follows

$$\frac{1}{R_{f,t}} \text{var}_t^* R_T = \frac{2}{S_t^2} \left[\int_0^{F_{t,T}} \text{put}_{t,T} (K) \, dK + \int_{F_{t,T}}^{\infty} \text{call}_{t,T} (K) \, dK \right]$$

where S_t is the stock price at t, $F_{t,T} = E_t^* (S_T)$ is the forward stock price, and K denotes the strike price. On any date, it is therefore possible to extract a lower bound estimate for each available maturity of expiring options. Consistent with Martin (2017), we use linear interpolation to calculate constant maturity lower bounds, which post 2006 allows estimates out to two years and 6 months.

To account for changes in (forward) equity risk premia past year two, we first run factor analysis on the constant maturity 1, 2, 3, 6, 12, 18, 24 and 30 month equity risk premia, extracting the first two factors and also the corresponding factor loadings. We then fit the factor loadings as a function of maturity. Guided by the data, for the first factor (on which loadings are all positive) we use a Box–Cox regression, transforming the factor loading y_i and regressing it on maturity τ_i as follows

$$\frac{y_i^\lambda - 1}{\lambda} = \alpha + \beta \tau_i + \epsilon_i$$

with λ, α and β estimated by maximum likelihood. For the second factor, we estimate the following relation by nonlinear least squares

$$y_i = \alpha + \beta \left(\frac{1 - e^{-\lambda \tau_i}}{\lambda \tau_i} \right) + \epsilon_i$$

The functional form used for the second factor is the same as typically used for the slope factor in the literature on the term structure of riskless rates. With the estimated functional forms of loadings against maturity, we can predict the loadings of longer-dated unobserved risk premia, and finally estimate longer-dated risk premium themselves.

Figure 2.3 summarises the results of the above factor analysis. Row one presents the time series of the two factors. The factor analysis uses standardized inputs (mean zero, unit standard deviation risk premia). Row two presents loadings on the factors across the observed risk premium maturities (up to 2.5 years). Note that the factor loadings in row two are reminiscent of the loadings on the well known level and slope factors in the interest rate term-structure literature. All maturities load similarly (close to one) on the first *level* factor, while short (long) maturities loading positively (negatively) on the second *slope* factor. Factor loadings on the first (most persistent) factor start to fall for the highest observed maturities.

To model non-standardized risk premia, we multiply the factor loadings for a given maturity by the standard deviation of the risk premium for that maturity. These rescaled factor loadings are shown in row 3 and the above-described modeling of factor loadings is done using these as inputs. The figures in row 3 include (solid lines) the predicted values from our factor modeling. For both factors, the estimated functional forms provide a close fit. Row four presents the extrapolated factor loadings up to 10 year maturity. To avoid extrapolating far past the range of available maturities, we only use extrapolated risk premia out to year 5 in the return decomposition.

2.4 Relating the true change in the equity premium to the change in the Martin lower bound

A central element of the implementation of our return decomposition is that we use the Martin (2017) methodology to calculate equity risk premium estimates. This is an alternative to estimate a VAR to decompose discount rate news and cash flow news, an approach that is often sensitive to which variables are included and relies on the strong assumption that relations are stable over time. Since our return decomposition relies on *changes* in the risk premium, it is essential to know how the *change* in the bound relates to the *change* in the true risk premium. We therefore state the general condition for when the change in the lower bound is smaller than the true change in the equity risk premium and test this condition in data for 1996-2020 finding supportive evidence.

We supplement this empirical evidence with theoretical analysis for the log-normal case and the CRRA log-normal case. In particular, we show that for the CRRA log-normal case, the same parameters that ensure that the lower bound is in fact a lower bound (Martin's negative correlation condition) also ensure that the change in the lower bound is smaller than the change in the true risk premium. To the extent the lower bound is not right, our return decomposition will thus tend to understate the role of risk premium changes.

2.4.1 The tightness of the Martin lower bound

Martin documents an average lower bound over the 1996-2012 period of about 5%, close to the equity premium estimates obtained by Fama and French (2002) using average realized dividend (or earnings) growth rates as an estimate of ex-ante expected capital gains. Martin also tests whether the lower bound is a good predictor of the realized excess return. He estimates the relation

$$\frac{1}{T-t}\left(R_T - R_{f,t}\right) = \alpha + \beta \times \frac{1}{T-t}\frac{1}{R_{f,t}}\text{var}_t^* R_T \tag{2.9}$$

and cannot reject the null of $\beta = 1$, $\alpha = 0$ for any horizon from 1 month to 1 year.[10]

---

[10]Martin's defines a variable $SVIX_{t\to T}^2 = \frac{1}{T-t}\text{var}_t^*\left(\frac{R_T}{R_{f,t}}\right)$ and his regressor is thus expressed as $R_{f,t}SVIX_{t\to T}^2$.

2.4.2 The change in the lower bound

Suppose an underlying state variable s_t changes and that s_t is signed such that $. \frac{\partial \left[\frac{1}{R_{f,t}} \text{var}_t^* R_T \right]}{\partial s_t} > 0$. Then

$$
\begin{aligned}
\frac{\partial \left[E_t R_T - R_{f,t} \right]}{\partial s_t} &= \frac{\partial \left[\frac{1}{R_{f,t}} \text{var}_t^* R_T \right]}{\partial s_t} - \frac{\partial cov_t \left(M_T R_T, R_T \right)}{\partial s_t} \\
&\geq \frac{\partial \left[\frac{1}{R_{f,t}} \text{var}_t^* R_T \right]}{\partial s_t} \text{ iff } \frac{\partial cov_t \left(M_T R_T, R_T \right)}{\partial s_t} \leq 0
\end{aligned}
$$

It follows that the change in the lower bound is, on average, equal to the true change in the risk premium if the regression coefficient β in (2.9) equals one. If instead $\beta > 1$ that would suggest that the regressor is positively correlated with the omitted variable $-cov_t \left(M_T R_T, R_T \right)$ implying that $\frac{\partial cov_t (M_T R_T, R_T)}{\partial s_t} < 0$ and thus that the true change in the risk premium is larger than the change in the lower bound. Martin's regression coefficients are below one for the shortest horizon (1 month) but above one for the 6 month and 1-year horizons, but standard errors are large. We expand Martin's the sample from 1996-2012 to 1996-2020.

Over the longer sample, we find that β is higher than one for most horizons, though not significantly so (see Table 2.1). We cannot reject that changes in the lower bound are unbiased estimates of changes in the true risk premium. The β estimates above one imply that the true risk premium change exceeds that of the change in the lower bound. It is possible, however, that realized excess returns exceeded expected returns over this particular time period, more so in times of stress (high values of the risk-neutral variance). Fama and French (2002) argue that realized returns exceeded expected returns even over a sample as long as 1951-2000. Cieslak, Morse and Vissing-Jorgensen (2019) argue that over the post-1994 period, unexpectedly accommodating monetary policy has contributed to much of the realized excess return on the stock market. If the unexpected positive component of realized returns is sufficiently correlated with risk-neutral variance, then an estimated β above one may not imply that changes in the lower bound are conservative. Given the lack of conclusive empirical evidence on this issue, we next consider what theory says about whether it is likely that $\frac{\partial cov_t (M_T R_T, R_T)}{\partial s_t} \leq 0$.

2.4.3 The log-normal case

Assume conditional log-normality as follows:

$$M_T = e^{-r_{f,t} + \sigma_{M,t} Z_{M,T} - \frac{1}{2}\sigma^2_{M,t}}$$

$$R_T = e^{\mu_{R,t} + \sigma_{R,t} Z_{R,T} - \frac{1}{2}\sigma^2_{R,t}}$$

where $Z_{M,t}$ and $R_{M,t}$ are (potentially correlated) standard normal random variables.

$E_t(M_T R_T) = 1$ implies that

$$cov_t(M_T R_T, R_T) = E_t(M_T R_T^2) - E(R_T)$$

Consider each term on the right hand side separately.

$$
\begin{aligned}
\ln E_t\left(M_T R_T^2\right) &= E_t\left(\ln M_T + 2\ln R_T\right) + \frac{1}{2}V_t\left(\ln M_T + 2\ln R_T\right) \\
&= -r_{f,t} - \frac{1}{2}\sigma^2_{M,t} + 2\left(\mu_{R,t} - \frac{1}{2}\sigma^2_{R,t}\right) + \frac{1}{2}\left(\sigma^2_{M,t} + 4\sigma^2_{R,t} - 4\left(\mu_{R,t} - r_{f,t}\right)\right) \\
&= r_{f,t} + \sigma^2_{R,t}
\end{aligned}
$$

$$
\begin{aligned}
\ln E_t\left(R_T\right) &= E_t\left(\ln R_T\right) + \frac{1}{2}V_t\left(\ln R_T\right) \\
&= \mu_{R,t}
\end{aligned}
$$

Combining these two expressions

$$cov_t\left(M_T R_T, R_T\right) = e^{r_{f,t} + \sigma^2_{R,t}} - e^{\mu_{R,t}}$$

Therefore, the NCC holds iff the conditional Sharpe ratio exceeds the conditional standard deviation:

$$cov_t\left(M_T R_T, R_T\right) \le 0 \text{ iff } e^{r_{f,t} + \sigma^2_{R,t}} \le e^{\mu_{R,t}} \iff \sigma_{R,t} \le \frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}$$

$E_t(M_T R_T) = 1$ furthermore implies that

$$
\begin{aligned}
\ln\left(E_t\left(M_T R_T\right)\right) &= E_t\left(\ln M_T + \ln R_T\right) + \frac{1}{2}V_t\left(\ln M_T + \ln R_T\right) \\
&= \left(\mu_{R,t} - r_{f,t} - \frac{1}{2}\sigma^2_{M,t} - \frac{1}{2}\sigma^2_{R,t}\right) + \frac{1}{2}\left(\sigma^2_{M,t} + \sigma^2_{R,t} + 2cov_t\left(\ln R_T, \ln M_T\right)\right) \\
&= \mu_{R,t} - r_{f,t} + cov_t\left(\ln R_T, \ln M_T\right) = 0
\end{aligned}
$$

and thus

$$\mu_{R,t} - r_{f,t} = -cov_t\left(\ln R_T, \ln M_T\right) \tag{2.10}$$

The above results for the log-normal case are known from Martin (2017). The following result adds conditions that relate the true change in the risk premium to the change in the lower bound.

Result 4: The true change in the risk premium is larger than the change in the lower bound iff $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$. Under log-normality, it is sufficient for $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$ that

(1) The NCC holds: $cov_t(M_T R_T, R_T) \leq 0 \iff \frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} \geq \sigma_{R,t}$, and

(2) $\frac{\partial}{\partial s_t} \left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} \right] \geq \frac{\partial \sigma_{R,t}}{\partial s_t}$.

Proof: See appendix.

2.4.4   The CRRA log-normal case

In addition to log-normality, assume CRRA utility,

$$
\begin{aligned}
M_T &= \beta \left(\frac{C_T}{C_t} \right)^{-\gamma} \\
&= e^{\ln \beta - \gamma \ln(C_T/C_t)}
\end{aligned}
$$

with $\ln(C_T/C_t)$ is normal $\mu_{c,t}$, $\sigma_{c,t}^2$ conditional on information known at t. We can map this assumption to the above log-normal framework

$$
\begin{aligned}
M_T &= e^{-r_{f,t} + \sigma_{M,t} Z_{M,T} - \frac{1}{2} \sigma_{M,t}^2} \\
R_T &= e^{\mu_{R,t} + \sigma_{R,t} Z_{R,T} - \frac{1}{2} \sigma_{R,t}^2}
\end{aligned}
$$

where $Z_{M,t}$ and $R_{M,t}$ are (potentially correlated) standard normal random variables. To link the two assumptions for M_T, equate the two expressions:

$$
-r_{f,t} + \sigma_{M,t} Z_{M,T} - \frac{1}{2} \sigma_{M,t}^2 = \ln \beta - \gamma \ln(C_T/C_t) \tag{2.11}
$$

Calculating the variance of each side of (2.11) we get

$$
\sigma_{M,t}^2 = \gamma^2 \sigma_{c,t}^2.
$$

Taking expectations in (2.11) implies

$$
\begin{aligned}
r_{f,t} &= -\ln \beta - \frac{1}{2} \sigma_{M,t}^2 + \gamma E_t \ln(C_T/C_t) \\
&= -\ln \beta - \frac{1}{2} \gamma^2 \sigma_{c,t}^2 + \gamma E_t \ln(C_T/C_t).
\end{aligned}
$$

Thus

$$
\begin{aligned}
Z_{M,T} &= \frac{1}{\sigma_{M,t}} \left[\ln \beta - \gamma \ln \left(C_T/C_t \right) + r_{f,t} + \frac{1}{2}\sigma_{M,t}^2 \right] \\
&= \frac{1}{\gamma \sigma_{c,t}} \left[\gamma E_t \ln \left(C_T/C_t \right) - \gamma \ln \left(C_T/C_t \right) \right] \\
&= \frac{1}{\sigma_{c,t}} \left[E_t \ln \left(C_T/C_t \right) - \ln \left(C_T/C_t \right) \right]
\end{aligned}
$$

We can thus exploit (2.10) to get

$$
\begin{aligned}
\mu_{R,t} - r_{f,t} &= -cov_t \left(\ln R_T, \ln M_T \right) \\
&= \gamma cov_t \left(\ln R_T, \ln \left(C_T/C_t \right) \right).
\end{aligned}
$$

This implies,

$$
\begin{aligned}
\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} &= \gamma \frac{cov_t \left(\ln R_T, \ln \left(C_T/C_t \right) \right)}{\sigma_{R,t}^2} \sigma_{R,t} \\
&= \gamma \beta_t^C \sigma_{R,t}
\end{aligned}
$$

where β_t^C is the (potentially time-varying) beta of $\ln \left(C_T/C_t \right)$ with respect to $\ln R_T$. Thus, $\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} \geq \sigma_{R,t}$ (the NCC holds) iff $\gamma \beta_t^C \geq 1$. Furthermore,

$$
\frac{\partial}{\partial s_t} \left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} \right] = \gamma \beta_t \frac{\partial \sigma_{R,t}}{\partial s_t}
$$

This implies that

$$
\frac{\partial}{\partial s_t} \left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}} \right] \geq \frac{\partial \sigma_{R,t}}{\partial s_t} \text{ iff } \gamma \beta_t \geq 1.
$$

Therefore, the same condition that ensures the NCC holds, $\gamma \beta_t^C \geq 1$, also ensures that the true change in the risk premium is larger than the change in the lower bound. In the CRRA log-normal case, the NCC is thus sufficient to ensure that the true change in the risk premium is larger than the change in the lower bound. Martin (2017) argues that the NCC is very likely to hold in the log-normal case since the Sharpe ratio based on realized returns has substantially exceeded the realized standard deviation.

2.5   Empirical Results

2.5.1   The risk premium

Figure 2.4 shows our estimated risk premia (the lower bounds) over 2020, by maturity. All risk premia shown in the figure are annualized. The top graph shows that 5-day risk premia peaked above 100 percent in March. For a one-year horizon, the risk premium

increases from around 3 percent at the start of the year to about 15 percent on March 18. Annualized risk premia for longer horizons rise less.

As a supplementary way to describe the term structure of risk premia, Figure 2.5 graphs the cumulative equity risk premium by maturity for the beginning and end of the year as well as for March 18, they day risk premia peak. Higher annualized risk premia at shorter horizons translate in to a concave cumulative equity risk premium. At the peak of the crisis, investors required a risk premium of 4.1 percent to invest for a 30-day period and a risk premium of 15.7 percent to invest over the next year.

Figure 2.6 illustrates the time series for the (annualized) risk premia for various 6-month periods. The red line in the left graph shows the dramatic increase and subsequent reversal of the risk premium for the month 1-6 horizon. By contrast, the forward equity premium for the subsequent 6-month period increases much less during the initial months of the year, about 2 percentage points, and stays largely flat after that. The figure to the right compares the series for months 7-12 to that for months 19-24. The latter increases more gradually but also remains higher at the end of the year than at the beginning. We infer from these facts that near-term risk premia increased sharply during the COVID crisis, as they did during the financial crisis as documented in Martin (2017) but that investors standing in March expected a lot of the uncertainty generated by COVID to be resolved over the first six-month period.

Given that there is some increase in the risk premium even for months 19-24 it is likely that risk premia increase to some extent even past this horizon. As described in Section 2.3.2.3.3, we therefore use factor analysis to estimate the perceived persistence of risk premium changes from the maturity structure at a given point in time in order to account for changes to the risk premium past year 2. Figure 2.7 shows (demeaned) estimated forward risk premia out to 5 years maturity. Although longer-dated forward risk premia move less than shorter-date forward risk premium, the five year forward still increased by about 100 basis points over 2020. To avoid issues with over-extrapolation, we assume there is no change in forward risk premia at maturities past five years.

