Three Essays in Financial Economics

by

Zhen Yan

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2019

Doctoral Committee:

Professor Stefan Nagel, Co–Chair
Professor Uday Rajan, Co–Chair
Assistant Professor Serhiy Kozak
Assistant Professor Heng Liu
Assistant Professor Yesim Orhun

Zhen Yan

zachyan@umich.edu

ORCID iD: 0000-0002-3493-5213

To my wife Lai Wei.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my deepest gratitude to my primary adviser Professor Stefan Nagel. It goes without saying that without his advice, encouragement, and support this dissertation would not have been possible. Over the last five years, I have been benefited tremendously from our discussions through emails, via Skype, and in person. I am also particularly grateful to Stefan's generous hospitality at the University of Chicago, which have made the main chapter of this dissertation possible even after his departure from the University of Michigan.

I would also take this opportunity to thank my co-chair Professor Uday Rajan, who was always ready to have long discussions on my papers and presentations, among many others. His feedback is always constructive and invaluable. I am also grateful to Professor Serhiy Kozak, who helped me sharpen my dissertation to a great extent. I would also like to thank Professor Heng Liu and Yesim Orhun, who dedicated many hours to this dissertation.

I would like to express my great appreciation to friends, fellow PhD students, and faculty members both within and outside the University of Michigan, who have provided many insightful suggestions and comments on various versions of the chapters in my dissertation. I have benefited from discussions with Matteo Crosignani, Zhiguo He, Jiacui Li, Ulrike Malmendier, Indrajit Mitra, Paolo Pasquariello, Amiyatosh Purnanandam, Martin Schmalz, Tyler Shumway, Zhengyang Xu, Stefan Zeume, and participants at 17th Trans-Atlantic Doctoral Conference (London Business School), Michigan Ross

Brown-bag seminar, University of Florida, and University of Illinois at Chicago. I would also thank the doctoral program at Ross School of Business for providing the financial support.

Finally, I would like to thank my wife Lai Wei. Her understanding, support, and belief in me are among the most important reasons that I could complete the doctoral studies. I have been fortunate to meet her at Michigan and to share with her both good times and difficult phases of my PhD study. I dedicate my dissertation and my doctorate to her.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation examines i) how market structure and the beliefs of market participants impact U.S. corporate bond mutual funds as well as the underlying bond market, and ii) whether personal experiences affect central bankers' belief formation and their decision making.

In the first chapter of the dissertation, entitled "Returns to Scale among Corporate Bond Mutual Funds", I document a (within-fund) hump-shaped relation between fund size and subsequent fund alpha among U.S. corporate bond mutual funds. When funds are small, they exhibit *increasing* returns to scale but when they become large, they exhibit *decreasing* returns to scale. This sharply contrasts with the previous finding of decreasing returns to scale among equity mutual funds. Further, I show that the nature of trading costs in the corporate bond market—in particular, a U-shaped relation between trade size and unit trading cost at the corporate bond level—is relevant for explaining hump-shaped returns to scale. Interpreting these empirical patterns is not straightforward, though. In a rational expectations framework, we expect a fund's net alpha always to be zero and hence, there should be no time-series relation between fund size and subsequent fund alpha. To help interpret the empirical findings, I propose a dynamic model in which investors learn about a fund's ability to manage its trading cost from the fund's past returns. The evolution of investors' beliefs provides a source of variation in fund size and further, in fund alpha over time.

In the second chapter of the dissertation, entitled "The Making of Hawks and

Doves" and co–authored with Ulrike Malmendier and Stefan Nagel, we argue that central bankers' personal inflation experiences significantly alter their inflation forecasts, votes, and speeches. First, we show that inflation experiences have a direct impact on Federal Open Market Committee members' inflation forecasts in their semi-annual Monetary Policy Reports to U.S. Congress. Second, members with higher inflation experiences are significantly more likely to cast a hawkish dissent. Over the FOMC's voting history since March 1951, an increase in a member's experience-based inflation forecast by one within-meeting standard deviation raises the probability of a hawkish dissent by about one third, and decreases the probability of a dovish dissent also by about one third. Third, higher inflation experiences also predict a significantly more hawkish tone in speeches. Finally, aggregating over all FOMC members present at a meeting, the average experience-based forecast helps predict the federal funds target rate, over and above conventional forward-looking Taylor rule components. Our findings indicate strong and long-lasting effects of personal inflation experiences even among monetary-policy experts, and point to the importance of FOMC members' selection.

In the third chapter of the dissertation, entitled "Electronic Trading in the U.S. Corporate Bond Market", I study how the addition of electronic trading affects the U.S corporate bond market overall, which primarily depends on conventional voice trading. Motivated by the empirical trade-off between price improvement and execution risk of electronic trading, I develop a model of strategic trading-platform selection by investors and show that i) the equilibrium market share of electronic trading decreases as the capacity of dealers' balance sheets becomes smaller, and further ii) the inclusion of electronic trading may not benefit every single dealer—those with small balance sheet capacity would prefer the market structure with voice trading only.

CHAPTER I

Returns to Scale among Corporate Bond Mutual Funds

1.1 Introduction

Understanding the nature of returns to scale is important among mutual funds. First, it helps resolve one of the most important puzzles in the mutual fund literature: why do equity fund investors chase fund returns even though they are not persistent (Chevalier and Ellison (1997)). It turns out that investors can chase fund returns rationally because past returns are informative signals about unobserved fund skill but, in the meantime, equity funds that have higher past returns (and are therefore growing) are expected to perform less impressively in the future due to decreasing returns to scale (Berk and Green (2004)). Moreover, understanding returns to scale helps us evaluate a fund's skill. Given that a fund's performance depends on both its skill and size, to evaluate the fund's skill, we need to first understand the relation between fund size and fund performance (Pastor, Stambaugh, and Taylor (2015)).

Given its importance, returns to scale for equity mutual funds has received much attention since Berk and Green (2004); a lot of theoretical work is now built on the premise of decreasing returns to scale and many empirical studies have attempted to verify whether such an assumption holds. More recently, empirical evidence suggests that equity mutual funds indeed face decreasing returns to scale at the fund level (see, e.g., Harvey and Liu (2017) and McLemore (2018)). For corporate bond mutual funds, however, we know surprisingly little about their fund-level returns to scale. This is particularly striking given that they are the second-largest investors in the underlying corporate bond market.[1] More importantly, our understanding of equity funds cannot be directly applied to corporate bond funds due to the sharp difference in market structure between underlying stocks and corporate bonds (i.e., exchanges vs. the over-the-counter (OTC) market). In this paper, I aim to fill this gap.

I start by documenting the empirical relation between fund size and subsequent fund performance among U.S. corporate bond mutual funds. One challenge when estimating fund-level returns to scale is the endogeneity of fund size; funds with better skills tend to be larger and deliver higher returns. The ordinary least squares (OLS) estimate of the size-performance relation thus tends to suffer from an omitted-variable bias as fund skill is not observable. Including fund fixed effects in the regression addresses this bias as long as the unobserved fund skill is time-invariant but, in the meantime, it introduces another bias—known as the finite-sample bias—due to the positive contemporaneous correlation between fund size and unexpected fund alpha (Stambaugh (1999)). To account for both the omitted-variable and the finite-sample biases, I apply a recursive demeaning procedure to an otherwise standard fund fixed effects model, following Pastor, Stambaugh, and Taylor (2015) and Zhu (2018).

Based on a sample of 707 such funds from 1991 to 2017, I find that for a given fund, the next-period fund's benchmark-adjusted return (alpha) first *increases* and then *decreases* in its size. Point estimates from a quadratic specification suggest a

[1] As of year-end 2017, U.S. corporate bond mutual funds collectively manage over \$2.1 trillion in assets (vs. \$7.4 trillion for U.S. domestic equity mutual funds) and hold 16% of corporate and foreign bonds (Investment Company Institute 2018).

turning point at fund's total net assets (TNA) of about \$619 million. Using a piecewise linear specification, I further find that when fund size is below that turning point, a one-standard-deviation (42%) increase in the fund's TNA *raises* the annualized gross alpha by 173 bps. Once fund size becomes larger than \$619 million, however, a one-standard-deviation (34%) increase in the fund's TNA *lowers* the annualized gross alpha by 107 bps. This hump-shaped pattern is robust to various specifications as well as to different measures of fund performance. Note that *hump-shaped* returns to scale among U.S. corporate bond funds contrasts with the *decreasing* returns to scale among U.S. equity funds documented in the literature.

Next, I examine whether the nature of trading cost in the underlying bond market is relevant for explaining such a hump-shaped relation among corporate bond funds. Often, decreasing returns to scale among equity funds is motivated by a *positive* relation between trade size and unit trading cost at the underlying stock level (e.g., Pastor, Stambaugh, and Taylor (2017)). But this positive relation no longer holds for corporate bonds, due to their different market structure compared to stocks. The key difference between the two market structures is that search friction matters a lot for corporate bond trading (OTC market) but not so much for equity trading (exchanges). As a consequence, the relation between trade size and unit transaction cost is *U-shaped* in the corporate bond market (see, e.g., Randall (2015) and Bessembinder et al. 2018).[2] Therefore, given that a *positive* relation between trade size and unit transaction cost at the stock level gives rise to *decreasing* returns to scale among equity funds, would such a *U-shaped* relation at the bond level imply *hump-shaped* returns to scale among corporate bond funds?

[2]Specifically, one friction that could lead to the *decreasing* part of such a U shape is the fixed component of search cost in the OTC market. Because it takes certain time and efforts for dealers to find counterparts to trade regardless of the trade size, a larger trade would thus lower the unit trading cost when trade size is relatively small. When trade size becomes sufficiently large, standard frictions for equity trading become relevant in the corporate bond market as well (e.g., information asymmetry and inventory costs) and therefore, we would then expect the *increasing* part of the U shape.

To tackle this question, I begin by examining whether such a U-shaped relation holds for each fund. So far, the evidence almost all focuses on the *cross-section* rather than on the *individual* bond investor. Presumably, larger bond investors (funds) tend to trade in larger sizes and, due to their larger bargaining power over the dealer, have lower unit transaction costs; this is dubbed the "clientele effect." Thus, a cross-sectional U-shaped relation may only reflect this clientele effect, but does not preclude the possibility that for each bond investor, unit transaction cost may still increase with trade size, as in the equity market. If so, such a *cross-sectional* U-shaped relation would no longer account for *within-fund* hump-shaped returns to scale. Using additional proprietary data which include bond investors' (anonymous) identities, I demonstrate that such a U-shaped relation remains *after* conditioning on investors. This evidence thus ensures that it is trade size itself—rather than the potential clientele effect for which trade size may serve as a proxy—that affects the unit transaction cost for any given investor in the underlying corporate bond market.

Further, I examine the fund-level returns to scale for U.S. Treasury bond funds (274 funds in the sample) because both Treasuries and corporate bonds have the same market structure, i.e., the OTC market. But search frictions for Treasuries trading tend to be much less severe than those for corporate bond trading. As a result, unit transaction cost for Treasuries, unlike that for corporate bonds, still increases in trade size (Fleming, Mizrach, and Nguyen (2017)). Therefore, if such a U-shaped relation is relevant, we should not expect hump-shaped returns to scale among U.S. Treasury bond funds. Indeed, I find no evidence of them.

At first glance, the finding of hump-shaped returns to scale among corporate bond funds may seem natural given the presence of a U-shaped relation between trade size and unit transaction cost at the underlying bond level. From the perspective of competitive fund investors with rational expectations (who know the objective probability they face in equilibrium), however, this finding—in particular the increasing part—is rather puzzling. After all, why, in the equilibrium, would they settle for a fund size that still rests in the region of increasing returns to scale when they could be (at least weakly) better off by simply putting more money into the fund? More broadly, as investors are competing with each other, we would expect that, for a given fund, its net alpha always to be zero and hence, no time-series relation between fund size and next-period fund alpha. How, then, should we interpret any empirical (within-fund) returns-to-scale relation with competitive fund investors who have rational expectations?

The key to interpret the empirical findings is perhaps to figure out what drives the variation in fund alpha in equilibrium over time. Here, I focus on the role of fund investors' behavior, and show that the evolution of investors' beliefs could be one driver. In particular, I develop a dynamic equilibrium model in which competitive fund investors are uncertain about the fund's ability to manage its trading cost and thus have to form beliefs about it from fund past returns. In the model, I distinguish between the *subjective* belief, perceived by fund investors in real time, and the *objective* belief, from the perspective of econometricians who analyze the data ex-post. Under the subjective belief, as we already know, the expected fund net alpha is always zero. But, under the objective belief, I show that the expected fund net alpha generally differs from zero and varies over time. This time-series variation comes from investors' beliefs updating over time. Once fund alpha is realized, investors update their beliefs accordingly. Given investors' updated beliefs, the fund size is then determined in the equilibrium, which, in turn, affects the next-period fund alpha.

In the model, the *subjective* returns to scale—the relation that matters for investors' decisions on capital allocation—is strictly negative when evaluating at the equilibrium fund size. This explains how the mutual fund market reaches its equilibrium: Competitive fund investors keep investing until the expected fund net alpha reaches zero under their subjective beliefs, given that the expected fund alpha decreases in fund size at that point. In the meantime, consistent with the empirical finding of hump-shaped returns to scale among corporate bond funds, the model also implies that such a hump-shaped relation holds under the *objective* belief. The intuition is: When investors' perceived fund's ability is sufficiently lower than the objective one, the equilibrium fund size becomes so small that, under the objective belief, it still remains in the increasing-returns-to-scale region, given a U-shaped unit cost function (this is not the case under the subjective belief, though).[3] Otherwise, the equilibrium fund size would be large enough that it always fell into the decreasing-returns-to-scale region under both the subjective and the objective beliefs.

Clearly, the model's implication for hump-shaped returns to scale depends on the assumption that investors' subjective beliefs generally differ from the objective one. Such a wedge can arise for at least two reasons: First, investors' learning may lead to the wedge.[4] As investors have to infer the fund's ability from a limited sample in *real time*, their inference may be different from that of econometricians, who analyze full-sample data *ex-post*. As an illustration, I characterize how such a wedge emerges in equilibrium, assuming investors are Bayesian learners. But note that the model's implication for hump-shaped returns to scale does not depend on the exact way investors learn. Investors can learn with the standard Bayes rule in a fully rational and frictionless world. Or, perhaps more realistically, investors may deviate from Bayesian updating due to certain frictions; for examples, they may learn from their personal experience (Malmendier and Nagel (2011)) or extrapolate from past returns (Barberis,

[3]To reflect the unique market structure of underlying corporate bonds as discussed above, I assume that the unit trading cost for a corporate bond fund is U-shaped in fund size (rather than a monotonically increasing one that is often assumed in the literature to capture the market structure for equity trading). More details can be found in section 1.4.1.

[4]Evidence on mutual fund investors learning includes Brown and Wu (2016), Choi, Kahraman, and Mukherjee (2016), and Franzoni and Schmalz (2017).

Greenwood, Jin, and Shleifer (2015)).[5] Second, investors' misperceptions of fund alpha could also contribute to the wedge because their inferences on the fund's ability depend on their estimates of fund alpha in each period. Indeed, empirically, I find that corporate bond fund flows respond not only to the fund's past alpha, which they should, but also to its past "benchmark-related" returns (its exposures to the benchmark $\times$ excess returns of the benchmark), which they shouldn't (because investors could replicate the "benchmark-related" returns themselves with a much lower cost by investing in the index funds). This evidence suggests that fund investors are confused about the fund alpha.[6]

To summarize, this paper has two main contributions. The first is to empirically document a novel (time-series) hump-shaped relation between fund size and subsequent fund alpha for U.S. corporate bond mutual funds. To the best of my knowledge, this is the first paper to provide such evidence for fixed-income mutual funds. Existing studies mostly focus on equity funds: Recent evidence suggests that U.S. equity mutual funds face decreasing returns to scale at the fund level (see, e.g., Harvey and Liu (2017), Pastor, Stambaugh, and Taylor (2017), McLemore (2018), and Zhu (2018)). Outside the realm of mutual funds, Dyck and Pomorski (2014) find increasing returns to scale at the fund level for private equity (PE) funds.[7]

Another contribution is to provide an equilibrium framework to help interpret the empirical findings. As discussed above, it is difficult to interpret any observed pattern

[5]The consequence of these frictions is that investors would place different weights on their priors relative to the Bayesian learning case. But the wedge between the subjective and the objective belief remains.

[6]Details on the behavior of corporate bond fund flows can be found in Appendix Section B.7. Similar evidence has also been documented with respect to equity fund flows (see Barber, Huang, and Odean (2016) and Berk and van Binsbergen (2016)).

[7]Note that the underlying mechanism that results in an increasing returns to scale among PE funds is different from the one that I emphasize in this paper. Dyck and Pomorski (2014) argue that as PE funds become larger, they rely more on direct investments rather than on more costly intermediaries (such as fund-of-funds). In addition, the authors do not address the equilibrium question of why investors would settle for a fund size that still rests in the increasing-returns-to-scale region.

of returns to scale within a standard rational expectations framework. This difficulty is not limited to corporate bond funds but applies to equity funds as well. Therefore, the framework here would also be useful in understanding the returns-to-scale relation for equity funds. In that regard, this paper relates to the broad and emerging literature on mutual funds in equilibrium. In a pioneering work in this strand of literature, Berk and Green (2004) demonstrate how a rational expectations framework can be applied to the mutual fund industry to resolve important puzzles in the literature. In this paper, I depart from Berk and Green (2004) in two important ways: First, I assume the unit cost function at the fund level is U-shaped in fund size and second, I distinguish between the subjective and the objective beliefs. This allows me to interpret otherwise puzzling hump-shaped returns to scale among corporate bond funds. More recently, Pastor, Stambaugh, and Taylor (2017) study the equilibrium relations among fund size, expense ratio, turnover, and portfolio liquidity among equity mutual funds, but their focus is not on returns to scale. As a complement to their paper, mine studies how hump-shaped returns to scale may emerge in the equilibrium.

The rest of the paper is organized as follows. Section 1.2 describes the data and the empirical measurement of fund performance. Section 1.3 examines the empirical relation between fund size and subsequent fund performance and investigates the role of trading cost in the underlying bond market. Section 1.4 provides a model with a focus on fund investors' behavior to interpret the empirical findings. Section 1.5 concludes. Appendix A provides proofs. Appendix B contains additional empirical results.

1.2 Data and Empirical Measurement

1.2.1 Sample Construction

Data come from Center for Research in Security Prices (CRSP) survivorship-bias-free mutual fund database. In this study, I focus on U.S. corporate bond mutual funds based on the objective codes provided by CRSP.[8] The sample period is from January 1991, when the CRSP database starts to report monthly data, to March 2017.

I first exclude funds with a history of less than three years as I require three years of returns to estimate fund alpha (as described in the next subsection).[9] Then I drop funds whose total net assets (TNA) never exceed $10 million. To address the potential incubation bias (Evans 2010), I drop fund-month observations whose date is prior to the reported starting date of the fund. I further eliminate observations before funds first reach the $5 million threshold for their TNA. For funds with multiple share classes, I aggregate all subclasses up to the fund level because they share the same underlying portfolio.[10] The resulting sample includes 707 unique funds.

1.2.2 Measurement of Fund Performance

As standard in the literature, I use fund alpha—the intercept from fund-by-fund 36-month rolling regressions of excess fund returns on excess returns of a benchmark portfolio—to measure the performance for each fund. Following the suggestion of Berk and van Binsbergen (2015), I construct the benchmark using a number of Vanguard

[8] I include funds with (i) CRSP objective codes starting with "IC" or (ii) Lipper objective codes in the set ("IID", "SID", "SII"). Then I exclude index funds, exchange traded funds (ETFs), and exchange traded notes (ETNs).

[9] For funds that switched fund category at some point in the history—e.g., from Fixed-income to Balanced or to Money Market—I require at least 3 years of history in the Fixed-income category to be included in the sample.

[10] I calculate fund-level TNA by summing up the TNAs of each subclasses. For the fund-level expense ratio, net return, turnover, I take the TNA-weighted average across share classes.

bond market index funds.[11] Unlike traditional non-tradable risk factors, returns on index funds reflect the real-time investment opportunities available to investors and incorporate actual transaction costs.[12] Thus, they are preferable to non-tradable ones when used as the benchmark.

In particular, I include the Vanguard Total Bond Market Index Fund (VBMFX) as the aggregate bond market factor. To further account for some mechanical yield curve strategies, I also include the Vanguard Short-Term Bond Index Fund (VBISX), Vanguard Intermediate-Term Bond Index Fund (VBIIX), and Vanguard Long-Term Bond Index Fund (VBLTX) as additional term structure factors. The choice of Vanguard index funds but not index funds from other asset management firms is somewhat arbitrary; the primary reason is that Vanguard index funds are often one of the largest index funds within their sectors and have a long history of fund performance.

1.2.3 Summary Statistics

Table 1.1 presents summary statistics. Panel A describes various fund characteristics. On average, a fund receives monthly inflow of 0.67%, reflecting the overall growth of the corporate bond fund industry during the sample period. A median fund holds \$157 million total net assets, but the dispersion in fund size is large: the interquartile range is \$455 million (= \$510 − \$55). The median fund age is about eight years and the median annual expense ratio is 72 bps.[13] Notably, the median turnover is 100% per year, which is similar to that of a typical equity fund. This high turnover ratio suggests that corporate bond funds do trade and trade frequently and therefore, transaction cost

[11] Recent studies on bond mutual funds have began to use bond index funds as the benchmark as well (e.g., Goldstein et al. 2017).

[12] Another advantage is that (net) returns on index funds also reflect a small yet non-zero cost for fund investors when investing passively (bond index funds typically charge about 10-20 bps fee per year).

[13] Note that the sample is tilted toward old funds since I require at least 3 years of history to estimate fund alpha.

should be a relevant concern for them.

Panel B illustrates fund performance. The average benchmark-adjusted return is about 104 bps per year before fees ("gross alpha") and 41 bps per year after fees ("net alpha"). As evidenced by a high adjusted-R^2 (78.5% on average), the four Vanguard bond market index funds that I include in the benchmark perform well in explaining the variation of excess fund returns.

1.3 Empirical Relation between Fund Size and Subsequent Fund Performance

1.3.1 Main Results

Estimating returns to scale at the fund level is not trivial because fund size is not assigned randomly; indeed, funds with better skills tend to manage larger amounts of capital and achieve better performance (e.g., Berk and Green (2004) and Pastor, Stambaugh, and Taylor (2015)). Thus, simply running a panel regression of a fund's benchmark-adjusted return (alpha) on its lagged size would result in an omitted-variable bias. One way to eliminate this bias is to introduce fund fixed effects (FE), which control for the unobserved fund skills as long as they are time-invariant. But including fund FE would introduce another bias—namely the finite-sample bias—because a positive shock in fund alpha tends to increase the contemporaneous fund size and therefore, the strict exogeneity condition is violated (see more discussion in Stambaugh (1999) and Pastor, Stambaugh, and Taylor (2015)). To account for both the omitted-variable and the finite-sample bias, I apply the recursive demeaning (RD) procedure to an otherwise standard fund FE model, following Pastor, Stambaugh, and Taylor (2015) and Zhu (2018).[14]

[14] For completeness, I describe the details of the recursive demeaning procedure in Appendix B.1.

Table 1.1: Summary Statistics

The sample includes 707 corporate bond mutual funds from January 1991 to March 2017 (excluding money market funds, index funds, ETFs, and ETNs). The unit of observation is the fund-month. Turnover is defined as $\frac{min\left(Buys, Sells\right)}{FundSize}$. Excess return equals fund returns minus three-month T-bill returns. *Net Alpha*, α_j, is estimated fund by fund from the following 36-month rolling regression:

$$R_{j,t} = \alpha_j + \sum_{i=1}^{4} \beta_{j,i} R_{i,t} + \epsilon_{j,t}$$

where $R_{j,t}$ is the excess fund returns (after fees) and $R_{i,t}$ is the excess new returns for each of the four benchmark Vanguard index funds; i.e., total bond market (VBMFX), short-term (VBISX), intermediate-term (VBIIX), and long-term (VBLTX). *Gross Alpha* equals *Net Alpha* plus the expense ratio.

	N	Mean	S.D.	p10	p25	p50	p75	p90
Panel A: Fund characteristics (fund-month obs.)								
Monthly fund flow (%)	92763	0.67	4.38	-2.91	-1.22	0.03	1.69	4.75
Fund size ($ millions)	92763	817	3157	22	55	157	510	1524
Fund age (years)	92763	9.7	7.3	2.0	4.2	8.1	13.6	19.2
(Annual) expense ratio (%)	92740	0.77	0.35	0.42	0.55	0.72	0.92	1.18
(Annual) turnover	92763	1.7	1.9	0.3	0.5	1.0	2.2	3.9
Panel B: Fund performance (annualized) (fund-month obs.)								
Excess return (%)	92763	2.34	14.29	-12.03	-4.11	2.46	9.37	17.14
Net alpha (%)	67296	0.41	2.17	-1.18	-0.56	0.09	1.04	2.53
Gross alpha (%)	67235	1.04	2.14	-0.47	0.07	0.66	1.63	3.20
Adjusted R^2	67296	0.785	0.231	0.432	0.682	0.882	0.955	0.979

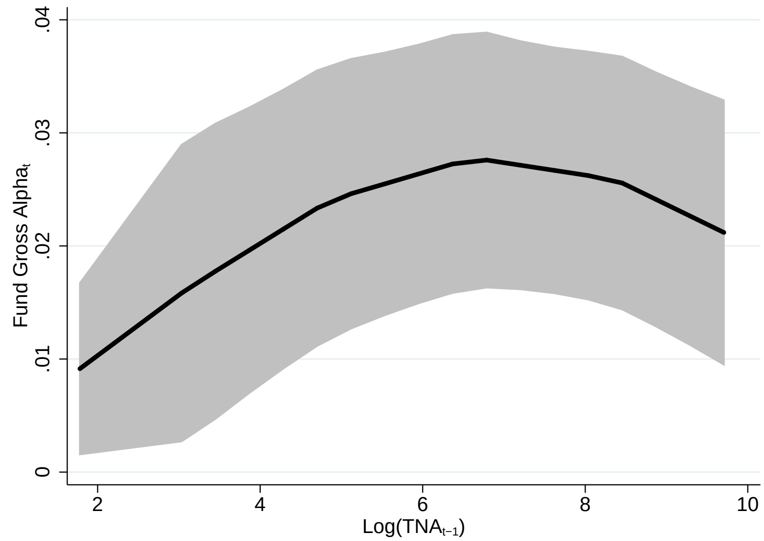

Figure 1.1: Spline Estimates of Fund-level Returns to Scale. This figure reports the estimated shape of fund-level returns-to-scale relation among U.S. corporate bond mutual funds using a linear spline regression with 4 knots equally placed at each quintile of $log(TNA_{t-1})$. The shaded area represents the 95% confidence intervals.

Figure 1.1 provides "semi-parametric" estimates on the fund-level returns-to-scale relation using a linear spline regression with four knots equally placed at each quintile of $FundSize$—the log of the fund's total net assets (TNA).[15] Clearly, the relation between fund size and subsequent fund gross alpha is hump-shaped. Moreover, the figure suggests a turning point when the fund's TNA is about $665 million ($e^{6.5} \approx 665$).

Next, I turn to parametric regressions to check whether such a hump-shaped pattern is significant and robust. Table 1.2 reports the results. I start with a quadratic specification. As shown in Column (i) of Panel A, I do find a statistically significantly hump-shaped relation between fund size and subsequent fund gross alpha. Interestingly, the turning point implied by the point estimates on $FundSize$ and $FundSize^2$ is when the fund's TNA is about $619 million ($e^{\frac{0.00905}{2\times 0.00069}} \approx 619$), which is very close to the one implied by the linear spline regression above.

[15]Note that the level of y-axis is not informative because the fund fixed effects are included in the estimation.

To further understand the economic magnitude with respect to both the increasing and decreasing parts of the "hump-shape," I estimate a piecewise linear regression with a turning point at \$619 million for a fund's TNA, implied by the quadratic specification. The result, shown in Column (i) of Panel B, again confirms a hump-shaped pattern of returns to scale. Moreover, such a hump-shaped pattern is economically meaningful: When the fund size is below the turning point, a one-standard-deviation increase in $FundSize$ (a 42% increase in the fund's TNA) *raises* the annualized gross alpha by 173 bps ($= 0.00344 \times 0.42 \times 12$). Once the fund size becomes larger than \$619 million, a one-standard-deviation increase in $FundSize$ (a 34% increase in the fund's TNA) *lowers* the annualized gross alpha by 107 bps ($= 0.00262 \times 0.34 \times 12$).

Note that an almost identical hump-shaped pattern emerges when fund performance is instead evaluated by net alpha (rather than gross alpha), as shown in Column (ii) of Panels A and B. I also find, as shown in Appendix B.2, a hump-shaped pattern of fund-level returns to scale among U.S. high-yield bond funds, where a U-shaped relation between trade size and unit transaction cost also holds at the underlying high-yield bond level.

1.3.2 Robustness: Alternative Benchmarks

In the above main analysis, I include four Vanguard bond market index funds in the benchmark to measure fund performance both before and after fees. As robustness checks, I consider several alternative compositions of the benchmark. First, I only include the aggregate bond market index fund, VBMFX, in the benchmark. Second, I add the aggregate stock market factor, using the returns on Vanguard 500 Index Fund (VFINX) as a proxy. Third, I include all five index funds: four bond market index funds (as used in the main analysis) plus one stock market index fund. Lastly, instead of using index funds, I adopt conventional risk factors as the benchmark; namely, short-

Table 1.2: Returns to Scale

This table shows returns-to-scale relations for U.S. corporate bond mutual funds from January 1991 to March 2017 (excluding index funds, ETFs, and ETNs). The dependent variable is either $GrossAlpha$ or $NetAlpha$, both of which are estimated over the previous 36 months. $FundSize$ is the log of the fund's total net assets (TNA) while $IndustrySize$ is the sum of the TNA for all corporate bond funds scaled by the amounts outstanding of all corporate bonds. $FundSize\_Small_{t-1}$ and $FundSize\_Large_{t-1}$ are linear splines constructed from $FundSize_{t-1}$ with a knot implied by the corresponding quadratic specification. In particular,

$$FundSize\_Small_{t-1} = \min(FundSize_{t-1}, k)$$
$$FundSize\_Large_{t-1} = \max(FundSize_{t-1}, k) - k$$

where k is the corresponding knot. Time-varying fund-level controls include: lagged log of fund age (in years), lagged fund expense ratio, lagged turnover ratio, and lagged realized return volatility in the past 12 months. Following Pastor, Stambaugh, and Taylor (2015), I first forward-demean the variables on both the left- and right-hand sides and then instrument any variables that involve $FundSize$ by their backward-demeaned counterparts. In the parentheses, I report the standard errors clustered by both month and fund.

	$GrossAlpha_t$ (i)	$NetAlpha_t$ (ii)
Panel A: Quadratic specification		
$FundSize_{t-1}$	0.00905	0.00910
	(0.00209)	(0.00213)
$(FundSize_{t-1})^2$	-0.00069	-0.00069
	(0.00020)	(0.00021)
$IndustrySize_{t-1}$	0.00043	0.00042
	(0.00009)	(0.00009)
Observations	67209	67246
Panel B: Piecewise linear specification		
$FundSize\_Small_{t-1}$	0.00344	0.00345
	(0.00064)	(0.00065)
$FundSize\_Large_{t-1}$	-0.00262	-0.00261
	(0.00129)	(0.00132)
$IndustrySize_{t-1}$	0.00039	0.00038
	(0.00009)	(0.00009)
Observations	67209	67246

rate factor (three-month Treasury-bill rate), slope factor (ten-year Treasury rate − one-year Treasury rate), curvature factor (two-year Treasury rate + ten-year Treasury rate − 2×five-year Treasury rate), and default risk factor (BAA-rated corporate bond yield − AAA-rated corporate bond yield).[16]

I reestimate the relation between fund size and subsequent fund alpha among corporate bond funds using a quadratic specification with the above four alternative benchmarks. The results are shown in Appendix Table B.2. As can be seen, the estimated patterns of fund-level returns to scale under these alternative benchmarks are all hump-shaped, as in the baseline case (where the benchmark includes four Vanguard bond market index funds). I therefore conclude that the finding of a hump-shaped pattern between fund size and subsequent fund alpha is robust to different choices of the benchmark.

1.3.3 Transaction Costs for Corporate Bond Funds

Transaction cost is often considered as the key driver in shaping the pattern of fund-level returns to scale (see, e.g., Pastor, Stambaugh, and Taylor (2017)). For transaction cost to empirically matter, its magnitude needs to be non trivial. However, empirical estimates of transaction cost for corporate bond mutual funds are scarce. Thus, before proceeding to investigate the relevance of transaction cost in explaining the hump-shaped returns to scale documented above, I first estimate the amount of transaction cost that a corporate bond fund would incur when adjusting its portfolio.[17]

[16]Note that the explanatory power of these four conventional risk factors on fund excess returns is much smaller than that of index funds, evaluated by the adjusted R^2 from fund-by-fund regressions. With above conventional risk factors, the average adjusted R^2 is merely 31.5% with an interquartile range from 20.3% to 42.9%. As a comparison, using four index funds as the benchmark (the baseline specification used in the main analysis), the average adjusted R^2 is 78.6% with an interquartile range from 68.4% to 95.5%.

[17]Note that the transaction cost I consider here only captures the implicit cost, i.e., the effective one-way trade execution cost. What I do not capture is the explicit cost, including commissions, taxes, and fees. According to Busse, Chordia, Jiang, and Tang (2018), for equity mutual funds, the explicit cost accounts for about 40% of the total transaction cost.

Ideally, one may first estimate the transaction costs trade by trade and then add them up for each fund by looking at the fund's transaction history. Unfortunately, such granular data is not publicly available.[18] As an imperfect but reasonable alternative, I estimate the fund total transaction cost as the product of its turnover and average unit transaction cost. While the fund's turnover is directly observable, I have to infer its average unit transaction cost from its average trade size, which I then estimate from the monthly mutual fund holding data from CRSP. In particular, the average trade size is estimated under two assumptions: (i) no trade occurred if a fund's position in a specific bond remains the same in two consecutive months, and (ii) exactly one trade occurred if a fund's position in a specific bond changes between two consecutive months—in other words, funds do not split one trade into multiple smaller pieces.[19] It is worth noting that with these two assumptions, I may underestimate the number of trades and thus overestimate the average trade size. Therefore, my estimates on the total transaction cost tend to be conservative, given that most trade sizes belong to the decreasing part of a U-shaped relation between trade size and unit transaction cost (Randall (2015)).

It turns out that for funds whose total net assets (TNA) are below 150 million, the median of the average trade size is about \$90K. According to the estimates in Bessembinder et al. (2018), unit transaction cost for customer trades whose size is less than \$100K is about 0.70%. Along with a median of reported turnover at 0.9 per year for those funds, the estimated transaction cost is thus 63 bps ($= 0.70\% \times 0.9$) per year.

[18] Recently, some studies have begun to use propitiatory data on actual transactions of institutional investors to investigate trading cost. For example, Busse, Chordia, Jiang, and Tang (2018) study equity mutual funds and Frazzini, Israel, and Moskowitz (2018) focus on one large institution. No such study, however, examines trading costs for fixed-income mutual funds.

[19] Although not ideal, these two assumptions are reasonable. Regarding the first assumption, given the nature of mutual funds and the illiquidity of the corporate bond market, it seems less likely that a fund manager would trade multiple times per month on any specific bond. Regarding the second assumption, corporate bond funds do have the incentive to trade in a larger size due to the U-shaped relation between trade size and unit transaction cost and hence, are less likely to split their trades into smaller pieces. Details on constructing the average trade size can be found in Appendix B.4.

A 63 bps annualized transaction cost is substantial, given that the average annualized gross alpha for those funds is only 86 bps. Interestingly, for larger funds whose TNA are above \$150 million, the median of the average trade size is about $\$210K$. Again, based on the estimates in Bessembinder et al. (2018), the corresponding unit transaction cost is about 0.30%. Thus, the estimated transaction cost for these larger funds is 28.5 bps per year, given a median of reported turnover at 0.95 per year. Based on the above back-of-the-envelope calculation, it appears that savings on transaction cost could be meaningful if a fund can increase its trade size (by managing more TNA).

1.3.4 Within-fund Relation between Trade Size and Unit Transaction Cost

Having established the importance of transaction cost for corporate bond funds, I now examine whether a U-shaped relation between trade size and unit transaction cost holds for each bond investor (fund). While such a U-shaped pattern is widely documented in the literature, most studies focus on the cross-section rather than on the individual bond investor (see, e.g., Randall (2015) and Bessembinder et al. 2018). The concern is that large bond investors (funds) tend to trade in large sizes and, due to their large bargaining power over the dealer, have lower unit transaction costs ("clientele effect"). Thus, a cross-sectional U-shaped relation may pick up this fixed clientele effect but still does not preclude the possibility that for each bond investor, the relation between trade size and unit transaction cost may still be positive, as in the equity market. If so, such a U-shaped relation would no longer account for within-fund hump-shaped returns to scale.

Ideally, one may address this concern by including investor fixed effects when estimating the relation between trade size and unit transaction cost. However, this approach is generally not feasible because even the most comprehensive TRACE dataset do not include investors' identifiers.[20] One imperfect—but at least feasible—way is to focus on a subsample of trades that were executed via electronic platforms on which investors' (anonymous) identifiers are available. Following this approach, I examine the relation between trade size and unit transaction cost using the transaction-level data provided by MarketAxess—a leading electronic platform provider in the U.S. corporate bond market.[21]

The results are reported in Appendix Table B.4. First, consistent with the literature, I document a U-shaped relation between trade size and unit transaction cost using a quadratic specification when only time fixed effects (FE) are included. The point estimates on the linear and quadratic terms suggest a turning point at a trade size of \$2.45 million, which is fairly large given that the average trade size is about \$600K. To account for the potential clientele effect, as discussed above, I first add investor FE. It turns out that, with both time and investor FE, a U-shaped relation between trade size and unit transaction cost remains. Further, by replacing "week + investor" FE with "week $\times$ investor" FE, I find that, for a given investor in a given week, the unit transaction cost again first decreases and then increases in trade size. Taken together, this evidence suggest that it is trade size itself, rather than the potential clientele effect for which trade size may serve as a proxy, that matters for the unit transaction cost for any given investor in the corporate bond market.

[20]TRACE stands for Trade Reporting and Compliance Engine, which, according to a Financial Industry Regulatory Authority (FINRA) mandate, records all secondary-market transactions for U.S. corporate bonds. The most comprehensive version of TRACE, made available by FINRA, includes identifies for registered dealers but not for customers (investors).

[21]Over the sample period, January 2014 to December 2015, MarketAxess accounts for about 30% of total trading volume (including both voice and electronic trades recored in TRACE). More details on MarketAxess data and empirical specifications can be found in Appendix B.5.

1.3.5 Returns to Scale among U.S. Treasury Bond Funds

In this subsection, I move on to examine whether such a U-shaped relation in trading cost would be relevant for hump-shaped returns to scale. One way to tackle this question is to look at the returns-to-scale relation for U.S. Treasury bond funds because both Treasuries and corporate bonds are traded in the OTC market. But search friction for Treasury trading is much less severe than that for corporate bond trading. As a result, unit trading cost for Treasuries, unlike that for corporate bonds, still increases in trade size (Fleming, Mizrach, and Nguyen (2017)). Therefore, if such a U-shaped relation is indeed relevant, we should not expect hump-shaped returns to scale among U.S. Treasury bond funds.

Table 1.3 presents the results. As shown in Column (i) of Panel A, estimates on $FundSize$ and $FundSize^2$ are statistically insignificant when four Vanguard bond market index funds are included in the benchmark (the same as the baseline case in the analysis of corporate bond funds), offering no evidence of a hump-shaped relation between fund size and subsequent fund alpha. Indeed, we cannot reject the null hypothesis of a constant returns-to-scale relation. This result is robust to various choices of the benchmark composition, as Columns (ii) to (v) show. The same conclusion holds when fund net alpha is the dependent variable, as Panel B shows.

To summarize, I find no evidence of a hump-shaped relation between fund size and subsequent fund alpha among U.S. Treasury bond mutual funds, which suggests that a U-shaped relation in trading cost is relevant for hump-shaped returns to scale.

1.4 A Model with Subjective Beliefs

The previous section documents a hump-shaped pattern of fund-level returns to scale among U.S. corporate bond mutual funds and further indicates that a U-shaped

Table 1.3: Returns to Scale among Treasury Bond Funds

This table replicates fund-level returns-to-scale relations as shown in Table 1.2, but with U.S. Treasury bond mutual funds (excluding money market funds, ultra-short duration funds, index funds, ETFs, and ETNs). The sample includes 273 such funds, from January 1991 to March 2017. The performance of Treasury bond funds is measured exactly the same way as that of corporate bond funds. "Four-factor" is the baseline case, referring to the benchmark with four bond market index funds (VBMFX, VBISX, VBIIX, and VBLTX); "one-factor" is the benchmark with only VBMFX; "two-factor" is the benchmark with both VBMFX and VFINX; "five-factor" is the benchmark with five index funds (VBMFX, VFINX, VBISX, VBIIX, and VBLTX); and "risk-factor" is short-rate factor (three-month Treasury bill rate), slope factor (ten-year Treasury rate − one-year Treasury rate), curvature factor (two-year Treasury rate + ten-year Treasury rate − 2×five-year Treasury rate), and default risk factor (BAA-rated corporate bond yield − AAA-rated corporate bond yield). $FundSize$ is the log of the fund's total net assets (TNA) while $IndustrySize$ is the sum of TNA of all Treasury bond funds scaled by the amounts outstanding of U.S. Treasury bonds. In Panel A, the dependent variable is $GrossAlpha$, in Panel B, it is $NetAlpha$. In the parentheses, I report the standard errors clustered by both month and fund.