2.5.2 The real rate

Figure 2.8, top graph, shows the evolution of the 10-year and 30-year real rates estimated from nominal Treasuries and inflation swaps. The bottom graph in Figure 2.8 illustrates the nominal yields and inflation swaps. Real yields fall dramatically over the year, with

a 119 bps decline in the 10-year real rate and an 85 bps decline in the 30-year real rate. The decline is interrupted by a sharp spike in real long yields from March 9 to March 18. Vissing-Jørgensen (2020) studies this spike which led to Federal Reserve purchases of over \$1T of Treasuries in 2020Q1 in order to stabilize Treasury markets. The spike is associated with heavy selling by bond mutual funds, foreign central banks and hedge funds and reverses as the Federal Reserve increases its daily purchases sharply starting on March 19. It is possible that the March spike in real Treasury yields is disconnected from the stock market in the sense that selling pressure affected Treasury yields without changing investors' view of the fundamental value of stocks. If so, our riskfree rate news component will overestimate the negative return effect of the spike on realized stock returns and will assign too small a role to declines in dividends past year 10 in explaining the market crash. This issue will not affect our decomposition for the full year, nor our assessment of the role of the risk premium for the crash, nor our estimate of the role of the real riskless rate outside of this short period of Treasury market dislocations.

Figure 2.9 seeks to determine whether our assumption of no changes to real rates past year 30 is realistic. We graph the real (annualized) 10-year forward rates for each of the next 3 decades. based on real rates from nominal Treasuries and inflation swaps (top left) or inflation-indexed bonds (top right). The real forward rate for the 3rd decade from now falls over the year, but less than the real forward rates for the first two decades. In the top left graph we illustrate the real forward rate for year 30 separately and even that appears to decline a bit (about 30 bps). It is thus possible that real rates change somewhat even past year 30. In the UK, inflation-indexed bonds are traded with 50-year maturity and as shown in the bottom graph, the real forward rate for years 31-50 falls about 40 bps for the year.

As a robustness check, we have therefore also calculated our main results assuming the change in the year-30 forward real risk-free rate is also the change in all longer-dated forward rates. However, despite the year-30 forward real risk-free rate falling by 33 bps over the year, the effect on the stock return over the year is approximately zero. This counter-intuitive result is due to the denominator in the right hand side of Result 1 not being a constant. If the right hand side tends to be higher on days with positive changes in long real rates than on days with negative changes of long real rates, then the net effect can be small even if long real rates decline overall for the year.

2.5.3 Dividends

Figure 2.10 shows the constant maturity expected dividend 2, 5 and 8 years ahead over the course of 2020. The left hand side figure shows nominal expected dividends and the right hand side shows real expected dividends. The year-2 expected real dividend fell by 36% from January 2nd to its lowest point on 03 April. It ended the year down 8% relative to the start of the year expectation. The moves in longer term dividends follow a similar, but less dramatic pattern. As the first 10 years of dividends make up less than 20% of total stock value, even these large moves in expected dividends have a small impact on the aggregate stock return.

2.5.4 Return decomposition results

Figure 2.1 reports the main result of our return decomposition based on the above-described inputs. The upward spike in risk premia in March generates a negative realized return effect which accounts for minus 14.3 percentage point of the realized return of minus 26 percent up to March 18. Although the risk premia news effect recovers somewhat from the height on crisis, it still ends the year negative. Our baseline specification only uses observed risk premia (to 2 years maturity), and shows that the increases in risk premia over 2020 generated a negative 4 percentage point return. By extrapolating forward risk premia to 5 years, we increase the estimated effect, with risk premia changes over 2020 generating a negative 7 percentage point return.

The fall in real riskless rates up to March 9 contributes a positive effect on realized stock returns as does the fall in the real riskless rates for the year as a whole. We estimate that the decline in the real riskless rates out to year 30 generate a plus 18.3 percent return component for the stock market for the year 2020. Changes to expected dividends out to year 10 play a minor role, consistent with most of the stock market value coming from later dividends. The expected return component contributes about 6 percent to the realized return for the year. We estimate this component from the 1-year real rate and the 1-year risk premium (compounded on a daily basis).

The top plot of Figure 2.11 presents the implied return from all of our observables combined. The residual (or unexplained) component of stock market return is then presented in the bottom plot. The residual captures the effect of dividends past year 10 and any changes to risk premia past year 2 and real riskless rates past year 30. We therefore call it long-term news. The long-term news component is large, contributing about 20

percentage points to the crash and a roughly equal amount to the recovery.

2.6 Conclusion

The paper contributes to answering a core question in asset pricing: what is the role of discount rate news versus cash flow news. We develop a simple return decomposition starting from the observation that a lot can be observed about how the real riskless rate and the risk premium evolves over time, with some additional observable information dividends out to year 10. We apply the decomposition to understand the evolution of the US stock market during the COVID crisis. Our findings highlight the role of discount rate variation in driving stock market variation. In particular, volatility in *short-term* (1 year to 2 year) equity risk premia had a substantial role in the market crash and rebound in March and April, while the fall in *long-term* (1 year to 30 year) real risk-free rates over the year was a key driver of the stock market ending the year with strong positive returns.

2.7 Figures

Figure 2.1: Decomposition of the S&P500 index return, 2020. This figure shows the cumulative return of the S&P500 index in 2020, along with the return contribution from real rate news, risk premium news, dividend news and also the realisation of the expected return. All components are extracted from observables. The effect of real rate news is estimated from changes in the real rates on index-linked government bonds (available up to a maturity of 30 years). The effect of risk premium news is estimated from changes in the Martin (2017) lower bound of equity risk premia (available up to a maturity of 2 years). The effect of dividend news is estimated from changes in the price of dividend futures (available up to a maturity of 10 years).

Figure 2.2: Cumulative sum of dividend weights by dividend maturity. This figure shows the cumulative weight of dividend prices relative to the overall stock market price. The cumulative sum weight at each maturity is the average weight across all days in 2020.

Figure 2.3: Estimating the equity risk premium term structure. This figure
summarises the factor analysis on the term structure of the equity risk premium. Row
one presents the time series of the two factors, row 2 presents loadings on the factors
across the observed risk premium maturities (up to 2.5 years), row 3 presents estimated
loadings across maturities, and row 4 presents the extrapolated factor loadings up to 10
year maturity.
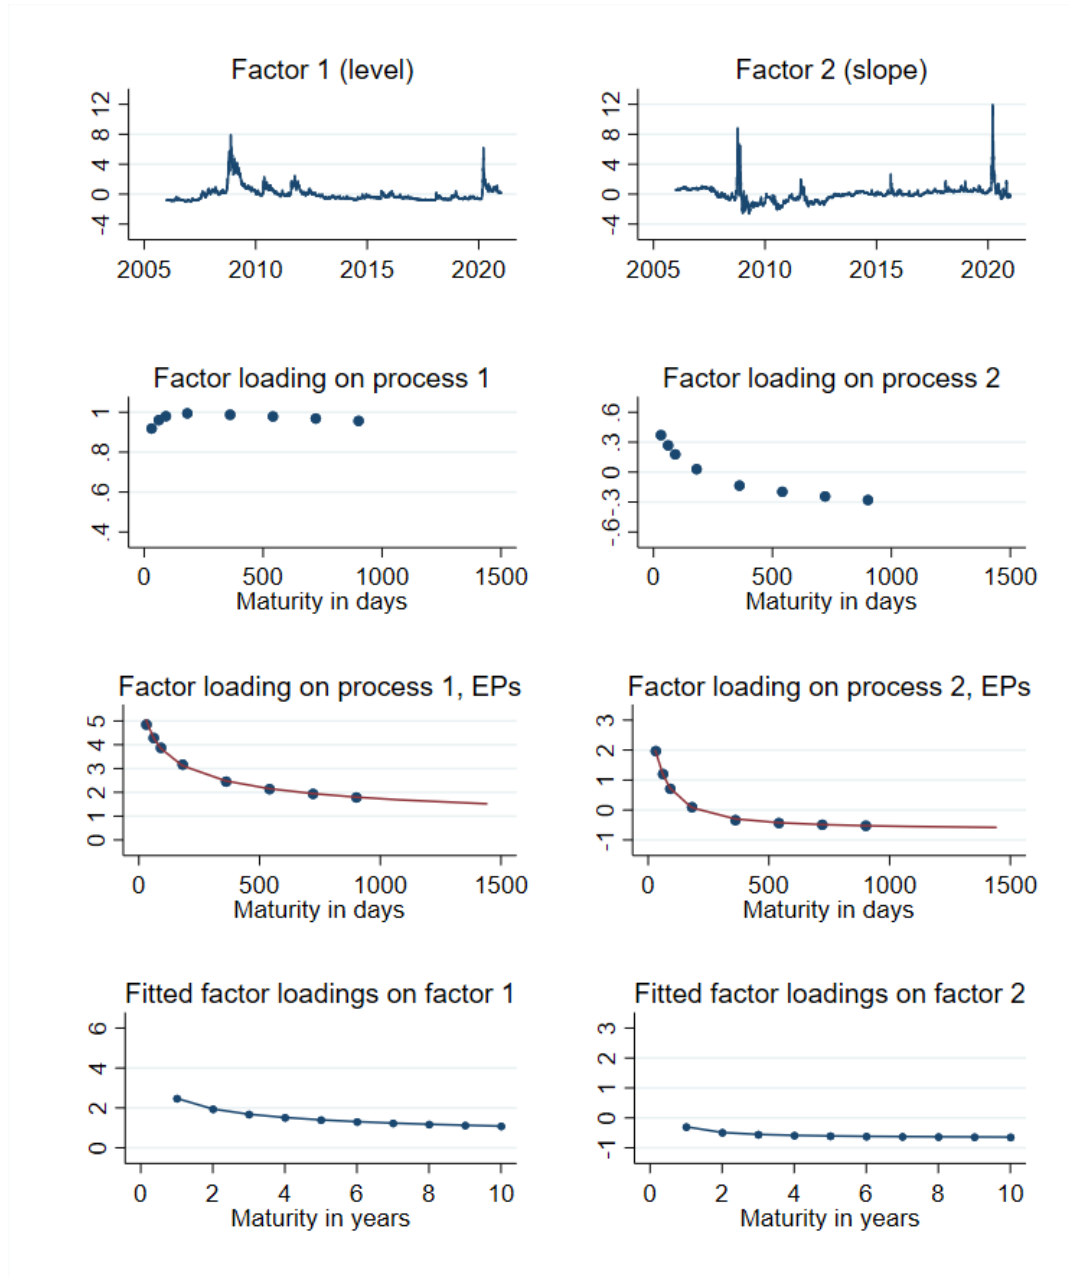
Figure 2.4: Equity risk premia by maturity (annualized). This figure shows annualized Martin (2017) equity risk premium estimates over 2020. The estimates are plotted for various observed maturities of the risk premium.
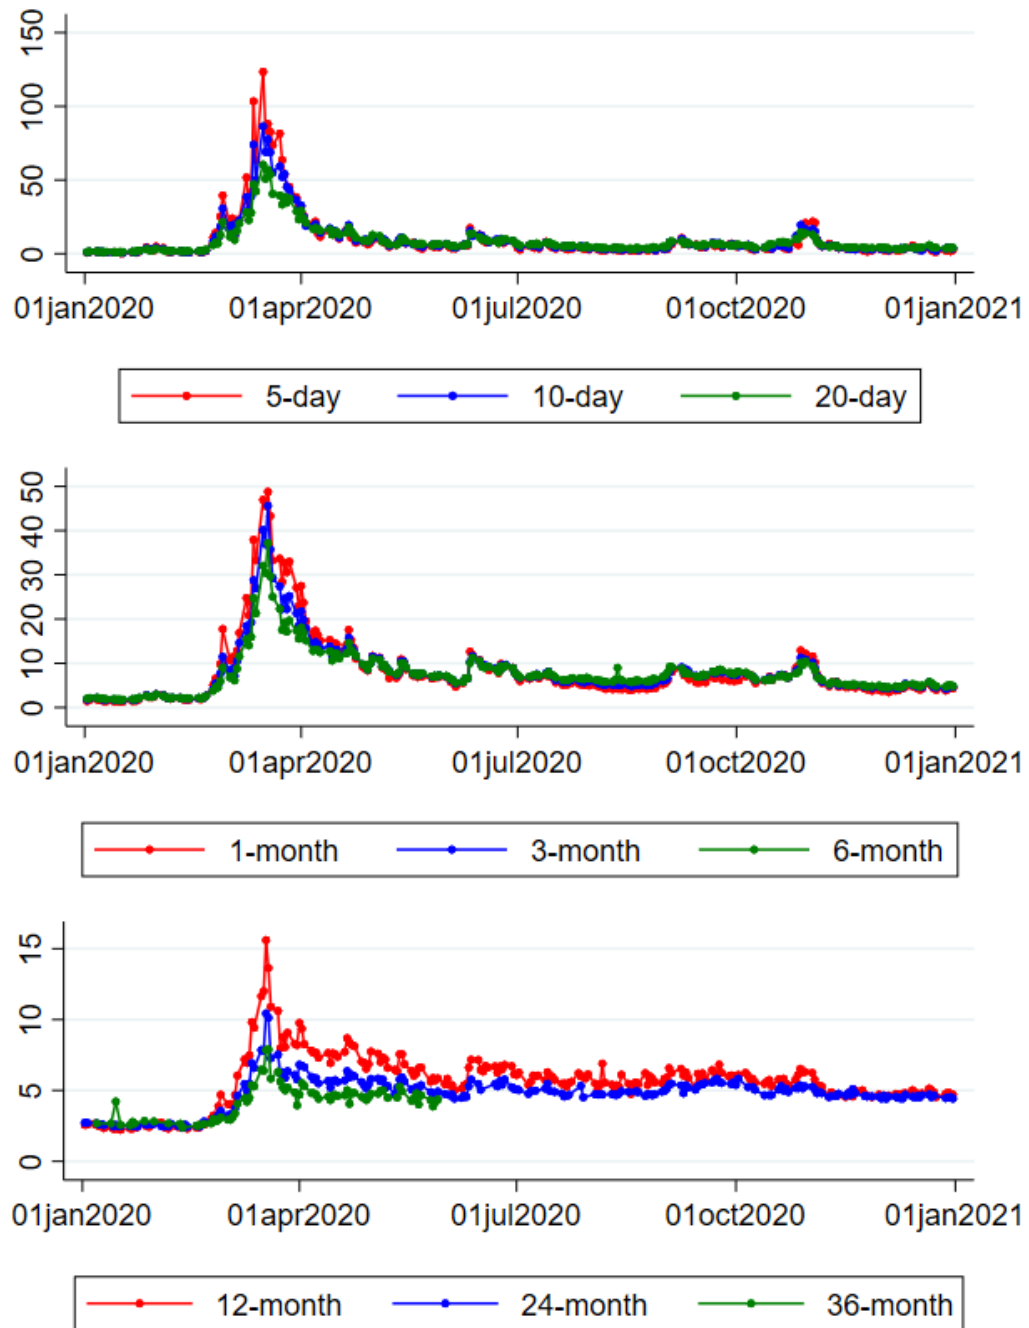
Figure 2.5: Cumulative equity risk premium. This figure shows the unanualised Martin (2017) equity risk premium estimates (day 0 to this day). We shot the cumulative risk premium for three separate days within the sample.

Figure 2.6: Observed forward equity risk premia (annualized). This figure shows the 6 month risk premium as well as the 6 months forward and 18 months forward 6 month risk premium.

Forward equity premia (annualized)

Figure 2.7: Estimated (demeaned) forward equity risk premia (annualized).
This figure shows the implied 1 year to 5 year forward 1 year risk premium, as estimated by the factor analysis. The left hand side shows the forward rates across the full sample (from 2006) and the right hand side focuses in on 2020.

Figure 2.8: Real risk-free rate and the underlying components. The top figure shows the 10 year and 30 year real yield over 2020. The bottom figures show the nominal yield and inflation swap rate (10 years on the left hand side and 30 years on the right hand side).
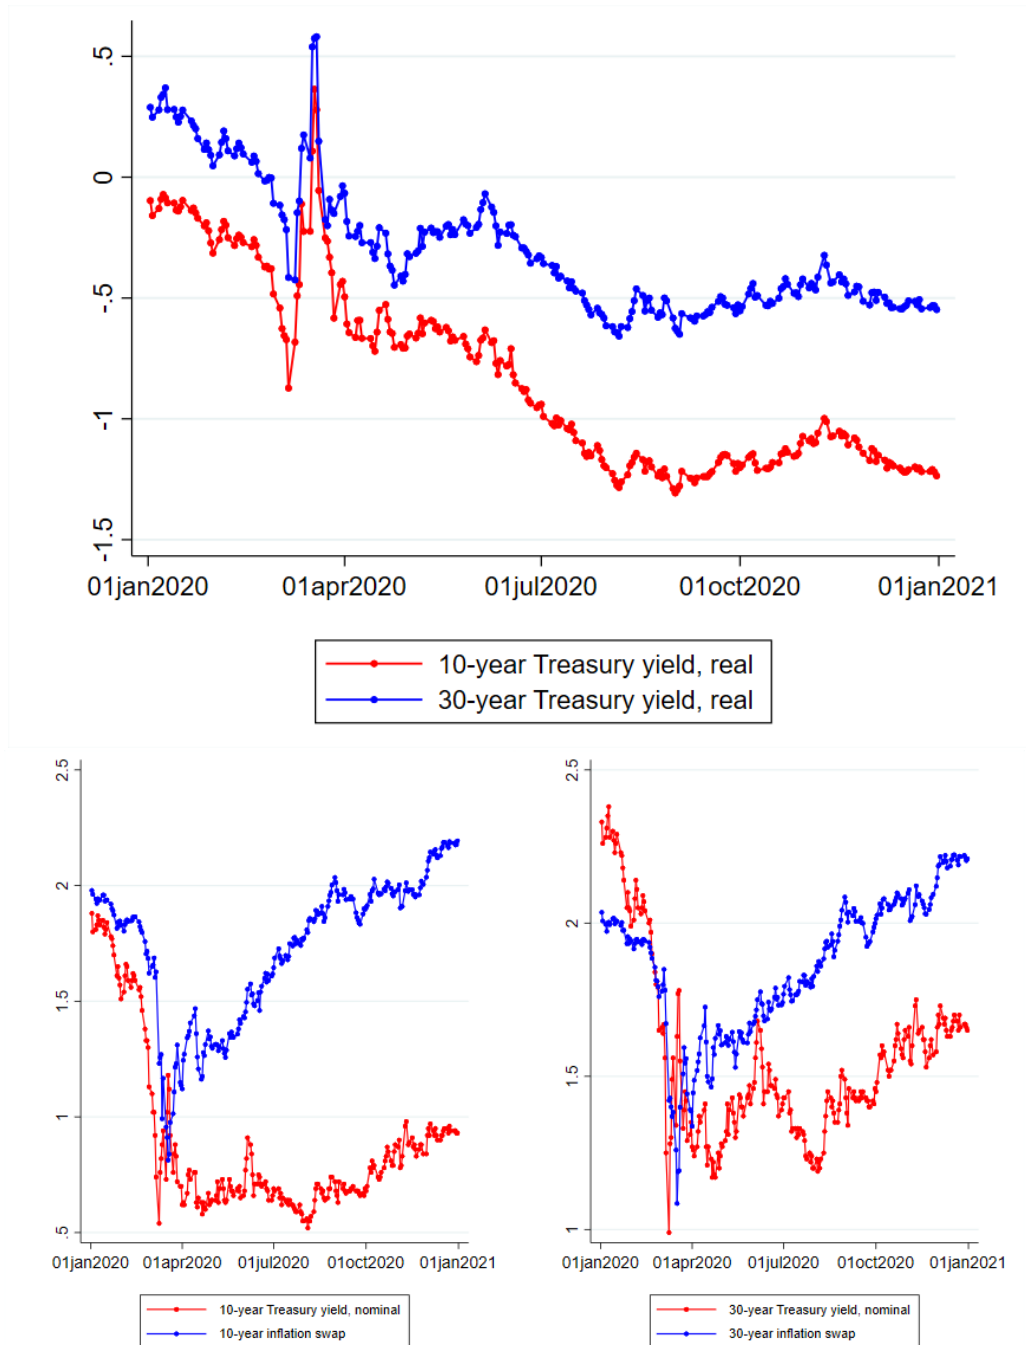
Figure 2.9: Real risk-free rate forwards This figure shows the real rate forwards
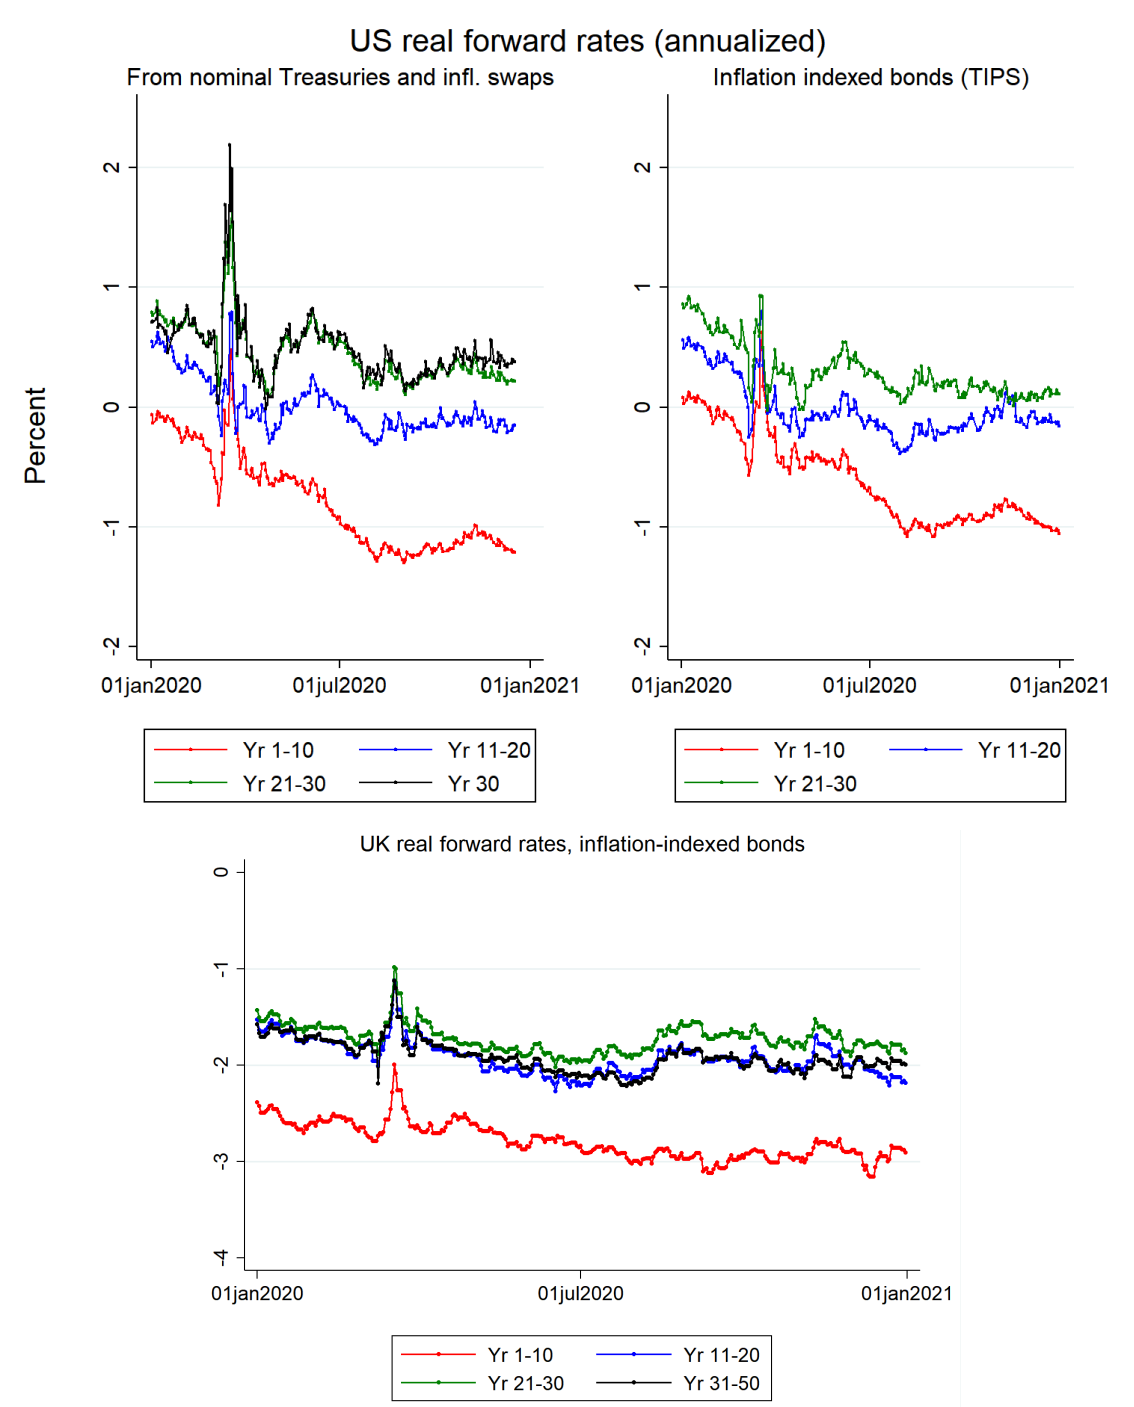
Figure 2.10: Constant maturity expected dividends These figure shows the constant maturity expected dividends 2, 5 and 8 years from each date. The left and right figures show nominal and real expected dividends respectively.

Figure 2.11: Stock market return from observables and long-term news. The top figure shows the cumulative return of the stock market return along with the implied return from our observables: news about real risk-free rates (1y-30yr), equity risk premium (1yr-5yr), expected dividends (1yr-10yr), and the realisation of the expected return. The bottom figure shows the residual stock return that is not explained by our observed variables. We call is long-term news.

2.8 Tables

Table 2.1: Risk premium estimate as a predictor variable.

This table reports the parameter estimate from the following time series regression: $\frac{1}{T-t}\left(R_T - R_{f,t}\right) = \alpha + \beta \times R_{f,t} \cdot \mathrm{SVIX}_{t \to T}^2 + \epsilon_T$ together with Newey-West standard errors with lag selection based on the number of overlapping observations. Columns refer to separate estimations with $T - t = 1$, 2, 3, 6 and 12 months respectively. The sample period is 1996-2020.

	Realised Return				
	1 month	2 month	3 month	6 month	1 year
Equity Risk Premium	1.26	1.33	1.33	1.79**	1.09
	(1.01)	(1.04)	(1.27)	(0.88)	(0.88)
Constant	0.01	-0.00	-0.00	-0.02	0.01
	(0.04)	(0.04)	(0.05)	(0.04)	(0.04)
R^2 (perc.)	0.96	1.68	2.05	4.64	1.92
Observations	6,254	6,235	6,216	6,151	6,026

2.9 Appendix

2.9.1 Proofs

Proof of Result 1

Write out the first few terms of the expression for P_t

$$P_t = \frac{E_t[D_{t+1}]}{E_t[1+R_{t+1}]} + \frac{E_t[D_{t+2}]}{E_t[1+R_{t+1}]E_t[1+R_{t+2}]} + \frac{E_t[D_{t+3}]}{E_t[1+R_{t+1}]E_t[1+R_{t+2}]E_t[1+R_{t+3}]} + ...$$

The effects of changes in the first three periods' expected returns are as follows

$$\frac{\partial P_t/P_t}{\partial E_t R_{t+1}} = -\frac{1}{E_t[1+R_{t+1}]}$$

$$\frac{\partial P_t/P_t}{\partial E_t R_{t+2}} = -\frac{1}{E_t[1+R_{t+2}]}\frac{E_t P_{t+1}/E_t[1+R_{t+1}]}{P_t}$$

$$= -\frac{1}{E_t[1+R_{t+2}]}\left[1 - \frac{E_t D_{t+1}/E_t[1+R_{t+1}]}{P_t}\right] = -\frac{1}{E_t[1+R_{t+2}]}[1-w_{t,1}]$$

$$\frac{\partial P_t/P_t}{\partial E_t R_{t+3}} = -\frac{1}{E_t[1+R_{t+3}]}\frac{E_t P_{t+2}/\{E_t[1+R_{t+1}][1+R_{t+2}]\}}{P_t}$$

$$= -\frac{1}{E_t[1+R_{t+3}]}\left[1 - \frac{E_t D_{t+1}/E_t[1+R_{t+1}]}{P_t} - \frac{E_t[D_{t+2}]/\{E_t[1+R_{t+1}][1+R_{t+2}]\}}{P_t}\right]$$

$$= -\frac{1}{E_t[1+R_{t+3}]}[1-w_{t,1}-w_{t,2}]$$

where we exploited that

$$P_t = \frac{E_t[D_{t+1}+P_{t+1}]}{E_t[1+R_{t+1}]} \iff \frac{E_t P_{t+1}/E_t[1+R_{t+1}]}{P_t} = 1 - \frac{E_t D_{t+1}/E_t[1+R_{t+1}]}{P_t}$$

and

$$P_t = \frac{E_t[D_{t+1}]}{E_t[1+R_{t+1}]} + \frac{E_t[D_{t+2}+P_{t+2}]}{E_t[1+R_{t+1}]E_t[1+R_{t+2}]} \iff$$

$$\frac{E_t P_{t+2}/\{E_t[1+R_{t+1}][1+R_{t+2}]\}}{P_t} = 1 - \frac{E_t D_{t+1}/E_t[1+R_{t+1}]}{P_t} - \frac{E_t[D_{t+2}]/\{E_t[1+R_{t+1}]E_t[1+R_{t+2}]\}}{P_t}.$$

The effect of changes to later expected returns follows from similar steps.

Proof of Result 4

From (2.12),

$$cov_t(M_T R_T, R_T) = e^{r_{f,t}+\sigma_{R,t}^2} - e^{\mu_{R,t}} \tag{2.12}$$

The derivatives with respect to a state variable s_t is

$$\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} = e^{r_{f,t}+\sigma_{R,t}^2}\left[\frac{\partial r_{f,t}}{\partial s_t} + 2\sigma_{R,T}\frac{\partial \sigma_{R,t}}{\partial s_t}\right] - e^{\mu_{R,t}}\left[\frac{\partial \mu_{R,t}}{\partial s_t}\right]$$

If the NCC holds, $cov_t\left(M_T R_T, R_T\right) \leq 0$ and thus $e^{r_{f,t} + \sigma^2_{R,t}} \leq e^{\mu_{R,t}}$. Therefore, it is sufficient for $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$ that $\frac{\partial r_{f,t}}{\partial s_t} + 2\sigma_{R,t}\frac{\partial \sigma_{R,t}}{\partial s_t} \leq \frac{\partial \mu_{R,t}}{\partial s_t}$. Rewrite this sufficient condition as follows

$$\left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\frac{1}{\sigma_{R,t}} - \frac{\partial \sigma_{R,t}}{\partial s_t} \geq \frac{\partial \sigma_{R,t}}{\partial s_t}$$

Consider now the change in the conditional Sharpe ratio (for log returns):

$$
\begin{aligned}
\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] &= \frac{1}{\sigma^2_{R,t}}\left[\left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\sigma_{R,t} - (\mu_{R,t} - r_{f,t})\frac{\partial \sigma_{R,t}}{\partial s_t}\right] \\
&= \left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\frac{1}{\sigma_{R,t}} - \frac{(\mu_{R,t} - r_{f,t})}{\sigma^2_{R,t}}\frac{\partial \sigma_{R,t}}{\partial s_t} \\
&\geq \left(\frac{\partial \mu_{R,t}}{\partial s_t} - \frac{\partial r_{f,t}}{\partial s_t}\right)\frac{1}{\sigma_{R,t}} - \frac{\partial \sigma_{R,t}}{\partial s_t}
\end{aligned}
$$

where the last line follows from the fact that $\frac{(\mu_{R,t} - r_{f,t})}{\sigma^2_{R,t}} \geq 1$ under the NCC and we are considering a state variable that increases risk ($\frac{\partial \sigma_{R,t}}{\partial s_t} \geq 0$). Thus, it is sufficient for $\frac{\partial cov_t(M_T R_T, R_T)}{\partial s_t} \leq 0$ that the change in the conditional Sharpe ratio is at least as large as the change in the conditional standard deviation

$$\frac{\partial}{\partial s_t}\left[\frac{\mu_{R,t} - r_{f,t}}{\sigma_{R,t}}\right] \geq \frac{\partial \sigma_{R,t}}{\partial s_t}.$$

2.9.2 Estimating the Martin lower bound

For 2020, we use option price data from CBOE to construct a time series of the Martin (2017) lower bound of the equity risk premium. The data includes the trading prices for every traded SP500 Index Option on each day (with intra-day data available), as well as each option's best bid and ask price, strike price, expiry date. The data also reports the underlying SP500 index price at the time of trade. We clean the initial data in several ways. First we delete all trades with a highest bid or ask of zero. Second, we delete trades where the trade price is greater (lower) than the highest ask (bid). Third, we delete all trades where the underlying index price is missing. Finally, from this selection of cleaned trades, we select the latest traded option for each date, expiry date, strike, option type (call/put) combination.

For each date, expiry date, strike combination in the dataset we then estimate the equity risk premium by discretizing the right-hand side of Martin (2017)'s lower bound given by

$$
\frac{1}{R_{f,t}} \mathrm{var}_t^* R_T = \frac{2}{S_t^2} \left[\int_0^{F_{t,T}} \mathrm{put}_{t,T}(K)\, dK + \int_{F_{t,T}}^{\infty} \mathrm{call}_{t,T}(K)\, dK \right].
$$

The forward price $F_{t,T}$ is the unique solution K of the equation $\mathrm{call}_{t,T}(K) = \mathrm{put}_{t,T}(K)$. We estimate the forward price by first interpolating trade prices across strikes for both calls and puts at each date and expiry date combination, and second identifying the intersection of these two interpolated series. We do not use the interpolated prices in discretization of the above equation.

Once equity risk premium estimates have been estimated for each date expiry date combination, we clean the data again. First, we only keep equity risk premium estimates where the number of strikes used in the estimation is greater than 15. Second, we delete estimates when the minimum call or put price is over 40% of the maximum trade price for that date and expiry date combination. These cleaning procedures delete estimates where the discretization is too coarse and where a large tail of options are missing, both of which cause biased estimates.

Finally, we generate constant maturity equity risk premiums by interpolating between those estimated at available expiry dates on any given date.[11] We also extrapolate to extend maturity. However, to avoid over extrapolation, we limit this extrapolation to a

---

[11]As an alternative to interpolation, we have also estimated the equity risk premium yield curve at each date using cubic spline methods. The results are very similar.

maximum of 150 days greater than the longest maturity option available at that date.