	Four-factor (i)	One-factor (ii)	Two-factor (iii)	Five-factor (iv)	Risk-factor (v)
Panel A: $GrossAlpha_t$ **as dependent variable**					
$FundSize_{t-1}$	0.00066	0.00173	0.00125	0.00054	0.00118
	(0.00079)	(0.00149)	(0.00115)	(0.00072)	(0.00146)
$\left(FundSize_{t-1}\right)^2$	-0.00004	-0.00013	-0.00010	-0.00005	-0.00014
	(0.00010)	(0.00018)	(0.00014)	(0.00009)	(0.00019)
$IndustrySize_{t-1}$	0.00011	0.00008	0.00005	0.00007	0.00018
	(0.00005)	(0.00006)	(0.00005)	(0.00004)	(0.00008)
Observations	26498	29287	29281	26426	29287
Panel B: $NetAlpha_t$ **as dependent variable**					
$FundSize_{t-1}$	0.00063	0.00161	0.00113	0.00050	0.00105
	(0.00075)	(0.00139)	(0.00105)	(0.00067)	(0.00138)
$\left(FundSize_{t-1}\right)^2$	-0.00004	-0.00012	-0.00009	-0.00005	-0.00013
	(0.00010)	(0.00017)	(0.00013)	(0.00009)	(0.00018)
$IndustrySize_{t-1}$	0.00011	0.00008	0.00005	0.00007	0.00018
	(0.00005)	(0.00006)	(0.00005)	(0.00004)	(0.00008)
Observations	26498	29287	29281	26426	29287

relation between trade size and unit transaction cost—a unique feature at the underlying bond level—is relevant for explaining such a hump-shaped relation. But what remains puzzling is how to interpret these empirical findings; in a standard rational expectations framework with competitive fund investors, we should not expect any time-series relation between the fund size and next-period fund alpha.

In this section, I propose an equilibrium framework to help interpret the empirical findings. I start by developing a dynamic model in which competitive fund investors are uncertain about the fund's ability to manage its trading cost and thus have to form beliefs about it from fund past returns over time. I show that the evolution of investors' beliefs could drive the variation of fund alpha over time. I then examine model implications for fund-level returns to scale under both the subjective and the objective beliefs. Assuming that investors are Bayesian learners, I further illustrate how the wedge between the subjective and the objective beliefs emerges by characterizing investors' belief updating process and the flow-performance sensitivity accordingly. Finally, I discuss the dynamics between fund size and subsequent fund alpha, using a numerical example.

1.4.1 Setting

The model is dynamic, with time discrete and infinite. There are two types of agent in the model: one risk-neutral fund manager and (many) competitive fund investors. In the model, I do not distinguish between fund and manager, abstracting away any potential agency friction. As in Berk and Green (2004), the model is partial equilibrium; the fund manager's trading activities do not affect prices of the underlying securities in her portfolio.

At each period t, the fund's benchmark-adjusted return before fees ("gross alpha") is given by

$$\alpha_t^{Gross} = \mu - c(A_{t-1}) + e_t \tag{1.1}$$

where μ captures the manager's ability to outperform certain passive benchmark in the absence of trading cost, $c(A_{t-1})$ stands for unit trading cost, which is a function of fund size A_{t-1}, and e_t represents idiosyncratic shock with mean zero.

I specify the unit trading cost function as $c(A_{t-1}) = c_0 + c_1 A_{t-1} - \theta A_{t-1}^{\beta}$, where $\beta \in (0, 1)$ and θ, c_0, and c_1 are positive constants.[22] These parameters can be interpreted as the fund's ability to manage its trading cost. Clearly, a fund's unit trading cost here first decreases and then increases in its size. The choice of this unit cost function is motivated by the empirical U-shaped relation between trade size and unit transaction cost at the underlying corporate bond level, as discussed previously. Presumably, corporate bond funds tend to trade in larger sizes when they become bigger.[23] Therefore, a U-shaped relation at the bond level naturally translates into a U-shaped unit cost function at the fund level. Note that this specification is different from a monotonically increasing function that is often assumed in the mutual fund literature (e.g., Berk and Green (2004) and Pastor, Stambaugh, and Taylor (2017)). I further assume $\mu \geq c_0$; that is, when fund size is zero, the expected gross alpha is at least zero.

As this paper focuses on how trading cost affects the fund-level returns to scale, for tractability, I consider the case where investors know the fund manager's ability to beat the benchmark, i.e., μ. Due to the complexity of trading in the OTC market,

[22] I assume that c_0 is sufficiently large that the unit trading cost $c_0 + c_1 A_t - \theta A_t^{\beta}$ is always non-negative. Note that the minimum of $c_1 A_t - \theta A_t^{\beta}$ is $(\frac{1}{c_1})^{\frac{\beta}{1-\beta}} (\theta\beta)^{\frac{1}{1-\beta}} (1 - \frac{1}{\beta})$.

[23] Corporate bond funds, unlike equity mutual funds, do have the incentive to trade in large sizes if they can, rather than splitting one trade into smaller pieces, given that most of the trades fell into the decreasing part of the U-shaped relation in trading cost. In the meantime, the ability that a corporate bond fund lump smaller trades into a larger one is also limited because otherwise the number of holdings in the fund's portfolio would be too small relative to the benchmark and thus render the fund's tracking error becoming too large.

however, I assume that investors are only aware of the shape of the unit cost function but do not know the value of all relevant parameters and thus have to learn about them from fund's past returns. Among the set of cost parameters $\{c_0,\ c_1,\ \beta,\ \theta\}$, for tractability, I further assume that investors only have to learn about θ.[24] This learning process, as will be discussed in detail below, drives the evolution of equilibrium fund size and the expected fund alpha.

The timing of the model is as follows: at the end of each period t, fund gross alpha is realized and investors update their beliefs about θ accordingly. Given their updated beliefs, fund size is then determined in the equilibrium, which in turn affects the next-period fund alpha (as in Equation 1.1), and so on.

1.4.2 Fund size and Expense Ratio in Equilibrium

The fund manager maximizes total fee revenue at each period, $f_t A_t$, by choosing the expense ratio (or fee) f_t. By doing so, the manager is fully aware that her choice of fee would affect the amount of capital, A_t, that investors are willing to supply. Because investors are competitive, in the equilibrium, the benchmark-adjusted return after fees ("net alpha") is zero in expectation. Further, since (i) only the subjective belief matters for investors' capital allocation decisions and (ii) the subjective belief may not always coincide with the objective one (due to investors' learning and/or misperception of fund alpha), the assumption of competitive investors implies a zero expected net alpha under the subjective belief but *not necessarily* under the objective one.

[24]In Appendix A.2, I show that the model's implication for fund-level returns to scale under both the subjective and the objective beliefs remain the same if investors instead only learn about c_1 (rather than θ).

Having this in mind, the fund manager's problem can be formulated as follows:

$$\begin{aligned} \max_{f_t} \;& f_t A_t \\ \text{s.t. } & \hat{\mathbb{E}}_t\big(\alpha_{t+1}^{Gross}\big) - f_t = 0 \\ & f_t \geq 0 \end{aligned} \tag{1.2}$$

where $\hat{\mathbb{E}}_t(\cdot)$ denotes the expectation under investors' subjective beliefs at time t. To conserve space, in the text below, I denote $\phi_t \equiv \hat{\mathbb{E}}_t(\theta)$. The following proposition characterizes the equilibrium fund size, A_t^*, and expense ratio, f_t^*:

Proposition I.1. *Assuming $\mu \geq c_0$, at each period t, the equilibrium fund size that satisfies the manager's problem (1.2), A_t^*, is given by the unique positive solution to the equation:*

$$\phi_t(\beta + 1)(A_t)^\beta - 2c_1 A_t + (\mu - c_0) = 0 \tag{1.3}$$

The equilibrium expense ratio, f_t^*, is given by $\frac{1-\beta}{1+\beta}c_1 A_t^* + \frac{\beta}{1+\beta}(\mu - c_0)$. Moreover, the equilibrium fund size, A_t^*, is monotonically increasing in ϕ_t.

Note that the equilibrium fund size A_t^* increases in investors' perception of θ; that is, ϕ_t. The intuition is that as ϕ_t becomes larger, the perceived fund's ability to manage its trading cost increases and thus, investors are willing to put more money in the fund in the equilibrium. It is also worth noting that it is the investors' perception of θ, rather than θ itself, that determines the equilibrium fund size and expense ratio at each point of time.

1.4.3 Time-varying Expected Fund Alpha

Before turning to returns to scale, it is useful to first examine what drives the variation of expected fund alpha under the objective belief. As shown in Appendix A, the objective expected net alpha can be written as follows:

$$\mathbb{E}_t(\alpha_{t+1}^{Net}) = \big(\theta - \phi_t\big)(A_t^*)^\beta \tag{1.4}$$

where $\mathbb{E}_t(\cdot)$ denotes the expectation under the objective belief at time t. Given that econometricians are able to estimate the fund's unit cost function with high precision from large samples ex-post, I assume (for simplicity), under the objective belief, that econometricians know the value of all parameters, including θ.

As we can see, as long as ϕ_t differs from θ (which is generally the case), the expected fund net alpha is no longer *zero* under the objective belief, as it is under the subjective belief. When investors are optimistic about the fund's ability—that is, when ϕ_t is larger than θ—the objective expected net alpha is negative; otherwise, it is positive.

Moreover, as shown in Appendix A, the expected fund gross alpha under the objective belief also varies over time as investors' beliefs changes. Thus, I conclude that the evolution of investors' beliefs drives the variation of expected fund alpha over time.

1.4.4 Returns to Scale

In this subsection, I examine model-implied returns to scale at the fund level. I start with the relation under the *subjective* belief—the one that matters for investors' capital allocation decisions. As shown in Appendix A, when evaluating at the equilibrium fund size, the *subjective* returns to scale, $\frac{\partial \hat{\mathbb{E}}(\alpha_{t+1}^{Net})}{\partial A_t}|_{A_t=A_t^*}$, is strictly negative. This explains how the mutual fund market equilibrates: Competitive fund investors keep providing capital to the fund until the expected net alpha becomes zero under their subjective beliefs, given that the expected fund alpha decreases in fund size at that point.

While examining the subjective returns to scale helps us understand how the equilibrium is reached, we, as econometricians, cannot observe investors' subjective beliefs in the data. What we can observe instead is the returns-to-scale relation under the *objective* belief. For a given fund, such a relation is constant over time, since the underlying parameters are not time-varying (see Equation 1.1). To trace out such a time-invariant relation, we rely on the time-series variation in the equilibrium fund size driven by the fluctuation of investors' beliefs: As investors' perception of θ evolves over time (perhaps due to learning), the equilibrium fund size adjusts accordingly, which in turn affects the next-period fund alpha. This dynamic effectively acts as a "natural experiment" by perturbing the equilibrium fund size, and thus traces out the relation between fund size and subsequent fund alpha under the objective belief.

The following proposition characterizes how the expected fund alpha varies as investors' perception of θ—that is, ϕ_t—evolves:

Proposition I.2. *For a given fund, both the expected gross and net alpha under the objective belief first* ***increase*** *and then* ***decrease*** *in ϕ_t. The turning point, $\tilde{\phi}$, is smaller than $\beta\theta$.*

Given a positive one-to-one mapping between ϕ_t and the equilibrium fund size (see Proposition I.1), it immediately follows:

Corollary I.3. *For a given fund, both the expected gross and net alpha under the objective belief first* ***increase*** *and then* ***decrease*** *in the equilibrium fund size A_t^*.*

Both a U-shaped unit cost function and a wedge between ϕ_t and θ are responsible for the emergence of hump-shaped returns to scale at the fund level: When ϕ_t is sufficiently smaller than θ, investors significantly underestimate the fund's ability to manage its trading cost and thus, the equilibrium fund size becomes too small. As a consequence, under the objective belief, the fund would still be able to benefit from the reduction in unit trading cost by further expanding its size, given a U-shaped unit cost function. When ϕ_t is sufficiently large, the equilibrium fund size would be big enough that it always fell into the decreasing-returns-to-scale region under the objective belief.

At this point, it is natural to draw the link between the model's implication and the empirical evidence documented in Section 1.3. Given that the model focuses on a single fund, to map it into the panel data, I assume that different funds have the same ability to manage their trading cost but I allow their ability to outperform the benchmark *before trading cost* to be different (i.e., in the empirical specification, this fund-specific ability is captured by fund fixed effects).[25] Further, I approximate the unit cost function specified in the model with either a quadratic or a piecewise linear specification when implemented empirically. With these identification assumptions, the model can be exactly mapped into the empirical specifications, as in Section 1.3. Therefore, the model's implication for a hump-shaped pattern of fund-level returns to scale under the objective belief is consistent with the empirical finding.

1.4.5 Beliefs Updating and Flow-performance Sensitivity

In this subsection, I complete the model by studying how investors update their beliefs over time, assuming they are Bayesian learners, and then examining the flow-performance sensitivity accordingly. Note that the model's main implication for hump-shaped returns to scale does not depend on exactly how investors update their subjective beliefs. But assuming Bayesian investors allows me to provide a concrete example on how the wedge between the subjective and the objective beliefs may emerge with a tractable beliefs-updating process and flow-performance sensitivity.

[25] In particular, I assume the cost parameters c_0, c_1, θ, and β are the same across funds but μ are different across funds. Note that it is mathematically equivalent to assume that c_1, θ, and β are the same across funds, but $\mu - c_0$ vary across funds.

As Bayesian learners, investors update their beliefs on the true parameter θ by observing the most recent fund net alpha along with past fund size and expense ratio. In particular, suppose investors receive a noisy signal x_t of θ at each time t: $x_t \equiv \theta + \epsilon_t$, where ϵ_t stands for the noise term. Rearranging the definition of fund net alpha, we can explicitly write out the expression of the signal: $x_t = (A^*_{t-1})^{-\beta}\big(\alpha^{Net}_t - (\mu - c_0) + c_1 A^*_{t-1} + f^*_{t-1}\big)$.[26] Applying the Bayes rule to x_t leads to the following proposition, which characterizes the investors' belief updating process:

Proposition I.4. *Suppose investors' prior on θ follows $N(\theta_0, 1/\gamma)$ and the noise term, ϵ_t, is independently distributed through time and follows $N(0, 1/\omega)$, where γ is the precision of the prior and ω is the precision of the signal. The evolution of investors' perception on θ—that is, ϕ_t—is thus given by* [27]

$$\phi_t = \phi_{t-1} + \frac{\omega}{\gamma + t\omega}\Big(\frac{\alpha^{Net}_t}{(A^*_{t-1})^{\beta}}\Big) \tag{1.5}$$

Because investors always think the expected net alpha is zero under their subjective beliefs, when observing a *positive* (realized) net alpha, they would thus realize that they had *underestimated* θ, and would update their beliefs *upwards*. By the same token, investors would update their beliefs *downwards* when seeing a *negative* net alpha. As ϕ_t fluctuates over time, the wedge between ϕ_t and θ tends to be different from zero. Therefore, investors' real-time learning provides one plausible channel which could lead

[26] Recall that $\alpha^{Net}_t = \mu - c_0 + \theta(A^*_{t-1})^{\beta} - c_1 A^*_{t-1} - f^*_{t-1} + e_t$, which implies $(A^*_{t-1})^{-\beta}\big(\alpha^{Net}_t - (\mu - c_0) + c_1 A^*_{t-1} + f^*_{t-1}\big) = \theta + \epsilon_t$, where $\epsilon_t = (A^*_{t-1})^{-\beta} e_t$. Thus, one can interpret $(A^*_{t-1})^{-\beta}\big(\alpha^{Net}_t - (\mu - c_0) + c_1 A^*_{t-1} + f^*_{t-1}\big)$ as the signal x_t.

[27] The way investors update their beliefs, along with the flow-performance sensitivity, in my model are similar to those in Berk and Green (2004). This similarity is fully expected because, as in Berk and Green (2004), I assume that (i) fund investors are competitive and (ii) they are Bayesian learners. But I depart from Berk and Green (2004) in two important respects. First, they assume the unit cost that a fund faces is linearly increasing in its size, but I assume the shaped of unit cost function is instead U-shaped, motivated by the empirical evidence in the corporate bond market. Second, while Berk and Green (2004) do not distinguish between the subjective and the objective beliefs, I assume investors' subjective beliefs are in general different from econometricians' objective ones.

to the wedge between the subjective and objective beliefs.[28]

Now, let's turn to the flow-performance sensitivity. Given that the equilibrium fund size is uniquely determined by investors' perception of θ (see Proposition 1.2), as investors update their beliefs when observing the most recent fund net alpha, the equilibrium fund size adjusts accordingly. The following corollary summarizes how the equilibrium fund size reacts to the most recent fund net alpha:

Corollary I.5. *For a given fund, the percentage change in the equilibrium fund size is given by:*

$$\frac{A_t^* - A_{t-1}^*}{A_{t-1}^*} \approx \Phi_{1,t}\alpha_t^{Net} + \Phi_{2,t}\Big(\alpha_t^{Net}\Big)^2 \tag{1.6}$$

where

$$\Phi_{1,t} = \frac{(\beta+1)}{(1-\beta)2c_1A_{t-1}^* + \beta(\mu - c_0)}\Big(\frac{\omega}{\gamma + t\omega}\Big) > 0$$
$$\Phi_{2,t} = \frac{\beta(\beta+1)^2\Big((1-\beta)2c_1A_{t-1}^* + (\mu - c_0)(1+\beta)\Big)}{\Big((1-\beta)2c_1A_{t-1}^* + \beta(\mu - c_0)\Big)^3}\Big(\frac{\omega}{\gamma + t\omega}\Big)^2 > 0$$

Clearly, in the model, the percentage change in the equilibrium fund size is an increasing and convex function of fund's most recent net alpha.[29]

[28] One potential concern with Bayesian updating is that nothing can be learned in the limit and thus, ϕ_t converges to θ as t goes to ∞. To address this concern, one may instead adopt a perpetual learning rule (e.g., constant-gain learning) in which ϕ_t never converges to θ.

[29] In the Appendix, I further show that, in my framework, as long as fund investors update their beliefs on θ upwards (downwards) when seeing a positive (negative) net alpha (which does not necessarily have to be in line with Bayesian updating), an increasing and convex flow-performance sensitivity still holds.

Empirical flow-performance sensitivity Now, I examine whether this is indeed the case in the data. In particular, I estimate the following specification:

$$Flow_{j,t} = a_j + b_t + \gamma_1 NetAlpha_{j,t-1} + \gamma_2 (NetAlpha_{j,t-1})^2 + \theta X_{j,t-1} + e_{j,t} \tag{1.7}$$

The dependent variable, $Flow_{j,t}$, is defined as percentage growth of total net assets ($TNA_{j,t}$) for fund j in month t, net of internal growth $R_{j,t}$:

$$Flow_{j,t} = \frac{TNA_{j,t} - TNA_{j,t-1}(1 + R_{j,t})}{TNA_{j,t-1}}$$

To mitigate the impact of a number of errors in the CRSP database on flows (Elton, Gruber, and Blake (2001)), I winsorize fund flows at the 1% and 99% levels. I include fund fixed effects, a_j, to focus on the *within-fund* variation in flows. I include time fixed effects, b_t, to control for any unobserved aggregate force that may affect fund flows in general, such as investors sentiment. I include lagged fund flow, among others, as part of the time-varying fund-level controls, $X_{j,t-1}$, to account for potential sticky fund flows.[30]

Table 1.4 reports the results. Column (i) confirms an increasing and convex flow-performance sensitivity as predicted by Corollary I.5; point estimates on both $NetAlpha$ and $NetAlpha^2$ are positive and statistically significant.[31] As Columns (ii) to (v) show,

[30] The other controls are: lagged log of fund total net assets, lagged fund turnover, lagged log of fund age (in years), lagged fund expense ratio, and lagged realized return volatility in the past 12 months.

[31] Note that the convex flow-performance sensitivity documented here is in contrast with the concave relation documented in Goldstein, Jiang, and Ng (2017). I believe the main reason that they obtain a concave flow-performance sensitivity is their choice of fixed effects (FE). While they have both investment-grade (IG) and high-yield (HY) funds in their analysis, they include time FE, rather than time×sector FE (i.e. IG vs. HY sector), in their main specification (Table 2 in their paper). Given potential sector-specific forces that may affect fund flows across funds within the sector (for example, investors' overall preference on IG funds), it seems sensible to include time×sector FE. As I show in Appendix Table B.5, when time×sector FE are included, flow-performance sensitivity are no longer concave under the original specification of Goldstein, Jiang, and Ng (2017). Moreover, as shown in Appendix Table B.6, when I estimate the flow-performance sensitivity separately in the subsample of

Table 1.4: Flow-performance Sensitivity

This table shows the flow-performance sensitivity for U.S. corporate bond mutual funds from January 1991 to March 2017 (excluding money market funds, index funds, ETFs, and ETNs). The dependent variable is monthly net fund flow, defined as:

$$Flow_{jt} = \frac{TNA_{j,t} - TNA_{j,t-1}(1 + R_{jt})}{TNA_{j,t-1}}$$

where $TNA_{j,t}$ stands for the total net assets for fund j in month t and R_{jt} is the fund j's net return in month t. $NetAlpha^2$ is the quadratic term of $NetAlpha$. "Four-factor" is the baseline case, referring to the benchmark with four bond market index funds (VBMFX, VBISX, VBIIX, and VBLTX); "one-factor" is the benchmark with only VBMFX; "two-factor" is the benchmark with both VBMFX and VFINX; "five-factor" is the benchmark with five index funds (VBMFX, VFINX, VBISX, VBIIX, and VBLTX); and "risk-factor" is short-rate factor (three-month Treasury bill rate), slope factor (ten-year Treasury rate − one-year Treasury rate), curvature factor (two-year Treasury rate + ten-year Treasury rate − 2×five-year Treasury rate), and default risk factor (BAA-rated corporate bond yield − AAA-rated corporate bond yield). The time-varying fund-level controls include lagged log of fund total net assets, lagged fund turnover, lagged log of fund age (in years), lagged fund expense ratio, lagged realized return volatility in the past 12 months, and lagged fund flow. In the parentheses, I report the standard errors clustered by both month and fund.

	Four-factor (i)	One-factor (ii)	Two-factor (iii)	Five-factor (iv)	Risk-factor (v)
$NetAlpha_{t-1}$	2.73	3.42	3.09	2.97	2.01
	(0.35)	(0.38)	(0.53)	(0.38)	(0.29)
$(NetAlpha_{t-1})^2$	37.46	41.92	43.83	45.87	21.69
	(8.61)	(6.30)	(9.00)	(8.15)	(8.44)
Fund FE	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y
Observations	65109	69725	69717	65109	69725
Adjusted R^2	0.204	0.200	0.198	0.204	0.195

the finding of an increasing and convex flow-performance sensitivity is robust to various measures of fund performance.

The evidence of such an increasing and convex flow-performance sensitivity suggests that investors do learn about the fund over time in a way consistent with the model.

1.4.6 Dynamics

Having established (i) the pattern of returns to scale in each period and (ii) how investors update their beliefs on θ over time, I conclude this section by discussing the dynamics between fund size and next-period fund alpha with a numerical example.

Figure 1.2 illustrates the evolution of expected fund net alpha as a function of (equilibrium) fund size. The blue solid line describes the time-series variation between the fund size and the expected net alpha under the objective belief (which can be observed by econometricians). For each point on the solid blue line, there is a corresponding curve (dashed line) describing the same relation but under investors' subjective beliefs (which is not observable in the data). Each dashed line corresponds to a certain value of investors' perception of θ; that is, ϕ_t. As we can see from the figure, in the equilibrium, the subjective expected net alpha is always zero (red markers on the zero horizontal line) and the subjective returns to scale is always negative. The expected net alpha under the subjective and the objective beliefs coincide at zero when investors' perception of θ happens to be true (point D in the figure).

Now, let's consider the dynamics after some shocks to fund alpha. First suppose that, after a sequence of *bad* shocks, investors become quite pessimistic on θ and thus, the (equilibrium) fund size becomes fairly small (say, point A in the figure). At that point, since the (observed) next-period fund net alpha (which is the same as the expected net alpha under the objective belief because no further shock hits) is positive,

IG and HY funds, I do not find such a concave relation in either subsample.

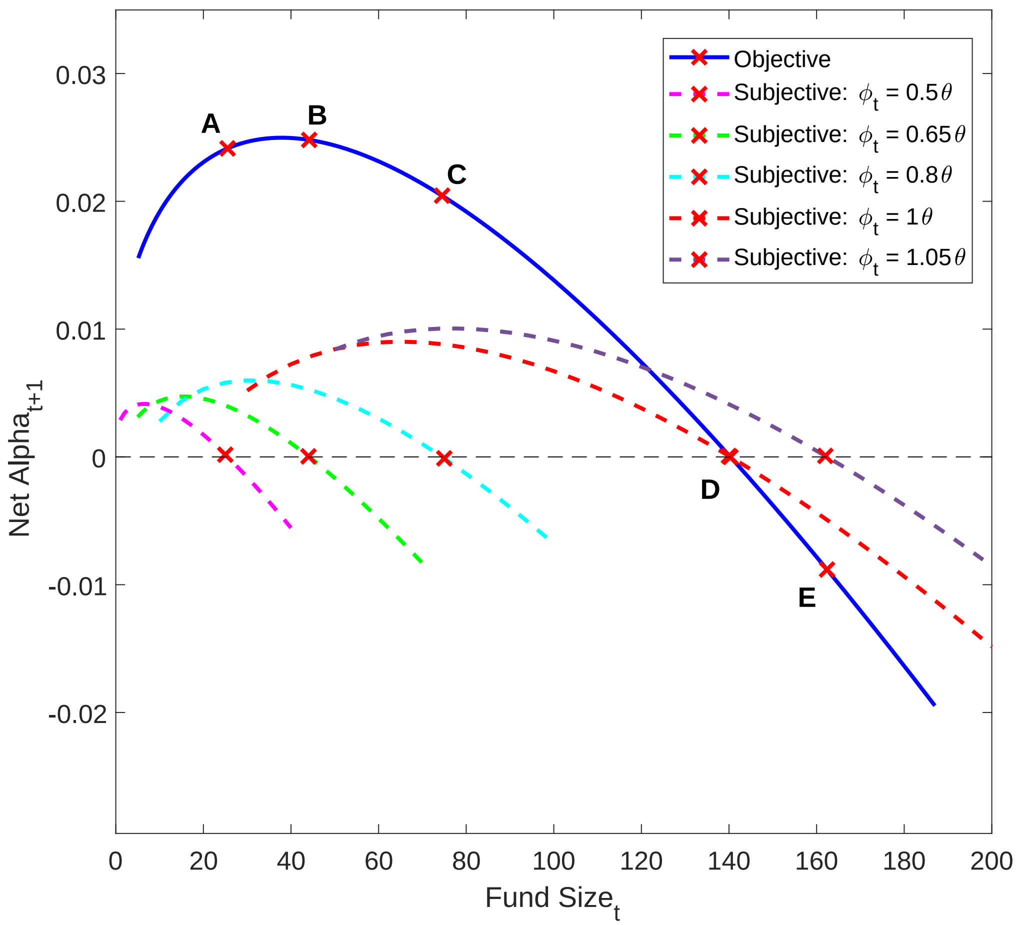

Figure 1.2: Dynamics between Fund Size and Expected Next-period Fund Alpha. This figure illustrates how expected next-period net alpha evolves as fund size varies under both investors' subjective and the objective beliefs. The parameter values are set as: $\mu - c_0 = 0.01$, $c_1 = 0.001$, $\theta = 0.005$, $\beta = 0.7$

investors would update their beliefs on θ upwards (see Equation 1.5). As a result, the subsequent fund size would become larger, which further leads to an increase in the future fund net alpha due to *increasing* returns to scale. Once ϕ_t passes the turning point but still remains smaller than θ (say, from point A to point B in the figure), the next-period fund net alpha remains positive. Under this scenario, both ϕ_t and the (equilibrium) fund size keep expanding but the subsequent fund net alpha now decreases rather than increases because of *decreasing* returns to scale (from point B to point C in the figure). Eventually, when ϕ_t reaches θ, the next-period fund net alpha becomes zero and hence, no more updating would occur in the absence of any further shock (point D in the figure).

The dynamics after a sequence of *good* shocks to fund alpha operate in a similar fashion but in the opposite direction. Suppose ϕ_t is now greater than θ and hence, the (equilibrium) fund size is quite large (say, point E in the figure). In this situation, the (observed) next-period fund net alpha is negative and investors would therefore update their beliefs downwards, resulting in a smaller subsequent fund size but a higher future fund net alpha (due to decreasing returns to scale). Once ϕ_t retreats to θ (from point E to point D in the figure), the next-period fund net alpha stays at zero and no further updating will occur if there are no more shocks.

1.5 Conclusion

In this paper, I study fund-level returns to scale for U.S. corporate bond mutual funds. In contrast to the decreasing returns to scale among equity mutual funds documented in the literature, I find a novel hump-shaped relation between fund size and subsequent fund alpha among corporate bond funds. Further, I show that a U-shaped relation between trade size and unit transaction cost at the bond level—a unique feature in the corporate bond market—is relevant for explaining such a hump-shaped pattern. To interpret the observed pattern of returns to scale (which is otherwise difficult in a standard rational expectations framework), I propose a model with a focus on fund investors' behavior and show that the evolution of investors' beliefs could drive the variation of fund size and further in fund alpha over time and thus, traces out the empirical (within-fund) returns-to-scale relation.

Although my proposed source of variation—the wedge between the subjective and objective beliefs about the fund's trading skill—is intuitive and plausible, it remains open what is the underlying mechanism that leads to the wedge. I suggest that investors' real-time learning as well as their misperception of fund alpha are potential channels. It will be interesting to explore in future research the extent to which these two plausible mechanisms can quantitatively match the empirical pattern of returns to scale and the flow-performance sensitivity.

CHAPTER II

The Making of Hawks and Doves

2.1 Introduction

Members of central-bank committees, such as the Federal Open Market Committee (FOMC) or the European Central Bank (ECB) Governing Council, often disagree on future inflation rates and whether to loosen or tighten monetary policy. Why do these highly educated and well-informed experts differ in their forecasts and recommendations when they have access to the same data and tools? Why do they deviate in their expectations from forecasts produced by their staff as documented by Romer and Romer (2008)?

Existing macroeconomic models of optimal monetary policy do not offer much of an explanation. Monetary policy makers, if modeled at all, assign the same weights to inflation and output stabilization, based on private-sector agent preferences and objective data, when maximizing social welfare (see, e. g., Rotemberg and Woodford (1999)). Even in models with learning, such as Sargent (1999), policy makers form beliefs based on objective historical data, which leaves no room for subjective disagreement.[1]

These modeling approaches are hard to square with the discussions among prac-

[1] Outside of macroeconomics, research on group decision-making has explored sources of heterogeneity among monetary policy committee members, including variation in preferences such as career-concerns, and differential information. For an overview, see Sibert (2006).

titioners and in the media classifying central bankers as 'hawks' or 'doves.' Debates about new appointments and their policy implications typically refer to appointees' background and personal experiences. For example, when Charles Plosser and Richard Fisher resigned as the Philadelphia and Dallas Federal Reserve Bank Presidents in 2014, much of the news coverage was about 'the generational shift rooted in personal inflation experiences: *"Annual inflation in the United States has averaged 3.8 percent during Mr. Plosser's adult life. By contrast, inflation has averaged just 2.5 percent during the adult life of Narayana Kocherlakota, president of the Federal Reserve Bank of Minneapolis, who at 50 is the youngest member of the policymaking committee and who has become the most outspoken proponent of expanding the Fed's stimulus campaign."*[2]

In this paper, we argue that personal experiences exert a measurable and statistically significant longterm influence on FOMC members. Whether and at what age they experienced, say, the Great Inflation or other inflation realizations, affects their stated beliefs about future inflation, their monetary-policy decisions, and the tone of their speeches on monetary-policy issues. We further show that time-variation in the average inflation experiences of all FOMC members present at a given meeting helps explain deviations of the federal funds rate from a conventional forward-looking Taylor rule.

Our research hypothesis and design build on a growing literature on *experience effects*. Personal experiences of macro-finance, labor-market, or political outcomes appear to be a strong determinant of individual attitudes and willingness to take risks in these areas in the long run. For example, prior experiences of stock market returns predict stock-market investment, prior experiences with IPOs predict future participation in IPOs, and prior experiences in the bond market predict future bond investment.[3]

[2] See "Charles Plosser and Richard Fisher, Both Dissenters, to Retire From Fed," by Binyamin Appelbaum, New York Times Sept. 22, 2014, www.nytimes.com/2014/09/23/business/fed-official-critical-of-policies-set-to-retire-in-march.html.

[3] Cf. Vissing-Jorgensen (2003), Kaustia and Knüpfer (2008), Chiang, Hirshleifer, Qian, and Sher-

Evidence in line with experience effects is also found among college students who graduate in recessions, among consumers who live through economic booms or busts, and in the political realm in terms of the long-term consequences of living under communism, its surveillance system, and propaganda.[4] Most closely related, Malmendier and Nagel (2016) show that life-time experiences of inflation significantly affect beliefs about future inflation, and that this channel explains the substantial disagreement between young and old individuals in periods of highly volatile inflation, such as the 1970s.

The monetary-policy setting in this paper is different. FOMC members are presumably highly educated and well informed about macroeconomic history, and monetary policy is generally considered a technocratic and model-driven area of economic policy. Experience effects may thus seem much less plausible than for the consumers and individual investors examined in earlier studies. Nevertheless we find a robust influence of personal experiences on FOMC members' stated beliefs and decisions, consistent with views in the media about generational origins of 'hawkishness.'

For our empirical analyses, we employ a model of experience-based learning that maps the history of each member's experienced inflation into a perceived long-run mean and persistence of inflation. Members' perceived-inflation dynamics follow a seasonal AR(1) process where experienced lifetime data is weighted with (roughly) linearly declining weights as estimated in Malmendier and Nagel (2016). The parameter estimates of the inflation process are updated each period. We then construct an experience-based inflation forecast for each FOMC member at each point in time as the

man (2011), Malmendier and Nagel (2011), and Strahilevitz, Odean, and Barber (2011). There is similar evidence for the housing market (Malmendier and Steiny (2017), Botsch and Malmendier (2016)), and the insurance markets (Gallagher (2014)).

[4]Cf. Kahn (2010) and Oreopoulos, von Wachter, and Heisz (2012) for labor markets; Malmendier and Shen (2017) for consumption expenditures (controlling for financial constraints and wealth); and Alesina and Fuchs-Schündeln (2007), Lichter, Löffler, and Siegloch (2016), Fuchs-Schündeln and Schündeln (2015), or Laudenbach, Malmendier, and Niessen-Ruenzi (2018) for political experiences. For example, Fuchs-Schündeln and Schündeln (2015) argue that the amount of time a person has lived under a democratic system determines her political preferences for democracy.

inflation forecast that is implied by these parameter estimates. These forecasts differ not only across cohorts in each period, but also change within each cohort over time as beliefs are updated in response to new inflation realizations. Hence, the identifying variation that we rely on to explain FOMC member behavior is not spanned by fixed age, time, and cohort effects.

As our first outcome variable, we analyze the inflation forecasts FOMC members submit for the semi-annual Monetary Policy Reports (MPRs) to Congress. The individual forecasts are made available with a 10-year lag, starting in 1992. We relate each member's experience-based forecast at a given time directly to their MPR forecast at that time. Despite the limited sample period, our estimation provides robust evidence that members put a substantial weight—37% or more, depending on the specification—on their experience-based forecasts. Hence, differences in members' lifetime experiences of inflation explain an economically significant portion of the differences in their inflation forecasts.

This first finding helps explain the puzzling time-series evidence in Romer and Romer (2008) that the central tendency of FOMC members' inflation expectations often deviates from the Federal Reserve staff's Greenbook forecast, even though their deviations *reduce* forecast accuracy. Our results imply that, to a large extent, the deviations are explained by reliance on personal inflation experiences. Hence, while our research design emphasizes between-member differences in experiences and outcomes, the estimates are also useful to understand why FOMC members as a group deviate from objective benchmarks.

Next, we turn to differences in decision-making. We study FOMC votes, which allow us to study clearly defined policy decisions over a sample period spanning several decades, from March 1951 to January 2014. The FOMC meets at least four (and typically eight) times per year. Members share their assessments of the economic situation and propose targets and policy measures for the upcoming inter-meeting time window, as well as the long-run. To analyze whether FOMC members' voting decisions are influenced by the inflation experiences they have accumulated during their lifetimes, we have to map their experience-based forecasts from the first step of our analysis into a voting decision. For this second step, we link the experience-based inflation forecasts to the desired level of nominal interest rates using a subjective version of the Taylor (1993) rule. We allow FOMC members to differ, based on their personal characteristics, in their weights on the inflation and output stabilization objectives as well as in their views about the appropriate inflation and output targets and the natural interest rate. Most importantly, we allow for the possibility that they evaluate deviations from the inflation target in terms of their own experience-based inflation forecasts. We estimate a highly significant relationship between inflation experiences and voting decisions. A one within-meeting standard-deviation increase in the experience-based inflation forecast raises the probability of a hawkish dissent by about one third, and it lowers the probability of a dovish dissent also by about one third, relative to the unconditional dissent probabilities.

The voting outcome is a clear indication that experiences significantly affect FOMC members' behavior; but it is also coarse, given the well-known reluctance of FOMC members, in particular governors, to formally cast a dissenting vote. To tease out more subtle differences in desired interest rate changes, we analyze, in a third step, the opinions FOMC members express in their speeches. We construct a data set of all "Speeches and Statements" from the Federal Reserve Archival System for Economic Research (FRASER) as well as hand-collected speeches from the websites of the regional Federal Reserve Banks (FRBs). We classify the language in these speeches and discussions as hawkish or dovish using the automated search-and-counts-approach of Apel and Grimaldi (2014). Applied to our sample, their *Net Index* of hawkishness reveals that FOMC members use a significantly more hawkish tone when their lifetime experiences imply a higher experience-based inflation forecast.

Finally, we turn from the cross-sectional analysis of individual behavior to the time series of the federal funds rate target. Traditionally, the FOMC implements monetary policy by setting a target for the federal funds rate, i. e., the interest rate at which banks lend overnight to each other. Within the forward-looking Taylor rule framework, we show that the federal funds rate target is tilted away from the Federal Reserve Board staff's Greenbook forecast of inflation and towards the experience-based inflation forecasts of the voting members present at the FOMC meeting. This result is robust to including the lagged federal funds rate in the interest-rate rule to account for interest-rate smoothing as part of the Federal Reserve's policy. Moreover, the strength of the tilt that we estimate here is broadly consistent with the tilt away from the staff forecast and towards personal experiences in the initial analysis of inflation forecasts. We quantify the implied effect in a rough calculation that abstracts from the equilibrium consequences of a different interest-rate path. We find that, relying only on the staff forecast and *not* on members' own inflation experiences, a counterfactual FOMC would have chosen a similar interest-rate path in the late 1980s and 1990s, but 50 to 100 basis points lower in the 2000s.

The four sets of empirical results can be parsimoniously explained by a model of experience effects, in which personal inflation experiences affect subjective beliefs about future inflation. Under such a model of *experience-based learning*, individuals overweight realizations of past inflation that they have experienced in their lives so far, consistent with earlier evidence on experience effects in individual inflation expectations (Malmendier and Nagel (2016)). In addition, there might be a preference-based link between inflation experiences and aversion to inflation. A preference-based explanation does not suffice, though, to explain all of our findings for at least two reasons. First, the preference channel does not easily explain the link between inflation experiences and FOMC members' stated beliefs in their MPR forecasts. While it is possible that the MPR forecasts reflect members' inflation preferences rather than their beliefs, this is not the standard interpretation of these data (e.g., Romer and Romer (2008)). Second, it is not clear why experience-based forecasts generated by an adaptive learning rule, which our empirical analysis employs, would be a good way to summarize FOMC members' inflation preferences. Ultimately, pinning down the precise channel is not essential for the validity of our findings. Irrespective of the preferred explanation, our findings show that heterogeneity in lifetime experiences has significant explanatory power for the heterogeneity in monetary-policy views and for the decisions of the experts on the FOMC.

Our findings add to a growing literature that studies experience-related heterogeneity in economic decisions and macroeconomic expectations. Relative to the macro and finance literature on experience effects cited above, our analysis stands out in that it is the first paper to provide evidence of personal experiences affecting policy experts. At the same time, it is not the first to find experience effects among professional agents, and complements prior evidence on experience effects among mutual fund managers who experienced the stock market boom of the 1990s (Greenwood and Nagel (2009)), among CEOs who grew up in the Great Depression (Malmendier and Tate (2005), Malmendier, Tate, and Yan (2011)), and even among lenders in 18th century Amsterdam (Koudijs and Voth (2016)).

Our results provide a new perspective on macroeconomic models in which monetary policy makers learn about the economy's stochastic processes (see Sargent (1999), Cho, Williams, and Sargent (2002), and Primiceri (2006), among others). A common assumption in these models is that policy makers update their beliefs (e. g., about the natural rate of unemployment, the slope of the Philips curve, or inflation persistence) using a constant-gain updating scheme that leads to perpetual learning with exponential downweighting of data in the past. Primiceri (2006) shows that this constant-gain assumption helps the model explain macroeconomic dynamics, including the high inflation in the 1970s and the subsequent disinflation. However, it is unclear why policymakers would update beliefs with a constant gain. One (standard) explanation is structural change in the stochastic processes agents learn about. Our findings point to an alternative: Data in the distant past carries low weight because policy makers overweight personal experience relative to objective historical data. In fact, Malmendier and Nagel (2016) show that the average experience-based belief of a group of individuals can be closely approximated by a constant-gain learning rule, and hence experience effects can provide an approximate "microfoundation" for constant-gain learning.

In addition, our results highlight sources of belief heterogeneity that the standard representative policy-maker approach in the literature would miss: the age distribution of the policy committee, as well as the differences in such age effects over time. As such, the evidence in this paper sheds light on the likely consequences of choosing specific individuals as central bankers—a topic much discussed in practice. Romer and Romer (2004) provide narrative evidence that the Federal Reserve chairs are heterogeneous in their views about the workings of the macroeconomy and the potency of monetary policy. They argue that this heterogeneity affects policy choices. Accordingly, Reis (2013) suggests that the choice of a central banker shapes the effective objective function for the central bank. Our evidence suggests that heterogeneity in macroeconomic experiences influence the beliefs that enter as inputs into this objective function.

Our evidence on the role of inflation experiences also adds a new dimension to a prior literature that links monetary policy decisions to the personal characteristics of FOMC members. Chappell, Havrilesky, and McGregor (Chappell, Havrilesky, and McGregor (1993), Chappell, Havrilesky, and McGregor (1995)) and Chappell and Mc-

Gregor (2000) document that a number of characteristics, including the role of regional Federal Reserve president versus Federal Reserve governor, are associated with differences in voting. While this earlier literature views policy maker characteristics as determinants of their preferences or incentives, our approach is motivated by a subjective beliefs channel. In support of this channel, we show that lifetime experiences explain FOMC members' stated beliefs about future inflation.

Finally, our analysis of the tone in FOMC members' speeches relates to the literature on textual analysis in monetary policy. Apel and Grimaldi (2014) measure the tone of the Swedish central bank minutes and use it to predict policy rate decisions. Numerous other text-mining approaches have recently been employed, for example by Hansen and McMahon (Hansen and McMahon (2016a), Hansen and McMahon (2016b)). Lucca and Trebbi (2011) analyze FOMC communication using an automated linguistics-based method to predict long-term Treasury yields. Most studies in this area focus on the role of transcripts, minutes, and statements of official meetings in predicting macro variables. We focus on how personal experiences explain tone differences across FOMC members' speeches outside their meetings.

The rest of the paper is organized as follows. In the next section, we lay out the methodology underlying our empirical approach and specify FOMC members' learning rule. We show that the resulting experience-based forecasts of inflation help predict the MPR inflation forecasts of FOMC members. In Section 2.3, we map the experience-based inflation forecasts into desired interest rates and show that they help explain dissenting votes. In Section 2.4, we perform a similar analysis for FOMC members' speeches. Section 2.5 relates the average inflation experiences of all FOMC members at each meeting to the federal funds rate decision, and Section 2.6 concludes.

2.2 Inflation Experiences and Inflation Forecasts

We start our analysis by examining the stated inflation expectations of FOMC members in the Semiannual Monetary Policy Report (MPR). This data set provides us with an inflation forecast for each individual FOMC member twice a year during the period from 1992 to 2004. We test whether we can detect experience-related heterogeneity in inflation expectations, even among the highly educated and professionally trained individuals on the FOMC: Does their personal lifetime experience of more or less inflationary environments affect their stated beliefs about future inflation? Do they attach higher weights to past realizations of inflation if they happen to have personally lived through those times?

2.2.1 Learning from Experience

Experience-based learning is a variant of adaptive learning where economic agents have a perceived law of motion for the variable they want to forecast, which may be a simple approximation of some unknown true law of motion. The agents estimate the parameters of this law of motion based on observed data and then use the estimated model to construct forecasts. As new observations arrive, they update the parameter estimates and forecasts. (See, e.g., Bray (1982), Marcet and Sargent (1989), Sargent (1993), and Evans and Honkapohja (2001).) The key modification of the standard approach that introduces learning from experience is that we allow the learning gain, i. e., the strength of updating in response to surprise inflation, to depend on age. Young individuals react more strongly to an inflation surprise than older individuals who already have accumulated a longer data set of lifetime observations. As a result, experience-based forecasts at a given point in time are heterogeneous by age (or, equivalently, across cohorts). Moreover, since individuals update their beliefs in response to new observations, experience-based forecasts vary within person, and hence within cohort.

There are no fixed cohort effects.

We utilize the learning-from-experience model of Malmendier and Nagel (2016) to generate FOMC members' experience-based inflation forecasts based on their experienced inflation histories, which we then compare with FOMC members' actual inflation forecasts. In the learning-from-experience framework of Malmendier and Nagel (2016), individual consumers perceive inflation as an AR(1) process, and use data on experienced inflation to estimate the AR(1) parameters and construct their forecasts. As they experience new inflation realizations, they update the AR(1) parameters and revise their forecasts. Intuitively, the AR(1) assumption implies that experienced inflation is summarized in terms of long-run mean and the persistence of shocks.

We modify this framework in a minor way to address seasonality. Especially towards the end of our sample period, the seasonal component of inflation accounts for a substantial share of its variance,[5] and we expect experts to be aware of the pattern. While the seasonality adjustment is not material for the results, it avoids seasonality-induced volatility in experienced-based forecasts in the later part of the sample, which plays a bigger role in the analysis here than in the Malmendier and Nagel (2016) sample that reached back to the 1950s. Hence, we model their perceived law of motion as a mixed seasonal AR(1) process,

$$\pi_{t+1} = \alpha + \phi_1 \pi_t + \phi_4 \pi_{t-3} - \phi_5 \pi_{t-4} + \eta_{t+1}, \tag{2.1}$$

where the $t-3$ and $t-4$ lags capture a four-quarter seasonal pattern.[6]

[5] Bryan and Cecchetti (1995) show that the relative variance share of the seasonal component rose as inflation became more stable after 1982, and Gospodinov and Wei (2015) note a strong seasonal component since the financial crisis in 2008.

[6] With the restriction $\phi_5 = \phi_4\phi_1$, this is an $ARIMA(1,0,0) \times (1,0,0)_4$ model, a special case of the seasonal ARIMA model discussed, e.g., in Box, Jenkins, Reinsel, and Ljung (2015). We do not impose this restriction in the learning algorithm (which does not affect consistency), so that the belief updating formulas still retain a recursive least-squares form. An alternative would be to use seasonally-adjusted data. However, seasonally-adjusted data is available only back to 1947. Moreover, standard seasonally-adjusted data suffers from a potential look-ahead bias as the seasonal adjustment

FOMC members use least-squares to estimate the vector b of parameters in (2.1), $b \equiv (\alpha, \phi_1, \phi_4, \phi_5)'$. Expressed recursively, the least-squares estimates of an FOMC member born in year s are updated every period as follows:

$$b_{t,s} = b_{t-1,s} + \gamma_{t,s} R_{t,s}^{-1} h_{t-1}(\pi_t - b'_{t-1,s} h_{t-1}), \tag{2.2}$$

$$R_{t,s} = R_{t-1,s} + \gamma_{t,s}(h_{t-1} h'_{t-1} - R_{t-1,s}), \tag{2.3}$$

where $h_t \equiv (1, \pi_t, \pi_{t-3}, \pi_{t-4})'$. Based on the newly revised estimates of $b_{t,s}$, members of cohort s form their subjective expectation of next period inflation as

$$\pi^e_{j,t+1|t} = b'_{t,s} h_t. \tag{2.4}$$

The sequence of gains $\gamma_{t,s}$ in (2.2) and (2.3) determines how strongly cohort s revises the parameter estimates when faced with an inflation surprise, $\pi_t - b'_{t-1,s} h_{t-1}$, at time t. Following Malmendier and Nagel (2016), we specify the gain as

$$\gamma_{t,s} = \begin{cases} \frac{\theta}{t-s} & \text{if } t - s \geq \theta, \\ 1 & \text{if } t - s < \theta. \end{cases} \tag{2.5}$$

That is, while the recursive least-squares set up follows standard implementations of adaptive learning (cf.; Evans and Honkapohja (2001)), the gain specification is different. In standard adaptive-learning models with decreasing gain, the gain is decreasing in the total size of available historical data and is the same for everybody. In contrast, the gain in (2.5) is decreasing in the size $t - s$ of the *lifetime* data of cohort s at time t. As a consequence, younger individuals have a higher gain and react more strongly to an inflation surprise than older individuals. Hence, the variation in gains is the

factors applied to the CPI time-series are estimated and retroactively updated by the Bureau of Labor Statistics using ex-post realized data over the full sample.

source of between-cohort heterogeneity in inflation forecasts, as well as within-cohort heterogeneity (over time), in our framework.

The parameter $\theta > 0$ is constant and determines how much weight the forecaster puts on recent data versus data in the distant past. For example, $\theta = 1$ implies equal weighting of recent data and data earlier in life, while $\theta > 1$ implies that recent data receives more weight than early experiences. Throughout the paper, we conduct our baseline estimation by setting $\theta = 3.044$, which is the value Malmendier and Nagel (2016) estimate from the data on inflation expectations in the *Michigan Survey of Consumers* (MSC). This value of θ implies that weights on past observations decline a little faster than linearly, going back from the current period to a weight of zero at birth.[7] By using this value of θ, we impose consistency with earlier evidence and tie our hands with regards to this parameter, rather than picking θ to best fit the FOMC member data. We test the robustness of our results to using a range of values around this point estimate. We also reestimate θ on the sample of college graduates in the MSC, which makes it plausibly more representative of the typical FOMC member. Our results are unaffected when we use the resulting parameter estimate of $\theta = 3.334$.

For a given θ, we calculate the experience-based inflation forecast $\pi^e_{j,t+1|t}$ of member j at time t based on inflation data since j's birth year. Our data source is the quarterly CPI series from Shiller (2005) that goes back to 1871Q1.[8] We measure inflation rates

[7]We find that the inflation forecast of an adult is not sensitive to the precise starting point of the experience accumulation for a fairly wide range of values around $\theta = 3.044$. In Malmendier and Nagel (2016), we stretch and compress the weighting function to include years before birth into the experience accumulation or start later (e.g., at the age of 18) without much effect, also because the initial years in an adult's lifetime carry relatively little weight.

[8]See the updated long-term stock, bond, interest rate and consumption data at `http://www.econ.yale.edu/~shiller/data.htm`. Shiller's inflation rate series is based on the CPI-U (Consumer Price Index-All Urban Consumers) published by the U.S. Bureau of Labor Statistics from 1913 onwards, and on the Warren-Pearson wholesale price index before 1913. Since the earlier price index is focused on commodities, it is more volatile. Appendix C.7 replicates key parts of our analyses excluding pre-1913 data, i.e., restricting the sample to FOMC members born after 1913. The results on voting remain essentially unchanged, as do the results on speech tone; the other two sets of analyses do not use pre-1913 data.

as annualized quarterly changes in the log CPI. As in Malmendier and Nagel (2016), we iterate on the perceived law of motion (2.1) at each cohort's quarter-t parameter estimates to construct experience-based forecasts of the average inflation rate over the relevant horizon (which is four quarters in most of our applications, unless otherwise noted).

In Appendix C.1, we illustrate the resulting heterogeneity in expectations and learning-from-experience dynamics in more details. There, we plot how the perceived persistence and long-run mean of inflation evolve over time, separately for different age groups. The graphs highlight the two key features of experience-based expectations formation. First, since individuals update their beliefs in response to new inflation observations, experience-based forecasts vary within person (and hence also within cohort) over time. Second, since younger individuals have a shorter life-time data set and place a higher weight on recent inflation surprises than older individuals, expectations are heterogeneous by age, but in a time-varying way. As a consequence, a linear combination of time, age, or cohort fixed effects cannot absorb experience-based expectations heterogeneity. For this reason, our approach to estimating experience effects is not subject to the age-time-cohort collinearity problem that plagues methods that are based on estimation of cohort fixed effects. (See Malmendier and Nagel (2016) for a more general discussion of this point.)

2.2.2 Inflation Forecast Data

We obtain individual inflation forecasts of FOMC members from the Semiannual MPR.[9] Twice a year, in February and July, the FOMC submits an MPR to Congress, which contains the FOMC members' inflation forecasts. In February, the forecasts concern the time period from Q4 of the previous year to Q4 of the current year. In July, two

[9] `www.philadelphiafed.org/research-and-data/real-time-center/monetary-policy-projection`

sets of forecasts are included in the report: one for Q4 of the previous year to Q4 of the current year, and another one for Q4 of the current year to Q4 in the next year.

We supplement the individual FOMC members' forecasts with forecasts in the "Greenbooks" that are prepared by Federal Reserve staff about a week prior to each FOMC meeting.[10] We use the Greenbooks for the February and July FOMC meeting and match them with the member forecasts from the MPR. As Romer and Romer (2008) discuss, the FOMC members have access to the Greenbook forecasts when they prepare their forecasts before the FOMC meeting that precedes the MPR. They also have an opportunity to revise their forecast after seeing other members' economic views and staff's summary of the other members' forecasts. Romer and Romer (2008) show that the central tendency of FOMC members' forecasts deviates from the staff forecast in the Greenbooks, and that this deviation from the staff forecasts reduces the forecast accuracy.

Our objective here is to test whether the deviations from staff forecasts reflect the influence of their personal inflation experiences. For this purpose, we extract the individual inflation forecasts contained in the MPRs (rather than the central tendency that Romer and Romer (2008) analyze) to construct a panel data set. The individual FOMC members' forecasts become available only with a 10-year lag, and the earliest ones available are from 1992. Hence, our sample runs from 1992 to 2004, covering 26 FOMC meetings. This data set of individual forecasts is introduced and described in Romer (2010).

2.2.3 Econometric specification

Our estimating equation relates FOMC members' deviation from the staff forecasts to their personal inflation experiences. We start from modelling FOMC member j's

[10] `www.federalreserve.gov/monetarypolicy/fomc_historical.htm`

forecast at time t, $\tilde{\pi}_{j,t+1|t}$, as a weighted average of j's experience-based forecast $\pi^e_{j,t+1|t}$ and the staff forecast $\tilde{\pi}_{t+1|t}$ reported in the most recent Greenbook:

$$\tilde{\pi}_{j,t+1|t} = \phi \pi^e_{j,t+1|t} + (1-\phi)\tilde{\pi}_{t+1|t}. \tag{2.6}$$

Subtracting $\tilde{\pi}_{t+1|t}$ on both sides, we obtain our estimating equation

$$\tilde{\pi}_{j,t+1|t} - \tilde{\pi}_{t+1|t} = a + \phi(\pi^e_{j,t+1|t} - \tilde{\pi}_{t+1|t}) + \varepsilon_t, \tag{2.7}$$

where we include a constant and a residual to account for other unobserved variables that could influence the FOMC members' forecasts.

One complication when estimating equation (2.7) is that the forecasted inflation variable switched in February 2000 from the consumer price index (CPI-U) to the price index for personal consumption expenditure (PCE). Our construction of $\pi^e_{j,t+1|t}$ is based on the history of the CPI, and from 2000 to the end of our sample in 2004, the average CPI inflation rate was about 0.40% higher than the PCE inflation rate. We take two approaches to address this discrepancy. First, we simply re-calculate $\tilde{\pi}_{j,t+1|t}$ post-1999 by adding the difference in CPI and PCE inflation rates over the 12 months prior to the meeting to the FOMC member forecast. Second, we estimate a version of equation (2.7) with time fixed effects. As long as views about the CPI-PCE discrepancy are similar among FOMC members, the effect of the discrepancy will be absorbed by the time fixed effects. In this case, the coefficient ϕ is identified purely from (time-varying) cross-sectional differences between FOMC members in their forecasts and their inflation experiences.

Another complication is that forecast horizons vary. To match the forecasts in the February MPR (from the end of the previous-year Q4 to the end of the current-year Q4), we construct the experience-based forecast using data until the end of previous-year Q4 and then iterate to construct a four-quarter-ahead forecast. To match the same (previous-year Q4 to current-year Q4) forecast in the July MPR, we average the two-quarter-ahead experience-based forecast (from end of Q2 to end of current-year Q4) and the realized inflation over the past two quarters (from end of last-year Q4 to end of Q2). To match the next-year forecast (from current-year Q4 to next-year Q4) in the July MPR, we subtract the same two-quarter-ahead experience-based forecast from the six-quarter-ahead experience-based forecast (from end of Q2 this year to end of Q4 next year).

Panel A in Table 2.1 reports summary statistics for the dependent and explanatory variables in (2.7), separately for each forecast horizon. The mean column shows that the FOMC members' actual MPR forecast exceeds the Greenbook forecast on average over the 1992-2004 sample period by between 0.17 to 0.32 percentage points. Interestingly, the same pattern, but at a greater magnitude, holds for FOMC members' experience-based forecast. This is a first hint that partial reliance on personal inflation experiences could be the reason why FOMC members deviate from the Greenbook forecast. The standard deviation column shows that actual and experience-based forecast deviations from the Greenbook have a standard deviation of around 0.50 percentage points for the February MPRs, and around 0.40 to 1.10 percentage points for the two July MPR forecasts. These means and standard deviations are large relative to the magnitudes of a typical federal-funds-rate target change of 0.25 percentage points that the FOMC might consider in a meeting.

The table also reports the within-member standard deviation of the actual and the experience-based forecast. This statistic reveals that member fixed effects do not absorb much of the variation. The much smaller within-meeting standard deviation in the next column indicates that much of the total standard deviation reflects time-series variation of the average members' deviation from the Greenbook forecast, rather

Table 2.1: Influence of FOMC Members' Inflation Experiences on their Inflation Forecasts

Panel A presents summary statistics for the dependent and explanatory variables in the estimations shown in Panel B. MPR fcst. - staff fcst. is the difference between i) FOMC members' stated inflation projection from the MPR and ii) the most recent Fed Staff's inflation forecast from the Greenbook prior to the February or July FOMC meeting. In February, the horizon of the members' MPR forecasts is over the four quarters until the end of the current year. In July, two horizons are available: four quarters until the end of the current year and the four quarters during next year. From February 2000 on, we add the difference between CPI and PCE inflation rate to each FOMC member forecast. The sample period runs from the first half of 1992 to the second half of 2004. In Panel B, MPR fcst. - staff fcst. is the dependent variable. The explanatory variable is the difference between the i) experience-based forecast $\pi^e_{j,t+1|t}$ for each FOMC member at each meeting, and ii) the Fed staff's inflation forecast. We calculate $\pi^e_{j,t+1|t}$ for each member at each meeting by recursively estimating a mixed seasonal AR(1) model using the member's lifetime history of inflation, as described in Section 2.2.1 (with $\theta = 3.044$). In parentheses we report the standard error based on clustering as described in the table.

Panel A: Summary statistics				
	Mean	S.D.	Within-Member S.D.	Within-Meeting S.D.
February MPR: Current-year forecast				
MPR fcst. - staff fcst.	0.26%	0.53%	0.44%	0.21%
Exp.-based fcst. - staff fcst.	0.66%	0.53%	0.43%	0.03%
July MPR: Current-year forecast				
MPR fcst. - staff fcst.	0.17%	0.44%	0.39%	0.18%
Exp.-based fcst. - staff fcst.	0.66%	1.09%	0.78%	0.03%
July MPR: Next-year forecast				
MPR fcst. - staff fcst.	0.32%	0.61%	0.50%	0.32%
Exp.-based fcst. - staff fcst.	1.16%	0.75%	0.61%	0.06%

Panel B: OLS regression	(i)	(ii)	(iii)	(iv)
Exp.-based fcst. - staff fcst.	0.37	0.40	0.81	0.82
	(0.10)	(0.12)	(0.37)	(0.39)
Member × fcst. horizon FE	No	Yes	No	No
Member FE	No	No	No	Yes
Meeting × fcst. horizon FE	No	No	Yes	Yes
Clustered s.e.	Member and Meeting	Member and Meeting	Member	Member
Observations	383	383	383	383
Adjusted R^2	34.7%	41.0%	77.7%	81.5%

than cross-sectional dispersion between members in a given FOMC meeting. This is a consequence of the fact that the sample period for these forecast data features relatively low and stable inflation rates. As a consequence, the heterogeneity in FOMC members' experience-based forecasts is limited. Our analysis of voting and speeches, which we turn to below, will instead cover the 1970s in its sample period, which bring in substantially greater dispersions in experience-based forecasts.

2.2.4 Estimation Results

The estimation results are in Panel B in Table 2.1. The panel reports the OLS estimates of the weight ϕ on the experience-based forecasts, relative to the staff forecasts, in equation (2.7). We find that the experience-based inflation forecast plays a significant role in explaining the variation of members' reported inflation forecasts. The specification in column (i) uses the total variation without fixed effects. The resulting estimate of 0.37 (s.e. 0.10) implies that FOMC members put about 37% weight on their experience-based forecast and 63% on the staff forecast. Figure 2.1 presents the scatter plot corresponding to this regression, comparing individual members' actual inflation forecast $\tilde{\pi}_{j,t+1|t}$ to their experience-based forecast $\pi^e_{j,t+1|t}$. The scatter plot illustrates the high R^2 of 34.7% in this regression.

The estimate of ϕ remains very similar when we add member$\times$forecast-horizon fixed effects, i. e., FOMC member dummies interacted with dummies for the three types of forecast in Panel A. As shown in column (ii), the coefficient estimate is now 0.40 (s.e. 0.12). This stability of the estimate implies that the results are not driven by cohort fixed effects (which are absorbed by the member fixed effects in this regression). Experience-based learners update their beliefs over time, and this time-variation in expectations is not captured by cohort fixed effects. Instead, the estimate is identified from variation in cross-sectional differences over time. The estimates in column (ii)

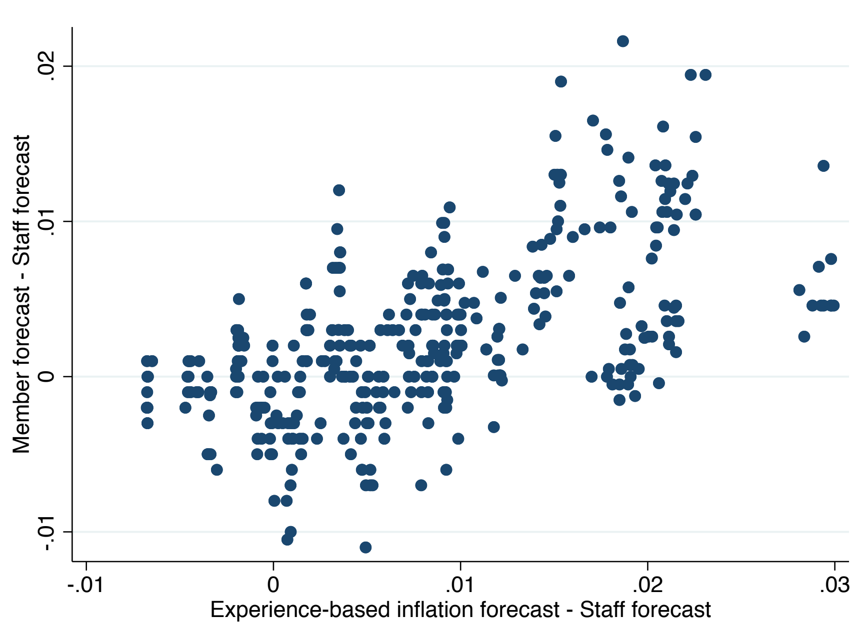

Figure 2.1: Relationship Between FOMC Member Inflation Forecasts in the MPR and their Experience-based Inflation Forecasts

also show that any alternative explanation based on fixed member characteristics (e.g., educational background) cannot explain the results.

The estimates so far largely reflect the time-series comovement of the average FOMC member's forecasts and experiences at a given meeting. Periods in which the average FOMC member submits an inflation forecast above the Greenbook forecast also tend to be periods in which the average FOMC member's experience-based forecast is above the Greenbook forecast. It is interesting that the time-series variation in these variables lines up so closely, as evident also from Figure 2.1. To rule out that that some omitted time-series factor is driving this co-movement, it is useful to focus on within-meeting variation. For this reason, we include meeting×forecast-horizon fixed effects in the estimations in columns (iii) and (iv). The magnitude of the ϕ estimate roughly doubles. However, only a small amount of variation remains after including this extensive set of fixed effects, and so the standard errors become fairly large. As a consequence, we cannot reject that the estimates are unchanged compared to those in column (i) and (ii). Nevertheless, even though pinning down the precise magnitude of the effect is difficult, it is reassuring that the results are qualitatively similar when we identify ϕ only from within-meeting variation.

Finally, we note that the estimates in column (iv) also include member fixed effects, on top of the meeting $\times$ forecast-horizon fixed effects. This estimation illustrates the point made earlier that the heterogeneity in experience-based inflation forecasts is not fully absorbed by time and member fixed effects. This dimension of identification constitutes the key difference between our approach and methods that try to capture experience effects through cohort fixed effects (which would be absorbed by the member fixed effects in column (iv)).

We conclude that the estimates are consistent with the view that heterogeneity in lifetime experiences of inflation results in significant heterogeneity in FOMC members' beliefs about future inflation. In terms of magnitude, while the focus on within-meeting variation in columns (iii) and (iv) is useful to achieve identification, independent of any correlated omitted time-series variables, the relevant variation for the assessment of experience effects and for counterfactual exercises is the total variation plotted in Figure 2.1, including the large between-meeting component. For example, to predict the policy stance of the committee, one may want to know by how much experience-based learning could shift the average member's inflation expectation away from the Greenbook forecast.

The large economic effect of personal inflation histories on FOMC members' stated beliefs has a similar order of magnitude as the effect estimated in the MSC. Among households surveyed in the MSC, Malmendier and Nagel (2016) find that that survey respondents put a weight of 0.67 on their experience-based forecasts. Considering the estimation uncertainty, it is difficult to make a precise comparison, but broadly, the weight put on personal experiences when forming inflation expectations appears quite similar across FOMC members and the households surveyed in the MSC.

In terms of interpretation, one potential concern specific to the FOMC setting is that strategic considerations might affect the forecasts stated in the MPR, including the desire to appear consistent or to send a message. This concern is somewhat muted because *individual* forecasts are actually not revealed in the MPR; they are made public only with a 10-year lag. The focus of public attention is usually on the published summary measures, especially the central tendency of the distribution of member forecasts. Also, as always with data on reported beliefs, it is important to keep in mind that it may not be possible to cleanly separate beliefs from preferences. Nevertheless, a direct effect of inflation experienced on beliefs about future inflation provides the most straightforward explanation of these results.

2.3 Inflation Experiences and Voting

Our first finding that FOMC members put substantial weights on their personal inflation experiences when forming inflation expectations raises the possibility that differences in experiences also give rise to differences in FOMC members' monetary policy stance. To find out, we examine how FOMC members' voting records relate to their inflation experiences. This analysis allows us to turn to actual monetary-policy decisions, and also to considerably expand the sample period backwards in time, compared to the relatively short sample period of MPR inflation expectations.

2.3.1 Policy Rule

In order to isolate the effects of inflation experiences on FOMC members' monetary-policy stance, we need a framework that allows us to map their beliefs about future inflation into their monetary-policy views. Such a framework should also allow for other sources of heterogeneity in policy preferences and incentives that could affect members' policy views.

We model monetary policy makers as following, explicitly or implicitly, an interest-rate rule that pins down their desired interest rates. We use the Taylor (1993) rule as a starting point, and augment it to allow for heterogeneity.

The standard Taylor rule implies a nominal interest rate

$$i_t^* = r + \pi^* + \lambda(\pi_t - \pi^*) + \gamma(y_t - y^*) \qquad \text{with } \lambda > 0,\ \gamma > 0, \tag{2.8}$$

where π_t is the inflation rate, π^* is the inflation target (assumed to be 2 percent by Taylor), y_t denotes output, y^* is potential output, and r is the "natural" real interest rate consistent with an output gap $y_t - y^*$ of zero. Orphanides (2003) shows that this rule explains well the evolution of the Federal Reserve's policy rate (federal funds rate) all the way back to the 1950s, with the exception of a few years in the early 1980s during the "Volcker disinflation." This does not mean that the FOMC explicitly followed such a rule; but its policy decisions are well described by this rule.

We augment the rule and introduce heterogeneity in two ways. First, we allow FOMC members to differ, relative to other members and over time, in their preferences for inflation versus output stabilization through different weights λ and γ, and in their views about the targets π^*, y^*, and about the natural rate r. Second, we introduce a subjective forward-looking element into the Taylor rule by allowing FOMC members to evaluate the deviations from the inflation target partly in terms of their own subjective inflation expectation.

With these sources of heterogeneity incorporated into the policy rule, FOMC member j's desired nominal interest rate at time t becomes

$$i_{j,t}^* = r_{j,t} + \pi_{j,t}^* + \lambda_{j,t}(\omega\pi_{j,t+1|t}^e + (1-\omega)\pi_t - \pi_{j,t}^*) + \gamma_{j,t}(y_t - y_{j,t}^*), \quad \text{where} 0 \le \omega \le 1. \tag{2.9}$$

The parameter ω represents the weight that FOMC members put on their own subjective expectation $\pi^e_{j,t+1|t}$ rather than the objective information π_t.

Note that one can go further and replace π_t and y_t with expectations of future inflation and output to make the Taylor rule fully forward-looking as in Clarida, Galí, and Gertler (2000). However, this does not change our estimating equation in our analysis of voting (and later our analysis of speeches) in a substantial way since we focus on cross-sectional heterogeneity between FOMC members. It will matter when we examine the time-series of the federal funds rate, and we will turn to a fully forward-looking specification there (at the expense of a much shorter sample, due to limited availability of forecast data).

We specify the heterogeneity of FOMC members in the following way:

$$\begin{aligned}
\lambda_{j,t} &= \lambda_0 + (x_{j,t} - \mu_x)'\lambda_1, \\
\gamma_{j,t} &= \gamma_0 + (x_{j,t} - \mu_x)'\gamma_1, \\
\pi^*_{j,t} &= \pi^* + (x_{j,t} - \mu_x)'\alpha_1, \\
y^*_{j,t} &= y^* + (x_{j,t} - \mu_x)'\alpha_2, \\
r_{j,t} &= r + (x_{j,t} - \mu_x)'\alpha_3,
\end{aligned} \tag{2.10}$$

where $x_{j,t}$ is a vector of characteristics of FOMC member j at time t with population mean μ_x. After substituting these expressions into equation (2.9), we perform a first-order Taylor approximation of $i_{j,t}$ as a function of $(\pi^e_{j,t+1|t}, x'_{j,t})$ around (π_t, μ'_x); cf. Appendix C.2. We obtain

$$i^*_{j,t} \approx a_t + \lambda_0\omega\pi^e_{j,t+1|t} + \kappa' x_{j,t} + \pi_t x'_{j,t}\lambda_1 + (y_t - y^*)x'_{j,t}\gamma_1, \tag{2.11}$$

where a_t is a time fixed effect and κ is a vector of constants. We use this version of the

Taylor rule to derive individual desired interest rates and corresponding policy views, whether expressed in voting decisions or speech tones.

2.3.2 Data on the FOMC Voting History

We study the FOMC voting history from March 1951 to January 2014. The starting point is dictated by the Treasury-Federal Reserve Accord of 1951, with which the Federal Reserve System regained its independence from the Department of Treasury after World War II.

The data comes from several sources. For meetings from January 1966 to December 1996, we use the data from Chappell, McGregor, and Vermilyea (2005). For meetings before January 1966 and after January 1997, we collect the data directly from FOMC meeting statements. Each statement reports all votes, typically followed by explanations of the dissenting opinions, if any. We exclude eight dissents that cannot easily be classified as hawkish or dovish.[11] Four FOMC members were both regional Fed presidents and governors at different points during their career, and we account for their varying roles in our empirical analysis.

We collect biographical information for each FOMC member from the Federal Reserve History Gateway[12] and the Who's Who database. The data includes the year and place of birth, gender, the highest degree earned, the program they graduated from, the role served in the Fed (board member or regional bank president), and the political party of the U.S president who was in office at the time of the member's first appointment.

We use these data to construct the vector $x_{j,t}$ of FOMC members' characteristics that we allow to influence the desired interest rate at meeting time t in equation (2.11). We include age to make sure the experience-based inflation forecast is not picking up

[11] Details on the construction of the voting data set are in Appendix C.3.

[12] `http://www.federalreservehistory.org/People`

Table 2.2: Summary Statistics

The table shows statistics for all FOMC meetings from 3/8/1951 to 1/29/2014. Details of the data construction are in Appendix C.3. The first column in Panel A reports the statistics for all FOMC members; and columns 2 to 4 report separately those for members who dissent towards monetary easing (*Dovish Dissent*), who consent (*Consent*), and who dissent towards monetary tightening (*Hawkish Dissent*). Panel B reports the pairwise correlations between voting record, experience-based inflation forecast, and member characteristics. We code *Vote* as 1 for a hawkish dissent, as 0 for a consent, and as -1 for a dovish dissent; *Fed Role* as 1 for regional Fed presidents and 0 for board members; *Party* as 1 if the member was first appointed while a Republican was U.S. president and 0 otherwise; and *Same Party* as 1 if the party of the U.S. president at the time of the appointment is the same as the party of the current president and 0 otherwise.

Panel A

	All	Dovish Dissent	Consent	Hawkish Dissent
#Meetings	659	109	659	178
#Votes	7,350	160	6,925	265
Avg. age	56.4	55.6	56.4	57.1
Avg. tenure (in days)	2,286	1,924	2,285	2,545
% w/ PhD	46.3	50.6	45.8	56.2
% studied Economics	67.5	70.6	67.0	78.9
% Male	93.9	83.1	93.9	100
% Regional Fed president	44.6	23.7	44.0	72.1
% Republicans	53.7	45.0	53.3	70.9
% Same party as current pres.	56.7	67.5	56.6	52.1
Expr.-based infl. fcst.: mean	3.4%	3.8%	3.4%	4.1%
std.dev.	1.8%	2.2%	1.8%	2.1%

Panel B: Pairwise Correlation

	Vote	Infl. fcst.	Male	Age	Fed role	Party	Same pty.
Vote	1.00	-	-	-	-	-	-
Exp.-based infl. fcst.	0.04	1.00	-	-	-	-	-
Male	0.08	-0.03	1.00	-	-	-	-
Age	0.02	-0.07	0.06	1.00	-	-	-
Fed role: Fed pres.	0.12	-0.01	0.10	-0.09	1.00	-	-
Party: Republican	0.07	0.15	-0.01	-0.02	0.10	1.00	-
Same Party	-0.03	0.05	-0.05	-0.18	0.03	0.12	1.00

an age effect, as well as other characteristics that the prior literature has found to be important determinants of FOMC voting (Chappell, Havrilesky, and McGregor Chappell, Havrilesky, and McGregor (1993), Chappell, Havrilesky, and McGregor (1995); Chappell and McGregor (2000)): gender, indicators for being a Regional Federal Reserve Bank President, for being appointed during the time a Republican U.S. president was in office, and for the U.S. president at the time of the first appointment being in the same party as the current president. For reasons we discuss below, we also include an interaction between the indicator for Regional Federal Reserve Bank President and an indicator for meeting times after November 1993.[13]

Table 2.2 presents the summary statistics. Our data covers 659 FOMC meetings with 7,350 votes. Overall, we have 160 dovish and 265 hawkish dissenting votes.

For the interpretation of the estimation results below, it is useful to keep in mind that the share of dovish and hawkish dissents is quite small, typically somewhere between 2.2% and 3.6%. These averages hide, however, a large degree of heterogeneity by role served and over time. Figure 2.2 shows the number of dissents in each FOMC meeting separately for Federal Reserve Board members (Panel a) and Regional Federal Reserve Presidents (Panel b). We can see that governors are much more likely to cast a dovish than a hawkish dissenting vote. The opposite holds for regional presidents, with a much higher fraction of hawkish dissents, as also indicated in Panel A of Table 2.2. Figure 2.2 also reveals a significant shift in voting behavior in November 1993, indicated by the red line. At that time, the Federal Reserve responded to pressure from Congress for more transparency and accountability, and agreed to publish lightly edited transcripts of the FOMC meetings with a five-year lag (Lindsey (2003)). Before 1993, the Federal Reserve published individual votes and summary minutes, but not

[13] In addition, we have checked the robustness to including further control variables and their interactions, such as tenure (as a possible control for expertise, cf. Hansen and McMahon (2016a)) and educational background. None of our results are affected if we include tenure, tenure squared, and controls for the school attended, the highest degree, and the field studied.

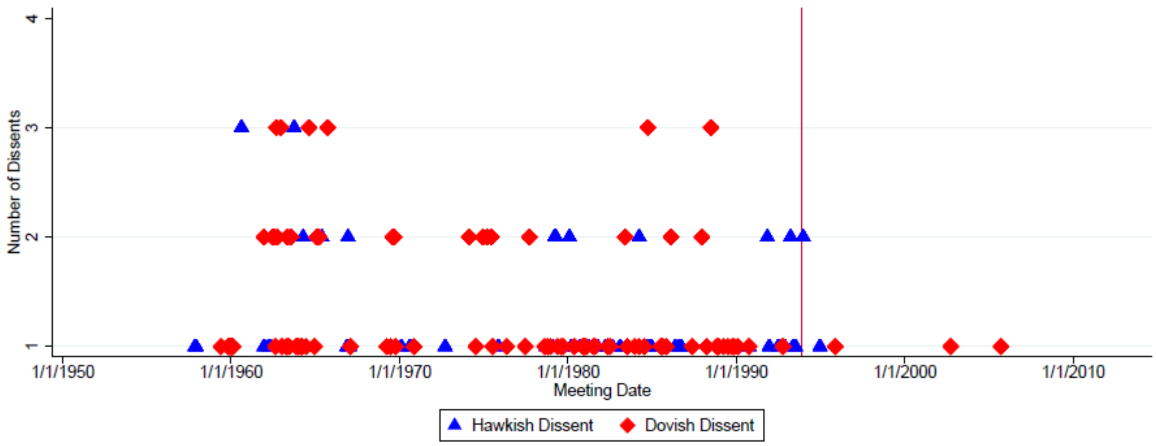

(a) Dissents by Federal Reserve Board Members

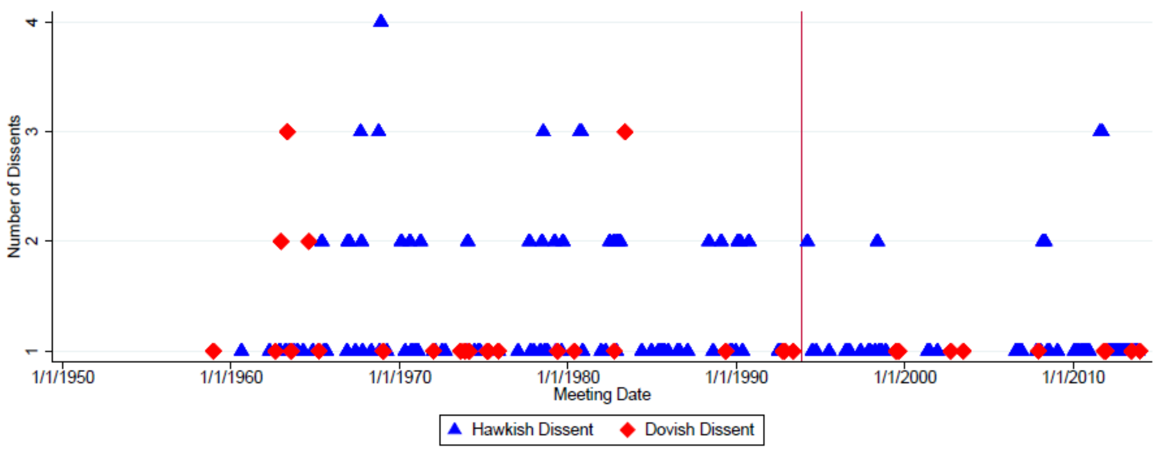

(b) Dissents by Regional Federal Reserve Presidents

Figure 2.2: Dissents in FOMC Meetings

Notes. The red vertical line is the time-stamp for November 1993, after which the FOMC agreed to make public its lightly-edited transcripts with a five-year lag.

the full transcripts. Meade and Stasavage (2008) find that this change reduced the willingness of FOMC members to verbally express dissents in the meetings. They also find a decrease in the propensity of Federal Reserve board members to dissent in formal voting, but the effect is not statistically significant in their sample until 1997. Figure 2.2, however, shows a fairly clear pattern. Dissents among Federal Reserve Board members became almost non-existent after the increase in transparency in 1993 (only 6 subsequent dissents). In contrast, dissents among regional Federal Reserve presidents remained quite common (71 subsequent dissents). Thus, the thresholds for FOMC members to voice dissent seems to have changed in 1993, and differently so for governors and presidents. This is an important feature of the data that we will need to accommodate in our econometric specification.

Returning to Panel A of Table 2.2, we see that hawkish dissenters are older, have a longer tenure on the FOMC, are more likely to have a PhD, to have studied economics, to be male, and to be appointed when the U.S. president in office was from a different party than the current U.S. president. (All differences other than the doctoral degree and field of study are statistically significant.) At the bottom of Panel A, we show the mean and standard deviation of FOMC members' experience-based forecasts $\pi^e_{j,t+1|t}$, calculated as described in Section 2.2.1. The average experience-based inflation forecasts for dovish dissenters is 3.8% while the average for hawkish dissenters is 4.1%, though the difference is not significant, and the average among consenters is even lower (3.4%).

Panel B shows the pairwise correlations between the key variables. We note again the positive relationship between the role of Fed president and votes leaning in a hawkish direction, and the same for being male, older, and Republican. Experience-based forecasts and hawkish voting are also positively correlated, and the correlation is significant. Our empirical analysis will test whether this relationship persists when analyzing the between-member variation in experiences after controlling for all other characteristics and their interaction effects, as implied by the policy rule (2.11).