Chapter 3

A Granular Identification of Monetary Policy

Benjamin Knox[1]

I document rich heterogeneity in business cycles across U.S. states. As a result, state-level Taylor rules imply very different optimal monetary policies across states. To exploit the cross-sectional variation, I present a Granular Instrumental Variables (GIV) approach to monetary policy identification. The intuition behind the approach is that idiosyncratic shocks to economic activity in one state can lead to changes in monetary policy, which are exogenous monetary policy shocks from the perspective of other states. Based on this approach, I find that the U.S. unemployment rate increases 1.80 percentage points following a 1 percentage point increase in the federal funds rate, with the peak in unemployment 15 months after the initial federal funds rate innovation. The impact is larger than existing estimates in the literature, although non-trivial standard errors mean that the difference may be due to noise.

3.1 Introduction

The effects of monetary policy on the economy is a central question in macroeconomics. However, while the conventional wisdom is that monetary policy matters qualitatively, the quantitative impact of monetary policy changes is still poorly understood. The main challenge is that monetary policy decisions are so clearly tied to expected future economic activity that it makes the identification of their effect on this same variable very hard to estimate. From December 2008 to December 2009, the Federal Reserve Board reduced the target federal funds rate from 4.25% to 1.00%. At the same time, real GDP growth in the U.S. fell from 1.97% p.a. to -2.75% p.a.. However, we do not conclude that the fall in growth was caused by the reduction in the federal funds rate. To the contrary, the general consensus is that economic growth would have fallen further if the Fed Board hadn't reduced rates. How much further, though, remains an open question.

In this paper, I present a new approach to the monetary policy identification problem that uses cross-sectional variation in the economy to make causal inference on the effect of monetary policy in the aggregate. I first document rich heterogeneity in business cycles across U.S. states and, to frame this economic variation in a monetary policy context, compute state-level Taylor rule residuals. These are the difference between the Taylor rule implied optimal federal funds rate for each state and the actual federal funds rate. Large state-level Taylor rule residuals indicate that monetary policy at regional levels are far from optimal, and instead influenced by the economic conditions in other regions. This variation offers an identification opportunity. Specifically, a shock to the economic growth of one state in the U.S. can lead to a change in monetary policy, with this change in the federal funds rate exogenous from the perspective of other states.

To implement this idea, I present a Granular Instrumental Variables (GIVs) approach to the monetary policy identification problem using econometric methods recently developed by Gabaix and Koijen (2020) (GK). I apply the approach in the U.S. using monthly state-level unemployment data over the sample period 1981-2009. Following a 1 percentage point increase in the fed funds rate, I find that the US unemployment rate increases 1.80 percentage points, with the peak in unemployment 15 months after the initial federal funds rate innovation. The estimated impact is statistically significant, and the magnitude is larger than existing estimates in the literature that are based on a variety of alternative identification techniques (see Ramey (2016) for a comprehensive summary). However,

large standard errors mean that the GIV estimate is relatively imprecise in my setting.

A uniting feature across all existing monetary policy identification techniques is that they aim to address the endogeneity problem by identifying shocks to monetary policy. There are various methods applied in the literature, including narrative methods using historical records (Friedman and Schwartz (1963)), monetary policy rate innovations estimated through VAR model's (Christiano et al. (1999)), and policy deviations away from typical policy responses to the economic outlook (Romer and Romer (2004) Coibion (2012)). More recently, high-frequency data has also been used to extract monetary policy shocks, with the changes in interest rates in narrow windows surrounding policy announcements treated as unexpected movements (Barakchian and Crowe (2013) Gertler and Karadi (2015) Nakamura and Steinsson (2018)).

However, regardless of the approach implemented, any method attempting to extract monetary policy shocks faces the same issue that central bank actions have become highly predictable over recent decades. This makes true monetary policy shocks smaller and less frequent, and existing identification techniques less well identified due to reduced statistical power. Ramey (2016) summarises the challenge as follows:

> "The breakdown of many specifications in the later sample is simply that we can no longer identify monetary policy shocks well. Monetary policy is being conducted more systematically, so true monetary policy shocks are now rare.... While this is bad news for econometric identification, it is good news for economic policy."

In this context, GIVs provides a very appealing alternative approach to monetary policy identification, with the method working even in the extreme case where monetary policy is completely predictable. Instead, the shocks used for identification are idiosyncratic shocks in the cross-section of the economy. The critical characteristic of the data for a successful GIV implementation is thus that there is sufficient cross-sectional variation in the economic variable of interest. In practice, any cross-section could potentially be used (by industry, for example), but I focus on a geographical cross-section and at the state-level of granularity.

To motivate the analysis, I begin by documenting rich heterogeneity in business cycles across states. For example, in 2001, which was a recession at an aggregate level, I show that 40% of states were still in positive growth territory. In fact, the top 20% of states

had an average economic growth of around 4% that year. As well as there being variation in economic activity across states, I also show that there is also non-trivial dispersion in prices across states using the data provided in Hazell et al. (2020). The average cross-sectional standard deviation of state-level inflation rates is in excess of 1%. The same products are consistently set at meaningfully different prices across the U.S. at any given point in time.

To frame this economic variation in a monetary policy context, I next compute state-level Taylor rule residuals. As a proof of concept, I first show that the correlation between the actual the fed funds rate and a simple Taylor rule rate using aggregate U.S. data has a correlation of over 90%. In other words, the Taylor rule is a good overall fit and broadly describes the Federal Reserve Board's monetary policy behaviour. However, the picture is very different looking at the optimal Taylor rule rates implied by state-level data. The average correlation between these and the actual policy rate is less than 50%. State-level Taylor rule residuals are consistently large, exceeding 2% on average, and can also reach double digit figures in the extreme. Such striking heterogeneity across states offers the potential for GIVs to identify causal evidence of monetary policy.

In the subsequent implementation, I focus on unemployment rates given that they are available at a state-level on a monthly frequency. To understand the dynamics of U.S. unemployment after innovations in the federal fund rates, I implement a Jordà (2005) local projection approach combined with GIV methods. I find that an increase in the federal funds rate leads to an increase in the unemployment rate, the peak of which is approximately 15 months after the monetary policy innovation. Critically, the change in the federal fund rates is instrumented by the GIVs instrument, which supports a causal interpretation of the impulse response of the unemployment rate to monetary policy.

The GIVs in my setting is the size-weighted sums of idiosyncratic shocks to state-level unemployment rates. The intuition behind the validity of the GIVs is twofold. Firstly, the construction from purely idiosyncratic shocks means the instrument is exogenous and therefore exclusion restrictions hold. Secondly, the size-weightings generate statistical power and help the instrument to predict aggregate monetary policy (the IV relevance condition). Intuitively, the larger a state, the more its idiosyncratic shocks effect aggregate monetary policy, the more relative weight one wants it to have in the instrument of multiple idiosyncratic shocks.

The main threat to identification of the GIVs approach is that the instrument contains

shocks common across states as well as the desired idiosyncratic shocks local to each state. Including common shocks in the instrument means exclusion restrictions do not hold and IV estimate becomes biased. To control for this, I follow procedures recommended in Gabaix and Koijen (2020) and control for common factors estimated through principle component analysis. However, this significantly reduces the variation in the instrument (i.e. a large fraction of cross-sectional variation is driven by common factors), and thus reduces the estimations first stage power. Reducing first stage power increases second stage standard errors.

The above analysis reveals the tension in the GIV procedure. By ensuring identification and controlling for common factors, first stage power is sacrificed. To evaluate the strength of the instrument, I therefore report the Cragg-Donald Wald F-statistic rank test from the first stage estimation. The F-statistic in the main specification is 10.08, slightly above the rule of thumb threshold of 10 that is required to alleviate weak instrument concerns ((Stock and Yogo (2005) Andrews et al. (2019))). Nevertheless, a stronger first stage would help reduce standard errors and improve the precision of the GIVs estimate in my setting.

There is an extensive literature that has studied the causal impact of U.S. monetary policy on economic variables (a none-exhaustive list includes Christiano et al. (1999) Romer and Romer (2004) Boivin et al. (2010) Coibion (2012) (Barakchian and Crowe (2013)) Gertler and Karadi (2015) Nakamura and Steinsson (2018)). However, as far as I am aware, this is the first paper to apply a GIV approach to the monetary policy identification problem.

Conceptually, the underlying idea behind the identification technique is related to Ioannidou et al. (2015) and Jordà et al. (2015). These papers study how countries with currency pegs to the U.S. in some sense import U.S. monetary policy. U.S. monetary policy changes in response to U.S. specific economic conditions can result in monetary policy changes in other countries that are exogenous. Ioannidou et al. (2015) focus on the case of Bolivia and Jordà et al. (2015) apply the idea more broadly across 17 economies. In effect, the GIVs approach also looks at how monetary policy is imported, but rather than looking internationally with currency pegs, it studies importation within the same currency union. This makes it especially useful to policymakers, as it identifies the causal impact of their monetary policy decisions within their own jurisdiction. Further, the economies of countries with a currency peg to the U.S. can be quite different to the U.S. itself, and one

might therefore expect the transmission of monetary policy to also be different.

Other papers have also implemented elements of the intuition underlying GIVs in a less formal setting. For example, Jiménez et al. (2012) treat the euro area monetary policy as exogenous from the perspective of Spain, given their small contribution to output and relatively uncorrelated business cycle. The paper does not control for common factors between Spain and the euro area, however, which is a threat to identification. The treatment of common shocks is an important component of the GIV method. Nevertheless, this example highlights the potential to implement GIVs in a monetary policy identification context in many settings, including the euro area, and perhaps even globally. My contribution to the literature is implementing GIVs in a monetary policy context for the first time in the U.S. using a state-level of granularity.

3.2 Data

3.2.1 Monetary policy rates

For the main measure of the monetary policy rate I use the effective federal funds rate. A monthly time series is available from the Board of Governors of the Federal Reserve System (Federal Reserve Board) since 1952. In some specifications, I also use the 1 year and 2 year constant maturity treasury rates. Longer maturity rates such as these can be used to capture the forward guidance that the Federal Reserve Board provides in respect to future policy rates (Romer and Romer (2000) Leombroni et al. (2021)). The constant maturity rates are also provided by the Federal Reserve Board, and are available from 1962.

3.2.2 State-level unemployment rates (monthly)

The U.S. Bureau of Labor Statistic (BLS) provides seasonally adjusted unemployment rates at a state-level on a monthly basis. This covers the 50 states plus the District of Columbia. The unemployment rate represents the number of unemployed as a percentage of the labor force. Labor force data in each state is restricted to people 16 years of age and older who do not reside in institutions (e.g., penal and mental facilities, homes for the aged), and who are not on active duty in the Armed Forces. Unemployment rate data is available at a state-level from 1976.

At the national level, I also use the natural rate of unemployment provided by the

Congressional Budget Office. It is the estimated rate of unemployment arising from all sources except fluctuations in aggregate demand. This data is available quarterly.

3.2.3 State-level price indexes (quarterly)

For state-level inflation data, I use the Hazell et al. (2020) state-level price indexes that are constructed based on the micro-price data the BLS collects for the purpose of constructing the CPI. For each state, Hazell et al. (2020) provide a quarterly time series of non-tradeables and tradeables price index. The sample period is 1978 to 2018 (with a 26 month gap in 1986-1988 due to missing micro-data). For aggregate US inflation, I use the *Consumer Price Index for All Urban Consumers: All Items* and also the *Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average.* Both are provided by the BLS from 1947.

3.2.4 State-level Gross Domestic Product (annual)

Gross Domestic Product by State is available at an annual frequency from the U.S. Bureau of Economic Analysis. The sample starts in 1963 and is split by industry. There is a discontinuity in the GDP by state time series at 1997, where the data change from SIC industry definitions to NAICS industry definitions. Since 2005, the data has also been available at a quarterly frequency.

3.2.5 Descriptive Statistics

Table 3.1 presents descriptive statistics for the state-level economic variables. The last three columns explore the sources of variation in the panel data. The first column, $\sigma(x_{it})$, reports the standard deviation of variable x across the whole panel. Column $\sigma(\bar{x}_t)$ reports the time-series standard deviation of the panel's cross-sectional averages. The final column reports the time-series average of the cross-sectional standard deviation. For every time observation t, the cross sectional standard deviation $\sigma(x_i)$ is calculated first, with the column presenting the mean of these standard deviations across all t.

The final two columns capture the time-series variation and cross-sectional variation of the economic variables respectively. They show there is a significant amount of cross-sectional variation in the variables. In fact, in the case of unemployment and output growth, there is more cross-sectional volatility than time-series volatility.

3.3 Motivating Evidence

3.3.1 Variation in economic activity across states

Figure 3.1 presents evidence on the variation in economic activity across states in the US. States are sorted into five portfolios in each time period conditional on an economic variable of interest. The portfolio average with respect to the conditioning variable are then plotted over time. Panels A-D present the portfolio time series using the conditioning variables real GDP growth per capita, output gap, the demeaned unemployment rate, and inflation rate respectively.[2]

There is a wide dispersion across states for all the economic variables considered. For example, the difference between portfolio 5 and portfolio 1 for real GDP growth per capita at any given point is in excess of 5% across the full sample. In terms of a relative magnitude for this cross-sectional variation, the time-series average standard deviation of real GDP growth per capita at the aggregated US-level is 2.2%. The cross sectional variation is therefore economically significant.

A closer look at aggregate recessions periods is also revealing. For example, in 2001, 40% of states were in positive real GDP growth territory. This growth is despite the year being a NBER defined recession period. In fact, the top 20% of states had an average economic growth of around 4% that year. It is thus clear that U.S. states can be in quite varying business cycle conditions during aggregate recessions.

One may think that the same states are always placed in the same portfolio due to state fixed effects. However, Panels B and C, which present output gap and unemployment rates demeaned, show that the variation holds even controlling for long-run state averages. For example, following the financial crisis, the 20% of worst effected states saw unemployment rates 5% above their average unemployment rate. At the same time, there was another 20% of states with unemployment rates less than 1% over their long run average.

As well as there being variation in economic activity across states, Panel D shows there is also none-trivial dispersion in prices across states. Inflation rates are computed from state-level price indexes based on micro-price data,[3] and the variation shows that the prices on the same individual products change significantly across US regions. The average

---

[2]The state-level output gap is estimated from the cyclical component of the Hodrick and Prescott (1997) filter on real output. I have also computed output gap with Hamilton (2018) proposed methods. Results are qualitatively similar, with even more variation using Hamilton (2018).

[3]see Hazell et al. (2020) for more details.

cross sectional standard deviation (final column of Table 3.1) of the resulting state-level inflation rates is in excess of 1%.

3.3.2 Taylor rule residuals at a state-level

What does this variation mean in a monetary policy context? To explore this question, I calculate the optimal federal funds rate for each state, using a simple Taylor rule

$$r_{i,t}^* = 2 + \pi_{i,t} + 0.5(\pi_{i,t} - 2) + (u_t^* - u_{i,t}) \tag{3.1}$$

where $\pi_{i,t}$ is the inflation rate of state i at time t, $u_{i,t}$ is the unemployment rate of state i at time t, and u_t^* is the long-run natural rate of unemployment for the US at time t. The rule follows the Bernanke (2015) suggested adjustment to the original Taylor (1993) rule with the coefficient on the unemployment gap set at 1 rather than 0.5.[4]

Figure 3.2 presents analysis on the variation of the optimal federal funds rate across states. In Panel A, states are sorted into five portfolios at each date conditional on their optimal rate. The portfolio average optimal rates are then plotted over time. I also plot the actual federal fund rate for comparison. Given the dispersion in the economic variables behind the Taylor rule, it is not surprising to see that there is also very large dispersion in the optimal federal fund rates.

Each state's Taylor rule residual $\epsilon_{i,t}^{TR} = r_{i,t}^* - r_t$ is the difference between the optimal rate at time t and the actual federal funds rate at time t. Figure 3.2 Panel B presents variation in this variable across states. In the 1990s, when the median optimal fed fund rate almost exactly matches the actual fed fund rate (i.e. the Taylor rule fits very well in the aggregate), the difference between portfolio 1 and portfolio 5 averages at 5%. This shows there were economically meaningful Taylor rule residuals at the local level, even in a period where there was close to no residual at the aggregate level.

Table 3.2 presents more analysis on the Taylor rule residuals over 1989-2009.[5] The first row implements the Taylor rule in equation (3.1) at the US aggregate-level as a proof of concept. The first column shows the correlation between the US Aggregate Taylor rule rate and actual fed funds rate is over 90% in this period. In other words, the Taylor rule

---

[4]In a similar exerise, Malkin and Nechio (2012) apply the Taylor rule to four U.S. regions: Northeast, Midwest, South, and West. However, they do not study the more granular cross-section of state-level Taylor rule implied rates.

[5]I have remove the dates post 2009 as the lower bound period restricted the flexibility of the Federal Reserve Board for reduce interest rates.

is a good overall fit and broadly describes the Federal Reserve Board's monetary policy behaviour.

The remaining rows, which look at various states within the US, show a much worse fit at the state-level. In the best case, South Carolina has a correlation of 75%, while in the worst case Mississippi ("MS") has a correlation of only 8%. The other columns look at the variation in Taylor rule residuals. The minimum and maximum show that, in the extreme, states can have federal fund rates which are double digits away from what their local economy requires at that time. On top of this, the mean of the absolute Taylor rule residual (final column) shows states are consistently a long distance from their optimal rates. While for the US aggregate the mean distance is less than 1%, individual states are typically over 2% away from their Taylor rule implied optimal rates.

As a case study, Figure 3.3 plots the Taylor rule optimal rate for Florida and North Carolina. These are two states that are relatively close together geographically within the US. Nevertheless, there are large deviations in their optimal policy rate over time. In the 1990s, Florida had high unemployment and thus an optimal policy rate close to zero. At the same time, North Carolina was enjoying much stronger economic conditions, and had a higher optimal federal funds rate. In fact, at times during the 1990s the federal funds rate was too low from North Carolina's perspective. Strong economic conditions for states like North Carolina prevented the Federal Reserve Board from cutting rates for the benefit of Florida, and therefore meant that Florida had too low an interest rate. The roles reverse in 2000s, however, with Florida (North Carolina) requiring a higher (lower) federal funds rate given their economic outlook in this decade.

3.3.3 Summary of motivating evidence

This section has documented strikingly large cross sectional variation in economic conditions across U.S. states. The variation is economically meaningful in monetary policy context, as demonstrated with Taylor rule residuals. These findings motivate the remainder of the paper, which explores how this heterogeneity can be used to identify the effects of monetary policy. Specifically, it will use GIV methods, which have been designed to exploit local idiosyncratic shocks in the cross section to make causal identification in the aggregate.

3.4 Empirical Framework

This section begins with some basic notation. It then introduces monetary policy identification with GIVs methods in a simple economy. This provides the econometric intuition for the empirical approach. The section concludes with one extension to the simple example that is important for the implementation.

3.4.1 Notation

It is useful to set out how we denote panel data with N panel members and T time periods. For any variable x, the size-weighted average at time t is denoted

$$x_{S,t} = \sum_{i=1}^{N} S_i x_{i,t} \tag{3.2}$$

where S_i is the relative size of member i with $\sum_{i=1} S_i = 1$. The equal-weighted average at time t is denoted

$$x_{E,t} = \sum_{i=1}^{N} E_i x_{i,t} \tag{3.3}$$

where E_i is the equal weight such that $\sum_{i=1} E_i = 1$.[6] The difference between the size-weighted and equal-weighted averages is denoted

$$x_{\Gamma,t} = \sum_{i=1}^{N} \Gamma_i x_{i,t} \tag{3.5}$$

where $\Gamma_i = S_i - E_i$ such that $\sum_{i=1} \Gamma_i = 0$ and $x_{\Gamma,t} = x_{S,t} - x_{E,t}$.

3.4.2 A simple economy with endogenous monetary policy

Consider a simple economy where there are $i = 1, ..., N$ regions in a currency union. There are two variables of interest, the short-term interest rate r and economic growth y. The dynamics of the economy, assuming variables are demeaned, is described by

$$y_{i,t} = b^r r_t + \eta_t + \nu_{i,t} \tag{3.6}$$

$$y_{S,t} = b^r r_t + \epsilon_t^{y_S} \tag{3.7}$$

---

[6]The basic equal weight is $E_i = \frac{1}{N}$. An alternative is the inverse-variance weight

$$E_i = \frac{\sigma_i^{-2}}{\sum_{j=1}^{N} \sigma_j^{-2}} \tag{3.4}$$

where σ_i is the standard deviation of x_i over the sample period. This E_i is used such that each panel member contributes equal variance to $x_{E,t}$ in the case that x_i are independent but heteroskedastic.

$$r_t = b^y y_{S,t} + \epsilon_t^r \tag{3.8}$$

where economic growth for region i at time t is $y_{i,t}$, the size-weighted economic growth across all regions in the union at time t is $y_{S,t}$, and the short-term interest rate set a time t by the central bank is r_t. Economic growth disturbances in region i are assumed to be the sum of a common shock η_t and an idiosyncratic shock $\nu_{i,t}$ local to that region. Aggregate economic growth disturbances $\epsilon_t^{ys} = \eta_t + \nu_{S,t}$ in equation (3.7) are thus the sum of the common shock η_t plus the size-weighted average of idiosyncratic shocks $\nu_{S,t}$. Finally, the monetary policy shock at time t is denoted ϵ_t^r. All shocks u_{it}, η_t and ϵ_t^r are i.i.d across dates, with the regional shocks ν_{it} also i.i.d across i's.

The parameter b^r, which is assumed to be homogeneous across regions, is the sensitivity of economic growth to the short-term interest rate. The parameter b^y is the central bank's standard response function to aggregate economic growth. In a typical economy $b^y > 0$ and $b^r < 0$. When economic activity is above its long run average, the central bank increases the short-term interest rate, which reduces economic growth and prevents an overheating of the economy as the central bank desired.

3.4.3 Traditional monetary policy identification in the simple economy

The objective of the monetary policy identification literature is to estimate, b^r, the sensitivity of economic growth to the short-term interest rate. However, estimating equation (3.7) by itself leads to endogeneity due to simultaneity. Solving for r_t results in

$$r_t = \frac{b^y}{1 - b^y b^r} \epsilon_t^{ys} + \frac{1}{1 - b^y b^r} \epsilon_t^r \tag{3.9}$$

and thus a regression of $y_{S,t}$ on r_t produces a biased estimate of b^r with the independent variable correlated with the error term

$$\mathbb{E}\left[r_t \epsilon_t^{ys}\right] = \frac{b^r}{1 - b^y b^r} \mathbb{E}\left[\epsilon_t^{ys} \epsilon_t^{ys}\right] \neq 0. \tag{3.10}$$

The traditional solution to the endogeneity problem is to extract monetary policy shocks ϵ_t^r directly. These can be used to identify b^r in equation (3.7) as they are, by definition, not correlated with the error term ϵ_t^{ys}. There are various techniques in the literature for extracting monetary shocks. For example, using narrative methods with historical case studies (Friedman and Schwartz (1963)), moneraty policy rate innovations estimated through VAR model's (Christiano et al. (1999)), and policy deviations away from typical (i.e. the Taylor rule) policy responses (Romer and Romer (2004) Coibion (2012)). More

recently, high-frequency data has also been used to extract monetary policy shocks. The idea is that the changes in interest rates in narrow windows surrounding policy announcements can be treated as unexpected movements (Gertler and Karadi (2015) Nakamura and Steinsson (2018)).

However, as Ramey (2016) stresses, in recent decades monetary policy changes have been well anticipated by market participants ahead of the policy announcements. The predictability of policy changes means that true monetary policy shocks ϵ_t^r are very small. This is a challenge for all the existing approaches to monetary policy identification. By attempting to extract and use ϵ_t^r, the econometrician replaces an endogeneity problem with an issue of statistical power.

3.4.4 GIVs monetary policy identification in the simple economy

GIVs provides an alternative monetary policy identification strategy. The approach is to apply an instrumental variables estimation to

$$y_{E,t} = b^r r_t + \epsilon_t^{y_E} \tag{3.11}$$

where $y_{E,t}$ for equal-weighted economic growth across regions at time t and the error term $\epsilon_t^{y_E} = \eta_t + \nu_{E,t}$ is the sum of the common economic growth shock η_t and the equal-weighted average of idiosyncratic growth shocks $\nu_{E,t}$. The "granular" instrumental variable for the estimation is

$$z_t = y_{\Gamma,t} \tag{3.12}$$

where $y_{\Gamma,t}$ is the difference between the size-weighted and equal-weighted averages of economic growth. For z_t to be a valid instrument for the estimation of equation (3.11), both the relevance condition $\mathbb{E}[z_t r_t] \neq 0$ and the exclusion restriction condition $\mathbb{E}[z_t \epsilon_t^{y_E}] = 0$ must hold. Below the intuition of the instrument with respect to these conditions is explained. The econometric proofs are provided in Appendix 3.9.1.

For the exclusion restriction condition to hold, the instrument must be uncorrelated with the estimation error term $\epsilon_t^{y_E} = \eta_t + \nu_{E,t}$. The crucial characteristic of the instrument from this perspective is that it is made up of idiosyncratic shocks only. The factors common across regions $\{r_t, \eta_t\}$ cancel when calculating the difference between $y_{S,t}$ and $y_{E,t}$ and hence we are left with $z_t = y_{\Gamma,t} = \nu_{\Gamma,t}$. The instrument is therefore uncorrelated with the common shock η_t in the estimation error term. Further, given the regional shocks are i.i.d

across i, the gamma-weighted average is uncorrelated with the equal-weighted average, and thus the instrument $\nu_{\Gamma,t}$ is uncorrelated with the $\nu_{E,t}$ in the estimation error term.[7]

For the instrument to be relevant, it must also have predictive power for the short-term interest rate. The crucial characteristic of z_t from this perspective is the weight $\Gamma_{i,t}$ on the idiosyncratic shocks $\nu_{i,t}$. By placing greater weight on the shocks of the larger regions, the instrument correlates with aggregate economic growth. The short-term interest rate is set in response to aggregate economic growth, and it therefore moves when there are shocks to large regions. This relationship is captured in first term of equation (3.9).

3.4.5 GIVs: Key intuition and required characteristics in the data

To summarise so far, a GIVs approach to identification is to harness idiosyncratic shocks in the cross-section of the economy and use them to make causal inference on the effect of monetary policy in the aggregate. Shocks to the economic growth of one region in the currency union leads to changes in monetary policy, with these changes in the short-term interest rate being exogenous from the perspective of the other regions.

GIVs methods therefore requires two features in the data. First, there needs to be sufficiently large idiosyncratic shocks in the cross section. Rich heterogeneity in business cycles and taylor rule residuals across states (documented in section 3.3) help validate the application of GIVs methods in a US monetary policy setting from this perspective. Second, there needs to be size variation across regions. The larger a region, the more its idiosyncratic shocks effect aggregate monetary policy, the more powerful it's shock is as a component of the granular instrumental variable.[8]

As shown by GK, the excess Herfindahl index

$$h = \sqrt{-\frac{1}{N} + \sum_{i=1}^{N} S_i^2} \tag{3.13}$$

is a useful measure of expressing the panel's size-weight variation in the context of a GIVs approach. Figure 3.4 presents the excess Herfindahl index for US states with S_i defined as state i's fraction of the aggregated US population. The size variation has been increasing over our sample and today $h = 16\%$.

---

[7] see appendix for details on the importance of the independence assumption for this result

[8] Indeed, if all regions are equal sized then $z_t = \nu_{S,t} - \nu_{E,t} = \nu_{E,t} - \nu_{E,t} = 0$.

3.4.6 Dealing with heterogeneous loadings on the common factor

One extension of the model that is important in practice is a relaxation of the assumption that all regions load equally on the common factor η_t. In this section, I denote the loading of region i on the common shock as λ_i. In the previous sections, the assumption was $\lambda_i = 1$ for all i. With this adjustment to the model, equation (3.6) becomes

$$y_{i,t} = b^r r_t + \lambda_i \eta_t + \nu_{i,t} \tag{3.14}$$

and the error terms for equations (3.7) and (3.11) are now $\epsilon_t^{yS} = \lambda_S \eta_t + \nu_{S,t}$ and $\epsilon_t^{yE} = \lambda_E \eta_t + \nu_{E,t}$ respectively. The GIV in this setting is

$$z_t = y_{\Gamma,t} = \lambda_\Gamma \eta_t + \nu_{\Gamma,t} \tag{3.15}$$

the sum of a common shock component $\lambda_\Gamma \eta_t$ and the gamma-weighted idiosyncratic shocks $\nu_{\Gamma,t}$. The common shock component, which results from the heterogeneous loadings λ_i, means the instrument is correlated with ϵ_t^{yE}. The exclusion restriction therefore no longer holds.

To resolve this issue, GK recommend to first compute the difference between regional economic growth and equal-weighted economic growth

$$y_{i,t} - y_{E,t} = (\lambda_i - \lambda_E) \eta_t + (\nu_{i,t} - \nu_{E,t}) \tag{3.16}$$

with this variable in essence removes date fixed effects (including the endogenous r_t) from the panel data of $y_{i,t}$'s. A factor model, such as Principal Component Analysis (PCA), can be run on this newly created variable to estimate the latent factor. The estimated factors are then used as control variable in the instrumental variables estimation.

There are many further extensions possible of the GIVs approach, with the reader referred to GK for details.

3.5 Empirical Results

3.5.1 Estimation Procedure

I implement a GIV approach to monetary policy identification using state-level monthly unemployment data. The estimation is based on an economic model that is a generalisation of the economy described in section 3.4. It can be summarised with two structural

equations

$$u_{S,t} = \alpha_u + b^r r_t + \sum_{f=1}^{F} \lambda_S^f \eta_t^f + \epsilon_t^{u_S} \tag{3.17}$$

$$r_t = \alpha_r + b^u u_{S,t} + b^\pi \pi_t + \epsilon_t^r \tag{3.18}$$

where $u_{S,t}$ is the size-weighted (or aggregate) unemployment rate in the US, r_t is the federal funds rate, η_t^f's are F common factors that affect unemployment across states, λ_i^f is states i's loading on factor f, and π_t is the US inflation rate. Equation (3.18) is a version of the Taylor rule, with the Federal Reserve Board responding to economic activity (in this case unemployment) and the inflation rate.

For the estimation setup, I follow the methodology recommended in GK. I first compute E_i equal-weights for U.S. states using the inverse-variance equal-weights defined in equation (3.4).[9] I then estimate a panel regression with state and date fixed effects

$$u_{i,t} = \alpha_i + \gamma_t + \check{u}_{i,t}$$

using E_i as regression weights, and construct the $\check{u}_{i,t}$ as the residuals. Finally, motivated by section 3.4.6, I extract estimated principal components of $\check{u}_{i,t}$ using PCA. These capture latent common factors within the residuals, and are denoted $\bar{\eta}_t$.

Following the above setup, the results are generated from the estimation of the equal-weighted unemployment rate $u_{E,t}$ response to monetary policy

$$u_{E,t} = \alpha_u + b^r r_t + \sum_{f=1}^{F} \lambda_E^f \bar{\eta}_t^f + \epsilon_t^{u_E} \tag{3.19}$$

where the GIVs $z_t = u_{\Gamma,t}$ is used as an instrument for the risk-free rate r_t, and the controls are the common factors $\bar{\eta}_t$ estimated via PCA. The instrument z_t is computed using time-varying size weights $S_{i,t-1}$ that are based on states i's faction of the aggregate population in the proceeding month $t-1$.

Previous studies of the effects on monetary policy on economic activity highlight the reaction of the economy over time. In my main results, I therefore implement a Jordà (2005) local projection approach combined with GIV methods to understand the dynamics of US unemployment following innovations in the federal fund rates. Jordà et al. (2015) is an example of a paper that implements an empirical strategy combining the local projection

---

[9]Inverse-variance equal-weights are used instead of simple equal-weights $E_i = 1/N$ so that less weight is placed on volatile states. This improves the GIV estimation efficiency, with GK showing that the inverse-variance equal-weight is the optimal GIV estimator in terms of precision.

approach with IV methods for the purpose of monetary policy identification. As far as I am aware, this is the first paper to combine the local projection approach with GIV methods.

3.5.2 Unemployment rate dynamics in response to an increase in the federal funds rate

I estimate the path of US unemployment after innovations in the federal fund rate using the following local projection specification

$$\Delta_h u_{E,t+h} = \alpha_h + \beta_h^r \Delta r_t + \sum_{j=1}^{J} \beta_{h,j}^{u_E} \Delta u_{E,t-j} + \sum_{f=1}^{F} \beta_{h,f}^{\eta} \Delta \bar{\eta}_t^f + \beta_h^{\pi} \Delta \pi_t + \epsilon_{t+h}, \quad h = 1, \ldots, H \tag{3.20}$$

where $\Delta_h u_{E,t+h}$ is the change in the equal-weighted unemployment rate from month t to $t+h$, α_h is a constant and Δr_t is the change in the fed funds rate at time t. The fed funds rate r_t is instrumented by the granular instrumental variable $z_t = u_{\Gamma,t}$, which has been computed with time-varying size weights $S_{i,t-1}$ that are based on states i's faction of the aggregate population in the proceeding month $t-1$.

As is standard in the Jordà (2005) local projection framework, I control for lags of the dependent variable $\Delta u_{E,t-j}$. The main specification includes lags $J = 3$. I also control for estimated unemployment factors $\bar{\eta}_t^f$ that are common across US states. The main specification includes $F = 3$ common factors. Finally, I control for the change in annualised inflation $\Delta \pi_t$, which is an important variable in the IV's first stage.

Figure 3.5 presents the impulse response of the US unemployment rate following innovations in the federal funds rate given by the estimated sequence of coefficients $\{\beta_h^r\}$ from equation (3.20). The estimated sample period is 1981-2009, and results are shown out to $H = 40$ months. An increase in the federal funds rate leads to an increase in the unemployment rate. The peak in the change in unemployment is approximately 15 months after the monetary policy innovation.

In terms of magnitudes, a one percentage point increase in the federal funds rate predicts a 1.80 percentage points increase in unemployment. Subsequent to the peak, the unemployment rate mean reverts, and is back to it's initial rate roughly 36 months after the innovation in the federal funds rates first occurred. The peak increase in unemployment is statistically significant, with the dashed lines represent 90% confidence intervals

computed using Newey and West (1987) standard errors with lag selection h.[10]

Ramey (2016) provides a detailed summary of monetary policy identification estimates that have been published in recent years. The GIVs approach estimates an impact of monetary policy on unemployment that is larger than existing estimates based on a variety of alternative identification techniques. For example, using Romer and Romer (2004) policy shocks, Coibion (2012) finds one of the larger estimates. Even in this case, unemployment increases by 0.95 percentage points following a 1 percentage point increase in the federal funds rate, with the peak is estimated to be 24 months following the federal funds innovation. The findings of the GIVs approach therefore suggest that the effect of monetary policy is more powerful than previously thought.