In order to illustrate the identifying variation in our estimations, we plot two measures of the cross-sectional differences in experience-based inflation forecasts. Panel (a) of Figure 2.3 shows the learning-from-experience forecasts $\pi^e_{j,t+1|t}$ of the youngest and oldest FOMC members at each meeting, both net of the forecast of the median-age member. The differences range from 0 to 1.5 percentage points, with the biggest differences occurring during the high-inflation years of the late 1970s and early 1980s. At that time, younger members' inflation experiences are dominated by the high and persistent inflation of the 1970s, more so than those of older members, and young members have the highest experience-based forecasts. From the mid-1980s onwards, younger members adapted more quickly to the now low rates of inflation and the relatively low persistence, and the lines cross. The perception of a low inflation persistence among younger members also contributes to the spike around 2010, when young members' learning-from-experience forecast is temporarily much higher than the median: When faced with the recession-driven low inflation rates at the time, young members expected a faster reversion of inflation rates up (towards the mean of slightly above 2%) than older members.

As a second measure of the heterogeneity in experience-based inflation forecasts, Panel (b) plots the time-series of the within-meeting standard deviation of $\pi^e_{j,t+1|t}$. There is a lot of variation in this dispersion measure over time. A typical value would be around 0.1 percentage points (the full-sample within-meeting s.d. is 0.10 pp). It is useful to keep these magnitudes in mind for the interpretation of our empirical results below. Overall, the within-meeting dispersion of the experience-based forecasts is higher than in our earlier 1992-2004 sample of FOMC member inflation expectations.

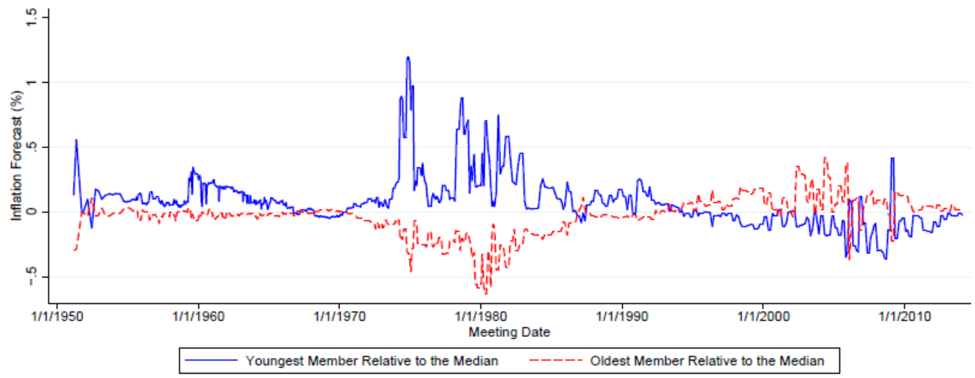

(a) Experience-based inflation forecasts of the youngest and the oldest FOMC member, relative to the median-age member's forecast

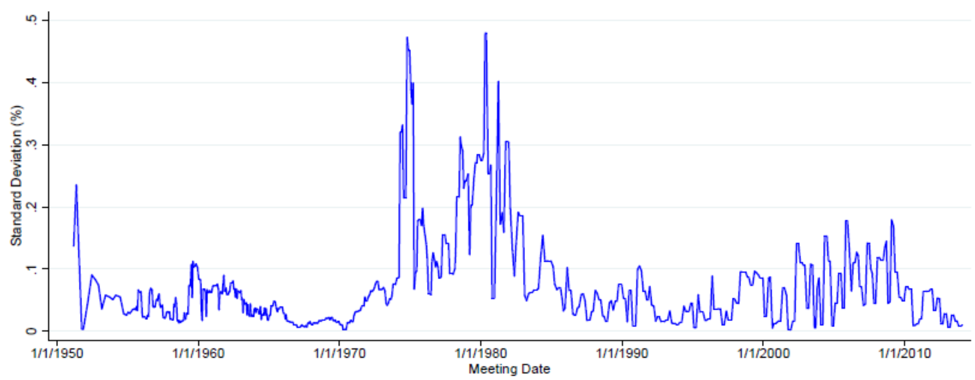

(b) Standard deviation of members' experience-based inflation forecasts

Figure 2.3: Dispersion of Experience-based Inflation Forecasts in Each FOMC Meeting

2.3.3 Econometric Specification

At each FOMC meeting, all current voting members cast a vote to either support or dissent from the proposal of the Fed chairperson. We classify the vote $V_{j,t}$ of member j in the meeting at time t as falling into one of three categories, $V_{j,t} \in \{-1, 0, 1\}$, for dovish dissent, no dissent, and hawkish dissent, respectively. We express the probability of being in one of these three categories as a function of the desired interest rate from equation (2.11) via the following ordered probit model: For $k \in \{-1, 0\}$,

$$\begin{aligned} &P(V_{j,t} \leq k | \pi^e_{j,t+1|t}, x_{j,t}, \pi_t, y_t) \\ &= \Phi[\delta_{k,j,t} - a_t - \lambda_0 \omega \pi^e_{j,t+1|t} - \kappa' x_{j,t} - \pi_t x'_{j,t} \lambda_1 - (y_t - y^*) x'_{j,t} \gamma_1], \end{aligned} \tag{2.12}$$

where $\Phi(.)$ denotes the standard normal cumulative distribution. We normalize $a_1 = 0$, and we suitably scale all variables so that the latent residual has unit standard deviation.[14] The main variable of interest in estimating equation (2.12) is the experience-based forecast $\pi^e_{j,t+1|t}$.

The model in equation (2.12) generalizes the ordered-probit model because we allow the dissent thresholds $\delta_{k,j,t}$ to vary with the characteristics of the FOMC member and over time, especially across the transparency regime change in 1993. The most important concern motivating this generalization is that regional Fed presidents may have different dissent thresholds than Federal Reserve Board governors. As we illustrated in Figure 2.2, this concern is particularly relevant since the November 1993 change in transparency. To accommodate the possibility of threshold-heterogeneity among FOMC members, we let the thresholds in equation (2.12) depend on the FOMC member characteristics $x_{j,t}$, including an interaction between indicators for the role of Fed

[14]These normalizations are of no consequence for the estimated partial effects, and so we do not explicitly write them out.

President and for a meeting time after November 1993:

$$\delta_{k,j,t} = \delta_{0,k} + \delta'_{1,k} x_{j,t} \qquad \text{for } k \in \{-1, 0\}.$$

Note that coefficients of $\delta_{0,k}$ and $\delta_{1,k}$ are threshold-specific. With this threshold specification, we obtain a version of the generalized ordered probit model in Williams (2006). We estimate the model with maximum likelihood. As a robustness check, we also explore conventional fixed-threshold ordered probit specifications in Section 2.3.6.

2.3.4 Hyperinflation Experiences

One FOMC member in our data set, Henry Wallich, personally experienced hyperinflation.[15] Wallich was born in Germany in 1914 in a family of bankers, and lived through Germany's hyperinflation from 1921 to 1924. In the 1930s, he emigrated to the United States. He was Federal Reserve governor from 1974 to 1986. Mr. Wallich dissented 27 times during his tenure on the Federal Reserve Board, the highest number of dissents among all FOMC members in Federal Reserve history, according to Thornton and Wheelock (2014).[16]

The presence of Wallich in our sample poses the question of how to include hyperinflation experiences into a parametric belief-updating scheme that is designed for (and works well in) a regime in which inflation rates are at most a few percent per quarter. How can we adjust it to properly describe expectation formation from data

[15] Henry Wallich is the only FOMC member with personal hyperinflation experiences that we could identify. H. Robert Heller, another German-born Federal Reserve Board member in the 1980s was born in 1940, after the hyperinflation. Stanley Fischer, who was born in Zambia in 1943, spent time in Israel, but not during its hyperinflation. He is not included in our sample because he started his tenure as vice chairman of the Federal Reserve Board in June 2014 while our sample ends in January 2014.

[16] In our sample, we identify only 26 dissents by Wallich, 24 of which were hawkish. The difference to Thornton and Wheelock's classification could be Wallich's vote on the 2/6/1979. In this meeting he dissented regarding the adopted growth rates of the monetary aggregates (M1-M3), but not regarding the open market transactions that were authorized. In our sample, this vote is not counted as dissent.

that include inflation rates around one million percent per quarter? Note that early life experiences are heavily downweighted in the calculation of the experience-based forecast, and it therefore makes virtually no difference whether we use inflation rates of the U.S. or another country, in which an individual might have grown up as a teenager, in low-inflation environments (with, say, single digit inflation rates). This is different with hyperinflation experiences. For example, if we naively plug German inflation rates from the 1920s into Wallich's experienced inflation history, the outliers are so big that three or four quarterly observations in 1923 would completely determine the autoregressive coefficients for the rest of Wallich's life. The post-1923 history would be rendered irrelevant, which is unlikely to be a plausible representation of how hyperinflation experiences influence inflation expectations.

We implement two approaches. First, we take a non-parametric approach and augment the inflation experience-based forecast (using U.S. data) with an indicator variable that we label "Wallich Dummy." With the caveat that this variable captures the voting behavior of just one individual member, the corresponding coefficient estimate provides at least tentative evidence on the effects of a "hyperinflation" treatment, i. e., how the extreme experience of hyperinflation may influence monetary policy views. Second, we also explore experience-based expectations formation with a mixed inflation process that includes a hyperinflation regime. This approach allows us to integrate hyperinflation experiences within one parametric framework with qualitatively similar results, but at the cost of additional complexity. We show the corresponding estimation results in Appendix C.4.

2.3.5 Baseline Results

Table 2.3 presents the estimates of our baseline ordered probit specification (2.12) using data from 1951 to 2014. Our focus is on the coefficient estimate, and the corresponding marginal effect, of each member's experience-based inflation forecast $\pi^e_{j,t+1|t}$. The chairman's vote is excluded from the sample because he never dissented during our sample period.

Column (i) of Table 2.3 reports estimates for a specification where the dissent thresholds can vary with indicators for the type of FOMC member (governor versus regional president) and with an indicator for the post-November 1993 period, as well as their interaction. This allows the model to accommodate the dramatic shift towards fewer dissents among Federal Reserve Board members after November 1993 that we saw in Figure 2.2. The coefficient on the experience-based inflation forecast of 216.6 (s.e. 66.1) is significantly different from zero at conventional significance levels. The magnitude of the effect on the probability of dissent can be inferred from the average partial effects (APE) reported in the middle block of the table. An increase of 0.1 percentage points (pp) in the experience-based forecasts of an FOMC member—which, according to Figure **??**, is a typical within-meeting standard deviation of FOMC members experience-based inflation forecasts during much of the sample—translates into an increase in the probability of a hawkish dissent vote of 1.21 pp, which is a little less than a third of the unconditional probability of hawkish dissent ($265/6707 \approx 4.0\%$). The probability of a dovish dissent drops by 0.76 pp, which is approximately a third of the unconditional probability of dovish dissent ($160/6707 \approx 2.4\%$). Thus, the estimates imply an economically large impact of inflation experiences on voting behavior.

The APE of the Wallich dummy indicates that the "hyperinflation treatment" is associated with a very large reduction in the probability of dovish dissent, 5 pp, and increase in the probability of hawkish dissent, 8 pp. In other words, the effects associated with the Wallich dummy are roughly of the same magnitude as those associated with a 1.0 pp increase in an FOMC member's experience-based inflation forecast.

All results are virtually identical in column (ii) where we allow the dissent thresholds to also depend on the FOMC members' individual characteristics (age, gender, party of president at appointment indicator, and same party as current president indicator).

2.3.6 Robustness Checks

One potential concern with the estimates in columns (i) and (ii) in Table 2.3 is that the inclusion of meeting fixed effects in the ordered probit model might introduce an incidental parameters problem.[17] To address this concern, we estimate an alternative specification in which we omit the meeting fixed effects. Instead, we specify that the probabilities of dissent are driven directly by cross-sectional differences (against the incumbent chairperson) in inflation experiences and other personal characteristics. That is, we forgo the non-parametric controls for the time-specific determinants of voting behavior, but still remove some of their effect to the extent that it is captured by the time-varying values associated with the chairperson.

The results are in columns (iii) and (iv) of Table 2.3. The coefficient estimates of the experience-effect forecast variable and the Wallich dummy decrease, but these changes largely reflect the altered econometric specification. As the APE calculations reveal, the implied economic magnitudes remain similar to those in columns (i) and (ii). Both sets of estimates also remain statistically significant. We conclude that our findings are not generated by estimator inconsistencies due to the incidental parameter problem.

As a second robustness check, we test whether we still find experience effects if we employ a simple ordered probit model with fixed dissent thresholds and restrict the analysis to subsamples in which the fixed-threshold assumption is more likely to hold, i. e., prior to the decrease in dissents in November 1993 and for the votes of regional

[17] As T increases, the number of meeting fixed effects grows at the same rate as T. As a consequence, the probit estimator is inconsistent and standard formulas for the asymptotic distribution of the estimator may not provide a good approximation of its finite-sample properties.

Table 2.3: Experience-based Inflation Forecasts and FOMC Voting Behavior

The sample period is March 8, 1951 to January 29, 2014. The experience-based inflation forecast for each member at each meeting is calculated by recursively estimating a mixed seasonal AR(1) model using the member's lifetime history of inflation, as described in Section 2.2.1 (with $\theta = 3.044$). The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. Column (i) and (iii) report the results assuming that the thresholds depend on a) whether the member is a board member or regional president, and b) whether the meeting occurs after Nov. 1993 and the interaction of a) and b). Column (ii) and (iv) report the results assuming that the thresholds depends, in addition, on age, gender, party of president at appointment indicator, and same party as current president indicator. In parentheses we report the standard error based on two-way clustering by both member and meeting.

	Ordered Probit		Ordered Probit "de-chaired"	
	(i)	(ii)	(iii)	(iv)
Experienced-Based Forecast	216.6	214.4	97.2	98.5
	(66.1)	(67.8)	(39.5)	(39.0)
Wallich Dummy	1.43	1.39	1.05	1.05
	(0.36)	(0.36)	(0.17)	(0.17)
Meeting FE	Yes	Yes	No	No
Thresholds	Role $\times$ $I_{>93}$	All	Role $\times$ $I_{>93}$	All
Observations	6,707	6,707	6,707	6,707
Pseudo R^2	39.0%	39.1%	9.7%	10.0%
APE of Experienced-Based Forecast:				
Dovish Dissent	-7.6	-7.6	-5.1	-5.1
Consent	-4.4	-4.3	-2.5	-2.5
Hawkish Dissent	12.1	11.9	7.6	7.7
APE of Wallich Dummy:				
Dovish Dissent	-0.050	-0.050	-0.055	-0.055
Consent	-0.029	-0.028	-0.027	-0.027
Hawkish Dissent	0.080	0.077	0.082	0.082

Table 2.4: Experience-based Inflation Forecasts and FOMC Voting Behavior: Different Sample Periods with Fixed Ordered Probit Thresholds

The experience-based inflation forecast for each member at each meeting is calculated as in Table 2.3. The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. Column (i) reports the results with all FOMC members prior to November 1993. Column (ii) reports the results with regional Fed presidents only prior to November 1993. Column (iii) reports the results with regional Fed presidents only over the entire sample. Column (iv) reports the results with all FOMC members prior to November 1993 and regional Fed presidents only afterwards. In parentheses we report the standard error based on two-way clustering by both member and meeting.

	All Members pre-1993 (i)	Regional Pres. Only Full Sample (ii)	Regional Pres. Only pre-1993 (iii)	Mixed Members Full Sample (iv)
Expr.-Based Fcst.	230.0	379.2	495.5	230.9
	(80.0)	(103.9)	(155.9)	(68.9)
Wallich Dummy	1.49	-	-	1.51
	(0.37)	-	-	(0.37)
Meeting FE	Yes	Yes	Yes	Yes
Observations	5,123	3,275	2,467	5,931
Pseudo R^2	38.0%	45.3%	49.2%	38.3%
APE of Expr.-Based Fcst.:				
Dovish Dissent	-9.5	- 6.4	-8.0	-9.0
Consent	-3.5	-19.5	-21.0	-5.2
Hawkish Dissent	13.0	26.0	29.0	14.2
APE of Wallich Dummy:				
Dovish Dissent	-0.062	-	-	-0.059
Consent	-0.022	-	-	-0.034
Hawkish Dissent	0.084	-	-	0.093

presidents.

Table 2.4 presents the results of this exercise. The specification in column (i) employs the voting records of all members from November 1993 onwards. The estimated results turn out to be very close to our benchmark case with characteristics-dependent dissent thresholds. We estimate slightly larger average partial effects of -9.5 pp for dovish dissents and $+13.0$ pp for hawkish dissents, again measured as the response to an increase of 1.0 pp in FOMC member's experience-based forecasts. The APE of the Wallich dummy also become slightly larger in both directions in this subsample.

In column (ii) we restrict the sample to regional Fed presidents, but use the full sample period. This subsample exploits the fact that the November 1993 transparency change did not have much effect on the voting behavior of regional presidents, as we showed in Figure 2.2. We find that the estimated effects are even stronger.[18] In this subsample, the proper comparison for the APEs is the unconditional probability of dovish or hawkish dissent by Federal Reserve presidents. The estimated average partial effects (APE) of changes in experience-based inflation forecast on the voting behavior of regional presidents suggests that an increase of 0.1% in the experience-based forecast of regional Fed presidents translates into an increase in the probability of a hawkish dissent by roughly 2.6 pp, which is a bit less than one half of the unconditional probability of a hawkish dissent by regional Fed presidents ($191/3275 \approx 5.8\%$). Meanwhile, the probability of a dovish dissent drops by 0.6 pp, which is roughly half of the unconditional probability of dovish dissent by regional Fed presidents ($38/3275 \approx 1.2\%$). Comparing these numbers to our baseline case with all FOMC members, it appears that past inflation experience has a stronger effect on the votes of regional Fed presidents.

In column (iii), we further restrict the sample of regional presidents to include only

[18]Since Henry Wallich is not a regional Fed president, we cannot estimate the Wallich dummy coefficient in this case.

the pre-November 1993 periods. The estimated APEs remain very similar.

Finally, in column (iv), we analyze the union of the column (i) and column (ii) sub-samples, i. e., all members pre-November 1993 and only Fed presidents post-November 1993. The estimated effects are very similar to those in column (i), as well as to the benchmark case.

Appendix C.5 contains an additional set of results with fixed thresholds where we use the full sample of all members and meetings. These results, shown in Table C.2, are again very similar. This simplified specification also allows a straightforward interpretation of the effects of the member characteristics, $x_{j,t}$. We report the coefficients associated with these variables in Table C.3.

As a last robustness check, we employ variations in the gain parameter θ of the learning algorithm. So far we fixed θ at the point estimate of 3.044 from Malmendier and Nagel (2016). Relying on a prior estimate has the advantage that we credibly tied our hands, rather than picking θ to fit the voting behavior of FOMC members. We now check how the fit and the estimated APE change if we vary θ. That is, we reestimate the learning rule for each FOMC member over a range of plausible values of θ. We then rerun the estimation from column (i) of Table 2.3 with the corresponding alternative experience-based forecasts of inflation.

For our first alternative value, we reestimate the gain parameter using MSC data based on the same procedure as in Malmendier and Nagel (2016), but with the sample restricted to college graduates. This sub-sample is more comparable to the FOMC members in terms of educational background. We estimate $\theta = 3.334$ (with s.e. of 0.347). That is, the θ estimate for college grads is less than one standard error from the full-sample estimate. As column (i) of Table 2.5 shows, employing $\theta = 3.334$ rather than $\theta = 3.044$ does not alter our findings. The results remain very similar to our baseline estimates in column (i) of Table 2.3.

Table 2.5: Experience-based Inflation Forecast and FOMC Voting Behavior: Varying Weights on Past Experience

The sample period is from March 8, 1951 to January 29, 2014. The ordered probit specification is the same as in column (i) of Table 2.3, but here with different values of the gain parameter θ in the calculation of the experience-based inflation forecast. The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. We assume that the ordered probit thresholds depend on a) whether the member is a board member or regional president, and b) whether the meeting occurs after Nov. 1993 and the interaction of a) and b). In parentheses we report the standard error based on two-way clustering by both member and meeting.

	$\theta = 3.334$ (i)	$\theta = 2$ (ii)	$\theta = 2.5$ (iii)	$\theta = 3.5$ (iv)	$\theta = 4$ (v)
Experience-Based Forecast	183.8	218.2	256.7	165.4	117.6
	(61.2)	(68.4)	(74.3)	(58.0)	(48.5)
Wallich Dummy	1.42	1.45	1.46	1.41	1.39
	(0.36)	(0.36)	(0.36)	(0.36)	(0.36)
Meeting FE	Yes	Yes	Yes	Yes	Yes
Observations	6,707	6,707	6,707	6,707	6,707
Pseudo R^2	38.9%	38.9%	39.1%	38.8%	38.6%
APE of Experienced-Based Forecast					
Dovish Dissent	-6.5	-7.7	-9.1	-5.9	-4.2
Consent	-3.8	-4.5	-5.2	-3.4	-2.4
Hawkish Dissent	10.3	12.2	14.3	9.2	6.6
APE of Wallich Dummy					
Dovish Dissent	-0.050	-0.051	-0.052	-0.058	-0.050
Consent	-0.029	-0.030	-0.030	-0.029	-0.029
Hawkish Dissent	0.079	0.081	0.081	0.079	0.078

Second, we employ a range of θ values between $\theta = 2$ to $\theta = 4$ (in steps of 0.5). As shown in columns (ii) to (v) of Table 2.5, all results are qualitatively similar to our baseline estimates as in column (i) of Table 2.3. We conclude that our results are robust to variations over a broad range of plausible θ values.

In summary, we find that lifetime inflation experiences have an economically large and robust effect on FOMC members' voting behavior. When an FOMC members' lifetime experience suggests higher inflation going forward than the experience of their peers, they are more likely to dissent in a hawkish direction. The opposite holds for inflation experiences suggesting lower future inflation; they induce dovish dissents.

2.4 Inflation Experiences and the Tone of FOMC Members' Speeches

The seeming reluctance of governors to dissent, especially since November 1993, indicates that FOMC members may not always fully reveal their disagreement in their voting behavior. They might voice their monetary policy views in discussions or speeches, but ultimately refrain from casting a dissenting vote.

In this section, we test whether FOMC members' attitude towards monetary policy can be detected in the language, or tone, they use in their speeches. To categorize language as hawkish or dovish, we employ an automated search-and-count approach that closely builds on the analysis of Apel and Grimaldi (2014). Apel and Grimaldi (2014) examine the Swedish Riksbank minutes and test whether the tone of an Executive Board member conveys a policy inclination toward loosening or tightening monetary policy. We apply their classification of tone to the speeches of FOMC members, with some adjustments to the different context and sample, as described in detail below.

Our data consists of all 6,353 "Speeches and Statements" available from the Federal

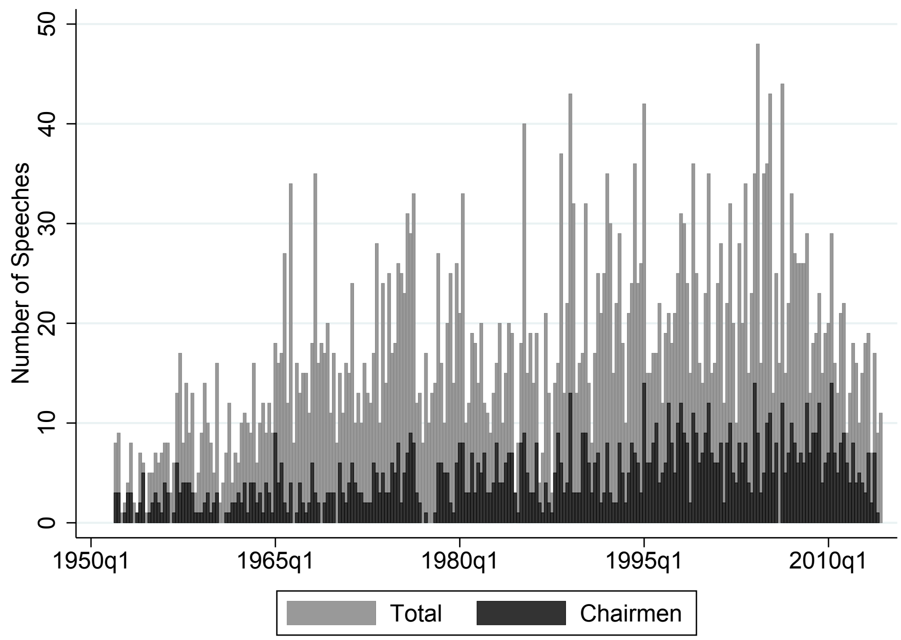

Figure 2.4: Number of FOMC Member Speeches Over Time

Reserve Archival System for Economic Research (FRASER), and additional 658 hand-collected speeches from the websites of the regional FRBs. To be consistent with the analysis of votes in the previous section, we focus on voting members and remove speeches delivered by the (rotating) non-voting regional Fed presidents. We also drop pdf files that could not be properly converted into text and for which the date of the speech cannot be determined. The final sample consists of 4,294 speeches for 86 FOMC members from the meeting on March 8th, 1951, to June 2014, with an average of 50 speeches per member. A quarter of the members have 15 or fewer speeches in the sample, while long-serving FOMC members, especially chairmen, tend to have more than 100 speeches. For example, our sample includes 482 speeches by Alan Greenspan and 264 by Ben Bernanke. Appendix C.6 details the construction of the data set.

Figure 2.4 shows the time series of the speeches in our sample. The total number increases over time. From 1965 onwards, the average number of speeches in a quarter is above 17, i.e., more than one speech per FOMC member per quarter. The share of speeches delivered by the chair increases only slightly over time and lies around 30%.

To classify the tone of these speeches, we follow Apel and Grimaldi (2014) and generate two-word combinations from two sets of words: nouns describing the *goals* of a central bank, and adjectives describing the *attitudes* of a central banker towards a goal. The list of goals in Apel and Grimaldi (2014) consists of "inflation," "cyclical position," "growth," "price," "wages," "oil price," and "development." In addition, we show estimation results after adapting the list to the FOMC context by adding "(un-)employment." Apel and Grimaldi had omitted this term because the Swedish Riksbank has price stability as a single goal, while the U.S. Federal Reserve System has a dual mandate. The list of attitudes consists of "decrease," "slow," "weak," and "low" on the dovish side, and "increase," "fast," "strong," and "high" for the hawkish counterpart. For unemployment, we swap the hawkish and the dovish adjectives.

For each mention of a *goal*, we check whether words from the *attitudes* list occur within a range (n-gram) of two words before and after the *goal*. While Apel and Grimaldi (2014) require the *attitude* word to appear directly before the *goal*, such two-word combinations do not generate sufficient variation between the speeches of FOMC members, possibly because the language is less formal and standardized than the Swedish central bank minutes, and the speeches of the FOMC members address a wider audience. We choose a range of two words before and after the goal (i.e., five-grams) in order to accommodate two-word goals such as "oil price," for which the *attitude* word is allowed to appear either one or two words before "oil" or one word after "price", as well as to accommodate different relative positions of the classification words. For example, an FOMC member might refer to "increasing prices" or mention that "prices are increasing." In addition, by centering the n-grams around the noun of interest, we avoid double-counting: Every word of the speech can occur in up to n

n-grams but is at most once in the center of an n-gram.

We drop n-grams containing more than one "goal" or "attitude" with different connotations. For example, the sequence "... low growth and unemployment ..." generates a five-gram centered around the *goal* 'growth' combined with the *attitude* 'low;' but the same five-gram also features another *goal*, unemployment. Since these two goals generate a dovish combination ("low growth") as well a hawkish one ("low unemployment"), we drop the five-gram from our analysis.

As in Apel and Grimaldi (2014), we then collapse the number of hawkish and dovish combinations in each speech into a single index:

$$\textit{Net Index} = \frac{\textit{Hawkish}}{\textit{Hawkish} + \textit{Dovish}} - \frac{\textit{Dovish}}{\textit{Hawkish} + \textit{Dovish}}. \tag{2.13}$$

The index ranges from -1 to $+1$, where -1 indicates that all of the tagged n-grams are dovish, and $+1$ that all tagged n-grams are hawkish. Hence, larger values of *Net Index* indicate greater hawkishness. If no hawkish or dovish n-grams can be found in the text, *Net Index* is set to zero.

Table 2.6 provides some summary statistics of *Net Index* and its components. On average, a speech contains 3,378 five-grams, but there is a large variation across speeches. A mean of 1.50 five-grams are tagged as hawkish, and 0.99 as dovish, when we use the original set of goals defined in Apel and Grimaldi (2014). By adding "employment/unemployment" to the goal list, we add an additional 0.29 hawkish and 0.22 dovish tags per speech. The average *Net Index* across speeches is about 0.10, irrespective of the specification of the goal list. The positive value indicates that the language used in our sample of speeches is slightly tilted towards a more hawkish wording, albeit with a large standard deviation of 0.55.

To develop our estimating equation, we assume that cross-sectional differences in

Table 2.6: Tone of Speeches: Summary Statistics

The sample includes voting FOMC members' speeches from March 1951 to June 2014. *Net Index* is an index of hawkishness calculated as described in equation (2.13). *Hawkish/Dovish Tags* is the average count of hawkish and dovish word combinations in a speech. *Hawkish/Dovish Tags for employment* counts the additional hawkish/dovish word combination per speech for the goal employment/unemployment.

	N	Mean	Std. Dev.	Min	Median	Max
5-grams per speech	4,294	3,378	2,098	10	3,058	23,891
Net Index excl. (un)empl.	4,294	0.10	0.55	-1	0	1
Net Index incl. (un)empl.	4,294	0.10	0.55	-1	0	1
Hawkish Tags excl. (un)empl.	4,294	1.50	3.05	0	0	68
Hawkish Tags for (un)empl.	4,294	0.29	0.85	0	0	16
Dovish Tags excl. (un)empl.	4,294	0.99	2.08	0	0	33
Dovish Tags for (un)empl.	4,294	0.22	0.72	0	0	12

Net Index between FOMC members map approximately linearly into differences in their desired interest rate according to equation (2.11). We obtain

$$\textit{Net Index}_{j,t} = \alpha_t + \beta_1 \pi^e_{j,t+1|t} + \beta_2' x_{j,t} + \pi_t x_{j,t}' \beta_3 + (y_t - y^*) x_{j,t}' \beta_4, \tag{2.14}$$

where the coefficients are multiples (by the same factor) of the corresponding coefficients in equation (2.11). As before in the voting analysis, we relate the outcome during quarter t to $\pi^e_{j,t+1|t}$, which is constructed based on the inflation history leading up to the end of quarter $t-1$. We also continue to focus on cross-sectional heterogeneity by employing time-fixed effects, α_t, to absorb common time-variation in the use of hawkish and dovish expressions.[19] The vector of member characteristics $x_{j,t}$ is the same as in the voting analysis (age, gender, party of president at appointment indicator, and same party as current president indicator), and it can influence the level of hawkishness as

[19] For example, in times of high unemployment, all FOMC members might be likely to employ the goal-attitude combination "high unemployment" in their five-grams.

well as the extent to which inflation or output gap increase or decrease hawkishness.

In addition, we also account for the fact that, differently from voting behavior, speech tone is likely subject to additional sources of heterogeneity. 'Speech style' and the choice of words can depend on other personal characteristics of the speaker, including education and prior professional experience. This heterogeneity adds noise and it could introduce correlated omitted variables. We use two approaches to account for these additional personal characteristics. First, we augment equation (2.14) with dummy variables that control for education and prior professional experience.[20] We generate indicator variables for having earned a PhD, a JD, an MBA, or a Master's degree as the highest degree. We also collect information on FOMC members' prior professional experience from the Fed's History Gateway and from the personal vitae of FOMC members. Using those sources, we generate indicator variables for prior experience in the financial industry, in non-finance industries, in other government organizations and agencies besides the Fed, and as an academic (i. e., having worked full-time in an academic department at some point prior to becoming an FOMC member). As a second approach to addressing heterogeneity in speech style, we absorb any time-invariant personal characteristics with member fixed effects. Under this approach, the coefficient of interest, β_1, is identified from within-member variation of speech tone as their inflation experience changes. The inclusion of member fixed effects is, on the one hand, most comprehensive in accounting for unobserved person-specific determinants of language use. On the other hand, it removes a substantial amount of variation coming from the differences in average experience-based inflation forecasts between FOMC members.

Table 2.7 presents the results. In columns (i) to (iii), we use the original *NetIndex* with the same list of goals as in Apel and Grimaldi (2014). In columns (iv) to (vi), we

[20]Details on the construction of both variables are at the end of Appendix C.6, including summary statistics in Appendix-Table C.4.

Table 2.7: Experience-based Inflation Forecasts and FOMC Members' Tone of Speeches

OLS regressions with the $NetIndex$ measure of speech hawkishness from equation (2.14) as the dependent variable. The experience-based inflation forecast for each member at each meeting is calculated as in Table 2.3. All estimations include the same controls and interactions with recent CPI inflation and unemployment as in Table 2.3. In addition, we include the controls for education and professional background detailed in the text, except for columns (3) and (6) where we instead employ member fixed effects. In columns (2) and (5), we drop speeches of chairmen. Standard errors, shown in parentheses, are calculated allowing for two-way clustering by FOMC member and year-quarter.

	Net Index excluding (un)empl.			Net Index including (un)empl.		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Experience-Based Fcst.	32.88	39.15	43.28	29.97	38.97	47.07
	(14.52)	(18.50)	(16.32)	(13.70)	(17.74)	(14.68)
Wallich Dummy	0.10	0.17	-	0.12	0.16	-
	(0.08)	(0.10)	-	(0.07)	(0.07)	-
Member FE	No	No	Yes	No	No	Yes
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Chair's speeches dropped	No	Yes	No	No	Yes	No
Industry expr. controls	Yes	Yes	No	Yes	Yes	No
Degree controls	Yes	Yes	No	Yes	Yes	No
Adjusted R^2	4.4%	4.7%	5.7%	3.9%	4.3%	5.1%
Observations	4294	3295	4294	4294	3295	4294

expand the index and add (un-)employment to the list of goals.

We estimate a significant effect of differences in inflation experiences on speech tone. In the baseline specification in column (i), the coefficient of 32.88 (s.e. 14.52) is significantly different from zero at the 5% level. An increase of 0.1 percentage points in the experience-based forecasts of an FOMC member—which is a typical within-meeting standard deviation—is associated with an increase of about 0.03 in the $NetIndex$, or about 1/16th of a standard deviation of $NetIndex$. This magnitude seems plausible for two reasons. First, the experience effects should be relatively subtle given the small age heterogeneity of FOMC members. Second, there is likely substantial measurement noise in $NetIndex$. This is apparent from the fact that the R^2 is only 4.4% despite the inclusion of time fixed effects, even though one would presumably expect substantial common time-variation in the *true* hawkishness of speeches.

The point estimate for the Wallich dummy suggests that hyperinflation experience predicts a 0.10 higher $NetIndex$ than that of other Fed governors with similar characteristics at the time; but given the standard error (0.08) it is not possible to rule out a zero effect at conventional significance levels in first specification. Nevertheless, it is noteworthy that the ratio of the point estimates for the experience-based forecasts and the Wallich dummy (about 200-300 here depending on the specification) is of the same order of magnitude as in the voting analysis in Table 2.3 (about 100-150).

In column (ii) we test the extent to which our estimation results are affected by the large number of speeches given by the chairperson. Speeches of the chair might systematically differ from the speeches of other FOMC member for at least two reasons. First, chairs might use a more balanced language for political reasons, especially given that they tend to attract more attention. Second, chairs might use the speeches to provide signals to financial markets, whereas the other FOMC member might primarily use the speeches to communicate their views between each other. When we drop the chair's speeches, we obtain a slightly larger coefficient of 39.15 (s.e. 18.50) which is also significant at the 5% level. In column (iii), we include both member fixed effects and speeches of the chair. The outcome remains almost unchanged.

In columns (iv) through (vi), we re-estimate the specifications from columns (i) through (iii) for the version of *Net Index* that includes (un-)employment as a goal. The results are very similar.

We conclude that the personal lifetime inflation experiences of FOMC members leave a significant imprint not only on their dissenting votes and the strong policy leanings expressed with those, but also on the more subtle expressions of attitudes towards monetary policy voiced in speeches.

2.5 Inflation Experiences and the Federal Funds Rate Target

Our analyses of cross-sectional differences in stated inflation expectations, voting decisions, and the tone of speeches all indicate that FOMC members rely, to a significant extent, on their own inflation experiences. We now test whether this partial reliance on personal experiences affects even the committee's ultimate decision about the Federal Funds target rate. That is, we test whether there is an incremental effect of FOMC members' experience-based inflation forecasts on the consensus decision, alongside conventional interest-rate determinants in a Taylor rule.

This last analysis has to overcome two additional difficulties. First, we aim to explain the time series of federal funds rates rather than cross-sectional differences in behavior. In the preceding analyses, we were able to identify the effects of inflation experiences from cross-sectional cohort-specific differences as well as from changes in those differences over time. Time dummies allowed us to absorb any potentially confounding time-series factors, including conventional determinants of monetary policy. Here, instead, we cannot absorb time-series factors but need to take a stand on a specific model of the time-series determinants of monetary policy decisions. We will focus on standard versions of the Taylor rule that have been proven successful in predicting the FOMC's federal funds rate policy in the recent empirical literature.

The second challenge is the limited data availability in the time-series dimension, relative to our earlier cross-sectional analyses. As we detail below, the need for output-gap forecast data and limitations of the forecast-based Taylor rule restrict our analysis to 1987Q3-2007Q2.

Because of these additional challenges, the time-series tests in this section should be viewed in conjunction with our earlier evidence from inflation forecasts, voting decisions, and the tone in speeches. The analysis in this section evaluates whether the federal funds rate moves over time in a way that is consistent with the evidence above.

In order to test whether we can detect the influence of FOMC members' personal experience in the fed funds rate target they set, we first have to aggregate the lifetime experiences of all members present at a given meeting, and hence their corresponding desired interest rates. We start from the linear approximation of the subjective Taylor rule in (2.11) that represents the desired federal funds rates of the individual FOMC members present at the meeting. In our baseline specification, we assume that the federal funds rate target decided at an FOMC meeting represents the average of the members' desired rate levels. (Alternatively, we use the median or the chairperson's desired rates instead; see Appendix C.8 for both robustness checks.) Averaging equation (2.11) across all FOMC members present at a meeting at time t, we obtain (as derived in Appendix C.2)

$$i_t^* = \beta_0 + \bar{z}_t + \beta_e \bar{\pi}^e_{t+1|t} + \beta_\pi \pi_t + \beta_y (y_t - y^*), \tag{2.15}$$

where $\bar{\pi}^e_{t+1|t}$ is the average of the FOMC members' experience-based inflation forecasts as of the meeting at time t, and $\bar{z}_t$ is the time-t average of

$$z_{j,t} = \kappa' x_{j,t} + \pi_t x'_{j,t} \lambda_1 + (y_t - y^*) x'_{j,t} \gamma_1. \tag{2.16}$$

With $\bar{z}_t = 0$ and $\beta_e = 0$ (the latter would follow from $\omega = 0$ in equation (2.11)), this reduces to the standard Taylor rule. Our earlier analyses suggest instead $\omega > 0$ and hence $\beta_e > 0$, i. e., that FOMC members rely to some extent on their experience-based inflation forecast, over and above the standard inflation- and output-gap components of the Taylor rule.

Turning to the empirical implementation, we aim to minimize the chance that $\bar{\pi}^e_{t+1|t}$ picks up the effects of measurement error in the objective macroeconomic information used by the FOMC. In order to do so, we need to use empirical measurements of π_t

and $(y_t - y^*)$ that are as close as possible to the information used by the FOMC. We do so in three steps. First, we build on Orphanides (Orphanides (2001), Orphanides (2003)), who shows that forecast-based variants of the Taylor rule provide a better empirical fit to the actual decisions about the federal funds rate target than a rule based on realized macroeconomic data. We follow Orphanides (2003) and replace, for every meeting in quarter t, π_t and $(y_t - y^*)$ with the Federal Reserve staff's Greenbook forecasts of inflation from quarter $t-1$ to $t+3$ and forecasts of the output gap in quarter $t+3$.[21] Second, we use the inflation index that the FOMC relies on primarily. Following Mehra and Sawhney (2010) and Bernanke (2010), we construct the time series of the staff's "core inflation forecast" from Greenbook forecasts of the core CPI inflation before the year 2000 and of the core PCE inflation thereafter. Third, we follow Coibion and Gorodnichenko (2012) and use one FOMC meeting per quarter (the one that is closest to the middle of the quarter). This ensures that the CPI information leading up to the end of the previous quarter, which is embedded in $\bar{\pi}^e_{t+1|t}$, is available to the FOMC. Moreover, obtaining data points that are almost equally spaced in time is useful when we include lagged interest rates.

We start the sample in 1987Q3 when the Federal Reserve's staff forecast of the output gap become available. As shown in Orphanides (2001), the Taylor rule, and its forecast-based variant in particular, then provides a good description of actual Federal Reserve policy. We end the sample in 2007Q2, just before the start of the financial crisis. Mishkin (2010) argues that starting in the summer of 2007, the FOMC reacted to information from financial markets that did not yet show up in inflation and output gap forecasts. As a result, the Taylor rule does not provide a good description of the

[21] In the earlier sample, the Greenbooks did not not explicitly include output gap forecasts, but the Board of Governors staff used them to construct wage and inflation forecasts. See www.philadelphiafed.org/research-and-data/real-time-center/greenbook-data/gap-and-financial-data-set.cfm for more details.