However, it must be noted that the standard errors of the main specification are relatively large, indicating a wide range of plausible coefficients from the GIVs approach. I explore the reasons for the large standard errors next.

3.5.3 Instrument power

This section provides first stage analysis of the GIVs estimation described in equation (3.20). Table 3.3 presents estimates from the first stage specification:

$$\Delta r_t = \alpha + \beta^z \Delta z_t + \sum_{j=1}^{J} \beta_j^{u_E} \Delta u_{E,t-j} + \sum_{f=1}^{F} \beta_f^{\eta} \Delta \bar{\eta}_t^f + \beta^\pi \Delta \pi_t + \epsilon_t, \quad F = 0, \dots, 6$$

where the parameter estimates are reported for β^z and β^π only. To evaluate the strength of the instrument, the Cragg-Donald Wald F-statistic rank test is also reported. Each column corresponds to an estimation with $F = 0, 1, \dots, 6$ common factors included in the controls variable.

The parameter estimates correspond to the response of the federal funds rate to changes in the unemployment rate instrument z_t (size-weighted minus equal-weighted unemployment) and annualised core inflation π_t. The reported coefficient are consistent with the traditional Taylor rule interpretation. When unemployment (inflation) increases the Federal Reserve Board increases (decreases) the federal funds rate.

The main estimation, as presented in Figure 3.5, uses $F = 3$. This corresponds to column 4 in Table 3.3. The fact that this is greater than 10 provides reassurances that the GIV estimation does not suffer from weak instruments (Stock and Yogo (2005) Andrews et al. (2019)) in the main specification. However, the weaker the instrument, the greater

---

[10]The standard errors must incorporate a correction for serial correlation of residuals ϵ_{t+h} over h.

the standard error. A more powerful first stage would therefore help to reduce the large second stage standard errors and produce more precise estimates.

A concern from the first stage analysis is the sensitivity of the first stage F-statistic to the number of common factors controlled for. Indeed, I have chose the specification that maximises the first stage F-statistic. Using either more or less than $F = 3$ factors would reduce the power of the first stage, bringing it below the threshold of 10 that is generally targeted in the literature.

3.5.4 Threats to identification

The concept behind the GIV is that it is constructed from idiosyncratic shocks. However, as explained in Section 3.4.6, it is likely that the basic GIV contains common factors due to heterogeneous loadings on the common factor across states. This is why the estimated common factors $\bar{\eta}$ are included in the estimation specifications. The main threat to identification is therefore that the common factors are not fully controlled for.

Figure 3.6 Panel A presents the time series of the instrument z_t as well as the underlying components $u_{S,t}$ and $u_{E,t}$. It is clear there is a cyclical component of the instrument, with the size-weighted unemployment rate increasing more in recessions relative to the equal-weighted employment rate. In the main estimation, I therefore include as many factors as possible while still maintaining a powerful instrument.

Figure 3.6 Panel B presents the residual of regressing the instrument z_t on the first three principle components η_t^1, η_t^2 and η_t^3. This is the instrument purged of the first three estimated common factors. The first three PCA absorb 32%, 26% and 10 % of variation in the panel of $\check{u}_{i,t}$ respectively. By controlling for common factors, we therefore remove a significant amount of variation.

This analysis highlights a trade-off in the construction of the GIVs. Unobserved common factors are a threat to identification. The more estimated common factors that are included in controls, the more likely the exclusion restriction holds and the instrument is exogenous. Indeed, the identifying assumption of the GIVs method is that the variable plotted in Panel B of Figure 3.6 contains local idiosyncratic shocks only. However, by increasing common factors, the variation on the instrument (purged of the common factors) is diminished, thus reducing the power of the first stage.

The question remains whether using $F = 3$ common factors successfully controls for endogeneity in the instrument. I therefore implement a series of tests for over-identifying

restrictions in the next subsections. The results are summarised in Figure 3.7.

Odd-even instruments

For the first over-identifying restrictions test, I follow the procedure in Gabaix and Koijen (2021) and sort states by size in each period, creating two instruments constructed purely on odd or even states respectively. Each state's idiosyncratic shock $\check{u}_{i,t}$ is a valid instrument, and therefore, by default, any portfolio of shocks is also a valid instrument. The idea here is therefore to create two instruments by arbitrary dividing the sample of idiosyncratic shocks into two subsets, whilst retaining instrument power by still weighting on size.[11]

The results from the new odd and even instruments are presented in Panel A and B of Figure 3.7 respectively. The coefficient on both are very close to the original coefficients presented in Figure 3.5. The consistent estimates provide support for the over identification test. The standard errors on the new instruments are larger than the original estimates, which is not surprising give 50% of the exogenous variation has been thrown out for each instrument.

Region fixed effects

A potential source of correlation between state shocks $\check{u}_{i,t}$ could be regional effects. For example, Crone (2005) shows how states in close geographical proximity exhibit similar business cycles. I therefore explicitly control for 4 regional fixed effects in the next test, with the states placed into their Census Bureau-designated regions.

Instead of just estimating latent common factors from the panel of $\check{u}_{i,t}$ with PCA, I first estimate a series of cross sectional regressions of $\check{u}_{i,t}$ on region fixed effects. The time series of estimated factor loadings are a time series of region effects that can be used as a control in the main IV specification. In addition, I also include PCA estimated common factors for further latent factors. These are estimated on the panel of residuals from the cross sectional regressions.

The result with region fixed effects is presented in Figure 3.7 Panel C. The estimated impact of monetary policy is an order of magnitude larger once region fixed effects are controlled for. The peak increase in unemployment is 2.72%, which occurs 15 months after

---

[11]As a reminder, weighting by size increases power of the instrument as large shocks are more likely to affect the aggregate.

the federal funds shock. Although this result is not completely consistent with the main specification, it is within the range of the confidence bands. The dynamics are also very consistent across specifications, with the peak in each specification exactly matching 15 months.

A synthetically created large state

In the final test, I group all states within one region together, to create an artificially large state. The four groups in the Census Bureau-designated regions are the Northeast, Midwest, South and West, and in the results presented in Figure 3.7 Panel D I have grouped the South states together. By creating a large region in the cross section of states, this has an added benefit of increasing the excess Herfindahl index as defined in equation (3.13). In theory this should improve the precision of the GIV. However, I find with this specification larger standard errors, and a peak impact of monetary policy on unemployment of close to 5%.

3.5.5 Sample Period

The main results have been estimated using the sample period 1981-2009. This removes the 1979-1981 period of reserves targeting and high inflation. Boivin et al. (2010) and Coibion (2012) have shown this period of unprecedented and unusual monetary policy distorts estimates of monetary policy effects. When this period is included in my sample, the first stage loses power and therefore any second stage estimates are unreliable.

My sample also removes the post great financial crisis period where the federal funds rate has been stuck at the lower bound. When this period is included, my results hold, with the effect of monetary policy slightly weaker. This is not surprising given their is no volatility in the key independent variable. On top of this, alternative monetary policy tools, such as quantitative easing, were implemented, which could distort results.

In sum, the sample period chosen for my main results is the most robust period for identifying a consistent estimate for the effect of monetary policy.

3.6 Conclusion

The paper contributes to answering a core economic question: how do monetary policy actions today affect economic variables tomorrow? By applying a GIV approach, the iden-

tification technique used does not require monetary policy shocks, which, while typically used by the literature, are very small and therefore have poor statistical power. Instead, a GIV approach harnesses idiosyncratic shocks in the cross-section of the economy to make causal inference on the effect of monetary policy in the aggregate.

Implementing the GIV approach in the US on monthly unemployment data that is available at a state-level, I find that the US unemployment rate increases 1.80 percentage points 15 months subsequent to a 1 percentage point increase in the federal funds rate. The estimated impact is larger than existing estimates in the literature, and points to monetary policy being more powerful than previously thought.

The potential application of GIVs in a monetary policy setting are quite ubiquitous, with alternative regions and alternative cross sections (i.e. by industry rather than by geography) also implementable. Further exploration of these options may prove fruitful in terms of identifying more precise estimates of monetary policy effects via GIV methods. These studies are left to future work.

3.7   Figures

Figure 3.1: Heterogeneous economies across the US. This figure presents state-level variation in economic activity. States are sorted into five portfolios conditional on an economic variable of interest in each time period. The portfolio averages with respect to the conditioning variable are then plotted over time. Each panel presents results using output growth per capita, output gap, unemployment rate and inflation rate respectively (panels (a)-(d)) The state-level output gap is estimated from the cyclical component of the Hodrick and Prescott (1997) filter on real output. State-level unemployment rates have been demeaned. All panels include NBER defined recession periods as shaded regions.

(a) Real Output Growth per Capita

(b) Real Output Gap

(c) (Demeaned) Unemployment Rate

(d) Inflation Rate (p.a.)

Figure 3.2: Heterogeneous optimal Taylor rules across the US. This figure presents variation in optimal (Taylor rule) fed fund rates across states. For each state, the optimal fed fund rate is defined as

$$r_{i,t}^* = 2 + \pi_{i,t} + 0.5(\pi_{i,t} - 2) - (u_{i,t} - u_t^*)$$

where $\pi_{i,t}$ is the inflation rate of state i at time t, $u_{i,t}$ is the unemployment rate of state i at time t, and u_t^* is the long-run natural rate of unemployment for the US at time t. The Taylor rule residual $\epsilon_{i,t}^{TR} = r_{i,t}^* - r_t$ is the difference between the optimal rate at time t and the actual federal funds rate at time t. Panel (a) [(b)] sorts states into five portfolios at each date conditional on their optimal fed funds rate [taylor rule residual], with the portfolio average optimal fed funds rate [taylor rule residual] plotted over time. The shaded regions are NBER defined recession periods.

(a) Optimal fed funds rate

(b) Taylor rule residuals

152

Figure 3.3: Optimal Taylor rules: Florida and North Carolina. This figure presents the optimal (Taylor rule) fed fund rate for Florida and North Carolina over 1989-2009. The optimal fed fund rate for each state is defined as

$$r^*_{i,t} = 2 + \pi_{i,t} + 0.5(\pi_{i,t} - 2) - (u_{i,t} - u^*_t)$$

where $\pi_{i,t}$ is the inflation rate of state i at time t, $u_{i,t}$ is the unemployment rate of state i at time t, and u^*_t is the long-run natural rate of unemployment for the US at time t. The shaded regions are NBER defined recession periods.

Figure 3.4: Variation in the size of US states. This figure presents a measure of the size variation across US states based on their populations. It plots the excess Herfindahl index

$$h_t = \sqrt{-\frac{1}{N} + \sum_{i=1}^{N} S_{i,t}^2}$$

where $S_{i,t}$ is state i's fraction of the aggregated US population at date t. $N = 51$ covers the 50 states plus the District of Columbia.

Figure 3.5: Unemployment dynamics following a federal funds rate innovation. This figure presents the impulse response of the US unemployment rate following innovations in the fed funds rate. It reports the estimated sequence of coefficients $\{\beta_h^r\}$ from equation

$$\Delta_h u_{E,t+h} = \alpha^h + \beta_h^r \Delta r_t + \sum_{j=1}^{3} \beta_{h,j}^{u_E} \Delta u_{E,t-j} + \sum_{f=1}^{3} \beta_{h,f}^{\eta} \Delta \eta_t^f + \beta_h^\pi \Delta \pi_t + \epsilon_{t+h}, \quad h = 1, \ldots, 40$$

where where $\Delta_h u_{E,t+h}$ is the change in the equal-weighted unemployment rate from month t to $t+h$, α is a constant and Δr_t is the change in the fed funds rate at time t. The fed funds rate r_t is instrumented by the granular instrumental variable $z_t = u_{\Gamma,t}$. Control variables include lags of the dependent variable $\Delta u_{E,t-j}$, estimated unemployment factors η_t^f that are common across US states, and annualised inflation $\Delta \pi_t$. The dashed lines represent 90% confidence intervals computed using Newey and West (1987) standard errors with lag selection h. The estimation period is 1981-2009.

Figure 3.6: GIVs Instrument Analysis. This figure first presents the the time series of the instrument z_t as well as the underlying components $u_{S,t}$ and $u_{E,t}$. The second figure presents the residual from a regression of the instrument z_t on the common factors η_t^1, η_t^2 and η_t^3. The identifying assumption of the GIV estimation is that this variable contains idiosyncratic shocks only.

(a) Raw instrument

(b) Instrument purged of common factors

Figure 3.7: Tests for over-identifying restrictions. This figure presents the over identification tests of the impulse response of the US unemployment rate following innovations in the fed funds rate. It reports the estimated sequence of coefficients $\{\beta_h^r\}$ from equation

$$\Delta_h u_{E,t+h} = \alpha^h + \beta_h^r \Delta r_t + \sum_{j=1}^{3} \beta_{h,j}^{u_E} \Delta u_{E,t-j} + \sum_{f=1}^{3} \beta_{h,f}^{\eta} \Delta \eta_t^f + \beta_h^\pi \Delta \pi_t + \epsilon_{t+h}, \quad h = 1, \ldots, 40$$

where where $\Delta_h u_{E,t+h}$ is the change in the equal-weighted unemployment rate from month t to $t+h$, α is a constant and Δr_t is the change in the fed funds rate at time t. The fed funds rate r_t is instrumented by the granular instrumental variable . Control variables include lags of the dependent variable $\Delta u_{E,t-j}$, estimated unemployment factors η_t^f that are common across US states, and annualised inflation $\Delta \pi_t$. The dashed lines represent 90% confidence intervals computed using Newey and West (1987) standard errors with lag selection h. The figures are robustness checks to the main specification estimated in Figure 3.5. In panels A and B, states are sorted by size in each period, creating two instruments constructed purely on odd or even states respectively. In panel C, 4 region fixed effects are included in the control vector η_t^f. In panel D, the South region states has been grouped together to construct a synthetically created large state.

(a) Odd-Even instruments (odd)

(b) Odd-Even instruments (even)

(c) Region fixed effects

(d) A synthetically created large state

3.8   Tables

Table 3.1: Descriptive statistics of state-level economic variables. This table provides summary statistics for state-level economic variables and explores the sources of variation. Unemployment rate data is available monthly (1976-2019) and is provided by the U.S. Bureau of Labor Statistic. Inflation data is available quarterly (1979-2019) and is taken from Hazell et al. (2020). Real GDP growth per capita is available annual (1967-2019) and is provided by the U.S. Bureau of Economic Analysis. All three economic variables are available for the 50 states plus the District of Columbia. To explore sources of variation in the panel data, column $\sigma(x_{it})$ first reports the standard deviation of variable x across the whole panel. Column $\sigma(\bar{x}_t)$ reports the time-series standard deviation of the panel's cross-sectional averages. The final column reports the time-series average of the cross-sectional standard deviation. For every time observation t, the cross sectional standard deviation $\sigma(x_i)$ is calculated first, with the column presenting the mean of these standard deviations across all t. Standard deviations are unannualised in all columns and rows of the table.

	observations	mean	min	max	$\sigma(x_{it})$	$\sigma(\bar{x}_t)$	$\overline{\sigma(x_i)}_t$
Unemployment rate	26,673	5.94	2.10	18.70	1.95	1.45	1.46
Inflation	4,615	2.68	-4.03	14.68	1.95	1.83	1.08
Real GDP growth per capita	2,466	1.34	-34.99	29.54	3.74	2.23	2.77

Table 3.2: State-level Taylor rule residuals (1989-2009). This table presents analysis of Taylor rule residuals at a state-level over the period 1989-2009. The Taylor rule residual $\epsilon_{i,t}^{TR} = r_{i,t}^* - r_t$ is the difference between the optimal rate at time t and the actual federal funds rate at time t. The optimal fed fund rate for state i at time t is defined as $r_{i,t}^* = 2 + \pi_{i,t} + 0.5(\pi_{i,t} - 2) - (u_{i,t} - u_t^*)$ where $\pi_{i,t}$ is the inflation rate, $u_{i,t}$ is the unemployment rate of state i, and u_t^* is the long-run natural rate of unemployment for the US. The columns present each state's: (i) optimal rate correlation with the actual rate, (ii) standard deviation of the Taylor rule residual, (iii) minimum of the Taylor rule residual, (iv) maximum of the Taylor rule residual, and (v) mean of the absolute of the Taylor rule residual. The first row presents these statistics for the USA as an aggregated economy, with the remaining rows presenting statistics of individual states. Results are preseneted for the 39 states that are included in the Hazell et al. (2020) dataset.

	$Corr\left(r_t, r_{i,t}^*\right)$	$\sigma\left(\epsilon_{i,t}^{TR}\right)$	$min\left(\epsilon_{i,t}^{TR}\right)$	$max\left(\epsilon_{i,t}^{TR}\right)$	$mean\left(|\epsilon_{i,t}^{TR}|\right)$	observations
USA Aggregrate	0.92	0.96	-2.54	1.48	0.88	80
AK	0.21	2.45	-4.29	5.54	2.55	80
AL	0.31	2.31	-6.60	3.98	1.82	80
AR	0.45	2.08	-4.97	4.87	1.71	80
CA	0.46	2.39	-3.88	5.33	2.07	80
CO	0.39	2.34	-9.15	2.72	2.11	80
CT	0.55	2.21	-4.78	4.65	1.83	80
DC	0.30	3.01	-4.73	8.19	2.49	80
FL	0.33	2.34	-4.65	3.99	1.95	80
GA	0.69	1.62	-3.78	2.81	1.44	80
HI	0.21	3.83	-9.66	6.54	3.52	80
IL	0.67	1.68	-4.43	3.29	1.36	80
IN	0.47	2.07	-6.05	3.83	1.76	80
KS	0.36	2.41	-7.09	5.99	2.15	80
LA	0.26	2.51	-6.61	6.47	2.10	80
MA	0.58	2.15	-7.55	2.96	1.83	80
MD	0.17	2.90	-11.90	3.17	2.15	80
MI	0.38	2.32	-3.55	6.11	1.94	80
MN	0.22	2.38	-7.23	4.60	2.39	80
MO	0.67	1.58	-3.95	2.37	1.37	80
MS	0.08	2.74	-6.66	4.20	2.22	80
NC	0.71	1.82	-4.59	4.32	1.50	80
NJ	0.18	3.01	-7.85	6.78	2.32	80
NY	0.49	2.18	-5.62	3.32	1.76	80
OH	0.61	1.71	-5.26	3.13	1.29	80
OK	0.17	2.92	-10.34	3.74	2.63	79
OR	0.73	1.60	-5.66	2.54	1.28	80
PA	0.44	2.13	-6.90	3.17	1.61	80
SC	0.75	1.63	-3.59	3.59	1.28	80
TN	0.46	2.02	-5.67	3.47	1.69	80
TX	0.26	2.47	-7.30	4.22	1.97	80
UT	0.67	1.79	-5.62	2.94	1.60	80
VA	0.40	2.35	-7.87	2.14	2.18	80
WA	0.60	1.87	-6.28	3.32	1.37	80
WI	0.46	1.91	-6.72	2.33	2.00	80

Table 3.3: Analysis of first stage results in the GIVs estimation. This table provides first stage analysis of the granular instrumental variables estimation described in equation (3.20). It presents estimates from the first stage specification:

$$\Delta r_t = \alpha + \beta^z \Delta z_t + \sum_{j=1}^{J} \beta_j^{u_E} \Delta u_{E,t-j} + \sum_{f=1}^{F} \beta_f^{\eta} \Delta \eta_t^f + \beta^{\pi} \Delta \pi_t + \epsilon_t, \quad F = 0, 1, \dots, 6$$

where the parameter estimates are reported for β^z and β^{π} only. To evaluate the strength of the instrument, the Cragg-Donald Wald F-statistic rank test is also reported. Each column corresponds to an estimation with $F = 0, 1, \dots, 6$ common factors included as controls variables.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Instrument (Z)	-1.541*	-2.344**	-3.698***	-4.300***	-4.002***	-4.051***	-3.926***
	(-1.78)	(-2.35)	(-3.34)	(-3.83)	(-2.80)	(-2.70)	(-2.60)
Inflation (π)	0.208**	0.161	0.163	0.161	0.158	0.158	0.160
	(2.12)	(1.58)	(1.61)	(1.61)	(1.57)	(1.57)	(1.58)
Adj R-sq	0.139	0.143	0.160	0.175	0.173	0.170	0.169
Observations	324	324	324	324	324	324	324
Cragg-Donald F-stat	2.353	3.254	8.423	10.082	4.375	4.833	4.619
Common factors	0	1	2	3	4	5	6

3.9 Appendix

3.9.1 Identifying assumptions

Stripped back to fundamentals, the insight of the econometrics behind GIVs is that the instrument, which is gamma-weighted local shocks ν_Γ, is correlated with size-weighted local shocks ν_S, but is *not* correlated with equal-weighted local shocks ν_E.

Note that ν i.i.d across i means $\mathbb{E}\left[\nu_i\nu_j\right] = \sigma_\nu^2$ if $i = j$, and $\mathbb{E}\left[\nu_i\nu_j\right] = 0$ if $i \neq j$. Using this result we see

Exclusion Restriction holds: $\mathbb{E}\left[\nu_{E,t}\nu_{\Gamma,t}\right] = 0$

$$
\begin{aligned}
\mathbb{E}\left[\nu_{E,t}\nu_{\Gamma,t}\right] &= \mathbb{E}\left[\left(\sum_i \frac{1}{N}\nu_{it}\right)\left(\sum_i \Gamma_i\nu_{it}\right)\right] \\
&= \mathbb{E}\left[\left(\sum_i \frac{1}{N}\nu_{it}\right)\left(\sum_i S_i\nu_{it} - \sum_i \frac{1}{N}\nu_{it}\right)\right] \\
&= \mathbb{E}\left[\left(\sum_i \frac{1}{N}\nu_{it}\right)\left(\sum_i S_i\nu_{it}\right)\right] - \mathbb{E}\left[\left(\sum_i \frac{1}{N}\nu_{it}\right)\left(\sum_i \frac{1}{N}\nu_{it}\right)\right] \\
&= \frac{1}{N}\sum_i S_i\sigma_\nu^2 - \frac{1}{N^2}\sum_i \sigma_\nu^2 \\
&= \frac{1}{N}\sigma_\nu^2\left(\sum_i S_i - \frac{1}{N}\sum_i 1\right) \\
&= 0
\end{aligned}
$$

Relevance condition holds: $\mathbb{E}\left[\nu_{S,t}\nu_{\Gamma,t}\right] \neq 0$

$$
\begin{aligned}
\mathbb{E}\left[\nu_{S,t}\nu_{\Gamma,t}\right] &= \mathbb{E}\left[\left(\sum_i S_i\nu_{it}\right)\left(\sum_i \Gamma_i\nu_{it}\right)\right] \\
&= \mathbb{E}\left[\left(\sum_i S_i\nu_{it}\right)\left(\sum_i S_i\nu_{it} - \sum_i \frac{1}{N}\nu_{it}\right)\right] \\
&= \mathbb{E}\left[\left(\sum_i S_i\nu_{it}\right)\left(\sum_i S_i\nu_{it}\right)\right] - \mathbb{E}\left[\left(\sum_i S_i\nu_{it}\right)\left(\sum_i \frac{1}{N}\nu_{it}\right)\right] \\
&= \sum_i S_i^2\sigma_\nu^2 - \frac{1}{N}\sum_i S_i\sigma_\nu^2 \\
&= \sigma_\nu^2\left(\sum_i S_i^2 - \frac{1}{N}\right) \\
&\neq 0
\end{aligned}
$$

Indeed, this final term shows $\mathbb{E}\left[\nu_{S,t}\nu_{\Gamma,t}\right]$ increases in σ_ν^2 and the variation in S_i. Consistent with the intuition, we can therefore see formally here that the instrument power is driven by the size of idiosyncratic shocks and by the size variation in the cross section.

3.9.2 Asymptotic Bias

Consider the structural equations

$$u_{S,t} = b^r r_t + \epsilon_t^{u_S} \tag{3.21}$$

$$r_t = b^u u_{S,t} + \epsilon_t^r \tag{3.22}$$

where $\epsilon_t^{u_S} = \lambda_S \eta_t + \nu_{S,t}$. Solving for r_t we have

$$r_t = b^u M \epsilon_t^{u_S} + M \epsilon_t^r \tag{3.23}$$

where $M = \frac{1}{1-b^u b^r}$.

OLS Bias

Bias of estimating equation (3.21) via OLS is

$$\begin{aligned}
\hat{b}_{OLS}^r - b^r &= \frac{Cov\left(r_t, \epsilon_t^{u_S}\right)}{Var\left(r_t\right)} \\
&= \frac{Cov\left(b^u M \epsilon_t^{u_S} + M \epsilon_t^r, \epsilon_t^{u_S}\right)}{Var\left(r_t\right)} \\
&= b^u M \frac{Var\left(\epsilon_t^{u_S}\right)}{Var\left(r_t\right)} \\
&< 0.
\end{aligned}$$

The inequality uses $b^r > 0$ (unemployment increases when fed funds increases) and $b^u < 0$ (fed funds down when unemployment up), which means $b^r b^u < 0$ and $0 < M < 1$ and therefore $b^u M < 0$. A positive shock to unemployment $\epsilon_t^{u_S}$ results in a lower interest rate endogenously, which means the estimated reaction of unemployment to interest rates is lower than the true relationship.

IV Bias without controlling for common factor

The IV estimation is regressing the equal-weighted average unemployment on the interest rate

$$u_{E,t} = b^r r_t + \epsilon_t^{u_E} \tag{3.24}$$

with the instrument $z_t = u_{\Gamma,t} = \lambda_\Gamma \eta_t + \nu_{\Gamma,t}$. If the common factor η_t is not controlled for in the regression, then the IV bias is

$$
\begin{aligned}
\hat{b}^r_{IV} - b^r &= \frac{Cov\left(z_t, \epsilon^{u_S}_t\right)}{Cov\left(z_t, r_t\right)} \\
&= \frac{Cov\left(\lambda_\Gamma \eta_t + \nu_{\Gamma,t}, \lambda_E \eta_t + \nu_{E,t}\right)}{Cov\left(\lambda_\Gamma \eta_t + \nu_{\Gamma,t}, b^u M \epsilon^{u_S}_t + M \epsilon^r_t\right)} \\
&= \frac{Cov\left(\lambda_\Gamma \eta_t + \nu_{\Gamma,t}, \lambda_E \eta_t + \nu_{E,t}\right)}{Cov\left(\lambda_\Gamma \eta_t + \nu_{\Gamma,t}, b^u M \left(\lambda_S \eta_t + \nu_{S,t}\right)\right)} \\
&= \frac{\lambda_\Gamma \lambda_E Var\left(\eta_t\right)}{b^u M \lambda_\Gamma \lambda_S Var\left(\eta_t\right) + b^u M Cov\left(\nu_{\Gamma,t}, \nu_{S,t}\right)} \\
&= \frac{\lambda_E}{b^u M} \left(\frac{1}{\lambda_S} + \frac{\lambda_\Gamma Var\left(\eta_t\right)}{Cov\left(\nu_{\Gamma,t}, \nu_{S,t}\right)}\right)
\end{aligned}
$$

The sign of the IV bias (unlike the OLS bias) is therefore undetermined and specification specific.

Bibliography

Adrian, T., E. Etula, and T. Muir (2014). Financial Intermediaries and the Cross-Section of Asset Returns. *Journal of Finance 69*(6), 2557–2596.

Andrews, I., J. H. Stock, and L. Sun (2019). Weak Instruments in Instrumental Variables Regression: Theory and Practice. *Annual Review of Economics 11*, 727–753.

Barakchian, S. M. and C. Crowe (2013). Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics 60*(8), 950–966.

Becker, B. and V. Ivashina (2015). Reaching for Yield in the Bond Market. *Journal of Finance 70*(5), 1863–1902.

Becker, B., M. M. Opp, and F. Saidi (2020). Regulatory Forbearance in the U.S. Insurance Industry: The Effects of Eliminating Capital Requirements. *Stockholm School of Economics Working Paper*.

Bernanke, B. S. (2015). The Taylor Rule: A benchmark for monetary policy? *The Brookings Institution*.

Boivin, J., M. T. Kiley, and F. S. Mishkin (2010). How has the monetary transmission mechanism evolved over time? *Handbook of Monetary Economics 3*(C), 369–422.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2020). Expectations of fundamentals and stock market puzzles. *Working Paper* (May).

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies 22*(6), 2201–2238.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review*.

Campbell, J. Y. (1991). A Variance Decomposition for Stock Returns. *The Economic Journal 101*(405), 157–179.

Campbell, J. Y. and J. Ammer (1993). What Moves the Stock and Bond Markets.

Campbell, J. Y. and R. J. Shiller (1988). The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies 1*(3), 195–228.

Chen, L. and X. Zhao (2009). Return decomposition. *Review of Financial Studies 22*(12), 5213–5249.

Chodorow-Reich, G., A. C. Ghent, and V. Haddad (2020). Asset Insulators. *Review of Financial Studies, forthcoming*.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Chapter 2 Monetary policy shocks: What have we learned and to what end? *Handbook of Macroeconomics 1*(PART A), 65–148.

Cochrane, J. H. (2011). Presidential Address: Discount Rates. *Journal of Finance 66*(4), 1047–1108.

Coibion, O. (2012). Are the effects of monetary policy shocks big or small? *American Economic Journal: Macroeconomics*.

Cox, J., D. L. Greenwald, and S. C. Ludvigson (2020). What Explains the COVID-19 Stock Market? *MIT Sloan Working Paper*.

Crone, T. M. (2005). An alternative definition of economic regions in the United States based on similarities in state business cycles. *Review of Economics and Statistics 87*(4), 617–626.

Dahlquist, M. and M. Ibert (2021). How Cyclical Are Stock Market Return Expectations? Evidence from Capital Market Assumptions. *Swedish House of Finance Research Paper No. 21-1* (April).

De La O, R. and S. Myers (2021). Subjective Cash Flow and Discount Rate Expectations. *Journal of Finance, Forthcoming* (0).

Diamond, D. W. and P. H. Dybvig (1983). Bank Runs, Deposit Insurance, and Liquidity. *Journal of Political Economy 24*(1), 14–23.

Ellul, A., C. Jotikasthira, A. V. Kartasheva, C. T. Lundblad, and W. Wagner (2020). Insurers As Asset Managers and Systemic Risk. *Indiana University Working Paper*.

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

Friedman, M. and A. J. Schwartz (1963). *A monetary history of the United States, 1867-1960*.

Froot, K. and P. O'Connell (1999). The pricing of US catastrophe reinsurance. *The Financing of Catastrophe Risk I*(January), 195–232.

Gabaix, X. and R. S. J. Koijen (2020). Granular Instrumental Variables. *Harvard Working Paper*.

Gabaix, X. and R. S. J. Koijen (2021). The Inelastic Markets Hypothesis. *Harvard Working Paper*.

Gao, C. and I. Martin (2019). Volatility, Valuation Ratios, and Bubbles: An Empirical Measure of Market Sentiment. *CEPR Discussion Papers*.

Ge, S. (2020). How Do Financial Constraints Affect Product Pricing? Evidence from Weather and Life Insurance Premiums. *Journal of Finance, Forthcoming*.

Ge, S. and M. S. Weisbach (2020). The Role of Financial Conditions in Portfolio Choices: The Case of Insurers. *Journal of Financial Economics, Forthcoming*.

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*.

Gilchrist, S. and E. Zakrajšek (2012). Credit spreads and business cycle fluctuations. *American Economic Reivew*.

Gormsen, N. J. and R. S. J. Koijen (2020). Coronavirus: Impact on Stock Prices and Growth Expectations. *The Review of Asset Pricing Studies*, 1–34.

Gormsen, N. J., R. S. J. Koijen, and I. Martin (2021). Implied Dividend Volatility and Expected Growth. *AEA Papers and Proceedings, forthcoming*.

Greenwald, D., M. Lettau, and S. Ludvigson (2019). How the Wealth Was Won: Factor Shares as Market Fundamentals. *NBER Working Paper Series* (510), 1–9.

Greenwood, R. and A. Vissing-Jorgensen (2018). The Impact of Pensions and Insurance on Global Yield Curves. *Harvard Business School Working Paper No. 18-109*.

Gron, A. (1994). Evidence of Capacity Constraints in Insurance Markets. *Journal of Law and Economics 37*(2), 349–377.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*(8), 2291–2304.

Hamilton, J. D. (2018). Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics 100*(5), 831–843.

Hanson, S. G., A. Shleifer, J. C. Stein, and R. W. Vishny (2015). Banks as patient fixed-income investors. *Journal of Financial Economics 117*(3), 449–469.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2020). The Slope of the Phillips Curve: Evidence from U.S. States. *NBER Working Paper Series*.

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

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

Hill, R. D. (1979). Profit Regulation in Property-Liability Insurance. *The Bell Jounral of Economics 10*(1), 172–191.

Hodrick, R. J. and E. C. Prescott (1997). Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking 29*(1), 1.

Ioannidou, V., S. R. G. Ongena, and J.-L. Peydro (2015). Monetary Policy, Risk-Taking and Pricing: Evidence from a Quasi-Natural Experiment. *Review of Finance 19*(1), 95–144.

Jiménez, B. G., S. Ongena, J.-l. Peydró, and J. Saurina (2012). Credit Supply and Monetary Policy : Identifying the Bank Balance-Sheet Channel with Loan Applications. *American Economic Reivew 102*(5), 2301–2326.

John Y. Campbell, G. B. T. (2003). Equity Volatility and Corporate Bond Yields. *Journal of Finance LVIII*(6), 2321–2349.

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *The American Economic Review 95*(1), 161–182.

Jordà, Ò., M. Schularick, and A. M. Taylor (2015). Betting the house. *Journal of International Economics 96*(S1), S2–S18.

Kashyap, A. K., R. Rajan, and J. C. Stein (2002). Banks as liquidity providers: An explanation for the coexistence of lending and deposit-taking. *Journal of Finance 57*(1), 33–73.

Koijen, R. S. and M. Yogo (2015). The cost of financial frictions for life insurers. *American Economic Review 105*(1), 445–475.

Kraus, A. and S. A. Ross (1982). The Determination of Fair Profits for the Property-Liability Insurance Firm. *The Journal of Finance 37*(4), 1015–1028.