FOMC's policy during this period.[22]

Column (i) of Table 2.8 provides a benchmark for the analysis. We replicate the standard Taylor rule findings without $\bar{z}_t$ and $\bar{\pi}^e_{t+1|t}$. The estimated coefficients on the output gap (0.67) and on the inflation variable (1.51) are consistent with typical findings in the literature. In column (ii), we include the average experience-based forecast, $\bar{\pi}^e_{t+1|t}$. We estimate a coefficient of 0.38 (s.e. 0.21) that is significantly different from zero at a 10% level. Hence, FOMC members' average experience-based inflation forecast has explanatory power for the federal funds rate target over and above the staff forecast of inflation and the output gap, albeit only marginally significant in this specification. Considering the coefficients on the two inflation variables together, the weight on the experience-based forecast in our experience-augmented Taylor rule (2.15) is about $0.38/(1.27 + 0.38) \approx 0.23$.

Column (iii) turns to the full specification (2.15) by including $\bar{z}_t$, which captures the effect of the changing characteristics of the FOMC members on interest-rate decisions. Through equation (2.16), $\bar{z}_t$ depends on parameters that we cannot credibly estimate purely from time-variation in the federal funds rate target. For this reason, we construct $\bar{z}_t$ from the estimates in our voting analysis. The fitted values of the latent desired interest rate of our ordered probit model (2.12) allow us to construct $z_{j,t}$ in equation (2.16) up to scaling by a constant. More precisely, we use the ordered probit specification with fixed thresholds, shown in the robustness tables in the Appendix in Table C.2. (With characteristics-dependent thresholds, we would not be able to separate the effect of characteristics on the thresholds from the effect on the latent desired interest rate.) Averaging the fitted $z_{j,t}$ across FOMC members each period yields $\bar{z}_t$. After adding $\bar{z}_t$ to the Taylor rule as an explanatory variable in column (iii) of Table

[22]Baxa, Horváth, and Vašíček (2013) provide empirical evidence consistent with this description of FOMC policy. They show that adding financial market variables to the Taylor rule equation matters significantly in 2008-09, over and above inflation and output gap information.

Table 2.8: Influence of FOMC Members' Inflation Experiences on the Target Federal Funds Rate

The sample period is from the 8/18/1987 to 6/28/2007. The dependent variable is the target federal funds rate set at the FOMC meeting closest to the middle of the quarter t. The experience-based forecast is the average of FOMC members' experienced-based 4-quarter forecast of inflation based on CPI data leading up to the end of quarter $t-1$, calculated as in Table 2.3. The staff's core inflation forecast is from end of quarter $t-1$ to end of quarter $t+3$ based on the core CPI before 2/1/2000 and the core PCE thereafter. The staff's output gap forecast at quarter t is the forecast for quarter $t+3$. The staff's forecasts of CPI/PCE and of the output gap are from the Philadelphia Fed Greenbook data set. Lagged fed funds rate target is the federal fund funds rate target from the previous quarter's meeting. Columns (i) to (iii) report the OLS estimates based on (2.15). Columns (iv) and (v) report the estimates of β_e, β_π, β_y, ρ, and c from non-linear least-squares regressions as specified in (2.18). Columns (iii) and (v) include a proxy for $\bar{z}_t$, the linear combination of five FOMC-member characteristics and their interaction with inflation and unemployment estimated from voting data as reported in the Appendix in Table C.2. In parentheses, we report Newey-West standard errors with six lags from column (i) to (iii), and zero lags in column (iv) and (v).

	(i)	(ii)	(iii)	(iv)	(v)
Experience-based inflation forecast	-	0.38	0.61	0.46	0.44
	-	(0.21)	(0.24)	(0.21)	(0.21)
Staff's core inflation forecast	1.51	1.27	1.44	1.27	1.25
	(0.13)	(0.23)	(0.23)	(0.17)	(0.20)
Staff's output gap forecast	0.67	0.69	0.46	0.98	1.00
	(0.06)	(0.06)	(0.10)	(0.08)	(0.15)
Lagged federal funds rate target	-	-	-	0.68	0.69
	-	-	-	(0.04)	(0.04)
Intercept	0.80	0.11	2.17	-0.03	-0.08
	(0.44)	(0.36)	(0.86)	(0.16)	(0.42)
Member characteristics	N	N	Y	N	Y
Method	OLS	OLS	OLS	NLS	NLS
Observations	80	80	80	80	80
Adjusted R^2	85.8%	86.5%	87.7%	97.6%	97.6%

2.8, we find that the coefficient on the experience-based inflation forecast increases to 0.61 (s.e. 0.24), which is now statistically highly significant.

Finally, in columns (iv) to (v), we check whether the experience variable might be picking up the effect of a lagged federal funds rate. Existing evidence from the literature on monetary policy rules, e. g., Clarida, Galí, and Gertler (2000) and more recently Coibion and Gorodnichenko (2012), indicates that the Federal Reserve's policy is best characterized by partial adjustment, where the actual federal funds rate target i_t is a weighted average of the desired federal funds rate i_t^* from equation (2.15) and the lagged actual federal funds rate target i_{t-1},

$$i_t = (1-\rho)i_t^* + \rho i_{t-1}. \tag{2.17}$$

To check whether accounting for partial adjustment of this form changes the conclusions regarding the experience effects, we combine the partial adjustment rule with equation (2.15):

$$i_t = c + (1-\rho)\left[\bar{z}_t + \beta_e \bar{\pi}^e_{t+1|t} + \beta_\pi \pi_t + \beta_y (y_t - y^*)\right] + \rho i_{t-1}. \tag{2.18}$$

Since the parameter of interest, β_e, is now interacted with $1-\rho$, we estimate (2.18) with non-linear least squares. We report the estimates of β_e, β_π, β_y, ρ, and c in columns (iv) and (v) for the specification without and with the $\bar{z}_t$ variable, respectively.

Column (iv) presents the version without the $\bar{z}_t$ variable. Consistent with the existing literature on federal funds rate inertia, the lagged target rate has a strong predictive power and absorbs a large portion of the residual. The coefficients on the inflation variables are not affected much, though. The estimate of β_e of 0.46 (s.e. 0.21) is now a bit higher than in column (ii), and significantly different from zero at the 5% level. The implied weight on experienced inflation relative to the staff forecast is now $0.46/(1.27+0.46) \approx 0.27$. Turning to the estimation with the $\bar{z}_t$ variable included

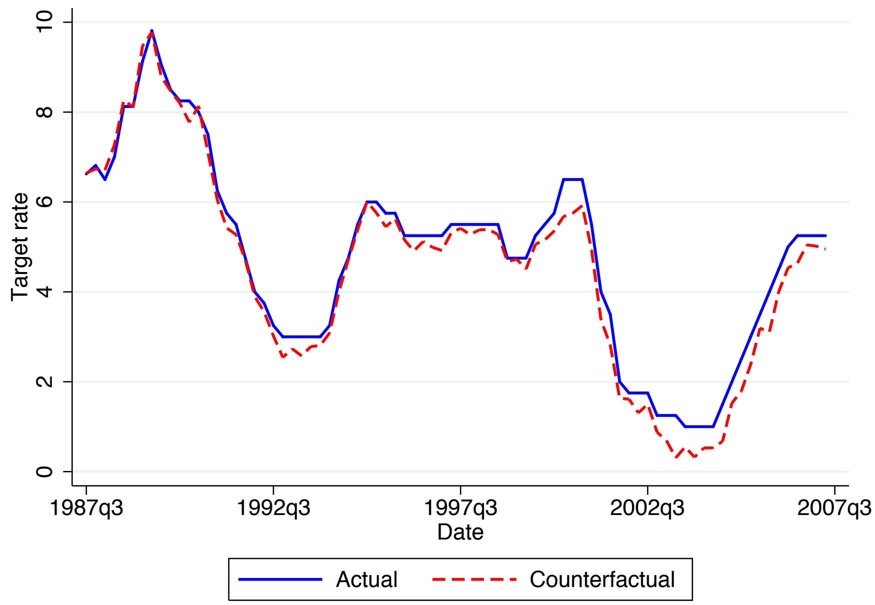

Figure 2.5: Counterfactual Federal Funds Rate Target (with experience effects removed)

in column (v), we find that adding $\bar{z}_t$ has very little effect on the estimates when the lagged federal funds rate target is included.

Overall, the evidence from the time-series of the target federal funds rate is consistent with the inflation experience effects that we identified in FOMC members' heterogeneous forecasts, voting decisions, and wording of speeches.

To assess the magnitude of this effect, we can compare these estimate to the those from the inflation forecast regressions in Table 2.1. There, we found that members put a weight of about 37-40% weight on their experience-based forecasts. It is reassuring that the weights obtained here, around 25%, are of very similar magnitude.

In Figure 2.5, we illustrate the magnitude of the effect by constructing a counterfactual federal funds rate target path that removes the estimated experience effects from the actual path. To construct the counterfactual path, we take the actual federal funds rate target and subtract the estimated β_e from column (ii) times the difference between FOMC members' average experience-based forecast and the Greenbook forecast of inflation. This counterfactual path represents the target that the FOMC would have chosen if its members had relied only on the staff forecast, not on their own inflation experiences—at least if we abstract from follow-on equilibrium effects.[23]

As the figure shows, the incremental effects of inflation experiences are substantial at times, but not unreasonably large. In the late 1980s and early 1990s, the effects were small. At the time, the average experience-based forecast remained very close to the staff's core inflation forecast. In contrast, in the 2000s the counterfactual federal funds rate target is often between 50 to 100 basis points lower than the actual federal funds rate.

2.6 Conclusion

We present novel evidence showing that personal lifetime experiences significantly affect the inflation forecasts, voting behavior, tone of speeches, and federal funds target rate decisions of FOMC members. Our findings suggest that heterogeneous inflation experiences generate heterogeneity in the desired policies and the macroeconomic outlook of FOMC members. Personal experiences exert this influence even though FOMC members are highly educated individuals and receive extensive decision-support from professional staff. In fact, experience effects help explain to a substantial extent why FOMC members deviate in their inflation forecasts from the forecasts prepared by Federal Reserve staff.

Our findings add to a growing literature on the role of experience-based heterogeneity in economic decisions and macroeconomic expectations. While existing studies focus on decisions and expectations of individual consumers and investors, this study

[23] If the FOMC had chosen a different target rate path, macroeconomic performance would presumably have been different. As a consequence, the inputs to the Taylor rule would have been different, which would in turn have affected the federal funds rate target. Our simple counterfactual analysis does not consider these equilibrium effects, but allows us to get a sense of the magnitude of the experience effects relative to the other drivers of the federal funds rate target.

is the first one to provide evidence of similar experience effects for policy makers.

The evidence in this paper also helps shed light on the behavioral origins of 'experience effects.' The overweighting of personal experiences by individual consumers documented in the earlier literature could perhaps be explained by informational frictions that restrict the availability of data they did not experience themselves. For sophisticated policy makers like the FOMC members in this study, such an explanation seems less plausible. Presumably, FOMC members are extensively exposed to historical macroeconomic data. Thus, there seems to be a deeper behavioral reason for why personal experiences get a relatively high weight in belief formation, even if historical information is easily accessible.

On the policy side, our results add a twist to the practical notion that the choice of a policy maker can have a long-lasting impact on policy outcomes: To predict a policy maker's leanings, it is helpful to look at the person's prior lifetime experiences. For a given outcome variable of interest, here inflation, we can calculate their weighted average experience with (roughly) linearly declining weights, and obtain a directional and quantitative prediction about their future decision-making. It will be interesting to explore in future research the extent to which such a model of experience-based learning is helpful in predicting policy makers' behavior in other policy areas.

CHAPTER III

Electronic Trading in the U.S. Corporate Bond Market

3.1 Introduction

Traditionally, trading in the U.S corporate bond market, a dealer-centric and over-the-counter market, has been primary done by voice where the investor contacts and negotiates with one dealer either over the phone or via electronic chat-room (such as Bloomberg). Over the last decade, however, the voice trading convention has gradually migrated to the use of electronic auction where the investor can now contact multiple dealers at the same time. According to Greenwich Associates, an industry consultancy, the market share of electronic trading (auction) in US investment-grade corporate bonds has grown from 5% in 2005 to about 20% as of 2015.[1] Despite the proliferation (and hence the growing importance) of electronic trading, we have little knowledge on how the addition of electronic trading would affect the corporate bond market overall. Specifically, how should an investor choose between voice and electronic platforms? Can both market structures coexist in the equilibrium? If so, what is the equilibrium market share of electronic trading? Further, would the introduction

[1] "The Continuing Corporate Bond Evolution" - Greenwich Associates (Q4 2015)

of electronic trading benefit all existing dealers?

This paper aims to shed some light on these questions through the lens of a model of strategic trading-platform selection by investors. In particular, I focus on the investor's choice between voice trading, which features a one-to-one negotiation between the investor and the dealer, and electronic auction, which allows the investor to contact multiple dealers simultaneously.

I start by motivating the theoretical model with two empirical observations of electronic trading in the US corporate bond market. On the one hand, I document that trading through the electronic auction has smaller transaction costs *ex-post*, comparing to the conventional voice trading. For instance, for investment-grade bonds, transaction cost on the electronic platform is about 30-70% of that through voice trading (depending on the trade sizes). On the other hand, however, I find that electronic trading entails a nontrivial execution risk *ex-ante*. On average, about one-quarter of the time, an investor receives zero response on her inquiry to the dealer community. Such trade delays are costly to investors. The execution delay costs are estimated to be about 8 basis points (0.08% of the notional value) for investment grade bond trades and about 19 basis points for high-yield ones, both of which are economically significant compared to their benchmark transaction costs.

Next, I develop a model of platform selection where an investor has to choose either voice or electronic trading.[2] In particular, I consider an environment in which the risk-averse investor and dealers share their endowment/inventory risks through trading. In voice trading, the investor negotiates with a dealer over both the price and quantity according to the standard Nash bargaining protocol. For electronic trading, I model it as a first-price-sealed-bid auction following the "Request-for-quote" (RFQ) protocol in

[2]In this paper, I only focus on customer-dealer trades and do not consider customer-customer or inter-dealer transactions because only customer-dealer trades take place on the electronic platform during the sample period.

practice.[3] This modeling choice builds on Biais (1993) but with one major difference motivated by the empirical observation: Dealers may not always participate in the electronic auction; instead, they decide whether to respond to an investor's electronic inquiry based on their inventory positions.

The key trade-off involved in this platform selection problem is price improvement versus execution risk of electronic trading. On the one hand, a potential price improvement comes from the observation that electronic auction, by allowing investors to contact multiple dealers simultaneously, increases the investor's probability of finding a dealer with a large opposite inventory position. Because holding (excess) inventory is costly,[4] dealers are more inclined to improve their quotes to induce orders from the public and thereby allow their inventory to revert to the desired level. Thus, trading through the electronic auction offers the investor a better chance to receive a more favorable price than (randomly) contacting a single dealer as in voice trading.[5]

On the other hand, however, trading through the electronic auction entails a non-trivial execution risk, which is costly to the investor. The risk of no execution is driven by dealers' positive participation costs. Given an investor buying inquiry, I show that only dealers with sufficiently large (positive) inventory positions would choose to participate in the auction because their expected gain-to-trade would then be large enough to compensate the participation cost. Therefore, the probability of investor receiving zero response from dealers is not negligible.

In the equilibrium, I confirm the intuition that investors tend to receive better prices through the electronic auction than through the voice market when dealer's inventory

[3]Although some other electronic trading protocols exist, according to Markets Committee (2016), RFQ remains the dominant protocol in the corporate bond market, accounting for more than 95% of total transaction volume.

[4]The inventory-holding cost exists due to dealer's limited risk-bearing capacity especially given the recently more stringent capital requirement on a risk-weighted basis.

[5]Note that reaching out to multiple dealers simultaneously is different from contacting them sequentially because of search costs. In the sequential contact case, a non-trivial search cost may well dampen the benefit of "shopping-around". More discussions on this are in Section 3.4.2.

position is large. Further, I show that investors with small relative bargaining powers would love to trade electronically because to them, the larger the bargaining power, the more to gain from the bilateral bargaining (voice trading) while, by definition, bargaining power does not matter for electronic trading. Relatedly, the equilibrium market share of electronic trading tends to decrease as i) dealer's (exogenous) inventory distribution becomes more concentrated around mean, and ii) dealers become less risk averse. When the dealer's inventory distribution becomes tighter, the investor's chance to trade with a dealer with a large opposite inventory becomes smaller and therefore, investor's benefit of trading electronically goes down. In the meantime, as dealers become less risk averse, they would be less worry about their own inventory risk and therefore, given the fixed participation cost for electronic auction, their participation threshold would thus be higher. Thus, the execution risk of electronic trading would be larger. Either way, the market share of electronic trading would decline accordingly.

In addition, as for dealer's welfare implication, I show that adding electronic auction on top of voice trading does not necessary improve each dealer's utility; dealers with relatively small balance sheet capacity would prefer the market structure with voice trading only. This is because when dealer's inventory position is not large enough, the probability to win the auction is essentially zero and therefore, dealer's expected gain from participating the auction would not compensate his relative loss in voice trading due to investors' self-selection.[6]

It is worth emphasizing that the key mechanism drives the results is dealer's inventory management motive. In the U.S. corporate bond market, inventory is a primary concern to dealers and hence, is of first-order importance in understanding the pricing and liquidity implications (see Friewald and Nagler (2016) and Dick-Nielsen and Rossi

[6] As mentioned earlier, with both voice and electronic trading, investors with large relative bargaining power would choose to trade in the voice market. Therefore, under this market structure with two platforms, the dealer's expected gain from voice trading would become smaller compared to the market structure with voice trading only. In that sense, dealers would incur a relative loss.

(2016)). This inventory management motive also complements the informational channel studied in the earlier literature on trading platform selection (see, e.g., Hendershott and Mendelson (2000) and Zhu (2014)).

The remaining of the paper proceeds as follows: Section 3.2 reviews the literature. Section 3.3 provides two motivating observations of electronic trading. Section 3.4 develops a model of strategic platform selection by investors and studies the equilibrium outcomes. Section 3.5 concludes. Appendix D contains additional figures. Appendix E provides proofs and related derivations.

3.2 Related Literature

This paper mainly connects to three streams of literature. First, it fits into the literature studying the role of dealer inventories in providing liquidity and their asset pricing implications. On the theory front, Ho and Stoll (1981, 1983) establish an optimal dynamic pricing rule for dealers to manage their inventory risks. Relatedly, Grossman and Miller (1988) model the role of dealers in providing the immediacy due to the temporary order imbalance. Empirically, Comerton-Forde et al. (2010) provide evidence that shocks to the market-maker's balance sheet and income statement affect daily stock market liquidity using a panel of New York Stock Exchange (NYSE) specialist inventory positions. More relatedly, equipped with transaction data in the US corporate bond market along with dealers' identities, Friewald and Nagler (2016) study the cross-sectional properties of dealer inventories and demonstrate that inventory considerations are important when it comes to the pricing of corporate bonds.

Second, this paper adds to the emerging literature that studies how relationship affects trading in the over-the-counter market. Theoretically, Seppi (1990) and Desgranges and Foucault (2005) study why dealers often offer price improvement to their clients and argue that the long-lasting investor-dealer relationship is the key for enforcing the no-informed trading agreement. More broadly, Duffie, Garleau, and Pedersen (2005, 2007) study the equilibrium allocation and price formation in the context of search-based friction and bargaining powers. Direct empirical evidence on the impact of trading relationship in the OTC market is relatively limited perhaps due to data limitations. Until very recently Hendershott et al. (2015) study the customer-dealer relationship in the US corporate bond market using data on insurer transactions and find that larger insurers trade with both large and small dealers and enjoy better execution while smaller insurers only trade with large dealers and suffer higher execution costs. As a complement to Hendershott et al. (2015), Maggio et al. (2016) focus on the relationship among dealers in the US corporate bond market and show that the inter-dealer market exhibits a clear core-periphery structure, with core dealers taking advantages of their connections to enjoy lower transaction costs at the expenses of smaller dealers and clients.

Third, this paper ties closely with the literature on comparing different market structures. For example, Biais (1993) and De Frutos and Manzano (2002) compare the centralized market, where quotes are visible to all market participants, and fragmented market, where the agents cannot observe the quotes of their competitors, with respect to equilibrium price and liquidity. Further, Yin (2005) extends Biais (1993) by taking the investor's search cost into consideration. More recently, Hendershott and Mendelson (2000) and Zhu (2014) study the impact of introducing a passive crossing network (dark pool) on top of an existing exchange (dealer market) with a focus on liquidity externalities (Hendershott and Mendelson (2000)) and price discovery (Zhu (2014)). While most existing studies in this stream either consider an anonymous market, where there is no room for trading relationship to play a role, or abstract away from the trading relationship, this paper explicitly studies the investor's platform selection problem, i.e., voice vs. electronic, with the consideration of bilateral relationship between investors and dealers.

Finally, perhaps the most related paper to this article is Hendershott and Madhavan (2015) (HM) who also compare the voice and electronic trading in the U.S corporate bond market. This article complements HM in one important way. In HM, they focus on what different bond characteristics matter for investor's platform choice and do not consider any heterogeneity among market participants. In this paper, however, I am particularly interested in understanding how heterogeneity among dealers (characterized by their different inventory positions) would affect investor's platform choice. My modeling approach allows me to tackle this question directly.

3.3 Motivating Evidence

The sample covers from January 2014 to December 2015. The electronic trading data are provided by MarketAxess, a leading electronic platform provider in the U.S corporate bond market.[7] It contains detailed information on each electronic inquiry submitted by investors, including the number of dealers quired, the inquiry size and direction, the quote from each responding dealer, an indicator of whether the inquiry leads to trade, and the transaction price if the trade occurs. One novelty of the MarketAxess data is that it provides the (anonymous) identities of bond investors. I further supplement the MarketAxess data with details on all customer-to-dealer corporate bond transactions, including both voice and electronic trades, from Trade Reporting and Compliance Engine (TRACE). Also, I obtain reference information on all corporate bonds, including ratings, coupon, maturity, and issue size, from Mergent FISD.

[7]According to Greenwich Associates (2015), MarketAxess is the dominate player in electronic trading of US investment-grade bonds with an estimated market share of 75%. The other players are Bloomberg FIT and Tradeweb with estimated market shares being 20% and 5%, respectively. For high-yield bonds, the market share of MarketAxess is even higher, amounting to nearly 90%. Therefore, it is reasonable to think of the data from MarketAxess is representative of overall electronic trading in the U.S corporate bond market.

Table 3.1: Summary Statistics

The sample period is from January 2014 to December 2015. Given that MarketAxess does not include inter-dealer trades, I restrict my attention to the customer-to-dealer transactions on TRACE accordingly.

	Electronic (MarketAxess)		Overall (TRACE)	
	Inv.-grade	High-yield	Inv.-grade	High-yield
Number of Bonds	10,812	3,808	15,677	16,751
Number of Trades	1,730,830	377,670	7,261,147	3,227,161
Client Buy	744,220	163,245	4,080,358	1,745,155
Client Sell	986,610	214,425	3,180,789	1,482,006
Average Daily Volume (in billions)	1.46	0.19	4.90	0.82
Client Buy	0.62	0.09	2.59	0.43
Client Sell	0.84	0.10	2.31	0.39
Average Trade Size (in thousands)	425	269	349	149
Client Buy	418	303	327	143
Client Sell	430	243	379	158

Table B.3 presents the summary statistics. During the sample period, around $16K$ ($17K$) different investment grade (high-yield) bonds were traded, among which about $11K$ ($4K$) bonds were covered by MarketAxess. In terms of number of trades, MarketAxess accounts for about 24% (12%) of all investment grade (high-yield) bonds transactions. As for trading volume, MarketAxess contributes roughly 30% (23%) of total volume of investment grade (high-yield) bonds.[8] Further breaking trades according to their directions, I find that the client selling volume is larger than the buying volume on MarketAxess, in contrast to the case found in the TRACE sample. Finally, as far as trade size is concerned, trades on MarketAxess appear to be larger than their counterparts in TRACE regardless of a bond's rating (either investment-grade or high-yield).

[8]Note that the trade size on TRACE is masked for so-called "mega" trades, i.e. \$5MM for IG and \$1MM for HY bonds. Thus, for the ease of comparison, I report the average trade volume only for those "non-Mega" trades on the MarketAxess. The same rule applies to the average trade size as well.

3.3.1 Transaction Cost: Voice vs. Electronic Trading

One natural way to compare different trading platforms is to look at its transaction cost. Measuring transaction cost in the OTC market is not trivial, though. Given the decentralized nature of corporate bond trading, high-quality intraday bid-ask quotes generally do not exist. Therefore, some commonly-used trading cost measures in the equity market, like effective bid-ask spreads, i.e. the difference between transaction price and mid-point of the quoted spread (see, e.g., Huang and Stoll (1996) and Hasbrouck (2009)), can not be directly applied in the corporate bond market. To address this problem, I instead define the transaction cost as the percentage difference between the trade price and the benchmark following Hendershott and Madhavan (2015), i.e.

$$\text{Transaction Cost} = \frac{\text{TradePrice} - \text{Benchmark}}{\text{Benchmark}} \times \text{Trade Direction}$$

where Trade Direction equals 1 if the investor is buying and -1 otherwise. Intuitively, the transaction cost defined above captures the premium that investors pay for the intermediation. The challenge here is how to find an appropriate benchmark empirically. Hendershott and Madhavan (2015) use the last interdealer price as the benchmark throughout. While this approach is straightforward, it has drawbacks. The main concern is that dealers may unwind the position obtained from their last customer trade with either other dealer(s) or customer(s) shortly after.[9] For these so-called "riskless principal trades",[10] the last interdealer price should not be considered as an appropriate benchmark because dealers tend to quote these inquiries based on the offsetting legs that show up *later* in the dataset instead of based on what happened *before*. This concern is more relevant when the last interdealer trade occurred days (or even weeks) before the current inquiry. To address this concern, for trades that I can identify

[9] About 20% of the trades fall into this category in the sample.

[10] More detailed discussions on riskless principal trades can be found in Harris (2015)

Table 3.2: Transaction Cost: MarketAxess vs. TRACE

The sample period is from January 2014 to December 2015. The calculation of transaction cost only involves "non-mega" trades as defined in the text. All numbers are in basis points.

	Electronic (MarketAxess)		Overall (TRACE)	
	Inv.-grade	High-yield	Inv.-grade	High-yield
All Trades				
Micro (1-100K)	13.1	26.5	46.9	80.8
Odd (100K-1MM)	13.3	15.8	19.2	30.9
Round (1-5MM)	10.6	-	9.6	-
Client Buy				
Micro (1-100K)	14.0	21.0	58.9	105.7
Odd (100K-1MM)	14.3	12.5	22.1	39.0
Round (1-5MM)	10.5	-	7.2	-
Client Sell				
Micro (1-100K)	12.5	29.1	30.8	49.6
Odd (100K-1MM)	12.4	19.0	15.4	21.6
Round (1-5MM)	10.7	-	12.3	-

their immediate offsets, I instead use the price on the offsetting leg to construct the benchmark.[11] Otherwise, I follow Hendershott and Madhavan (2015) and take the last interdealer price as the benchmark.

Table 3.2 presents the transaction costs by various trade-size bucket according to the industry convention.[12] Not surprisingly, transaction costs of investment grade bonds are smaller than that of high yield ones regardless of the trade-size bucket. For example, for an odd-lot investment grade bond trade on MarketAxess, the average transaction cost is about 13 basis points, or 0.13% of the notional value, while the cost of an odd-lot high-yield trade on MarketAxess is 16 basis points. Such transaction cost

[11] I treat the trade which occurred right after the original transaction (within the same day) with the same quantity and opposite direction as the offsetting leg. As far as how the benchmark is constructed, I directly take the trade price of the offsetting leg as the benchmark if the offsetting leg is an inter-dealer trade. Otherwise, I compute the benchmark as the average price of the pair of two offsetting trades.

[12] Given that the trade size for "mega" trades are masked on TRACE, it is difficult to find the exact match between MarketAxess and TRACE for those "mega" trades, and hence, obtaining a clean benchmark price becomes problematic. Therefore, I only focus on the "non-mega" trades for both MarketAxess and TRACE.

differences between investment grade bond and high yield ones are more prominent in the overall TRACE sample.

Comparing transaction costs between MarketAxess and TRACE, we can easily see that trades on MarketAxess tend to have more favorable prices to investors. Such price improvements on MarketAxess is more pronounced for micro trades and for high-yield bonds. However, it is important to keep in mind that a smaller *ex-post* transaction cost does not necessarily imply trading on MarketAxess is *ex-ante* cheaper. Indeed, as I will show in the next section, trading on MarketAxess involves nontrivial execution risks and such trade delays are costly to investors.

3.3.2 Execution Delay Cost

The risk of no execution on MarketAxess is not negligible.[13] At the bond level, about 10% of all inquiries on investment grade bonds receive no response from dealers while the zero-response rate for high-yield inquiries is even higher, which is about 39%. From the investor's perspective, on average, about 27% of the time, an investor receives zero response on her inquiry. Looking at a sub-sample of major investors, I find that the zero-response rate for the largest 100 investors is still as high as 16%.[14]

Execution risk matters for investors because trade delays are costly to them. Conceptually, for example, trade delay is costly to liquidity traders because they can not fulfill their liquidity needs immediately. To empirically estimate such execution delay cost on the electronic platform, I look at the difference in the transaction costs between repeated inquiries and normal ones. By repeated inquiries, I refer to those

[13] Unfortunately, direct evidence about the execution risk in the voice market is hard to obtain since data on those voice inquiries, such as phone tapes or Bloomberg chats, is not available to the public. But based on my conversations with corporate bond traders, at least anecdotally, practitioners believe that the execution risk of voice trading is much smaller than that of electronic market perhaps due to the reputation concerns of dealers.

[14] Among 800 investors in the sample, the largest 100 investors constitute about 85% of total trading volume.

electronic inquiries that initially received zero response but were sent later again and got filled electronically.[15] To identify such repeated inquiries, I restrict the initially unfilled inquiries and the later-filled ones to be on the same bond, with the same buy/sell direction and the same size, from the same investor, submitted within seven business days. By doing so, I am able to find $25,520$ such repeated inquiries out of $391,671$ zero-response ones.

Table 3.3 presents the results. As column (1) and (3) suggest, given an investor, the transaction cost of a repeated inquiry is significantly higher than that of a normal one. For investment grade bonds, the difference is about 8 basis points (0.08% of the notional value) while for high-yield ones, such difference is higher, which is about 19 basis points. Note that the differences in transaction costs between a repeated inquiry and a normal one, i.e., the execution delay cost, are economically large. As can be seen in Table 3.2, the average transaction cost for an investment grade (high-yield) bond trade is about 10 to 15 (15 to 30) basis points on MarketAxess.

Taken together, the empirical evidence suggests that, on average, trading on the electronic platform (MarketAxess) enjoys a smaller transaction cost *ex-post* but, *ex-ante*, it bears a nontrivial execution delay cost.

3.4 Model: Voice vs. Electronic Trading

This section presents a static model of trading platform selection between voice and electronic trading faced by investors, featuring the trade-off between potential price improvement and execution delay cost as observed in the data.

[15]Note that this definition of "repeated inquiries" is somewhat conservative. One may also think of a repeated inquiry as the one that was later executed in the voice market instead of in the electronic market again. However, if I were to apply this broader definition, any difference in the transaction cost between the normal inquiries (which are filled electronically) and the repeated ones (which are later filled in the voice market) could be partly due to the (perhaps unobserved) difference between voice and electronic platform.

Table 3.3: Execution Delay Cost

The sample period is from 1/1/2014 to 12/31/2015. The dependent variable is the transaction cost, defined as $\frac{\text{TradePrice - Benchmark}}{\text{Benchmark}} \times$Trade Direction, where Trade Direction equals 1 if the client is buying and -1 otherwise. $\mathbb{1}_{\text{Repeated inquiry}}$ is an indicator variable for a repeat inquiry as defined in the text. Client Buy equals one if the investor is buying from the dealer and zero otherwise. The order size is in million dollars. The bond-specific controls include time to maturity, time since initial offering, rating, and offering amount. In the parenthesis, I report the standard errors clustered at both bond and week level.

	Inv.-grade		High-yield	
	(1)	(2)	(3)	(4)
$\mathbb{1}_{\text{Repeated inquiry}}$	0.083	0.078	0.194	0.174
	(0.014)	(0.015)	(0.028)	(0.027)
Order size	-0.011	-0.011	-0.036	-0.100
	(0.002)	(0.002)	(0.008)	(0.013)
Client Buy	0.013	0.010	-0.093	-0.092
	(0.029)	(0.029)	(0.024)	(0.023)
Investor FE	Yes	No	Yes	No
Week FE	Yes	Yes	Yes	Yes
Bond-specific controls	Yes	Yes	Yes	Yes
Observations	1,432,787	1,432,787	239,695	239,695
Adjusted R^2	0.020	0.017	0.035	0.019

3.4.1 Setting

I consider the market of one risky asset with a random payoff v whose mean is μ and variance is σ^2. There are two types of agents participating in this market, i.e. risk-averse investors and dealers. Both investors and dealers are endowed with the mean-variance utility function with risk aversion parameter θ and γ, respectively. The investor receives an endowment shock e, which can be either a long position $+L$ or a short position $-L$ of the risky asset. Each dealer is endowed with some inventory I_j, following an exogenous distribution $F(\cdot)$. Further, the inventory position I_j is bounded between $-R$ and R, suggesting dealers can either long or short the risky asset.

The investor chooses either voice or electronic trading to share her endowment risk. In voice trading, the investor randomly contacts one of the dealers and then bargains with him over the price p and the quantity y according to standard Nash bargaining protocol. For electronic trading, the investor sends her inquiry to N dealers simultaneously.[16] Upon receiving the inquiry, each dealer first decides whether to respond. Once participating, dealers have to quote the entire size of the inquiry (rather than part of it). The dealer with the best quote wins the trade.

3.4.2 Assumptions

Before describing the model, I first discuss the key assumptions made throughout the analysis.

First, to focus on the effect of risk-sharing motive, I assume that there is no information asymmetry between the investor and dealers, i.e., both μ and σ^2 are known to each market participant. The trade, if any, is thus only motivated by risk-sharing between the investor and dealers.

Second, consistent with the real-world scenario, I assume that each dealer only

[16] For simplicity, I assume N is exogenously given.

knows his own inventory position I_j but not everyone else and the investor is not aware of the inventory position of any dealer. Yet, the inventory distribution, $F(\cdot)$, is common knowledge to both the investor and dealers.

Third, I assume that the investor would never contact more than one dealer (sequentially) in the voice market because the search cost is so high that it prevents her from contacting a second dealer (see Yin (2005)). Given that the electronic auction involves zero search cost, this assumption allows me to focus on the key trade-off of electronic trading, i.e., zero search cost vs. execution risk.[17]

Finally, for the electronic auction, I assume an exogenous reservation price $\bar{p}$ for the investor to prevent responding dealers from charging an arbitrarily high price.[18] I further assume an exogenous participation cost for dealers, $C > 0$. As we will see later, this assumption is critical since it drives dealers' endogenous participation strategy; each dealer responds if and only if his expected gain-to-trade is strictly greater than the participation cost. We may think of the participation cost as dealer's opportunity cost since most of the electronic inquiries are still manually priced by sell-side traders nowadays. Alternatively, we may interpret the participation cost as dealer's reputation cost. For many dealers, they would rather choose to pass the inquiry instead of quoting a non-competitive price, which may impair their relationship with clients (investors).[19]

[17]By reducing search frictions, the electronic auction allows the investor to contact multiple dealers simultaneously and thus, increases the investor's chance to trade with a dealer with a large opposite inventory position. As a result, the price that the investor would receive through the electronic auction tends to be more favorable than the one through the voice market where the investor (randomly) contacts and trades with one dealer.

[18]As shown in Figure D.1 in the appendix, any reasonable choice of $\bar{p}$ would not have a significant impact on the equilibrium relationship between different variables of interests

[19]Note that such non-competitive prices can be still away from investor's reservation prices.

3.4.3 Equilibrium Outcomes

3.4.3.1 Voice Trading

I first characterize the equilibrium in voice trading. Following Nash bargaining, both the investor and the (contacted) dealer choose price p and quantity y to maximize the product of their respective gain to trade weighted by their relative bargaining powers, i.e.

$$\max_{y,\,p} \left[U_{\text{gain}}^{MM}\right]^{\eta}\left[U_{\text{gain}}^{Inv}\right]^{1-\eta}$$

where,

$$\begin{aligned} U_{\text{gain}}^{MM} &= y(p-\mu) - \frac{1}{2}\gamma\sigma^2 y(y-2I) \\ U_{\text{gain}}^{Inv} &= y(\mu-p) - \frac{1}{2}\theta\sigma^2 y(y+2e) \end{aligned}$$

Solving for the above bargaining problem, the equilibrium trade size, y_{voice}, is given as follows (see details in Appendix E)

$$y_{\text{voice}} = \frac{\gamma}{\gamma+\theta}I - \frac{\theta}{\gamma+\theta}e \tag{3.1}$$

That is, through trading, the investor bears $\frac{\gamma}{\gamma+\theta}$ share of dealer's inventory risk and, in the meantime, offloads $\frac{\theta}{\gamma+\theta}$ share of her own endowment risk to the dealer. After trading, the investor ends up with $\frac{\gamma}{\gamma+\theta}$ share of the total risk, $I+e$, while the dealer bears the rest. Effectively, the investor and the dealer are thus sharing their inventory/endowment risk in proportion to their relative risk aversion.

Further, the equilibrium price, p_{voice}, can be shown as (also see Appendix E)

$$p_{\text{voice}} = \mu - \frac{1}{2}\Big[2 - \frac{\gamma}{\gamma+\theta} - \eta\Big]\gamma\sigma^2 I - \frac{1}{2}\Big[\frac{\gamma}{\gamma+\theta} + \eta\Big]\theta\sigma^2 e \tag{3.2}$$

Clearly, price goes down as the inventory level I becomes bigger since the coefficient of I is negative (assuming I is positive). The intuition is that given the mean-variance utility function, holding inventory is costly to dealers (which can be thought as a convex holding cost). Therefore, the larger amount of inventory the dealer holds, the more the dealer would like to offload by settling at a lower price. Further notice that the bigger the dealer's bargaining power η, the worse price the investor receives.[20]

Finally, given the equilibrium price p and trade size y, it is indeed incentive-compatible for both the investor and the dealer to trade since both parties' gain-to-trade are positive.[21] Put it differently, there is *no* execution risk in the voice trading.

3.4.3.2 Electronic Auction

The equilibrium in the electronic auction consists of dealers' participation and pricing strategies and the investor's trading strategy. By symmetry, I only consider the case where the investor receives a negative endowment shock $e = -L$. Assuming the

[20]This can be seen by re-arranging Eq. (3.2) as follows

$$p_{\text{voice}} = \mu - \frac{1}{2}[2 - \frac{\gamma}{\gamma+\theta}]\gamma\sigma^2 I - \frac{1}{2}[\frac{\gamma}{\gamma+\theta}]\theta\sigma^2 e + \frac{1}{2}[\gamma I - \theta e]\sigma^2\eta$$

Thus, as investor is buying ($y > 0$ or equivalently $\gamma I - \theta e > 0$), the price is increasing in η. As investor is selling ($y < 0$ or equivalently $\gamma I - \theta e < 0$), the price is decreasing in η. Taken together, the larger the η, the worse the price for the investor.

[21]Both investor and dealer's gain to trade are as follows

$$U^{\text{MM}}_{\text{gain}} = \frac{\eta\sigma^2}{2(\gamma+\theta)}\big(\gamma I - \theta e\big)^2 \geq 0$$
$$U^{\text{Inv}}_{\text{gain}} = \frac{(1-\eta)\sigma^2}{2(\gamma+\theta)}\big(\gamma I - \theta e\big)^2 \geq 0$$

investor wants to buy in this case,[22] I look for a symmetric equilibrium where a dealer would participate if and only if his inventory position is larger than some endogenous threshold.

Given the investor's inquiry size y and his own inventory position I_j, dealer's participation and pricing strategy are summarized as:

Proposition III.1. *Given that an investor receives a negative endowment shock, $e = -L$, each dealer chooses to participate in the electronic auction if and only if $I_j > \bar{I}^*$, where $\bar{I}^*$ uniquely solves*

$$F(\bar{I}^*)^{N-1} y \Big[\bar{p} - \tilde{p}(\bar{I}^*, y)\Big] = C \tag{3.3}$$

Once participating, dealer's optimal pricing strategy is given by

$$p(I_j, y) = \tilde{p}(I_j, y) + \gamma\sigma^2 \frac{\int_{\bar{I}^*}^{I_j} F(I)^{N-1} dI}{F(I_j)^{N-1}} + \left[\frac{F(\bar{I}^*)}{F(I_j)}\right]^{N-1} \Big[\bar{p} - \tilde{p}(\bar{I}^*, y)\Big] \tag{3.4}$$

where,

$$\tilde{p}(I_j, y) = \mu + \frac{1}{2}\gamma\sigma^2 (y - 2I_j)$$

The price is decreasing in the inventory position I_j and the number of dealers N. Further, dealer's expected trading surplus is increasing in his inventory position I_j.

The intuition of the dealer's participation strategy is as follows. Since the dealer's expected gain-to-trade increases in his inventory position, only dealers with sufficiently large inventory would love to participate (whose the expected gain would thus outweigh the exogenous participation cost). Further note that as the number of competing dealers becomes bigger (a larger N), the participation threshold $\bar{I}^*$ would increase

[22] I will confirm this conjecture after solving for the equilibrium

accordingly (see Eq. 3.3), and therefore, the participation probability for each dealer would decrease.

Next, I turn to describe dealer's pricing strategy, which consists of three components. The first component $\tilde{p}(I_j, y)$ is dealer's reservation price, which amounts to the sum of the fundamental value μ and the dealer's required risk premium $\frac{1}{2}\gamma\sigma^2(y - 2I_j)$. It turns out that the dealer's reservation price in the electronic auction equals the lowest possible price in the voice market given the same inquiry size y (in which case the dealer's relative bargaining power η is zero). The second component $\frac{\int_{\bar{I}^*}^{I_j} F(I)^{N-1} dI}{F(I_j)^{N-1}}\gamma\sigma^2$ is the standard first-price auction "adjustment term", which guarantees the seller (dealer) would obtain a strictly positive surplus in the case of winning the auction. This pure "first-price auction effect" is more pronounced when dealers are more risk averse (a larger γ) and the asset is riskier (a higher σ^2).

The key of dealer's pricing rule is the third component, $\left[\frac{F(\bar{I}^*)}{F(I_j)}\right]^{N-1}\left[\bar{p} - \tilde{p}(\bar{I}^*, y)\right]$, which can be interpreted as the threshold effect. When the dealer's inventory position I_j is higher than but still close to the threshold $\bar{I}^*$, $\left[\frac{F(\bar{I}^*)}{F(I_j)}\right]^{N-1}$ is close to one and thus, dealer's optimal quote remains close to the investor's reservation price $\bar{p}$. The intuition is that when the dealer's inventory position is not sufficiently high, the incentive for him to participate is low (as the chance to win remains close to zero and is not sensitive to the price change) and hence, he would quote a price close to the highest possible price he could charge, i.e., the investor's reservation price. As the inventory position gets larger and further away from the threshold, the dealer's incentive to participate increases and, as a result, the optimal quote becomes smaller.

Taken together, when deciding his optimal quote, the dealer starts from his reservation price $\tilde{p}$ and then raise the quote to capture the required surplus (the sum of the second and third element of the pricing rule). The more competing dealers, the greater chances that raising the quote would result in dealer missing the trade. Therefore, the surplus (as well as the quote) decreases in total number of dealers N. As N goes to infinity, the optimal quote converges to the reservation price, suggesting dealers can no longer earn any surplus when the number of competitors is infinitely large. Further notice that if the participation cost were zero, there would be no threshold anymore, i.e., $\bar{I}^* = -R$. Consequently, the pricing rule $p(I_j, y)$ would collapse to $\tilde{p}(I_j, y) + \gamma\sigma^2 \frac{\int_{-R}^{I_j} F(I)^{N-1} dI}{F(I_j)^{N-1}}$, which is the same as in Biais (1993).

Compared to dealer's strategies discussed above, the investor's strategy is relatively straightforward; she simply chooses an order size to maximize her expected gain to trade (by taking the execution risk into consideration). In the equilibrium, the dealer's endogenous participation threshold $\bar{I}^*$ and the investor's optimal order size y are jointly determined.[23] Unfortunately, neither the threshold nor the order size has a simple expression. As a consequence, the equilibrium price does not have a closed-form expression either, but I will illustrate the relation between the dealer's inventory position and the equilibrium price in the following numerical example.

3.4.3.3 A Numerical Example

Figure 3.1 illustrates how the price changes as the inventory position I varies in both voice and electronic trading. Here, I set the investor's reservation price $\bar{p}$ as 2μ and dealer's relative bargaining power η as 0.5.[24] For the ease of comparison, I plot the graph starting from $I = \bar{I}^*$.[25] As expected, both voice and electronic price decrease as dealer's inventory position gets larger. Yet, the electronic price tends to drop much faster than the voice one. This is mainly because as the inventory becomes further away from the threshold, the impact of the investor's reservation price on the electronic price decays dramatically, which is shown as the diminishing difference between the red solid

[23] Details on solving the equilibrium in the electronic auction are provided in Appendix E.

[24] As shown in Appendix figure D.2, the choice of $\bar{p}$ would not have a significant impact here.

[25] The price of voice trading exists for any inventory position while the price of electronic trading is only valid when $I > \bar{I}^*$.

line and red dash line in the graph. When the inventory position is relatively small (but still larger than the threshold), the electronic price can be higher than the voice one since the investor's reservation price still plays an important role in determining the electronic price in this case. As the inventory position gets large enough, however, the electronic price tends to be smaller than the voice one due to i) the competition among dealers and ii) the diminishing impact of investor's reservation price.

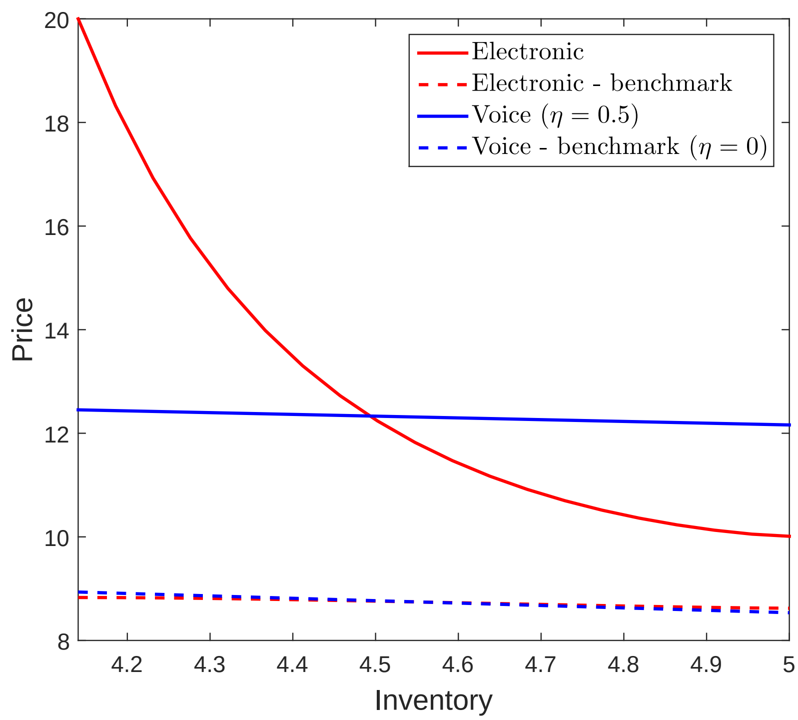

Figure 3.1: Dealer's Pricing Strategy: Voice vs. Electronic trading. The voice price (and voice benchmark) refers to Eq. (3.2). The electronic price corresponds to Eq. (3.4) with the electronic benchmark being the sum of first two components in the pricing rule, i.e. $\tilde{p}(I_j, y) + \gamma\sigma^2 \frac{\int_{I^*}^{I_j} F(I)^{N-1} dI}{F(I_j)^{N-1}}$. Model parameters: $\mu = 10$, $\gamma = 0.5$, $\theta = 3$, $\sigma^2 = 1$, $L = 4$, $R = 5$, $C = 5$, $N = 50$, and I has a scaled beta distribution with $a = b = 1.5$. Moreover, $\bar{p} = 2\mu$.

3.4.4 Equilibrium Market Share of Electronic Trading

So far, our analysis has focused on the equilibrium in voice and electronic trading separately for each investor. To study the aggregate effect, I now consider a continuum of investors characterized by their relative bargaining power over the dealer, i.e. $\delta \equiv 1-\eta \in [0,1]$, and each facing the choice between voice and electronic trading. Assume δ follows some exogenous distribution $H(\cdot)$. Further, to be consistent with the discussion above, I only consider the case where investors receive negative endowment shocks in this subsection.

To calculate the equilibrium market share of electronic trading, I first derive the investors' participation rate in the electronic auction. For each investor, the gain from voice trading, denoted as $V_{\text{voice}}(\delta)$, monotonically increases in her relative bargaining power δ while the gain from the electronic auction, denoted as V_{elec}, is independent of δ since bilateral negotiation never occurs in the electronic auction.[26] Therefore, if $V_{\text{voice}}(1) > V_{\text{elec}}$, there must exist a unique threshold δ^* such that $V_{\text{voice}}(\delta^*) = V_{\text{elec}}$. Thus, any investor whose bargaining power is smaller than δ^* would love to trade electronically and as a result, the investor's participation rate of electronic trading equals $H(\delta^*)$. Otherwise, if $V_{\text{voice}}(1) \leq V_{\text{elec}}$, the participation rate is simply 100%.

With the investors' participation rate, I now compute the equilibrium market share of electronic trading, i.e., the proportion of trading volume occurred through the electronic auction. The expected trading volume in voice and electronic trading can be written as follows, respectively,

$$\text{Vol}_{\text{voice}} = \Big(1 - H(\delta^*)\Big)\mathbb{E}(y_{\text{voice}})$$
$$\text{Vol}_{\text{elec}} = H(\delta^*)\Big(1 - F(\bar{I}^*)^N\Big)y_{\text{elec}}$$

[26]In fact, investor's gain from voice trading can be written as $\frac{\delta\sigma^2}{2(\gamma+\theta)}\mathbb{E}\big[(\gamma I + \theta L)^2\big]$.

where y_{elec} stands for the trade size in the electronic auction.

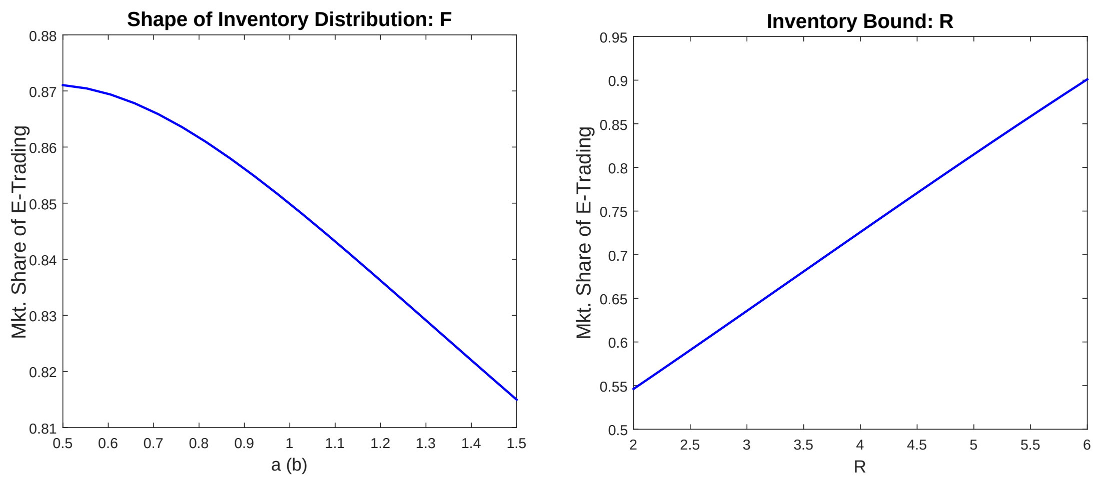

Figure 3.2: Market Share of E-trading with respect to the Inventory Distribution. Model parameters (unless otherwise stated): $\mu = 10$, $\gamma = 0.5$, $\theta = 3$, $\sigma^2 = 1$, $L = 4$, $R = 5$, $C = 5$, $N = 50$, δ has a uniform distribution between 0 and 1, and I has a scaled beta distribution with $a = b = 1.5$.

Figures 3.2 describes how the dealer's inventory distribution affects the equilibrium market share of electronic trading. The left-hand-side plot of Figure 3.2 focus on the shape of inventory distribution. In particular, I perturb the two parameters, a and b, of a scaled beta distribution, which, I believe, is flexible to capture different shapes of inventory positions across dealers. As I vary from $a = b = 0.5$ to $a = b = 1.5$, the inventory mass becomes more concentrated around the mean (see Figure D.1 in Appendix D). As a result, the chance that the investor would trade with a dealer with a large opposite inventory through the electronic auction becomes smaller. Therefore, the benefit for investors to trade electronically goes down, and the market share of electronic trading drops accordingly.

The right-hand-side plot of Figure 3.2 illustrates the effect of inventory bound R. When increasing the bound, one essentially shifts more inventory mass to the extreme. Therefore, similar intuition discussed above applies here as well; as R becomes larger, investors tend to benefit more from trading electronically since their chances to meet some dealer with a large opposite inventory position increase and consequently, the market share of electronic trading goes up.

The discussion above further generates an interesting empirical implication. That is, we can test whether dealer's inventory distribution matters for investor's trading-platform choice by running a cross-sectional regression of the (time-series) variance of dealer's inventory of each individual bond on its market share of electronic trading. A smaller variation in dealer's inventory over time corresponds to a more concentrated inventory distribution or a tighter bound R. Hence, my model suggests that, all else equal, bonds with smaller inventory variances have a lower market share of electronic trading compared to those with bigger variances.

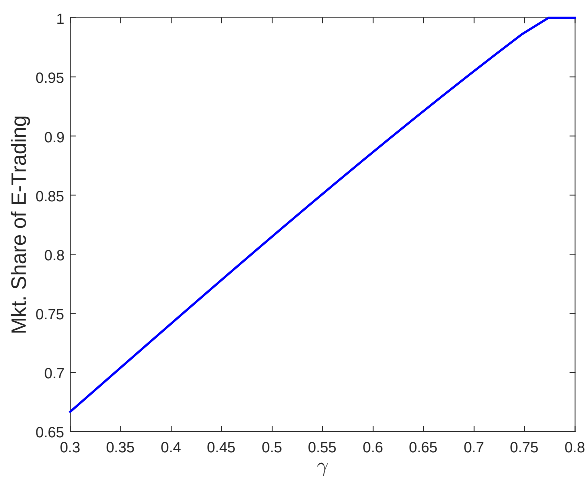

Figure 3.3: Market Share of E-trading with respect to Dealer's Risk Aversion γ. Model parameters: $\mu = 10$, $\theta = 3$, $\sigma^2 = 1$, $L = 4$, $R = 5$, $C = 5$, $N = 50$, δ has a uniform distribution between 0 and 1, and I has a scaled beta distribution with $a = b = 1.5$.

Figure 3.3 plots the equilibrium market share of electronic trading as a function of dealer's risk aversion γ. As dealers become more risk averse (a larger γ), their risk sharing motives become stronger and hence, their expected gain-to-trade increases.[27] Therefore, their participation threshold would be lower and thus result in a substantial

[27]Dealer's expected gain from the electronic auction can be written as $\gamma\sigma^2 y \int_{I^*}^{I_j} F(I)^{N-1} dI + C$

increase of trading volume through electronic trading. In the meantime, the volume through voice trading may only increase mildly or even decrease. Taken together, the market share of electronic trading would thus increase as dealer's risk aversion γ increases. Note that when γ is high enough, the investor's largest possible gain from voice trading, $V_{\text{voice}}(1)$, becomes smaller than the gain from electronic trading. In this case, investors would all choose to trade electronically and thus, the market share of electronic trading stays at 100% regardless of γ.

3.4.5 Welfare Implication for Dealers

I close the model section by discussing how the addition of electronic auction (on top of voice trading) would affect the welfare of each dealer.

With both voice and electronic trading available, as discussed above, investors with large relative bargaining powers would choose to trade by voice. Therefore, the dealer's expected gain from voice trading would be smaller in the "voice + electronic" case than that in the "voice only" case. Thus, the dealer's preference over the two market structures critically depends on whether such relative losses in voice trading can be compensated by the potential gain from electronic trading. If not, dealers would then prefer the market structure with voice trading only.

The left-hand-side plot of Figure 3.4 illustrates the dealer's expected surplus under two market structures, i.e., "voice only" and "voice + electronic", when the participation cost is strictly positive. First of all, if the dealer's inventory position is smaller than the threshold (which is not shown in the plot), he would never choose to participate in the electronic auction even if he receives an electronic inquiry. Therefore, in this case, the dealer would always prefer voice trading only. More interestingly, as shown in the plot, even if the inventory position is greater than the threshold, as long as it is not sufficiently large, the dealer would still prefer voice trading only. The intuition is that,

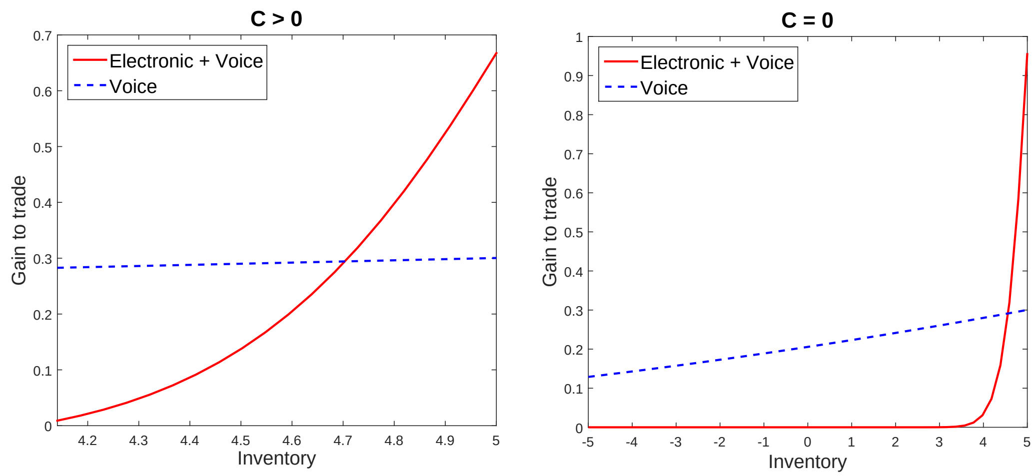

Figure 3.4: Dealer's Expected Gain under the Two Market Structures. Model parameters: $\mu = 10$, $\theta = 3$, $\sigma^2 = 1$, $L = 4$, $R = 5$, $C = 5$, $N = 50$, δ has a uniform distribution between 0 and 1, and I has a scaled beta distribution with $a = b = 1.5$.

with a close-to-threshold inventory position, the dealer's probability of winning the electronic auction remains close to zero. As a result, the dealer's expected gain from the electronic trading (which is close to zero) can not compensate his relative losses in the voice trading as discussed above. As the inventory position becomes sufficiently large, dealer's incentive to offload his inventory would be so strong that his chance to win the electronic auction would become much higher (by quoting a better price) and consequently, his expected gain from the electronic trading would be larger. Thus, dealers with large enough inventory position would prefer the market structure with both electronic and voice trading. The same conclusion holds even if the participation cost is zero, as shown in the right-hand-side plot of Figure 3.4.

3.5 Conclusion

Electronic trading in the U.S corporate bond market has become increasingly important over the past years. This paper studies the effect of adding the fast-growing electronic trading on top of the traditional voice trading, through the lens of a model of strategic trading-platform selection by investors. Empirically, I document that investors who trade electronically tend to have smaller transaction costs *ex-post* but, in the meantime, have to bear a nontrivial execution risk *ex-ante*, which is costly to them. Built on these empirical observations, I develop a theoretical model featuring the trade-off between price improvement and execution delay cost of the electronic auction.

Within the model, I find that investors tend to receive better prices through electronic auction than through voice trading when dealers' inventory position are large. Then, I show that electronic auction attracts investors with relatively small bargaining powers. Further, the equilibrium market share of electronic trading decreases as i) dealer's inventory distribution becomes more concentrated toward the mean and, ii) dealer becomes less risk-averse. Finally, for dealers' welfare implication, the model reveals that, through a mechanism of investors self-selection, adding electronic auction may not benefit every individual dealer; indeed, those with relatively small balance sheet capacity would prefer the market structure with voice trading only.

APPENDICES

APPENDIX A

Returns to Scale: Proofs

A.1 Proofs

Proof of Proposition I.1: The first-order-condition to the manager's problem (1.2) is

$$\phi_t(\beta+1)(A_t^*)^\beta - 2c_1 A_t^* + (\mu - c_0) = 0 \tag{A.1}$$

where $\phi_t \equiv \hat{\mathbb{E}}_t(\theta)$, i.e., investors' perception on θ at time t.

First, I show that if $\mu - c_0 \geq 0$, there exists a unique positive solution to Equation (A.1). Denote $f(A_t^*)$, a function of the equilibrium fund size A_t^*, as the left-hand side of Equation (A.1). When $\mu - c_0 = 0$, $A_t^* = \left(\frac{\phi_t(\beta+1)}{2c_1}\right)^{\frac{1}{1-\beta}} > 0$. When $\mu - c_0 > 0$, it is straightforward to see that (i) $\frac{\partial f(A_t^*)}{\partial A_t^*} < 0$, (ii) $f(0) > 0$, and (iii) $f(\infty) = -\infty$. Therefore, there exists a unique positive solution to equation (A.1).

Given that the second-order-condition to the manager's problem (1.2) is strictly negative, i.e.,

$$\phi_t\beta(\beta+1)(A_t^*)^{\beta-1} - 2c_1 = \frac{1}{A_t^*}\Big(-(1-\beta)2c_1 A_t^* - \beta(\mu - c_0)\Big) < 0$$

the unique solution to equation (A.1) is indeed the maximum as wanted. Plugging the equation (A.1) into $f_t^* = (\mu - c_0) - c_1 A_t^* + \phi_t (A_t^*)^\beta$, we have $f_t^* = \frac{1-\beta}{1+\beta} c_1 A_t^* + \frac{\beta}{1+\beta}(\mu - c_0)$, which is always strictly positive.

Next, I show that A_t^* increases in ϕ_t, i.e.,

$$\frac{\partial A_t^*}{\partial \phi_t} = -\frac{\frac{\partial f(\cdot)}{\partial \phi_t}}{\frac{\partial f(\cdot)}{\partial A_t^*}} = \frac{(\beta+1)(A_t^*)^{\beta}}{2c_1 - \phi_t\beta(\beta+1)(A_t^*)^{\beta-1}} = \frac{(\beta+1)(A_t^*)^{\beta+1}}{(1-\beta)2c_1A_t^* + \beta(\mu - c_0)} > 0 \qquad \text{(A.2)}$$

Proof of Negative Subjective Returns to Scale: The subjective returns to scale evaluated at the equilibrium fund size is given by,

$$\begin{aligned}
\frac{\partial \hat{\mathbb{E}}(\alpha_{t+1}^{Net})}{\partial A_t}|_{A_t=A_t^*} &= \phi_t\beta(A_t^*)^{\beta-1} - c_1 \\
&= \frac{\beta}{1+\beta}\Big(2c_1 - \frac{\mu - c_0}{A_t^*}\Big) - c_1 \\
&= -c_1\frac{1-\beta}{1+\beta} - \frac{\beta}{1+\beta}\frac{\mu - c_0}{A_t^*} < 0
\end{aligned}$$

Proof of Proposition I.2: First, I show that the expected fund net alpha first increases and then decreases in ϕ_t under the objective belief. The (objective) expected net alpha can be written as

$$\begin{aligned}
\mathbb{E}_t\big(\alpha_{t+1}^{Net}\big) &= \mu - (c_0 + c_1A_t - \theta A_t^{\beta}) - f_t^* \\
&= \theta(A_t^*)^{\beta} - \phi_t(A_t^*)^{\beta} \\
&= (\theta - \phi_t)(A_t^*)^{\beta}
\end{aligned}$$

The second equation follows because $\hat{\mathbb{E}}_t\big(\alpha_{t+1}^{Net}\big) = \mu - (c_0 + c_1A_t - \phi_tA_t^{\beta}) - f_t^* = 0$.

Thus, we can write:

$$\begin{aligned}
\frac{\partial \mathbb{E}(\alpha_{t+1}^{Net})}{\partial \phi_t} &= -(A_t^*)^{\beta} + (\theta - \phi_t)\beta(A_t^*)^{\beta-1}\frac{\partial A_t^*}{\partial \phi_t} \\
&= (A_t^*)^{\beta-1}\Big((\theta - \phi_t)\frac{\beta(\beta+1)(A_t^*)^{\beta}}{2c_1 - \phi_t\beta(\beta+1)(A_t^*)^{\beta-1}} - A_t^*\Big)
\end{aligned}$$

It follows:

$$\begin{aligned}
\frac{\partial \mathbb{E}(\alpha_{t+1}^{Net})}{\partial \phi_t} > 0 &\Leftrightarrow \theta > \frac{1}{\beta}\phi_t + \frac{\mu - c_0}{\beta(1+\beta)(A_t^*)^{\beta}} \\
&\Leftrightarrow \beta\theta > \underbrace{\phi_t\Big(\frac{2c_1A_t^*}{2c_1A_t^* - (\mu - c_0)}\Big)}_{h(\phi_t)}
\end{aligned}$$

Now, I prove $h(\phi_t)$ is invertible by showing that $h(\phi_t)$ is strictly increasing in ϕ_t:

$$\begin{aligned}\frac{\partial h(\phi_t)}{\partial \phi_t} &= \frac{2c_1 A_t^*}{2c_1 A_t^* - (\mu - c_0)} + \phi_t \frac{\partial\left(\frac{2c_1 A_t^*}{2c_1 A_t^* - (\mu - c_0)}\right)}{(\partial \phi_t)} \\ &= \frac{\phi_t (1+\beta)(A_t^*)^{\beta+1}}{2c_1 A_t^*(1-\beta) + (\mu - c_0)\beta} \frac{(1-\beta)\phi_t(1+\beta)(A_t^*)^{\beta}}{\left(2c_1 A_t^* - (\mu - c_0)\right)^2} > 0\end{aligned}$$

Therefore,

$$\frac{\partial \mathbb{E}(\alpha_{t+1}^{Net})}{\partial \phi_t} > 0 \Leftrightarrow \phi_t < h^{-1}(\beta\theta)$$

Because $h(\phi_t) > \phi_t$, $h^{-1}(\beta\theta) < \beta\theta$.

Next, I show the same pattern holds when net alpha is replaced with gross alpha. The (objective) expected gross alpha can be written as $\mathbb{E}(\alpha_{t+1}^{Gross}) = \mu - c_0 + \theta(A_t^*)^{\beta} - c_1 A_t^*$. Obviously, it is a function of ϕ_t and hence, varies with ϕ_t.

It follows that

$$\begin{aligned}\frac{\partial \mathbb{E}(\alpha_{t+1}^{Gross})}{\partial \phi_t} &= \theta\beta(A_t^*)^{\beta-1}\frac{\partial A_t^*}{\partial \phi_t} - c_1 \frac{\partial A_t^*}{\partial \phi_t} \\ &= \frac{\partial A_t^*}{\partial \phi_t}\Big(\theta\beta(A_t^*)^{\beta-1} - c_1\Big) \\ &= \frac{\partial A_t^*}{\partial \phi_t}\Big(\frac{\theta}{\phi_t}\frac{\beta}{1+\beta}(2c_1 - \frac{\mu - c_0}{A_t^*}) - c_1\Big)\end{aligned}$$

Therefore,

$$\frac{\partial \mathbb{E}(\alpha_{t+1}^{Gross})}{\partial \phi_t} > 0 \Leftrightarrow \beta\theta > \underbrace{\phi_t\left(\frac{(1+\beta)c_1 A_t^*}{2c_1 A_t^* - (\mu - c_0)}\right)}_{g(\phi_t) \equiv \frac{1+\beta}{2} h(\phi_t)}$$

As $h(\phi_t)$ is invertible, $g(\phi_t) \equiv \frac{1+\beta}{2} h(\phi_t)$ is also invertible. Thus,

$$\frac{\partial \mathbb{E}(\alpha_{t+1}^{Gross})}{\partial \phi_t} > 0 \Leftrightarrow \phi_t < g^{-1}(\beta\theta)$$

Because $g(\phi_t) > \phi_t$ and $g(\phi_t) < h(\phi_t)$ (as $\frac{1+\beta}{2} < 1$), $h^{-1}(\beta\theta) < g^{-1}(\beta\theta) < \beta\theta$.

Proof of Corollary I.3: Proposition I.2 along with the fact that the equilibrium fund size increases in ϕ_t (see Equation A.2) immediately leads to Corollary I.3.

Proof of Proposition I.4: Recall that $x_t \equiv \theta + \epsilon_t$ where $x_t = (A^*_{t-1})^{-\beta}\big(\alpha^{Net}_t - (\mu - c_0) + c_1 A^*_{t-1} + f^*_{t-1}\big)$. Suppose investors' prior on θ follows $N(\theta_0, 1/\gamma)$ and the noise term, ϵ_t, is independently distributed through time and follows $N(0, 1/\omega)$, where γ is the precision of prior and ω is the precision of signal uncertainty. Applying the Bayes rule to x_t leads to the investors' belief updating process as follows:

$$
\begin{aligned}
\phi_t &= \frac{\gamma}{\gamma + t\omega}\theta_0 + \frac{t\omega}{\gamma + t\omega}\bar{x}_t \\
&= \frac{\gamma + (t-1)\omega}{\gamma + t\omega}\phi_{t-1} + \frac{\omega}{\gamma + t\omega}x_t \\
&= \phi_{t-1} + \frac{\omega}{\gamma + t\omega}\Big(\frac{\alpha^{Net}_t}{(A^*_{t-1})^{\beta}}\Big) \qquad \text{(A.3)}
\end{aligned}
$$

The first equality directly follows Theorem 1 in DeGroot (1970, page 167). The second equality follows by substituting $\phi_{t-1} = \frac{\gamma}{\gamma+(t-1)\omega}\theta_0 + \frac{(t-1)\omega}{\gamma+(t-1)\omega}\bar{x}_{t-1}$. Finally, the last equality follows because $\mu - c_0 + \phi_{t-1}(A^*_{t-1})^{\beta} - c_1 A^*_{t-1} = f^*_{t-1}$ given the zero subjective expected net alpha condition i.e. $\hat{\mathbb{E}}(\alpha^{Net}_{t-1}) = 0$.

Proof of Corollary I.5: We can write the percentage change in the equilibrium fund size as follows:

$$
\begin{aligned}
\frac{A^*_t - A^*_{t-1}}{A^*_{t-1}} &\equiv \frac{A(\phi_t) - A(\phi_{t-1})}{A(\phi_{t-1})} \\
&\approx \frac{A(\phi_{t-1}) + A'(\phi_{t-1})(\phi_t - \phi_{t-1}) + \frac{A''(\phi_{t-1})}{2}(\phi_t - \phi_{t-1})^2 - A(\phi_{t-1})}{A(\phi_{t-1})} \\
&= \frac{A'(\phi_{t-1})}{(A^*_{t-1})^{\beta+1}}\Big(\frac{\omega}{\gamma + t\omega}\Big)\alpha^{Net}_t + \frac{A''(\phi_{t-1})}{2(A^*_{t-1})^{2\beta+1}}\Big(\frac{\omega}{\gamma + t\omega}\Big)^2\Big(\alpha^{Net}_t\Big)^2 \qquad \text{(A.4)}
\end{aligned}
$$

The approximation follows by the second-order Taylor expansion at ϕ_{t-1}. Note that

$$
\begin{aligned}
A'(\phi_{t-1}) &= \frac{(\beta+1)(A^*_{t-1})^{\beta+1}}{(1-\beta)2c_1 A^*_{t-1} + \beta(\mu - c_0)} > 0 \\
A''(\phi_{t-1}) &= \frac{\beta(\beta+1)(A^*_{t-1})^{\beta}(\frac{\partial A^*_{t-1}}{\partial \phi_{t-1}})\Big((1-\beta)2c_1 A^*_{t-1} + (\mu - c_0)(1+\beta)\Big)}{\big((1-\beta)2c_1 A^*_{t-1} + \beta(\mu - c_0)\big)^2} > 0
\end{aligned}
$$

Plugging $A'(\phi_{t-1})$ and $A''(\phi_{t-1})$ into equation (A.4), we have

$$\Phi_1 = \frac{(\beta+1)}{(1-\beta)2c_1A^*_{t-1}+\beta(\mu-c_0)}\Big(\frac{\omega}{\gamma+t\omega}\Big) > 0$$
$$\Phi_2 = \frac{\beta(\beta+1)^2\Big((1-\beta)2c_1A^*_{t-1}+(\mu-c_0)(1+\beta)\Big)}{\Big((1-\beta)2c_1A^*_{t-1}+\beta(\mu-c_0)\Big)^3}\Big(\frac{\omega}{\gamma+t\omega}\Big)^2 > 0$$

That is, the percentage change in fund size in positive on both α_t^{Net} and $(\alpha_t^{Net})^2$. Further notice that given that $A'(\phi_{t-1}) > 0$ and $A''(\phi_{t-1}) > 0$, as long as $\phi_t - \phi_{t-1}$ increases in α_t^{Net} (which does not have to be in line with Bayesian updating though), we still obtain an increasing and convex flow-performance relation.

A.2 Predictions on Returns to Scale when Investors Learn about c_1

In this subsection, I derive predictions on fund-level returns to scale under both the subjective and objective beliefs when investors learn about the parameter c_1 rather than θ, and show that the predictions remain the same as investors learn about θ.

Solving the equilibrium: The first-order-condition to the manager's problem (1.2) now becomes:

$$\theta(\beta+1)(A_t^*)^\beta - 2\psi_tA_t^* + (\mu-c_0) = 0 \tag{A.5}$$

where $\psi_t \equiv \hat{\mathbb{E}}(c_1)$.

First, I show that if $\mu - c_0 \geq 0$, there exists a unique positive solution to Equation (A.5). Denote $f(A_t^*)$, a function of fund size A_t^*, as the left-hand side of Equation (A.5). When $\mu - c_0 = 0$, $A_t^* = \big(\frac{\theta(\beta+1)}{2\psi_t}\big)^{\frac{1}{1-\beta}} > 0$. When $\mu - c_0 > 0$, it is easy to check that (i) $f(A)$ is monotonically decreasing in A, (ii) $f(0) > 0$, and (iii) $f(\infty) = -\infty$. Therefore, there exists a unique positive solution to equation (A.5).

It follows that $f_t^* = \frac{1-\beta}{1+\beta}\psi_tA_t^* + \frac{\beta}{1+\beta}(\mu-c_0) > 0$. Further, the second-order-condition is strictly negative, i.e.,

$$\theta(\beta+1)\beta(A_t^*)^{\beta-1} - 2\psi = -2\psi_t(1-\beta) - \frac{\mu-c_0}{A_t^*}\beta < 0$$

Therefore, the unique solution to equation (A.5) is indeed the maximum as wanted.

Next, I show the equilibrium fund size A_t^* decreases in investors' perception of c_1—that is, ψ_t:

$$\frac{\partial A_t^*}{\partial \psi_t} = -\frac{\frac{\partial f(\cdot)}{\partial \psi_t}}{\frac{\partial f(\cdot)}{\partial A_t^*}} = -\frac{2(A_t^*)^2}{(\mu - c_0)\beta + 2\psi_t A_t^*(1-\beta)} < 0 \tag{A.6}$$

Subjective Returns to Scale: The subjective returns to scale evaluated at the equilibrium fund size is strictly negative, i.e.,

$$\begin{aligned}\frac{\partial \hat{\mathbb{E}}(\alpha_{t+1}^{Net})}{\partial A_t}|_{A_t = A_t^*} &= \theta\beta(A_t^*)^{\beta-1} - \psi_t \\ &= -\frac{\mu - c_0}{2A_t^*} - \frac{1-\beta}{2}\theta\beta(A_t^*)^{\beta-1} < 0\end{aligned}$$

Objective Returns to Scale: First, I show that the expected fund net alpha first increases and then decreases in ψ_t under the objective belief. The (objective) expected net alpha can be written as

$$\begin{aligned}\mathbb{E}_t\big(\alpha_{t+1}^{Net}\big) &= \mu - (c_0 + c_1 A_t - \theta A_t^{\beta}) - f_t^* \\ &= (\psi_t - c_1)A_t^*\end{aligned}$$

The second equation follows because $\hat{\mathbb{E}}_t\big(\alpha_{t+1}^{Net}\big) = \mu - (c_0 + \psi_t A_t - \theta A_t^{\beta}) - f_t^* = 0$.

Thus, we can write:

$$\begin{aligned}\frac{\partial \mathbb{E}(\alpha_{t+1}^{Net})}{\partial \psi_t} &= A_t^* + (\psi_t - c_1)\frac{\partial A_t^*}{\partial \psi_t} \\ &= A_t^*\Big(1 + (c_1 - \psi_t)\frac{2A_t^*}{2\psi_t A_t^*(1-\beta) + (\mu - c_0)\beta}\Big)\end{aligned}$$

It follows:

$$\frac{\partial \mathbb{E}(\alpha_{t+1}^{Net})}{\partial \psi_t} > 0 \Leftrightarrow \frac{c_1}{\beta} > \underbrace{\psi_t - \frac{\mu - c_0}{2A_t^*}}_{h(\psi_t)}$$

Now, I prove $h(\psi_t)$ is invertible by showing that $h(\psi_t)$ is strictly increasing in ψ_t:

$$\begin{aligned}\frac{\partial h(\psi_t)}{\partial \psi_t} &= 1 + \frac{\mu - c_0}{2(A_t^*)^2}\frac{\partial A_t^*}{\partial \psi_t} \\ &= \frac{(1-\beta)\theta(\beta+1)(A_t^*)^\beta}{2\psi_t A_t^*(1-\beta) + (\mu - c_0)\beta} > 0\end{aligned}$$

Therefore,

$$\frac{\partial \mathbb{E}(\alpha_{t+1}^{Net})}{\partial \psi_t} > 0 \Leftrightarrow \psi_t < h^{-1}\left(\frac{c_1}{\beta}\right)$$

Note that $h^{-1}\left(\frac{c_1}{\beta}\right) > \frac{c_1}{\beta} > c_1$ as $\psi_t > h(\psi_t)$.

Next, I show the same pattern holds when gross alpha replaces net alpha. The (objective) expected gross alpha can be written as $\mathbb{E}(\alpha_{t+1}^{Gross}) = \mu - c_0 + \theta(A_t^*)^\beta - c_1 A_t^*$. It follows that

$$\begin{aligned}\frac{\partial \mathbb{E}(\alpha_{t+1}^{Gross})}{\partial \psi_t} &= -c_1\frac{\partial A_t^*}{\partial \psi_t} + \theta\beta(A_t^*)^{\beta-1}\frac{\partial A_t^*}{\partial \psi_t} \\ &= \frac{\partial A_t^*}{\partial \psi_t}\Big(\theta\beta(A_t^*)^{\beta-1} - c_1\Big) \\ &= \underbrace{\frac{\partial A_t^*}{\partial \psi_t}}_{<0}\Big(\frac{2\beta}{1+\beta}(\psi_t - \frac{\mu - c_0}{2A_t^*}) - c_1\Big)\end{aligned}$$

It follows:

$$\begin{aligned}\frac{\partial \mathbb{E}(\alpha_{t+1}^{Gross})}{\partial \psi_t} > 0 &\Leftrightarrow \frac{1+\beta}{2\beta}c_1 > \underbrace{\psi_t - \frac{\mu - c_0}{2A_t^*}}_{h(\psi_t)} \\ &\Leftrightarrow \psi_t < h^{-1}\Big(\frac{1+\beta}{2\beta}c_1\Big)\end{aligned}$$

Note that $h^{-1}\left(\frac{c_1}{\beta}\right) > h^{-1}\left(\frac{1+\beta}{2\beta}c_1\right) > \frac{1+\beta}{2\beta}c_1 > c_1$.

Given that A_t^* decreases in ψ_t, i.e. $\frac{\partial A_t^*}{\partial \psi_t} < 0$, it follows that both expected gross and net alpha first increase and then decrease in the equilibrium fund size. (When ψ_t is small, A_t^* is large and meanwhile, A_t^* decreases in ψ_t. Thus, a *positive* relation between ψ_t and expected fund alpha when ψ_t is small translates into a *negative* relation between A_t^* and expected fund alpha when A_t^* is large. By the same token, a *negative* relation between ψ_t and expected fund alpha when ψ_t is large translates into a *positive* relation between A_t^* and expected fund alpha when A_t^* is small. Putting together, the expected fund alpha is hump-shaped in the equilibrium fund size.)

APPENDIX B

Returns to Scale: Additional Results

B.1 Details on Recursive Demeaning Estimators

In this subsection, I describe the recursive demeaning (RD) procedure, following Pastor, Stambaugh, and Taylor (2015) and Zhu (2018).

Let's start with the standard fixed effects model,

$$Alpha_{j,t} = a_j + \beta FundSize_{j,t-1} + \theta IndSize_{t-1} + \gamma X_{j,t-1} + e_{j,t} \tag{B.1}$$

where a_j stands for fund fixed effects, $FundSize_{t-1}$ is the log of fund's total net assets (TNA), and $IndSize_{t-1}$ is the sum of TNA for all relevant bond funds scaled by amounts of outstanding of all corresponding bonds. $X_{j,t-1}$ represents fund-level time-varying controls, which include log of fund age (in years), fund expense ratio, turnover ratio, and realized return volatility in the past 12 months.

As in Pastor, Stambaugh, and Taylor (2015), for a given variable $x_{j,t}$, I define its backward-demeaned counterpart, $\underline{x}_{j,t}$, for $t = 2, \cdots, T_j$, as

$$\underline{x}_{j,t} \equiv x_{j,t} - \frac{1}{t-1} \sum_{s=1}^{t-1} x_{j,s}$$

Similarly, forward-demeaned variables are defined as

$$\bar{x}_{j,t} \equiv x_{j,t} - \frac{1}{T_j - t + 1} \sum_{s=t}^{T_j} x_{j,s}$$

First applying the recursively forward-demeaning procedure to variables on both sides of equation (B.1), we end up with the following,

$$\overline{Alpha}_{j,t} = \beta\overline{FundSize}_{j,t-1} + \theta\overline{IndSize}_{t-1} + \gamma\overline{X}_{j,t-1} + \bar{e}_{j,t} \tag{B.2}$$

Then, I estimate equation (B.2) with an instrumental variable approach where I use $\underline{\text{FundSize}}_{j,t-1}$ as the instrument for $\overline{FundSize}_{j,t-1}$ because $\overline{FundSize}_{j,t-1}$ and $\bar{e}_{j,t}$ are correlated by construction.[1]

B.2 Returns to Scale among High-yield Bond Funds

In this subsection, I examine the pattern of fund-level returns to scale among U.S. high-yield bond funds whose Lipper objective code is "HY". The sample includes 245 such funds, covering the period between January 1991 and March 2017. Since none of the four bond market index funds (VBMFX, VBISX, VBIIX, and VBLTX) used to construct the benchmark for corporate bond funds contain high-yield bonds, I instead choose Vanguard High-Yield Corporate Fund (VWEHX) as the benchmark for high-yield bond funds (the "one-factor" benchmark). For robustness, I also include both VWEHX and Vanguard 500 Index Fund (VFINX) in the benchmark (the "two-factor" benchmark). Other than the benchmark composition, the empirical specification here is identical to the one used in Table 1.2.