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

Landier, A. and D. Thesmar (2020). Earnings Expectations in the COVID Crisis. *MIT Sloan Working Paper*, 1–26.

Leombroni, M., A. Vedolin, G. Venter, and P. Whelan (2021). Central Bank Communication and the Yield Curve. *Journal of Financial Economics, Forthcoming*.

Malkin, I. and F. Nechio (2012). U.S. and Euro-Area Monetary Policy by Regions. *FRBSF Economic Letter 06*.

Martin, I. (2017). What is the expected return on the market? *Quarterly Journal of Economics 132*(1), 367–433.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown (1999). New evidence on the money's worth of individual annuities. *American Economic Review 89*(5), 1299–1318.

Modigliani, F. and M. H. Miller (1958). The Cost of Capital, Corporation Finance and Theory of Investment. *The American Economic Review 48*(3), 261–297.

NAIC (2018). Accounting Practices and Procedures Manual. *Statutory Issue Paper No . 53 Property Casualty Contracts - Premiums* (53), 1–8.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics 133*(3), 1283–1330.

Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*.

Novy-Marx, R. and J. Rauh (2011). Public pension promises: How big are they and what are they worth? *Journal of Finance 66*(4), 1211–1249.

Odomirok, K. C., L. M. Mcfarlane, G. L. Kennedy, and J. J. Brenden (2014). Financial Reporting Through the Lens of a Property/Casualty Actuary. *Casualty Actuarial Society (CAS)*.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics 2*, 71–162.

Rauh, J. D. (2016). Hidden Debt, Hidden Deficits How Pension Promises are Consuming State and Local Budgets. *Hoover Institution Essay*.

Romer, C. D. and D. H. Romer (2000). Federal reserve information and the behavior of interest rates. *American Economic Review 90*(3), 429–457.

Romer, D. H. and C. D. Romer (2004). A New Measure of Monetary Shocks : Derivation and Implications. *The American Economic Review 94*(4), 1055–1084.

Schultz, P. (2001). Corporate bond trading costs: A peek behind the curtain. *Journal of Finance 56*(2), 677–698.

Stein, J. C. (2012). Monetary policy as financial stability regulation. *Quarterly Journal of Economics*.

Stock, J. H. and M. Yogo (2005). Testing for weak instruments in Linear Iv regression. In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Confer. Series on Public Policy 39*(C), 195–214.

van Binsbergen, J. H. (2020). Duration-Based Stock Valuation: Reassessing Stock Market Performance and Volatility. *Wharton Working Paper*.

van Binsbergen, J. H., W. Hueskes, R. S. J. Koijen, and E. B. Vrugt (2013). Equity Yields. *Journal of Financial Economics 110*(3), 503–519.

Vissing-Jørgensen, A. (2020). Bond Markets in Spring 2020 and the Response of the Federal Reserve. *Berkeley Haas Working Paper*.

Vuolteenaho, T. and J. Y. Campbell (2004). Bad Beta, Good Beta. *American Economic Reivew 94* (5), 1249– 1275.

*A case study of the Fashion and
Design Branch of the Industrial District
of Montebelluna, NE Italy*

12. Mikkel Flyverbom
*Making the Global Information Society
Governable*
*On the Governmentality of Multi-
Stakeholder Networks*

13. Anette Grønning
Personen bag
*Tilstedevær i e-mail som inter-
aktionsform mellem kunde og med-
arbejder i dansk forsikringskontekst*

14. Jørn Helder
One Company – One Language?
The NN-case

15. Lars Bjerregaard Mikkelsen
*Differing perceptions of customer
value*
*Development and application of a tool
for mapping perceptions of customer
value at both ends of customer-suppli-
er dyads in industrial markets*

16. Lise Granerud
Exploring Learning
*Technological learning within small
manufacturers in South Africa*

17. Esben Rahbek Pedersen
*Between Hopes and Realities:
Reflections on the Promises and
Practices of Corporate Social
Responsibility (CSR)*

18. Ramona Samson
*The Cultural Integration Model and
European Transformation.*
The Case of Romania

2007
1. Jakob Vestergaard
Discipline in The Global Economy
*Panopticism and the Post-Washington
Consensus*

2. Heidi Lund Hansen
Spaces for learning and working
*A qualitative study of change of work,
management, vehicles of power and
social practices in open offices*

3. Sudhanshu Rai
*Exploring the internal dynamics of
software development teams during
user analysis*
*A tension enabled Institutionalization
Model; "Where process becomes the
objective"*

4. Norsk ph.d.
Ej til salg gennem Samfundslitteratur

5. Serden Ozcan
*EXPLORING HETEROGENEITY IN
ORGANIZATIONAL ACTIONS AND
OUTCOMES*
A Behavioural Perspective

6. Kim Sundtoft Hald
*Inter-organizational Performance
Measurement and Management in
Action*
*– An Ethnography on the Construction
of Management, Identity and
Relationships*

7. Tobias Lindeberg
Evaluative Technologies
*Quality and the Multiplicity of
Performance*

8. Merete Wedell-Wedellsborg
Den globale soldat
*Identitetsdannelse og identitetsledelse
i multinationale militære organisatio-
ner*

9. Lars Frederiksen
*Open Innovation Business Models
Innovation in firm-hosted online user
communities and inter-firm project
ventures in the music industry
– A collection of essays*

10. Jonas Gabrielsen
*Retorisk toposlære – fra statisk 'sted'
til persuasiv aktivitet*

integration in the process of product innovation

19. Throstur Olaf Sigurjonsson
   Governance Failure and Icelands's Financial Collapse

20. Allan Sall Tang Andersen
   Essays on the modeling of risks in interest-rate and inflation markets

21. Heidi Tscherning
   Mobile Devices in Social Contexts

22. Birgitte Gorm Hansen
   Adapting in the Knowledge Economy Lateral Strategies for Scientists and Those Who Study Them

23. Kristina Vaarst Andersen
   Optimal Levels of Embeddedness The Contingent Value of Networked Collaboration

24. Justine Grønbæk Pors
   Noisy Management A History of Danish School Governing from 1970-2010

25. Stefan Linder
   Micro-foundations of Strategic Entrepreneurship Essays on Autonomous Strategic Action

26. Xin Li
   Toward an Integrative Framework of National Competitiveness An application to China

27. Rune Thorbjørn Clausen
   Værdifuld arkitektur Et eksplorativt studie af bygningers rolle i virksomheders værdiskabelse

28. Monica Viken
   Markedsundersøkelser som bevis i varemerke- og markedsføringsrett

29. Christian Wymann
   Tattooing The Economic and Artistic Constitution of a Social Phenomenon

30. Sanne Frandsen
   Productive Incoherence A Case Study of Branding and Identity Struggles in a Low-Prestige Organization

31. Mads Stenbo Nielsen
   Essays on Correlation Modelling

32. Ivan Häuser
   Følelse og sprog Etablering af en ekspressiv kategori, eksemplificeret på russisk

33. Sebastian Schwenen
   Security of Supply in Electricity Markets

2012
1. Peter Holm Andreasen
   The Dynamics of Procurement Management - A Complexity Approach

2. Martin Haulrich
   Data-Driven Bitext Dependency Parsing and Alignment

3. Line Kirkegaard
   Konsulenten i den anden nat En undersøgelse af det intense arbejdsliv

4. Tonny Stenheim
   Decision usefulness of goodwill under IFRS

5. Morten Lind Larsen
   Produktivitet, vækst og velfærd Industrirådet og efterkrigstidens Danmark 1945 - 1958

6. Petter Berg
   Cartel Damages and Cost Asymmetries

7. Lynn Kahle
   Experiential Discourse in Marketing A methodical inquiry into practice and theory

8. Anne Roelsgaard Obling
   Management of Emotions in Accelerated Medical Relationships

2016

1. Signe Sofie Dyrby
   Enterprise Social Media at Work

2. Dorte Boesby Dahl
   The making of the public parking attendant
   Dirt, aesthetics and inclusion in public service work

3. Verena Girschik
   Realizing Corporate Responsibility Positioning and Framing in Nascent Institutional Change

4. Anders Ørding Olsen
   IN SEARCH OF SOLUTIONS
   Inertia, Knowledge Sources and Diversity in Collaborative Problem-solving

5. Pernille Steen Pedersen
   Udkast til et nyt copingbegreb
   En kvalifikation af ledelsesmuligheder for at forebygge sygefravær ved psykiske problemer.

6. Kerli Kant Hvass
   Weaving a Path from Waste to Value: Exploring fashion industry business models and the circular economy

7. Kasper Lindskow
   Exploring Digital News Publishing Business Models – a production network approach

8. Mikkel Mouritz Marfelt
   The chameleon workforce: Assembling and negotiating the content of a workforce

9. Marianne Bertelsen
   Aesthetic encounters
   Rethinking autonomy, space & time in today's world of art

10. Louise Hauberg Wilhelmsen
    EU PERSPECTIVES ON INTERNATIONAL COMMERCIAL ARBITRATION

11. Abid Hussain
    On the Design, Development and Use of the Social Data Analytics Tool (SODATO): Design Propositions, Patterns, and Principles for Big Social Data Analytics

12. Mark Bruun
    Essays on Earnings Predictability

13. Tor Bøe-Lillegraven
    BUSINESS PARADOXES, BLACK BOXES, AND BIG DATA: BEYOND ORGANIZATIONAL AMBIDEXTERITY

14. Hadis Khonsary-Atighi
    ECONOMIC DETERMINANTS OF DOMESTIC INVESTMENT IN AN OIL-BASED ECONOMY: THE CASE OF IRAN (1965-2010)

15. Maj Lervad Grasten
    Rule of Law or Rule by Lawyers?
    On the Politics of Translation in Global Governance

16. Lene Granzau Juel-Jacobsen
    SUPERMARKEDETS MODUS OPERANDI – en hverdagssociologisk undersøgelse af forholdet mellem rum og handlen og understøtte relationsopbygning?

17. Christine Thalsgård Henriques
    In search of entrepreneurial learning – Towards a relational perspective on incubating practices?

18. Patrick Bennett
    Essays in Education, Crime, and Job Displacement

19. Søren Korsgaard
    Payments and Central Bank Policy

20. Marie Kruse Skibsted
    Empirical Essays in Economics of Education and Labor

21. Elizabeth Benedict Christensen
    The Constantly Contingent Sense of Belonging of the 1.5 Generation Undocumented Youth
    An Everyday Perspective

33. Thomas Jensen
*Shipping Information Pipeline:
An information infrastructure to
improve international containerized
shipping*

34. Dzmitry Bartalevich
*Do economic theories inform policy?
Analysis of the influence of the Chicago
School on European Union competition
policy*

35. Kristian Roed Nielsen
*Crowdfunding for Sustainability: A
study on the potential of reward-based
crowdfunding in supporting sustainable
entrepreneurship*

36. Emil Husted
*There is always an alternative: A study
of control and commitment in political
organization*

37. Anders Ludvig Sevelsted
*Interpreting Bonds and Boundaries of
Obligation. A genealogy of the emer-
gence and development of Protestant
voluntary social work in Denmark as
shown through the cases of the Co-
penhagen Home Mission and the Blue
Cross (1850 – 1950)*

38. Niklas Kohl
Essays on Stock Issuance

39. Maya Christiane Flensborg Jensen
*BOUNDARIES OF
PROFESSIONALIZATION AT WORK
An ethnography-inspired study of care
workers' dilemmas at the margin*

40. Andreas Kamstrup
*Crowdsourcing and the Architectural
Competition as Organisational
Technologies*

41. Louise Lyngfeldt Gorm Hansen
*Triggering Earthquakes in Science,
Politics and Chinese Hydropower
- A Controversy Study*

2018

1. Vishv Priya Kohli
*Combatting Falsifi cation and Coun-
terfeiting of Medicinal Products in
the E uropean Union – A Legal
Analysis*

2. Helle Haurum
*Customer Engagement Behavior
in the context of Continuous Service
Relationships*

3. Nis Grünberg
*The Party -state order: Essays on
China's political organization and
political economic institutions*

4. Jesper Christensen
*A Behavioral Theory of Human
Capital Integration*

5. Poula Marie Helth
Learning in practice

6. Rasmus Vendler Toft-Kehler
*Entrepreneurship as a career? An
investigation of the relationship
between entrepreneurial experience
and entrepreneurial outcome*

7. Szymon Furtak
*Sensing the Future: Designing
sensor-based predictive information
systems for forecasting spare part
demand for diesel engines*

8. Mette Brehm Johansen
*Organizing patient involvement. An
ethnographic study*

9. Iwona Sulinska
*Complexities of Social Capital in
Boards of Directors*

10. Cecilie Fanøe Petersen
*Award of public contracts as a
means to conferring State aid: A
legal analysis of the interface
between public procurement law
and State aid law*

11. Ahmad Ahmad Barirani
*Three Experimental Studies on
Entrepreneurship*

2019

1. Shihan Du
ESSAYS IN EMPIRICAL STUDIES BASED ON ADMINISTRATIVE LABOUR MARKET DATA

2. Mart Laatsit
Policy learning in innovation policy: A comparative analysis of European Union member states

3. Peter J. Wynne
Proactively Building Capabilities for the Post-Acquisition Integration of Information Systems

4. Kalina S. Staykova
Generative Mechanisms for Digital Platform Ecosystem Evolution

5. Ieva Linkeviciute
Essays on the Demand-Side Management in Electricity Markets

6. Jonatan Echebarria Fernández
Jurisdiction and Arbitration Agreements in Contracts for the Carriage of Goods by Sea – Limitations on Party Autonomy

7. Louise Thorn Bøttkjær
Votes for sale. Essays on clientelism in new democracies.

8. Ditte Vilstrup Holm
The Poetics of Participation: the organizing of participation in contemporary art

9. Philip Rosenbaum
Essays in Labor Markets – Gender, Fertility and Education

10. Mia Olsen
Mobile Betalinger - Succesfaktorer og Adfærdsmæssige Konsekvenser

11. Adrián Luis Mérida Gutiérrez
Entrepreneurial Careers: Determinants, Trajectories, and Outcomes

12. Frederik Regli
Essays on Crude Oil Tanker Markets

13. Cancan Wang
Becoming Adaptive through Social Media: Transforming Governance and Organizational Form in Collaborative E-government

14. Lena Lindbjerg Sperling
Economic and Cultural Development: Empirical Studies of Micro-level Data

15. Xia Zhang
Obligation, face and facework: An empirical study of the communicative act of cancellation of an obligation by Chinese, Danish and British business professionals in both L1 and ELF contexts

16. Stefan Kirkegaard Sløk-Madsen
Entrepreneurial Judgment and Commercialization

17. Erin Leitheiser
*The Comparative Dynamics of Private Governance
The case of the Bangladesh Ready-Made Garment Industry*

18. Lone Christensen
STRATEGIIMPLEMENTERING: STYRINGSBESTRÆBELSER, IDENTITET OG AFFEKT

19. Thomas Kjær Poulsen
Essays on Asset Pricing with Financial Frictions

20. Maria Lundberg
Trust and self-trust in leadership identity constructions: A qualitative exploration of narrative ecology in the discursive aftermath of heroic discourse

TITLER I ATV PH.D.-SERIEN

1992
1. Niels Kornum
   Servicesamkørsel – organisation, øko-nomi og planlægningsmetode

1995
2. Verner Worm
   *Nordiske virksomheder i Kina
   Kulturspecifikke interaktionsrelationer
   ved nordiske virksomhedsetableringer i
   Kina*

1999
3. Mogens Bjerre
   *Key Account Management of Complex
   Strategic Relationships
   An Empirical Study of the Fast Moving
   Consumer Goods Industry*

2000
4. Lotte Darsø
   *Innovation in the Making
   Interaction Research with heteroge-neous Groups of Knowledge Workers
   creating new Knowledge and new
   Leads*

2001
5. Peter Hobolt Jensen
   *Managing Strategic Design Identities
   The case of the Lego Developer Net-work*

2002
6. Peter Lohmann
   *The Deleuzian Other of Organizational
   Change – Moving Perspectives of the
   Human*

7. Anne Marie Jess Hansen
   *To lead from a distance: The dynamic
   interplay between strategy and strate-gizing – A case study of the strategic
   management process*

2003
8. Lotte Henriksen
   *Videndeling
   – om organisatoriske og ledelsesmæs-sige udfordringer ved videndeling i
   praksis*

9. Niels Christian Nickelsen
   *Arrangements of Knowing: Coordi-nating Procedures Tools and Bodies in
   Industrial Production – a case study of
   the collective making of new products*

2005
10. Carsten Ørts Hansen
    *Konstruktion af ledelsesteknologier og
    effektivitet*

TITLER I DBA PH.D.-SERIEN

2007
1. Peter Kastrup-Misir
   *Endeavoring to Understand Market
   Orientation – and the concomitant
   co-mutation of the researched, the
   re searcher, the research itself and the
   truth*

2009
1. Torkild Leo Thellefsen
   *Fundamental Signs and Significance
   effects
   A Semeiotic outline of Fundamental
   Signs, Significance-effects, Knowledge
   Profiling and their use in Knowledge
   Organization and Branding*

2. Daniel Ronzani
   *When Bits Learn to Walk Don't Make
   Them Trip. Technological Innovation
   and the Role of Regulation by Law
   in Information Systems Research: the
   Case of Radio Frequency Identification
   (RFID)*

2010
1. Alexander Carnera
   *Magten over livet og livet som magt
   Studier i den biopolitiske ambivalens*