Table B.1 presents the results. I stat with a quadratic specification. As shown in Column (i) of Panel A, I find a statistically significant hump-shaped relation between fund size and subsequent fund gross alpha. The point estimates on $FundSize$ and $FundSize^2$ suggest a turning point at a fund size of \$454 million ($e^{\frac{0.0245}{2\times 0.0020}} \approx 454$). Interesting, the turning point is close to the one for corporate bond funds. Turning a piecewise linear specification with a turning point at \$454 million implied by the quadratic specification, the result confirms a hump-shaped pattern of return to scale, as shown in Column (i) of Panel B.

Note that the finding of hump-shaped pattern of returns to scale are robust to the case where the fund alpha is calculated instead using the "two-factor" benchmark, as shown in Column (iii) and (iv). In addition, the same hump-shaped pattern holds when fund performance is evaluated by net alpha, as present in Panel B.

[1] Unlike in Pastor, Stambaugh, and Taylor (2015) where intercept is not included in their first-stage regression, I include both the intercept and $\underline{\text{FundSize}}_{j,t-1}$ in the first-stage regression as suggested in Zhu (2018).

In sum, the above findings suggest the pattern of fund-level returns to scale among high-yield bond funds is hump-shaped, similar to the one documented among corporate bond funds.

B.3 Returns to Scale: Alternative Benchmarks

In this subsection, I re-estimate the relation between fund size and subsequent fund alpha among corporate bond funds using a quadratic specification with four different benchmarks: (1) One-factor benchmark includes only VBMFX; (2) two-factor benchmark includes both VBMFX and VFINX; (3) five-factor benchmark contains VBMFX, VFINX, VBISX, VBIIX, VBLTX; and finally (4) risk-factor benchmark consists of short rate factor (three-month Treasury-bill rate), slope factor (ten-year Treasury rate - one-year Treasury rate), curvature factor (two-year Treasury rate + ten-year Treasury rate - 2×five-year Treasury rate), and default risk factor (BAA-rated corporate bond yield - AAA-rated corporate bond yield).

The results are reported in Table B.2. As can be seen, across different choices of the benchmark composition, the relation between fund size and subsequent fund performance is always significantly hump-shaped, consistent with the baseline case (where four bond market index funds are included in the benchmark).

B.4 Estimating the Average Trade Size

In this subsection, I describe how I estimate the average trade size for fund j in month t using the monthly holding data from CRSP. Denote $p_{i,t}$ as bond i's price in month t, and $n_{i,j,t}$ as fund j's position on bond i in month t. Further denote $\mathcal{I}$ as the set of trades in which case $n_{i,j,t}$ does not equal $n_{i,j,t-1}$, and N as the number of trades. Thus, given the two assumptions mentioned in the main text, the average trade size for fund j in month t is calculated as follows

$$AverageTradeSize_{j,t} = \frac{1}{N} \sum_{k \in \mathcal{I}} |p_{k,t} n_{k,j,t} - p_{k,t-1} n_{k,j,t-1}|$$

B.5 Details on Electronic Trading in Corporate Bond Market

In this subsection, I provide details on how I estimate the relation between trade size and the unit transaction cost.

Table B.1: Returns to Scale: High-yield Bond Funds

This table replicates the returns-to-scale results shown in Table 1.2, but using U.S. high-yield bond funds defined with Lipper objective codes being “HY” (excluding index funds, ETFs and ETNs). The sample includes 245 such funds, from January 1991 to March 2017. The dependent variable is either fund gross alpha or net alpha, estimated over the previous 36 months. “One-factor” is the benchmark only with VWEHX (Vanguard High-Yield Corporate Funds); “two-factor” is the benchmark with both VWEHX and VFINX. $FundSize\_Small_{t-1}$ and $FundSize\_Large_{t-1}$ are linear splines constructed from $FundSize_{t-1}$ with the knot implied by the corresponding quadratic specification. In particular,

$$FundSize\_Small_{t-1} = \min(FundSize_{t-1}, k)$$
$$FundSize\_Large_{t-1} = \max(FundSize_{t-1}, k) - k$$

where k is the corresponding knot. Time-varying fund-level controls include: lagged log of fund age (in years), lagged fund expense ratio, lagged turnover ratio, and lagged realized return volatility in the past 12 months. In the parentheses, I report the standard errors clustered by both fund and month.

	One-factor		Two-factor	
	$GrossAlpha_t$ (i)	$NetAlpha_t$ (ii)	$GrossAlpha_t$ (iii)	$NetAlpha_t$ (iv)
Panel A: Quadratic specification				
$FundSize_{t-1}$	0.0245	0.0239	0.0266	0.0261
	(0.0085)	(0.0084)	(0.0092)	(0.0090)
$(FundSize_{t-1})^2$	-0.0020	-0.0020	-0.0021	-0.0021
	(0.0008)	(0.0008)	(0.0008)	(0.0008
$IndustrySize_{t-1}$	-0.0003	-0.0003	-0.0004	-0.0004
	(0.0008)	(0.0007)	(0.0008)	(0.0008)
Observations	24579	24579	24579	24579
Panel B: Piece-wise linear specification				
$FundSize\_Small_{t-1}$	0.0082	0.0080	0.0089	0.0087
	(0.0025)	(0.0025)	(0.0027)	(0.0027)
$FundSize\_Large_{t-1}$	-0.0067	-0.0065	-0.0072	-0.0070
	(0.0036)	(0.0035)	(0.0040)	(0.0039)
$IndustrySize_{t-1}$	-0.0005	-0.0005	-0.0006	-0.0006
	(0.0007)	(0.0006)	(0.0007)	(0.0007)
Observations	24579	24579	24579	24579

Table B.2: Returns to Scale: Alternative Benchmarks

This table replicates the returns-to-scale results shown in Table 1.2, but using different benchmarks to estimate fund alpha. The dependent variable is either gross alpha (Panel A) or net alpha (Panel B), estimated over the previous 36 months. "One-factor" is the benchmark with only VBMFX; "two-factor" is the benchmark with both VBMFX and VFINX; "five-factor" is the benchmark with five index funds (VBMFX, VFINX, VBISX, VBIIX, and VBLTX); and "risk-factor" is short-rate factor (three-month Treasury bill rate), slope factor (ten-year Treasury rate − one-year Treasury rate), curvature factor (two-year Treasury rate + ten-year Treasury rate − 2×five-year Treasury rate), and default risk factor (BAA-rated corporate bond yield − AAA-rated corporate bond yield). In the parentheses, I report the standard errors clustered by both fund and month.

	One-factor (i)	Two-factor (ii)	Five-factor (iii)	Risk-factor (iv)
Panel A: $GrossAlpha_t$ **as dependent variable**				
$FundSize_{t-1}$	0.0070	0.0072	0.0081	0.0127
	(0.0018)	(0.0018)	(0.0019)	(0.0027)
$(FundSize_{t-1})^2$	-0.0005	-0.0006	-0.0007	-0.0009
	(0.0002)	(0.0002)	(0.0002)	(0.0003)
$IndustrySize_{t-1}$	0.0003	0.0002	0.0003	0.0007
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Observations	71905	71897	67209	71905
Panel B: $NetAlpha_t$ **as dependent variable**				
$FundSize_{t-1}$	0.0069	0.0072	0.0081	0.0126
	(0.0018)	(0.0018)	(0.0020)	(0.0027)
$(FundSize_{t-1})^2$	-0.0005	-0.0006	-0.0007	-0.0009
	(0.0002)	(0.0002)	(0.0002)	(0.0003)
$IndustrySize_{t-1}$	0.0003	0.0002	0.0003	0.0007
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Observations	71942	71934	67246	71942

The electronic trading data are provided by MarketAxess, a leading electronic platform provider in the U.S corporate bond market.[2] The sample covers the period from January 2014 to December 2015. The data contains detailed information on each electronic inquiry submitted by bond investors, including the number of dealers quired, the inquiry size and direction, the quote from each responding dealer, an indicator of whether the inquiry leads to trade, and the transaction price if so. One novelty of the MarketAxess data is that it contains investors' (anonymous) identities. I further supplement the MarketAxess data with details on all customer-to-dealer corporate bond transactions, including both voice and electronic trades, from Trade Reporting and Compliance Engine (TRACE). Also, I obtain reference information on all corporate bonds, including ratings, coupon, maturity, and issue size, from Mergent FISD.

I define the unit transaction cost as the percentage difference between the trade price and the benchmark following Hendershott and Madhavan (2015), i.e.

$$\text{Unit Transaction Cost} = \frac{\text{TradePrice} - \text{Benchmark}}{\text{Benchmark}} \times \text{Trade Direction}$$

where Trade Direction equals 1 if the investor is buying and -1 otherwise. Intuitively, the unit transaction cost defined above captures the premium that investors pay for the intermediation. For trades that I can identify their immediate offsets, I use the price on the offsetting leg to construct the benchmark.[3] Otherwise, I follow Hendershott and Madhavan (2015) and take the last inter-dealer price as the benchmark.

Table B.3 presents the summary statistics. Over the sample period, around 16 thousands different investment grade bonds were traded, among which about 11 thousand were covered by MarketAxess. In terms of market share, MarketAxess accounts for about 24% of total number of trades and roughly 30% of total trading volume.[4] For trading costs, as an example, the average unit transaction cost for an odd-lot trade on

[2]According to Greenwich Associates (2015), MarketAxess is the dominate player in electronic trading of US investment-grade bonds with an estimated market share of 75%. The other players are Bloomberg FIT and Tradeweb with estimated market shares being 20% and 5%, respectively. Therefore, it is reasonable to think of the data from MarketAxess is representative of overall electronic trading in the U.S corporate bond market.

[3]I treat the trade that occurred right after the original transaction (within the same day) with the same quantity and opposite direction as the offsetting leg. To construct the benchmark, I use the trade price of the offsetting leg as the benchmark if the offsetting leg is an inter-dealer trade. Otherwise, I compute the benchmark as the average price of the pair of two offsetting trades.

[4]Trade sizes from MarketAxess and TRACE are masked for trades whose sizes are $5 million and above ("Mega" trades). Thus, I report the average trade volume only for those "non-Mega" trades. The same rule applies to the average trade size.

Table B.3: Summary Statistics: MarketAxess vs. TRACE

The sample period is from January 2014 to December 2015. Given that MarketAxess only include customer-to-dealer trades, I restrict my attention to the customer-to-dealer transactions on TRACE as well. The unit transaction cost is defined as $\frac{\text{TradePrice - Benchmark}}{\text{Benchmark}}\times$Trade Direction$\times$100%, where Trade Direction equals 1 if the client is buying and -1 otherwise, and Benchmark equals the price on the offsetting leg, if any, and otherwise, the last inter-dealer price.

	Electronic (MarketAxess)	Overall (TRACE)
Number of Bonds	10,812	15,677
Number of Trades	1,730,830	7,261,147
Client Buy	744,220	4,080,358
Client Sell	986,610	3,180,789
Average Daily Volume (in billions)	1.46	4.90
Client Buy	0.62	2.59
Client Sell	0.84	2.31
Average Trade Size (in thousands)	425	349
Client Buy	418	327
Client Sell	430	379
Unit transaction cost (in bps)		
Micro (1-100K)	13.1	46.9
Odd (100K-1MM)	13.3	19.2
Round (1-5MM)	10.6	9.6

MarketAxess is about 13 bps, or 0.13% of the notional value.

Table B.4 reports the results on how the unit transaction cost changes as trade size varies. Consistent with the literature, as shown in Column (i), the relation between trade size and the unit transaction cost is U-shaped when only time fixed effects (FE) are included. The point estimates on $TradeSize$ and $TradeSize^2$ suggest a turning point at a trade size of \$2.45 million ($\frac{0.0319}{2\times0.0065}\approx 2.45$), which is fairly large given the average trade size is about \$425K. Then, to control for potential client effect, I first add investor fixed effects. As Column (ii) shows, with both time and investor FE, the relation between trade size and unit transaction cost remains U-shaped. Further, by replacing week + investor FEs with week $\times$ investor FE, I find that, given an investor in a given week, the unit transaction cost again first decreases and then increases in trade size, as is shown in Column (iii). Taken together, the above evidence suggests

Table B.4: Relation between Trade Size and the Unit Transaction Cost on MarketAxess

The sample period is from 1/1/2014 to 12/31/2015. The dependent variable is the unit transaction cost, defined as $\frac{\text{TradePrice - Benchmark}}{\text{Benchmark}} \times$Trade Direction$\times$100%, where Trade Direction equals 1 if the client is buying and -1 otherwise, and Benchmark equals the price on the offsetting leg, if any, and otherwise, the last inter-dealer price. Trade size is in million dollars. Time-varying bond-specific controls include time to maturity, time since initial offering, rating, and offering amount. In the parentheses, I report the standard errors clustered by both week and bond.

	(i)	(ii)	(iii)
Trade Size	-0.0319	-0.0339	-0.0296
	(0.0055)	(0.0053)	(0.0065)
(Trade Size)^2	0.0065	0.0068	0.0060
	(0.0013)	(0.0012)	(0.0014)
Week FE	Y	Y	N
Investor FE	N	Y	N
Week $\times$ Investor FE	N	N	Y
Observations	1,408,379	1,408,379	1,405,628
Adjusted R^2	0.017	0.020	0.051

that it is trade size itself—rather than the potential clientele effect that trade size may approximate—that matters for the unit transaction cost for any given corporate bond investor.

B.6 A Replication of Goldstein, Jiang, and Ng (2017) on Flow-performance Sensitivity

In this subsection, I aim to understand what might lead Goldstein, Jiang, and Ng (2017) to document a concave flow-performance sensitivity among corporate bond mutual funds. I start with replicating their main result, i.e., Column (i) of Table 2. Following Goldstein, Jiang, and Ng (2017), I include two index funds in the benchmark, i.e., aggregate bond market and aggregate stock market index funds, when estimating fund alpha. Also, as in Goldstein, Jiang, and Ng (2017), I adopt a piecewise linear specification with a kink at zero net alpha (instead of a quadratic specification used in the main analysis). Lastly, I include the same set of fund-level controls as in Goldstein, Jiang, and Ng (2017): lagged log of total net assets, lagged fund flow, lagged log of fund age (in years), lagged expense ratio, and an indicator variable that equals one if the fund charges rear loads in the last period and zero otherwise.

Column (i) of Table B.5 represents the closest possible specification relative to Goldstein, Jiang, and Ng (2017). Consistent with their finding, the interaction term, i.e. Alpha$\times \mathbb{1}_{Alpha<0}$, is positive and statistically significant, suggesting a concave flow-performance sensitivity. The same conclusion continues to hold when I further add fund fixed effects (FE) and cluster the standard errors at both fund and month level (both of which are reasonable modifications of the original specification in Goldstein, Jiang, and Ng (2017)), as Column (2) and (3) show.

Next, I modify the original specification in Goldstein, Jiang, and Ng (2017) by replacing month FE with "month × sector" FE—investment-grade (IG) and high-yield (HY) sector—to control for any sector-specific unobservable variables which may affect different bond funds within the sector (e.g., investors' overall preference for IG bond funds).

As shown in Column (4), when "month×sector" FE is included, the point estimate of the interaction term, i.e. Alpha$\times \mathbb{1}_{Alpha<0}$, becomes negative and statistically insignificant, indicating no evidence of a concave flow-performance sensitivity. The same conclusion holds when I add fund FE and cluster the standard errors at both fund and month level, as Column (5) and (6) show.

Moreover, when I repeat the above exercise for IG and HY bond funds separately, as shown in Table B.6, I do not find any evidence of a concave flow-performance sensitivity in either subsample, which is consistent with the above finding with month×sector FE.

B.7 Flow-performance Sensitivity with Return Decomposition

In this subsection, I decompose fund excess returns into two parts—benchmark-adjusted (alpha) and benchmark-related (fund's exposures to the benchmark × benchmark excess returns)—and examine how corporate bond fund flows respond to the two components. In particular, I run the following panel regression:

$$Flow_{j,t} = a + \theta NetAlpha_{j,t} + \psi \sum_{i=1}^{4} \beta_{j,i} r_{i,t-1} + \gamma X_{j,t} + \mu_t + \delta_j + e_{j,t} \tag{B.3}$$

where $\beta_{j,i}$ is estimated loading on each of the four index funds included in the benchmark, i.e. VBMFX, VBISX, VBIIX, and VBLTX, $r_{i,t-1}$ corresponds to the excess return of these four index funds, μ_t stands for time fixed effects, and δ_j is fund fixed effects. The time-varying fund-level controls, $X_{j,t}$, includes: lagged log of fund total net assets, lagged fund turnover, lagged log of fund age (in years), lagged fund expense ratio, lagged realized return volatility in the past 12 months, and lagged fund flow. In addition, I run a similar panel regression with four individual components, $\beta_{j,i} r_{i,t-1}$ for $i = 1, \cdots, 4$, rather than just one aggregate component $\sum_{i=1}^{4} \beta_{j,i} r_{i,t-1}$.

I report the results in Table B.7. As shown in Column (i), flows not only respond positively to the fund net alpha, which they should, but also to the benchmark-related returns, which they shouldn't. This evidence suggests that fund investors are confused the fund alpha when allocating their capital. In terms of which component(s) of the benchmark-related returns investors respond most strongly to, the total market component (VBMFX) and the short-term component (VBISX) appear to stand out, as shown in Column (ii).

Table B.5: A Replication of Goldstein, Jiang, and Ng (2017): Flow-performance Sensitivity

This table aims to replicate the main result in Goldstein et al. (2017), i.e. Column (1) of Table 2, with various specifications. The sample includes both investment-grade and high-yield bond mutual funds, from January 1991 to March 2017. The unit of observation is fund-month. The dependent variable is the monthly net fund flow, defined as

$$Flow_{jt} = \frac{TNA_{j,t} - TNA_{j,t-1}(1 + R_{jt})}{TNA_{j,t-1}}$$

where $TNA_{j,t}$ is the total net assets for fund j in month t and R_{jt} is the fund j's net return in month t. As in Goldstein et al. (2017), Alpha refers to net alpha estimated from 36-month rolling regressions with two factors—the aggregate bond market and the aggregate stock market factor, $\mathbb{1}_{Alpha<0}$ is an indicator variable whose value equals one if Alpha is negative and zero otherwise, and fund-level time-varying controls include lagged log of total net assets, lagged proportional fund flow, lagged log of fund age (in years), lagged fund expense ratio, and an indicator variable that equals one if the fund charges rear loads in the last period and zero otherwise. I report standard errors in the parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
Alpha	0.16	0.09	0.09	1.91	1.79	1.79
	(0.14)	(0.16)	(0.32)	(0.30)	(0.33)	(0.35)
Alpha$\times\mathbb{1}_{Alpha<0}$	0.92	1.12	1.12	-0.46	-0.23	-0.23
	(0.24)	(0.25)	(0.42)	(0.44)	(0.42)	(0.44)
$\mathbb{1}_{Alpha<0}$	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Month FE	Y	Y	Y	N	N	N
Month×Sector FE	**N**	**N**	**N**	**Y**	**Y**	**Y**
Fund FE	N	Y	Y	N	Y	Y
Cluster at Fund Level	Y	Y	N	Y	Y	N
Two-way Cluster at Fund & Month	N	N	Y	N	N	Y
Observations	93489	93479	93479	93489	93479	93479
Adjusted R^2	0.162	0.188	0.188	0.193	0.219	0.219

Table B.6: A Replication of Goldstein, Jiang, and Ng (2017): Flow-performance Sensitivity in IG and HY subsamples

This table aims to replicate the main results in Goldstein et al. (2017) in two subsamples: Investment Grade (IG) bond and High Yield (HY) bond funds. The sample is from January 1991 to March 2017. The unit of observation is fund-month. The dependent variable is the monthly net fund flow, defined as

$$Flow_{jt} = \frac{TNA_{j,t} - TNA_{j,t-1}(1 + R_{jt})}{TNA_{j,t-1}}$$

where $TNA_{j,t}$ is the total net assets for fund j in month t and R_{jt} is the fund j's net return in month t. As in Goldstein et al. (2017), Alpha refers to net alpha estimated from 36-month rolling regressions with two factors—the aggregate bond market and the aggregate stock market factor, $\mathbb{1}_{Alpha<0}$ is an indicator variable whose value equals one if Alpha is negative and zero otherwise, and fund-level time-varying controls include lagged log of total net assets, lagged proportional fund flow, lagged log of fund age (in years), lagged fund expense ratio, and an indicator variable that equals one if the fund charges rear loads in the last period and zero otherwise. I report standard errors in the parentheses.

	IG Bond Funds Only			HY Bond Funds Only		
	(1)	(2)	(3)	(4)	(5)	(6)
Alpha	1.85	1.73	1.73	2.07	1.95	1.95
	(0.50)	(0.59)	(0.62)	(0.35)	(0.37)	(0.40)
Alpha×$\mathbb{1}_{Alpha<0}$	-0.62	-0.18	-0.18	-0.35	-0.30	-0.30
	(0.66)	(0.75)	(0.77)	(0.51)	(0.50)	(0.53)
$\mathbb{1}_{Alpha<0}$	-0.01	-0.01	-0.01	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Month FE	Y	Y	Y	Y	Y	Y
Fund FE	N	Y	Y	N	Y	Y
Cluster at Fund-level	Y	Y	N	Y	Y	N
2-way Cluster at Fund & Month	N	N	Y	N	N	Y
Observations	69739	69730	69730	23750	23749	23749
Adjusted R^2	0.173	0.199	0.199	0.247	0.269	0.269

Table B.7: Flow-performance Sensitivity with Return Decomposition

This table shows the flow-performance sensitivity for U.S. corporate bond mutual funds from January 1991 to March 2017 (excluding money market funds, index funds, ETFs, and ETNs). The dependent variable is monthly net fund flow, defined as:

$$Flow_{jt} = \frac{TNA_{j,t} - TNA_{j,t-1}(1 + R_{jt})}{TNA_{j,t-1}}$$

where $TNA_{j,t}$ is the total net assets for fund j in month t and R_{jt} is the fund j's net return in month t. VBMFX component is defined as the product of estimated loading on VBMFX and realized excess return of VBMFX. Same definition applies to VBISX, VBIIX, and VBLTX. $\sum_{i=1}^{4} \beta_{j,i} r_{i,t-1}$ is the sum of these four components. The time-varying fund-level controls include lagged log of fund total net assets, lagged fund turnover, lagged log of fund age (in years), lagged fund expense ratio, lagged realized return volatility in the past 12 months, and lagged fund flow. In the parentheses, I report the standard errors clustered by both month and fund.

	(i)	(ii)
Net Alpha	2.25	2.31
	(0.59)	(0.68)
$\sum_{i=1}^{4} \beta_{j,i} r_{i,t-1}$	0.93	
	(0.30)	
VBMFX component		0.81
		(0.34)
VBISX component		1.22
		(0.31)
VBIIX component		0.46
		(0.29)
VBLTX component		0.25
		(0.30)
Fund FE	Y	Y
Month FE	Y	Y
Observations	65109	65109
Adjusted R^2	0.204	0.204

APPENDIX C

FOMC: Additional Results and Related Derivations

C.1 Evolution of Perceived Law of Motion Parameters

We illustrate the experience-based belief-updating mechanism by showing how individuals' estimates of the parameters of the perceived law of motion (2.1) evolve over time. Figure C.1 presents the estimates of persistence (autocorrelation) ϕ_1 and of the long-run mean inflation rate $\mu = \frac{\alpha}{1-\phi_1-\phi_4+\phi_5}$ obtained from the learning algorithm described in the main text with $\theta = 3.044$, separately for individuals of a few selected ages, 45, 60, and 75.

As the figure shows, the perceived mean rises until 1980 and then declines, while the path of perceived persistence is flatter but also increases around 1980 and then drops dramatically after 2000. Both graphs reveal that the assessments of younger individuals are more volatile than those of older individuals: In 1980s, younger individuals perceived a higher mean than older individuals, while after 2000, the perceived mean of younger individuals falls below that of older individuals. The same pattern also holds for the perceived persistence.

(a) Persistence (autocorrelation) ϕ_1

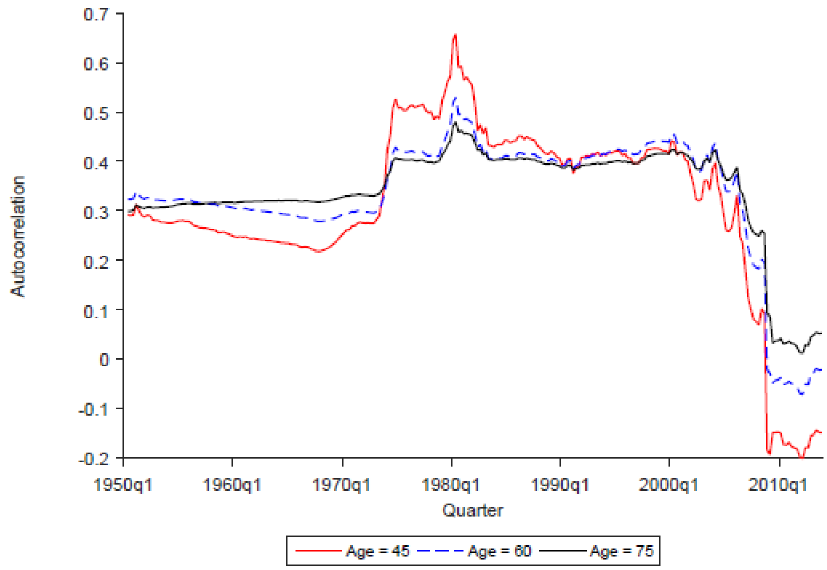

(b) Long-run mean μ

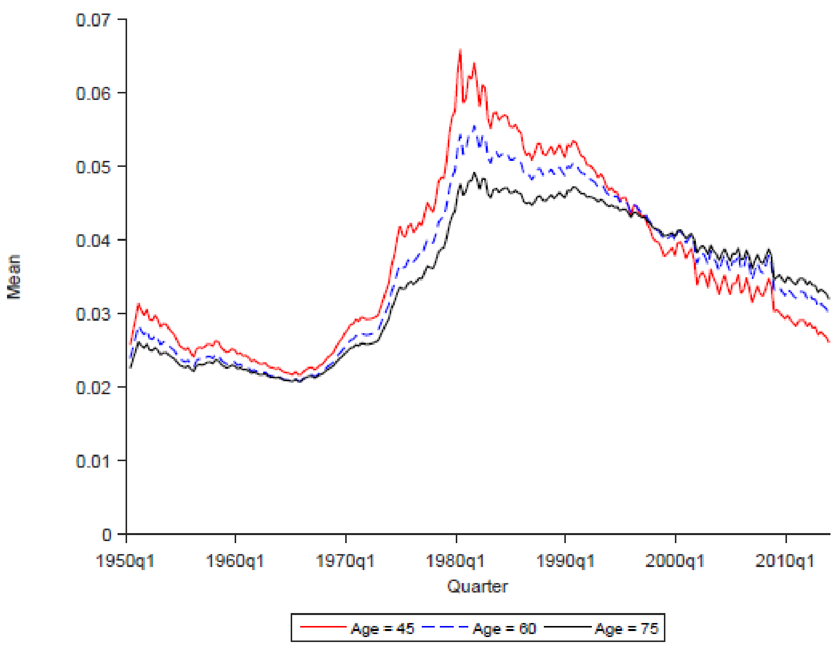

Figure C.1: Mixed Seasonal AR(1) Model Estimates (with $\theta = 3.044$ at ages 45, 60, and 75)

C.2 First-order Taylor approximation of the Subjective Taylor Rule

We start from the subjective Taylor rule in equation (2.9) and substitute the linear specifications in (2.10) to obtain

$$\begin{aligned} i_{j,t} = &\, r + (x_{j,t} - \mu_x)'\alpha_3 + \pi^* + (x_{j,t} - \mu_x)'\alpha_1 \\ &+ \big(\lambda_0 + (x_{j,t} - \mu_x)'\lambda_1\big)\big[\omega\pi^e_{j,t+t|t} + (1-\omega)\pi_t - \pi^* - (x_{j,t} - \mu_x)'\alpha_1\big] \\ &+ \big(\gamma_0 + (x_{j,t} - \mu_x)'\gamma_1\big)\big[y_t - y^* - (x_{j,t} - \mu_x)'\alpha_2\big]. \end{aligned} \tag{C.1}$$

We then perform a first-order Taylor approximation of $i_{j,t}$ as a function of $(\pi^e_{j,t+1|t}, x'_{j,t})$ around (π_t, μ'_x), which yields

$$\begin{aligned} i_{j,t} \approx &\, r + \pi^* + \lambda_0(\pi_t - \pi^*) + \gamma_0(y_t - y^*) + (\pi^e_{j,t+1|t} - \pi_t)\omega\lambda_0 \\ &+ (x_{j,t} - \mu_x)'\left[\alpha_3 + \alpha_1 - \lambda_0\alpha_1 - \gamma_0\alpha_2 + \lambda_1(\pi_t - \pi^*) + \gamma_1(y_t - y^*)\right]. \end{aligned} \tag{C.2}$$

We can rewrite this expression as

$$\begin{aligned} i_{j,t} \approx &\, a_0 + \left[\lambda_0(1-\omega) - \mu'_x\lambda_1\right]\pi_t + (\gamma_0 - \mu'_x\gamma_1)(y_t - y^*) \\ &+ \lambda_0\omega\pi^e_{j,t+1|t} + \kappa' x_{j,t} + \pi_t x'_{j,t}\lambda_1 + (y_t - y^*)x'_{j,t}\gamma_1, \end{aligned} \tag{C.3}$$

where

$$\begin{aligned} a_0 &= r + \pi^*(1-\lambda_0) - \mu'_x(\alpha_3 + \alpha_1 - \lambda_0\alpha_1 - \gamma_0\alpha_2 - \lambda_1\pi^*), \\ \kappa &= \alpha_3 + \alpha_1 - \lambda_0\alpha_1 - \gamma_0\alpha_2 - \pi^*\lambda_1. \end{aligned}$$

Denoting the first three terms on the right-hand side of (C.3) as a_t, we obtain equation (2.11) in the main text. Defining

$$\beta_0 = a_0, \qquad \beta_e = \lambda_0\omega, \qquad \beta_\pi = \lambda_0(1-\omega) - \mu'_x\lambda_1, \qquad \beta_y = \gamma_0 - \mu'_x\gamma_1, \tag{C.4}$$

and averaging across FOMC members at meeting time t yields equation (2.15) in the text.

C.3 Vote Sample Construction

Our sample of FOMC votes starts in 1951, after the official reinstatement of the Federal Reserve Bank's independence in the Treasury-Federal Reserve agreement of March 4, 1951. During our sample period from March 1951 to January 2014, eight Fed Chairmen lead the FOMC: McCabe (4/1948 to 4/1951), Martin (4/1951 to 1/1970), Burns (2/1970 to 3/1978), Miller (3/1978 to 8/1979), Volcker (8/1979 to 8/1987), Greenspan (8/1987 to 1/2006), and Bernanke (2/2006 to 1/2014).

The data set is constructed from two main sources. First, for meetings before January 1966 and after January 1997, we collect information on the votes from the FOMC meeting statements available at `http://www.federalreserve.gov/monetarypolicy/fomccalendars.htm`. Second, for meetings between January 1966 and December 1996, we use the data from Chappell, McGregor, and Vermilyea (2005), available at `http://professorchappell.com/Data/Book/index.htm`. In this latter data, we correct one coding error: In the meeting on 11/5/1985, governor Seger cast a dovish dissent (-1); the original data set had her vote coded as consent (0).

We also note several discrepancies between our sample and the data employed by Thornton and Wheelock (2014) in their analysis of votes in the Federal Reserve Bank of St. Louis Review:

- For the meeting on 10/3/1961, the Fed Review data records one dissent. We find no dissent reported in the meeting minutes.

- For the meeting on 2/9/1983, the Fed Review data records one dissent. We find four dissents reported in the minutes.

- Other discrepancies reflect dissents that occurred in conference calls (no separate Record of Policy Actions was released), which are not included in our sample. Our sample does include nine conference calls (94 total votes and 2 dissents), after which a separate Record of Policy Actions/Statement was available. We exclude those from the baseline sample. Including them does not alter the results.

We further exclude five votes by the two members who voted less than five times during their tenure with the FOMC, Paul Miller and Jamie Stewart. Mr. Miller only had one vote because he died in office (on Oct. 21, 1954), less than three month after he was appointed to the Board of Governors (on Aug. 13, 1954). Mr. Stewart cast four votes as the acting governor, when he was the first vice president of New York Fed, from

June through December 2003, during which the position of New York Fed president was vacant after McDonough resigned in 2003 and before his successor Geithner took place in Nov. 2003.

After the above corrections (and excluding votes from conference calls), our sample contains 160 dovish dissents, 265 hawkish dissents, and 8 un-codeable dissents between 3/8/1951 to 1/29/2014.[1] The eight un-codeable dissents are as follows:

- In the 12/19/1961 meeting, Robertson dissented with the reason explained as follows: *"While Mr. Robertson's analysis of the economic situation and the proper direction of policy was the same in its essentials as that of the majority, he voted against adoption of this directive on the grounds that it was undesirable to tie monetary policy to the bill rate."* See `www.federalreserve.gov/monetarypolicy/files/fomcropa19611219.pdf`.

- In the 7/30/1963 meeting, Bopp dissented with the reason explained as follows: *"Mr. Bopp stated that he had voted favorably on the policy directive at the July 9 meeting because it seemed to him that the use of the different instruments of monetary policy should be consistent and an increase in the discount rate was then imminent. Under such circumstances, it had seemed undesirable to reverse what had taken place in terms of yields only to reverse again. His vote, therefore, was essentially a vote on tactics. As to the future, it was still an open question whether short-term rates could be maintained at the new levels, and reserve availability at the old. Under these conditions, he agreed with the view that it would be desirable to maintain essentially an even keel for the time being, and to supply reserves through purchases of coupon issues, selling bills if necessary. In his opinion, emphasis should be placed on the availability of reserves."* See `www.federalreserve.gov/monetarypolicy/files/fomchistmin19630730.pdf`.

- In the 12/12/1967 meeting, Maisel dissented with the reason explained as follows: *"Mr. Maisel dissented from this action in part because he thought the directive was susceptible to an interpretation under which growth in member bank reserves and bank deposits would be slowed too abruptly, and perhaps succeeded by contraction. He favored seeking growth rates in reserves, deposits, and bank credit considerably below the average rates thus far in 1967, but still high enough*

[1] There are 13 additional dissents that occurred between 1936 and 1950, and two dissenting votes were cast during the nine conference calls in our sample. Neither are included in our data.

to facilitate expansion in GNP at a somewhat faster rate than had prevailed on average in the first three quarters of the year. He noted that whether or not interest rates would rise further under the course he advocated would depend upon the strength of market demands for funds in relation to the supplies that would be available under such a Committee policy. Mr. Maisel also thought that the statement of the Committee's general policy stance contained in today's directive had far too narrow a focus; in particular, he objected to the omission of reference to the basic policy goal of facilitating sustainable economic expansion. This omission resulted from the substitution of language stating that it was the Committee's policy "to foster financial conditions conducive to resistance of inflationary pressures and progress toward reasonable equilibrium in the country's balance of payments" for the language of other recent directives stating that it was the Committee's policy "to foster financial conditions, including bank credit growth, conducive to sustainable economic expansion, recognizing the need for reasonable price stability for both domestic and balance of payments purposes." See `www.federalreserve.gov/monetarypolicy/files/fomcropa19671212.pdf`.

- In the 1/11/1972 meeting, Brimmer dissented with the reason explained as follows: *"Mr. Brimmer shared the majority's views concerning broad objectives of policy at this time, and he indicated that he would have voted favorably on the directive were it not for the decision to give special emphasis to total reserves as an operating target during coming weeks. In his judgment the Committee should have had more discussion of the implications of that decision, and in any case it should have postponed the decision until after it had held a contemplated meeting to be devoted primarily to discussion of its general procedures with respect to operating targets."* See `www.federalreserve.gov/monetarypolicy/files/fomcropa19720111.pdf`.

- In the 7/17/1973 meeting, Francis dissented with the reason explained as follows: *"Mr. Francis dissented from this action not because he disagreed with the objectives of the policy adopted by the Committee but because he believed that—as had proved to be the case following other recent meetings—the objectives would not be achieved because of the constraint on money market conditions."* See `www.federalreserve.gov/monetarypolicy/files/fomcropa19730717.pdf`.

- In the 7/20/1976 meeting, Volcker dissented with the reason explained as follows:

"Mr. Volcker dissented from this action because in the present circumstances he would not wish to raise or lower the Federal funds rate by as much as 1/2 of a percentage point—a change that might be interpreted as a strong signal of a change in policy and that could have repercussions in financial markets—in response merely to short-term fluctuations in the monetary aggregates that might well prove transient." See `www.federalreserve.gov/monetarypolicy/files/fomcropa19760720.pdf`.

- In the 12/22/1981 meeting, Soloman dissented with the reason explained as follows: *"Mr. Solomon dissented from this action because he felt it was particularly important at the beginning of an annual target period that the Committee not formulate its directive in terms that conveyed an unrealistic sense of precision. In his view, the directive language referring to the November-to-March growth rates in M1 and M2 did seem to convey such a sense."* See `www.federalreserve.gov/monetarypolicy/files/fomcropa19811222.pdf`.
- In the 2/9/1983 meeting, Horn dissented with the reason explained as follows: *"Mr. Black and Mrs. Horn dissented from this action because they preferred to give more weight to M1 as a policy objective. While recognizing the difficulties in interpreting M1 currently, they believed that over time M1 was more reliably related to the Committee's ultimate economic objectives than were the broader aggregates and that it constituted a better basis for setting appropriate paths for reserve growth. They also favored reemphasizing M1 because they viewed it as a more controllable aggregate. In addition, Mr. Black indicated that he saw a need for lower target ranges, but he wanted to reduce monetary expansion gradually to avert dislocative effects."* See `www.federalreserve.gov/monetarypolicy/files/fomcropa19830209.pdf`. We record Black's vote as hawkish (+1).

As we note in the main text, four members of the FOMC were both regional Fed presidents and governors at some point, and we account for their varying roles in our empirical analysis. These four members are: Phillip Coldwell (Dallas Fed President from 2/68 to 10/74 and governor from 10/74 to 2/80), Oliver Powell (governor from 9/50 to 6/52 and Minneapolis Fed President from 7/52 to 3/57), Paul Volcker (NY Fed president from 5/75 to 8/79 and Fed Chairman from 8/79 to 11/87), and Janet Yellen (governor from 8/94 to 2/97, SF Fed president from 6/04 to 10/10, and then again governor since 10/2010, including her role as Fed Chairwoman).

C.4 Mixed Inflation Process with a Hyperinflation Regime

This section presents an alternative approach for integrating Henry Wallich's hyperinflation experiences into the estimation.

We assume that every period, inflation is drawn from the following mixed process with two regimes, one for hyperinflation, which takes place with probability p, and one for non-hyperinflationary periods

$$\pi_{t+1} = \mu + u_{t+1} \qquad \text{with probability} \quad p, \tag{C.5}$$

$$\pi_{t+1} = \alpha + \phi\pi_t + e_{t+1} \qquad \text{with probability} \quad 1 - p, \tag{C.6}$$

where $E_t[u_{t+1}] = 0$ and $E_t[e_{t+1}] = 0$. Therefore, μ is the expected value of π_{t+1} conditional on a hyperinflation occurring, and we can define

$$\mu_0 = \frac{\alpha}{1 - \phi} \tag{C.7}$$

as the expected value conditional on no hyperinflation. With known parameters, a forecast conditional on observed inflation would be

$$E_t[\pi_{t+1}] = p\mu + (1 - p)(\alpha + \phi\pi_t) = p(\mu - \mu_0) + \alpha + \phi\pi_t - p(\alpha + \phi\pi_t - \mu_0). \tag{C.8}$$

For small hyperinflation probabilities, the last term $p(\alpha + \phi\pi_t - \mu_0)$ is tiny relative to the others ($\mu - \mu_0$ is orders of magnitude bigger than to $\alpha + \phi\pi_t - \mu_0$). Thus, we can approximate,

$$E_t[\pi_{t+1}] \approx p(\mu - \mu_0) + \alpha + \phi\pi_t \tag{C.9}$$

i.e., the usual AR(1) forecast conditional on no hyperinflation plus an upward adjustment to the long-run mean to account for the fact that a hyperinflation might occur with probability p. This is the forecast we want to construct (in an experience-based way).

Parameters can now be estimated as follows: α and ϕ can be estimated in the usual way (the same way we do it for other FOMC members) from a sample excluding hyperinflation periods, for which we simply take US data only (mixing in some early German data would not make a difference as long as the hyperinflation years are excluded). To estimate $p(\mu - \mu_0)$, we can use the fact that the mean from sampling data from both regimes (i.e., German data for Wallich's youth years included) is

$$E[\pi_t] = p\mu + (1-p)\mu_0 \tag{C.10}$$

which implies

$$p(\mu - \mu_0) = E[\pi_t] - \mu_0 \tag{C.11}$$

We can estimate $E[\pi_t]$ as the simple mean estimate from mixed German-US data. And $\mu_0 = \alpha/(1-\phi)$ follows from the AR(1) estimates based on US data. Combining these gives us an estimate for $p(\mu - \mu_0)$ which we can then add to the no-hyperinflation AR(1) forecast $\alpha + \phi\pi_t$ to get $E_t[\pi_{t+1}]$ as in (C.9). For simplicity of exposition, we have illustrated the approach above with a simple AR(1) for the non-hyperinflation regime. But in our estimation, we instead use a mixed seasonal Ar(1) as in (2.1) in the main text.

Table C.1 reports the results. Apart from the use of the mixed inflation process and the absence of the Wallich dummy, everything else is the same as in Table 2.3 in the main text. As Table C.1 shows, there is still a strong and statistically highly significant effect on voting decisions. The APE show at the bottom of the table are somewhat smaller than in Table 2.3 in the main text, but with Wallich's hyperinflation experiences integrated through the mixed inflation process, the average within-meeting dispersion is now 0.15 percentage points (instead of the 0.10 that we had earlier). A one standard deviation change now translates into a change in the probability of hawkish or dovish dissent of about 1/6 of the unconditional dissent probabilities (compared with between 1/4 to 1/3 earlier).

C.5 Fixed-Threshold Ordered Probit Estimates

This section presents estimates from an ordered probit model as in (2.12), but with fixed dissent thresholds. Note that we use the fitted values from this estimation to construct the $\bar{z}_t$ variable in (2.15), which is the basis for the results on the Fed Funds Rate target presented in Table 2.8.

Table C.2 presents the ordered probit estimates. In column (i) we employ time fixed effects, and in column (ii) we express explanatory variables values as deviations from their values for the chairperson. The results are similar to the corresponding ones in Table 2.3 in the main text.

This fixed-threshold specification also offers the opportunity to examine the co-

Table C.1: Experience-based Inflation Forecasts and FOMC Voting Behavior

This table repeats the estimation from Table 2.3 in the main text, but with experience-forecasts for Henry Wallich calculated using the mixed inflation process with a hyperinflation regime.

	Ordered Probit		Ordered Probit "de-chaired"	
	(i)	(ii)	(iii)	(iv)
Experienced-Based Forecast	79.5	75.3	47.8	48.0
	(23.3)	(23.8)	(11.6)	(12.1)
Meeting FE	Yes	Yes	No	No
Thresholds	Role × $I_{>93}$	All	Role × $I_{>93}$	All
Observations	6,707	6,707	6,707	6,707
Pseudo R^2	0.394	0.396	0.108	0.112
APE of Experienced-Based Forecast:				
Dovish Dissent	-2.8	-2.7	-2.5	-2.5
Consent	-1.6	-1.5	-1.3	-1.2
Hawkish Dissent	4.4	4.1	3.7	3.7

efficients of the control variables. In the characteristics-dependent specification they are difficult to interpret because their effect on the dissent threshold is intertwined with their effect on the conditional mean of the latent variable and hence the voting decision. Table C.3 presents the coefficient estimates, including those for the interactions. Directionally, the results are broadly sensible. For example, FOMC members put more weight on current inflation and less weight on unemployment if they are older, are regional Fed presidents, male, appointed when a Republican U.S. president was in office, and are not in the same party as the current president. However, many of these estimates are statistically not significantly different from zero. To interpret the direct effect of the characteristics, we need to add the interacted terms evaluated at particular values of CPI inflation (e.g., 2%) and unemployment (e.g., 6%). Doing so reveals that there is a fairly strong association of hawkishness with regional president role and appointment while a Republican president was in office, while female gender is associated with a more dovish voting behavior.

Table C.2: Experience-based Inflation Forecasts and FOMC Voting Behavior: Simple Ordered Probit without Characteristics-Dependent Thresholds

The sample period is from March 8, 1951 to January 29, 2014. The experience-based inflation forecast for each member at each meeting is calculated by recursively estimating a mixed seasonal AR(1) model using the member's lifetime history of inflation with $\theta = 3.044$, as described in Section 2.2.1. The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. In parentheses, we report the standard error based on two-way clustering by both member and meeting.

	Ordered Probit	Ordered Probit "de-chaired"
	(i)	(ii)
Experienced-Based Forecast	192.2	89.7
	(60.0)	(36.1)
Wallich Dummy	1.6	1.2
	(0.4)	(0.2)
Meeting FE	Yes	No
Observations	6,707	6,707
Pseudo R^2	37.0%	8.2%
APE of Experienced-Based Forecast:		
Dovish Dissent	-7.0	-4.7
Consent	-4.1	-2.3
Hawkish Dissent	11.1	7.1
APE of Wallich Dummy:		
Dovish Dissent	-0.06	-0.06
Consent	-0.03	-0.03
Hawkish Dissent	0.09	0.09

Table C.3: Experience-based Inflation Forecasts and FOMC Voting Behavior: All coefficients

The sample period is from March 8, 1951 to January 29, 2014. The variables are defined as described in the main text. In parentheses we report standard errors based on two-way clustering by both member and meeting.

	Ordered Probit	Ordered Probit - “de-chaired”
Experienced-Based Forecast	192.24	89.66
	(60.04)	(36.12)
Wallich Dummy	1.57	1.16
	(0.37)	(0.18)
Age	-0.04	-0.03
	(0.03)	(0.01)
Fed Role	0.41	0.15
	(0.36)	(0.28)
Gender	0.01	0.09
	(0.87)	(0.58)
Party	1.09	0.47
	(0.46)	(0.29)
Same Party	-0.09	-0.42
	(0.43)	(0.25)
Fed Role $\times$ $\mathbb{1}_{\text{Post1993}}$	-0.11	-0.03
	(0.25)	(0.20)
CPI $\times$ Age	0.45	0.44
	(0.30)	(0.14)
CPI $\times$ Fed Role	4.23	5.42
	(3.88)	(1.96)
CPI $\times$ Gender	12.44	6.22
	(6.21)	(3.23)
CPI $\times$ Party	-5.83	-1.72
	(4.08)	(2.57)
CPI $\times$ Same Party	-0.88	-2.85
	(3.68)	(1.88)
Unemp. rate $\times$ Age	-0.67	-0.39
	(0.45)	(0.25)
Unemp. rate $\times$ Fed Role	-1.21	-2.25
	(5.90)	(4.89)
Unemp. rate $\times$ Gender	-9.87	-4.49
	(11.54)	(6.58)
Unemp. rate $\times$ Party	9.78	5.16
	(7.61)	(4.47)
Unemp. rate $\times$ Same Party	0.36	-7.43
	(7.60)	(4.31)
Meeting FE	Yes	No
Observations	6707	6707
Pseudo R^2	37.0%	8.2%

C.6 Speech Sample Construction

The FRASER economic history database at the Federal Reserve Bank of St. Louis maintains a digital library of speeches of past and current FOMC members. To construct our sample of speeches, we first download the HTML source code of the webpage listing the *Statements and Speeches of Federal Reserve Officials*. The source code contains a list of the FOMC members and their record IDs. (See the screenshot in Figure C.2a.) Each record ID uniquely identifies a webpage with the links to all speeches of the respective FOMC member. We use the record IDs to download the HTML source code of those webpages (see Figure C.2b), and then extract the so-called issue IDs of the individual speeches. The issue IDs, in turn, link to the webpages containing the metadata of the speeches, including the links to the pdfs (see Figure C.2c). We collect all links to the pdfs of the speeches in a single text document and parse the document to the *wget* function, which downloads the pdf files.[2] In addition, we hand-collected speeches from the websites of the regional FRBs for the regional presidents.

To search the speeches for hawkish and dovish language, the downloaded pdfs are converted to text format using a unix shell executable script. During this step, the speech text is cleaned of reference sections, typographic ligature, and duplicates of the speech header or title which is often repeated on every page of the pdfs. (Even though some of the speeches are photographs of the manuscript, the images are already translated into text and we do not have to run OCR for any of the cases.)

We restructure the text into sequences of five adjacent words, and then select the relevant subset of goal-centered five-grams. For example, words from the sentence "Inflation continued to be well behaved, and in fact with talk of lower oil prices there was even a whiff of deflation." said by Thomas Meltzer in a 1985 address to the Harry J. Loman Foundation, initially show up in twenty nine different five-grams. Only two of these five grams are kept and searched for words from the *attitudes* list: "[*two words from the previous sentence*]. Inflation continued to" and "of lower oil prices there". After searching for these attitude words, the second five-gram is tagged as dovish, because it contains the word "lower" from the *attitudes* list, and the first is not tagged at all.

There is a cluster of short speeches with around 500 n-grams. Checking these speeches by hand reveals that a large fraction are short opening remarks and introductions for other speeches, or short-hand notes for longer speeches instead of full

[2]We invoke the *wget* function from www.gnu.org/software/wget/Overview via OS X Terminal.

```html
                    <li id="record-905" class="issue-list-item">
      <p class="issue-list-item-firstline"><span> Statements and Speeches of Abbot Low Mills
</span></p>
      <input type="hidden" class="record-content-type" value="title">
      <input type='hidden' class='record-id' value="905">
    </li>
                      <li id="record-452" class="issue-list-item">
      <p class="issue-list-item-firstline"><span> Statements and Speeches of Alan Greenspan
</span></p>
      <input type="hidden" class="record-content-type" value="title">
      <input type='hidden' class='record-id' value="452">
    </li>
                      <li id="record-906" class="issue-list-item">
      <p class="issue-list-item-firstline"><span> Statements and Speeches of Alan S. Blinder
</span></p>
      <input type="hidden" class="record-content-type" value="title">
      <input type='hidden' class='record-id' value="906">
    </li>
                      <li id="record-907" class="issue-list-item">
      <p class="issue-list-item-firstline"><span> Statements and Speeches of Alice M. Rivlin
</span></p>
      <input type="hidden" class="record-content-type" value="title">
      <input type='hidden' class='record-id' value="907">
    </li>
                      <li id="record-463" class="issue-list-item">
      <p class="issue-list-item-firstline"><span> Statements and Speeches of Andrew F. Brimmer
</span></p>
```

(a) Step 1: HTML source code of the FRASER webpage for the *Statements and Speeches of Federal Reserve Officials*. The record IDs, highlighted by the box, identify the webpages with all speeches of the respective FOMC member.

```html
              <input type='hidden' class='issue-id' value="35277">
      </li>
          <li id="issue-35278" class="issue-list-item item-decade-1990">
        <p class='issue-list-item-firstline'>1997 |
                                                  <a href="/title/907#!35278">
          Sustaining CRA's Success
  : Remarks at the Annual Meeting of the National Community Reinvestment Coalition, Washington, D.C.
            </a>
        </p>
              <input type='hidden' class='issue-id' value="35278">
      </li>
          <li id="issue-35261" class="issue-list-item item-decade-1990">
        <p class='issue-list-item-firstline'>1997 |
                                                  <a href="/title/907#!35261">
          The Challenges of Macroeconomic Policy
  : Remarks at the Annual Meeting of the Eastern Economic Association, Washington, D.C.
            </a>
        </p>
              <input type='hidden' class='issue-id' value="35261">
      </li>
          <li id="issue-35262" class="issue-list-item item-decade-1990">
        <p class='issue-list-item-firstline'>1997 |
                                                  <a href="/title/907#!35262">
          Appropriate Monetary Policy and the Strong Economy
  : Testimony before the Committee on Banking and Financial Services, U.S. House of Representatives
            </a>
        </p>
              <input type='hidden' class='issue-id' value="35262">
      </li>
          <li id="issue-35263" class="issue-list-item item-decade-1990">
        <p class='issue-list-item-firstline'>1997 |
                                                  <a href="/title/907#!35263">
```

(b) Step 2: HTML code identified by the record ID obtained in the previous step. The issue IDs, highlighted by the boxes, identify the webpages with the metadata of the speeches of the respective FOMC member, including the links to the pdf files with the speeches.

```xml
- <modsCollection default:xsi="http://www.loc.gov/standards/mods/v3/mods-3-5.xsd" default:schemaLocation="">
  - <mods>
      <genre>speech</genre>
      <language>eng</language>
    - <location>
      - <url>
          https://fraser.stlouisfed.org/scribd/?item_id=808&filepath=/docs/historical/frbminn/presidents/corrigan/corrigan_19810701.pdf
        </url>
      </location>
    - <titleInfo>
        <title>Current Problems and Prospects for the Economy</title>
      - <subTitle>
          Statement before the Subcommittee on Conservation, Credit, and Rural Development, Committee on Agriculture, House of Representatives
        </subTitle>
        <titlePartNumber/>
```

(c) Step 3: Metadata of a speech, including a link to the pdf (highlighted by the box).

Figure C.2: FRASER Source Code to Obtain Speech PDFs

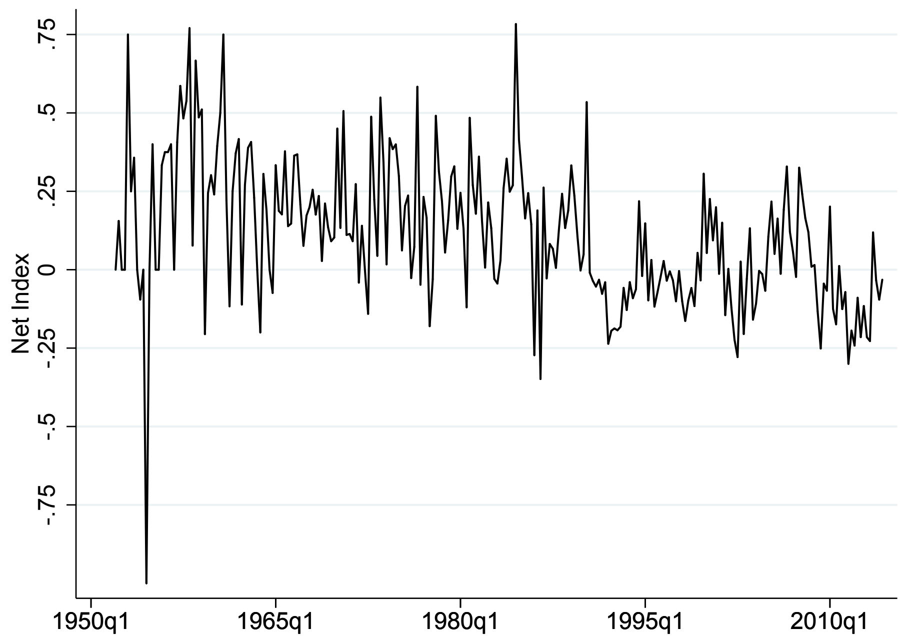

Figure C.3: Net Index Over Time

Notes. The graph depicts the average *Net Index* (using the expanded set of goals) of all speeches in year-quarter.

transcripts. Controlling for these short speeches by including an indicator variable for less than 750 n-grams has virtually no effect on the results.

In the main text, we describe the construction of the *Net Index* of speech hawkishness. Figure C.3 plots the time-series of the index. The index decreases slightly over time, especially after 1980. But overall there is fairly strong time-variation without much persistence. This may reflect a considerable amount of measurement noise in *Net Index*. The more muted amplitude of the *Net Index* in later sample years probably reflects the substantially larger number of speeches available, rather than a general trend towards a more neutral language, implying that the mean of *Net Index* contains less measurement error in later years.

As also discussed in the main text, our analysis of FOMC members' choice of words and tone of speeches might warrant further controls for personal characteristics to reduce noise and concerns about correlated omitted variables. We construct control variables for education and prior professional experience. Information on education, including degree type and degree granting institutions, is available from the member

Table C.4: Summary Statistics on FOMC Members' Educational Background

The table below shows statistics on the educational background for the 144 FOMC members who voted at least 5 times during the meetings from 3/8/1951 to 1/29/2014. Panel A shows every school that awarded the highest degree of at least three members, along with the number of bachelor's and PhD degrees awarded by those schools. Panel B shows the frequency with which each degree type was the highest degree awarded to an FOMC member. All data is from the Federal Reserve History Gateway.

Panel A: Most Common Schools

School	Highest Degree	PhD	Bachelors
Harvard University	24	10	8
University of Pennsylvania	10	6	4
MIT	7	7	1
University of Michigan	6	4	1
University of Missouri	6	1	3
Indiana University	5	4	2
University of Chicago	4	4	1
John Hopkins University	4	2	0
Stanford University	4	1	3
UCLA	3	3	0
University of Wisconsin	3	3	0
University of California, Berkeley	3	2	3
Yale	3	1	5
University of Virginia	3	1	3
Columbia University	3	1	2
Iowa State University	3	1	1
NYU	3	1	1
Georgetown University	3	0	1

Panel B: Highest Degree

School	Number of FOMC Members	Percentage
PhD	65	45.1%
JD	22	15.3%
Master's	20	13.9%
Bachelor's	17	11.8%
MBA	15	10.4%
Other	5	3.5%

biographies provided by the Fed on the Federal Reserve History Gateway website.

Table C.4 shows the summary statistics on the educational background for the 144 FOMC members in our sample.: 45.1% of members have a PhD as their highest degree, while 15.3% have a law degree, and 10.4% have an MBA. 24 of the 144 members hold their highest degrees from Harvard, ten from the University of Pennsylvania, seven from MIT, and six each from the University of Michigan and the University of Missouri. Harvard has also granted the most PhDs to FOMC members (ten). MIT follows with seven, six members have PhDs from the University of Pennsylvania, and four have PhDs from the Universities of Chicago, Michigan, and Indiana each. 67.4% have their highest degree in economics, or majored in it if their highest degree is a bachelors.

Also from the Federal Reserve History Gateway website, we collect mentions of FOMC members' industry experience prior to their first FOMC meeting. Members are classified as having had, or not had work experience in the financial industry, an academic department, the military, a government agency other than the Federal Reserve or the military, and other industries, e.g. manufacturing. 76 of the 144 members with at least three votes are classified as having financial industry experience, 74 as having worked at another government agency, 62 in academia, 53 in another industry, and 37 as having military experience.

C.7 Results without Members born before 1913

We replicate the results on voting and the tone of speeches including only FOMC Members born after 1913. These analyses address potential concerns about the methodological change in the inflation series in 1913. As can be seen below, our results remain the same. Our analyses of Fed Funds target rate and MPR inflation forecasts are not affected by this methodological change as they do not use pre-1913 data.

Voting The following three tables replicate the results as in Table 2.3 to 2.5 focusing on FOMC Members born after 1913.

Table C.5: Experience-based Inflation Forecasts and FOMC Voting Behavior: Only with Members who were Born after 1913

The sample period is from March 8, 1951 to January 29, 2014. The sample excludes FOMC Members who were born before 1913. The experience-based inflation forecast for each member at each meeting is calculated by recursively estimating a mixed seasonal AR(1) model using the member's lifetime history of inflation, as described in Section 2.2.1 (with $\theta = 3.044$). The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. Column (i) and (iii) report the results assuming that the thresholds depend on a) whether the member is a board member or regional president, and b) whether the meeting occurs after Nov. 1993 and the interaction of a) and b). Column (ii) and (iv) report the results assuming that the thresholds depends, in addition, on age, gender, party of president at appointment indicator, and same party as current president indicator. In parentheses we report the standard error based on two-way clustering by both member and meeting.

	Ordered Probit		Ordered Probit "de-chaired"	
	(i)	(ii)	(iii)	(iv)
Experienced-Based Forecast	265.3	289.8	126.6	138.0
	(72.6)	(78.9)	(42.5)	(45.3)
Wallich Dummy	1.4	1.3	1.0	0.8
	(0.4)	(0.3)	(0.2)	(0.2)
Meeting FE	Yes	Yes	No	No
Thresholds	Role × $I_{>93}$	All	Role × $I_{>93}$	All
Observations	4284	4284	4284	4284
Pseudo R^2	38.2%	39.4%	12.0%	13.5%
APE of Experienced-Based Forecast:				
Dovish Dissent	-9.1	-9.8	-6.2	-6.7
Consent	-7.8	-8.3	-4.7	-4.9
Hawkish Dissent	16.9	18.1	11.0	11.7
APE of Wallich Dummy:				
Dovish Dissent	-0.048	-0.042	-0.047	-0.040
Consent	-0.041	-0.036	-0.036	-0.029
Hawkish Dissent	0.089	0.079	0.083	0.069

Table C.6: Experience-based Inflation Forecasts and FOMC Voting Behavior: Different Sample Periods with Fixed Ordered Probit Thresholds and Only with Members who were Born after 1913

The sample excludes FOMC Members who were born before 1913. The experience-based inflation forecast for each member at each meeting is calculated as in Table 2.3. The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. Column (i) reports the results with all FOMC members prior to November 1993. Column (ii) reports the results with regional Fed presidents only prior to November 1993. Column (iii) reports the results with regional Fed presidents only over the entire sample. Column (iv) reports the results with all FOMC members prior to November 1993 and regional Fed presidents only afterwards. In parentheses we report the standard error based on two-way clustering by both member and meeting.

	All Members pre-1993 (i)	Regional Pres. Only Full Sample (ii)	Regional Pres. Only pre-1993 (iii)	Mixed Members Full Sample (iv)
Expr.-Based Fcst.	282.5	403.4	498.4	288.7
	(85.8)	(107.3)	(133.9)	(76.7)
Wallich Dummy	1.4	-	-	1.5
	(0.4)	-	-	(0.4)
Meeting FE	Yes	Yes	Yes	Yes
Observations	2700	2046	1238	3508
Pseudo R^2	35.3%	45.0%	50.5%	36.6%
APE of Expr.-Based Fcst.:				
Dovish Dissent	-13.0	- 7.7	-9.8	-11.5
Consent	-6.9	-24.5	-24.3	-10.2
Hawkish Dissent	19.9	32.2	34.2	21.7
APE of Wallich Dummy:				
Dovish Dissent	-0.065	-	-	-0.058
Consent	-0.035	-	-	-0.052
Hawkish Dissent	0.099	-	-	0.110

Table C.7: Experience-based Inflation Forecast and FOMC Voting Behavior: Varying Weights on Past Experience and Only with Members who were Born after 1913

The sample period is from March 8, 1951 to January 29, 2014. The sample excludes FOMC Members who were born before 1913. The ordered probit specification is the same as in column (i) of Table 2.3, but here with different values of the gain parameter θ in the calculation of the experience-based inflation forecast. The *Wallich Dummy* equals one if the member is Henry Wallich; 0 otherwise. The average partial effects (APE) reported at the bottom of the table are calculated by taking the partial derivative of the probability of a given voting category with respect to the experience-based inflation forecast at each sample observation and then averaging these partial derivatives across the whole sample. We assume that the ordered probit thresholds depend on a) whether the member is a board member or regional president, and b) whether the meeting occurs after Nov. 1993 and the interaction of a) and b). In parentheses we report the standard error based on two-way clustering by both member and meeting.

	$\theta = 3.334$	$\theta = 2$	$\theta = 2.5$	$\theta = 3.5$	$\theta = 4$
	(i)	(ii)	(iii)	(iv)	(v)
Experience-Based Forecast	246.9	150.5	231.5	230.6	182.5
	(71.3)	(68.0)	(76.5)	(69.6)	(60.6)
Wallich Dummy	1.4	1.4	1.4	1.4	1.4
	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)
Meeting FE	Yes	Yes	Yes	Yes	Yes
Observations	4284	4284	4284	4284	4284
Pseudo R^2	38.1%	37.7%	38.0%	38.1%	38.0%
APE of Experienced-Based Forecast					
Dovish Dissent	-8.5	-5.2	-8.0	-7.9	-6.3
Consent	-7.3	-4.5	-6.8	-6.8	-5.4
Hawkish Dissent	15.7	9.7	14.8	14.7	11.7
APE of Wallich Dummy					
Dovish Dissent	-0.048	-0.049	-0.049	-0.048	-0.049
Consent	-0.041	-0.042	-0.041	-0.042	-0.042
Hawkish Dissent	0.089	0.091	0.090	0.090	0.091

The Tone of FOMC Members' Speeches The following table replicates the results in Table 2.7 with an focus on FOMC members who were born after 1913.

Table C.8: Experience-based Inflation Forecasts and FOMC Members' Tone of Speeches: Only with Members who were Born after 1913

The sample excludes FOMC Members who were born before 1913. Dependent variable is the $NetIndex$ measure of speech hawkishness defined as in equation (2.14). The experience-based inflation forecast for each member at each meeting is calculated as in Table 2.3. All estimations include the same controls and interactions with recent CPI inflation and unemployment as in Table 2.3. In addition, we include controls for education and professional background as explained in the text, except for columns (iii) and (vi) where we instead employ member fixed effects. In columns (ii) and (v), we drop speeches from chairmen. The regressions are estimated with OLS. Standard errors, shown in parentheses, are calculated allowing for two-way clustering by FOMC member and year-quarter.

	Net Index excluding (un)empl.			Net Index including (un)empl.		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Experience-Based Fcst.	41.13	55.11	47.84	44.02	61.90	51.38
	(17.91)	(22.83)	(19.30)	(16.07)	(20.46)	(17.30)
Wallich dummy	0.14	0.13	-	0.16	0.14	-
	(0.11)	(0.12)	-	(0.08)	(0.09)	-
Member FE	No	No	Yes	No	No	Yes
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Chair's speeches dropped	No	Yes	No	No	Yes	No
Industry expr. controls	Yes	Yes	No	Yes	Yes	No
Degree controls	Yes	Yes	No	Yes	Yes	No
Adjusted R^2	3.7%	4.2%	4.6%	3.2%	3.5%	3.6%
Observations	3519	2639	3519	3519	2639	3519

C.8 Target Federal Funds Rate Regressions with Median and Chair's Experience Measures

The results on experience effects on the fed funds rate target in Table 2.8 use a measure of mean experiences across FOMC members. To address the concern that committee decisions do not necessarily reflect the average opinion of the committee's members, we show that our results are robust to using the median or the chairman's experience-based forecast, rather than the average. We also note that the concern is immaterial in our application as the difference between the average experience-based forecast at a meeting and the conventional, objective inflation-rate component of the Taylor rule tends to be substantially bigger than the differences between FOMC members. As a result, it does not matter much whether we use the average, the median, or even any specific FOMC member's experience-based forecast.

In columns (i) and (ii) of Table C.9, we use the median, and in columns (iii) and (iv) the chairman's experience-based forecast. As the table show, these changes result in only minor changes in the coefficient estimate compared with Table 2.8. The same is true when we add the lagged federal funds rate in columns (v) to (viii). The reason is that the time-series variation in the members' experience-based forecasts relative to the staff forecast is much greater than the dispersion between members' experience-based forecasts. These results imply that it does not matter much which measure of central tendency of the experience-based forecasts, or which individual experience-based forecast is used.

Table C.9: Influence of FOMC Members' Inflation Experiences on Target Federal Funds Rate: Median and Chair's Experienced Inflation

The sample period is from the 8/18/1987 to 6/28/2007. The dependent variable is the target federal funds rate set at the FOMC meeting closest to the middle of the quarter t. The experience-based forecast is the median (chair's) experienced-based CPI forecast from quarter $t-1$ to quarter $t+3$ at each meeting. The staff's core inflation forecast is from quarter $t-1$ to quarter $t+3$ and represents the core CPI before 2/1/2000 and the core PCE thereafter. The staff's output gap forecast at quarter t is the forecast for quarter $t+3$. The staff's forecasts of CPI/PCE and of the output gap are from the Philadelphia Fed Greenbook data set. Lagged fed funds rate target is the federal funds rate target from the previous quarter. Columns (i) to (iii) report the OLS coefficient estimates for the estimating equation in (2.15). Columns (iv) and (v) report the estimates of c, β_e, β_π, β_y, and ρ from non-linear least-squares regressions as specified in (2.18). In parentheses, we report Newey-West standard errors with six lags from column (i) to (iii), and zero lags in column (iv) and (v).

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
Exp.-based infl. fcst. (median)	0.39	0.62	-	-	0.47	0.46	-	-
	(0.21)	(0.24)	-	-	(0.21)	(0.21)	-	-
Exp.-based infl. fcst. (chair)	-	-	0.40	0.63	-	-	0.47	0.45
	-	-	(0.22)	(0.24)	-	-	(0.21)	(0.21)
Staff's core inflation forecast	1.27	1.44	1.26	1.44	1.26	1.25	1.26	1.25
	(0.23)	(0.23)	(0.23)	(0.23)	(0.17)	(0.20)	(0.17)	(0.20)
Staff's output gap forecast	0.69	0.46	0.70	0.46	0.98	1.00	0.98	1.00
	(0.06)	(0.10)	(0.06)	(0.10)	(0.07)	(0.15)	(0.07)	(0.15)
Lagged FFR target	-	-	-	-	0.68	0.69	0.68	0.69
	-	-	-	-	(0.04)	(0.04)	(0.04)	(0.04)
Intercept	0.10	2.16	0.10	2.19	-0.03	-0.08	-0.03	-0.08
	(0.35)	(0.86)	(0.36)	(0.86)	(0.16)	(0.42)	(0.16)	(0.42)
Member characteristics	N	Y	N	Y	N	Y	N	Y
Method	OLS	OLS	OLS	OLS	NLS	NLS	NLS	NLS
Observations	80	80	80	80	80	80	80	80
Adjusted R^2	86.6%	87.7%	86.6%	87.8%	97.6%	97.6%	97.6%	97.6%

APPENDIX D

E-trading in Corporate Bond Market: Additional Figures

D.1 Additional Figures

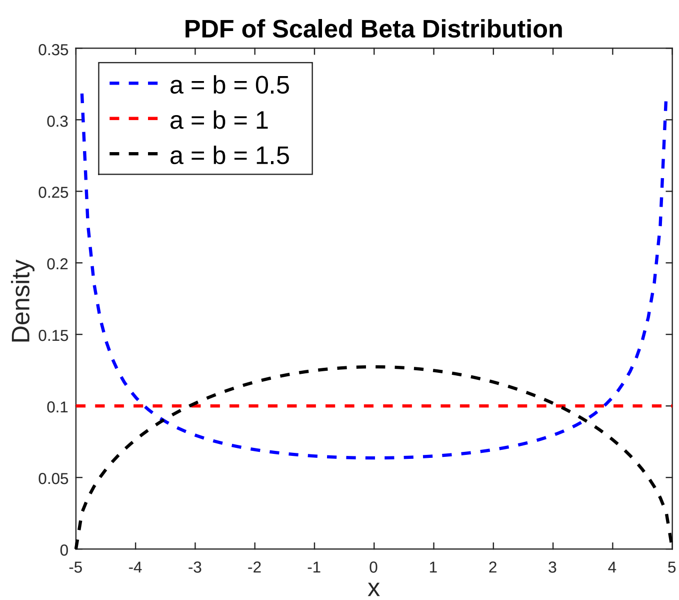

Figure D.1: Probability density function (PDF) of a Scaled Beta Distribution

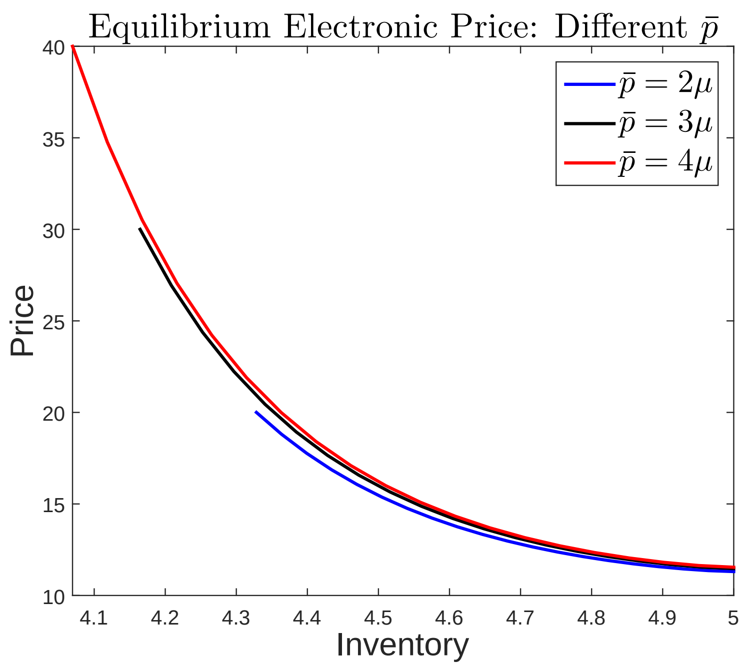

Figure D.2: Dealer's Pricing Strategy in E-trading with respect to Different Investor's Reservation Price $\bar{p}$. Model parameters: $\mu = 10$, $\gamma = 0.5$, $\theta = 3$, $\sigma^2 = 1$, $L = 4$, $R = 5$, $C = 5$, $N = 50$, and I has a scaled beta distribution with $a = b = 1.5$.

APPENDIX E

E-trading in Corporate Bond Market: Proofs and Derivations

E.1 Voice Trading

Following Nash bargaining, both investor and the dealer face the following optimization problem

$$\max_{y,p} \left[U^{MM}_{\text{gain}}\right]^{\eta} \left[U^{Inv}_{\text{gain}}\right]^{1-\eta}$$

where,

$$\begin{aligned} U^{MM}_{\text{gain}} &= y(p-\mu) - \frac{1}{2}\gamma\sigma^2 y(y-2I) \\ U^{Inv}_{\text{gain}} &= y(\mu-p) - \frac{1}{2}\theta\sigma^2 y(y+2e) \end{aligned}$$

Taking f.o.c with respect to p and y yields,

$$\frac{\eta}{1-\eta} = \frac{U^{MM}_{\text{gain}}}{U^{Inv}_{\text{gain}}} \tag{E.1}$$

$$\eta\left[p-\mu-\gamma\sigma^2 y+\gamma\sigma^2 I\right] = [1-\eta]\frac{U^{MM}_{\text{gain}}}{U^{Inv}_{\text{gain}}}\left[p-\mu+\theta\sigma^2 y+\theta\sigma^2 e\right] \tag{E.2}$$

Substituting Eq.(E.1) into Eq.(E.2), we have the optimal trade size as

$$y = \frac{\gamma I - \theta e}{\gamma + \theta} \tag{E.3}$$

Plugging Eq.(E.3) into Eq.(E.1), we obtain the trade price

$$p = \mu - \frac{1}{2}\gamma\sigma^2\Big[2 - \eta - \frac{\gamma}{\gamma+\theta}\Big]I - \frac{1}{2}\theta\sigma^2\Big[\frac{\gamma}{\gamma+\theta} + \eta\Big]e$$

After trading, investor holds $y+e = \frac{\gamma}{\gamma+\theta}[I+e]$ while the dealer holds $I-y = \frac{\theta}{\gamma+\theta}[I+e]$

E.2 Electronic Trading

Throughout, by symmetry, I only consider the case where the investor receives a negative endowment shock, i.e. $e = -L$. I conjecture, in this case, the investor would like to buy from the market maker and will confirm this when I solve for the equilibrium.

Dealer's problem Given the inquiry size y and its inventory level I_j, each dealer chooses the pricing function $p(\cdot)$ to maximize its expected gain to trade, i.e.

$$\max_{p(\cdot)} G(p)\Big[y\big(p - \tilde{p}(I_j, y)\big)\Big]$$

where

$$\tilde{p}(I_j, y) = \mu + \frac{1}{2}\gamma\sigma^2(y - 2I_j)$$

and,

$$G(p) = G(I_j) = \begin{cases} F(\bar{I}^*)^{N-1} & \text{if } I_j \leq \bar{I}^* \\ F(I_j)^{N-1} & \text{if } I_j > \bar{I}^* \end{cases}$$

F.O.C. yields

$$
\begin{aligned}
&\frac{\partial G}{\partial p}\Big[y(p-\tilde{p})\Big] + G(I_j)y = 0 \\
&\frac{\partial G}{\partial I}\frac{\partial I}{\partial p}(p-\tilde{p}) + G(I_j) = 0 \\
&\frac{\partial G}{\partial I}p + \frac{\partial p}{\partial I}G(I_j) = \tilde{p}\frac{\partial G}{\partial I} \\
&\frac{\partial Gp}{\partial I} = \tilde{p}\frac{\partial G}{\partial I}
\end{aligned}
\tag{E.4}
$$

Integrate over $[-R, I_j]$ on both sides of Eq.(E.4), we have

- When $I_j < \bar{I}^*$,

$$
\begin{aligned}
\int_{-R}^{I_j} \frac{\partial Gp}{\partial I} dI &= \int_{-R}^{I_j} \tilde{p}\frac{\partial G}{\partial I} dI \\
F(\bar{I}^*)^{N-1}\Big[p(I_j, y) - p(-R, y)\Big] &= \tilde{p} \times 0 \\
p(I_j, y) &= p(-R, y) \;\forall\; I_j
\end{aligned}
$$

- When $I_j > \bar{I}^*$,

$$
\begin{aligned}
&\int_{-R}^{\bar{I}^*} \frac{\partial Gp}{\partial I} dI + \int_{\bar{I}^*}^{I_j} \frac{\partial Gp}{\partial I} dI = \int_{-R}^{\bar{I}^*} \tilde{p}\frac{\partial G}{\partial I} dI + \int_{\bar{I}^*}^{I_j} \tilde{p}\frac{\partial G}{\partial I} dI \\
&F(\bar{I}^*)^{N-1}\Big[p(\bar{I}^*, y) - p(-R, y)\Big] + \Big[F(I_j)^{N-1}p(I_j, y) - F(\bar{I}^*)^{N-1}p(\bar{I}^*, y)\Big] = 0 + \int_{\bar{I}^*}^{I_j} \tilde{p}\frac{\partial G}{\partial I} dI \\
&\Rightarrow \cdots = \tilde{p}G(\cdot)\big|_{\bar{I}^*}^{I_j} - \int_{\bar{I}^*}^{I_j} G(\cdot) d\tilde{p} \\
&\Rightarrow \cdots = \tilde{p}(I_j, y)F(I_j)^{N-1} - \tilde{p}(\bar{I}^*, y)F(\bar{I}^*)^{N-1} + \gamma\sigma^2 \int_{\bar{I}^*}^{I_j} G(I) dI \\
&\Rightarrow \; p(I_j, y) = \tilde{p}(I_j, y) + \frac{\gamma\sigma^2 \int_{\bar{I}^*}^{I_j} F(I)^{N-1} dI}{F(I_j)^{N-1}} + \left[\frac{F(\bar{I}^*)}{F(I_j)}\right]^{N-1}\Big[p(-R, y) - \tilde{p}(\bar{I}^*, y)\Big]
\end{aligned}
\tag{E.5}
$$

Therefore, we can write down the dealer's expected gain as

$$
\begin{cases}
F(\bar{I}^*)^{N-1}y\Big[p(-R, y) - \tilde{p}(I_j, y)\Big] & \text{if } I_j \leq \bar{I}^* \\
y\Big[\gamma\sigma^2 \int_{\bar{I}^*}^{I_j} F(I)^{N-1} dI + F(\bar{I}^*)^{N-1}\big[p(-R, y) - \tilde{p}(\bar{I}^*, y)\big]\Big] & \text{if } I_j > \bar{I}^*
\end{cases}
$$

It follows that dealer's expected gain is increasing in inventory I_j. Further note that, when $I_j < \bar{I}^*$, the dealer would simply charge the highest possible price, i.e. the reservation price of the investor, since his probability to win the trade does not depend on his quote anymore. That is $p(-R, y) = \bar{p}$.

Given above two observations, dealer's indifference condition for the electronic auction given the participation cost C can be written as

$$F(\bar{I}^*)^{N-1} y \Big[\bar{p} - \tilde{p}(\bar{I}^*, y)\Big] = C \tag{E.6}$$

Note that the LHS of Eq.(E.6) is increasing in I and also $LHS = 0$ when $\bar{I}^* = -R$. Taken together, it implies there must exists some unique $\bar{I}^*$ which satisfy the indifference condition.

Now, let's consider the comparative statics of price p with respect to inventory I and the number of dealers N. First, note that $\frac{\partial p}{\partial I_j} = \frac{\frac{\partial G}{\partial I}(\tilde{p}-p)}{G(I_j)} < 0$ when $I_j > \bar{I}^*$, suggesting the price is decreasing in inventory.

Further, since

$$\frac{\partial p}{\partial N} = \gamma\sigma^2 \int_{\bar{I}^*}^{I_j} \underbrace{log\big(\frac{F(I)}{F(I_j)}\big) e^{(N-1)log(\frac{F(I)}{F(I_j)})}}_{<0} dI + \underbrace{log\big(\frac{F(\bar{I}^*)}{F(I_j)}\big)\big[\frac{F(\bar{I}^*)}{F(I_j)}\big]^{N-1}(\bar{p} - \tilde{p}(\bar{I}^*, y))}_{<0} < 0$$

when $I_j > \bar{I}^*$, the price is decreasing in the number of dealers. As N goes to infinity, p converges to $\tilde{p}(I_j, y)$.

Finally, when $C = 0$, there would be no participation threshold, i.e. $\bar{I}^* = -R$. Consequently, the pricing rule, i.e., Eq. (E.5), would collapse to

$$p(I_j, y) = \tilde{p}(I_j, y) + \gamma\sigma^2 \frac{\int_{-R}^{I_j} F(I)^{N-1} dI}{F(I_j)^{N-1}}$$

This is exactly the Biais (1993) result.

Investor's problem Knowing the dealer's pricing function, the investor faces the following problem

$$\max_y \mathbb{E}\Big[\Big(1 - F(\bar{I}^*)^N\Big)\Big(y\big(\mu - p(I_{max}, y)\big) - \frac{1}{2}\theta\sigma^2 y(y + 2e)\Big)\Big]$$

F.O.C is

$$\frac{\partial(1-F(\bar{I}^*)^N)}{\partial y}\mathbb{E}\Big[y(\mu-p(I_{max}))-\frac{1}{2}\theta\sigma^2 y(y+2e)\Big]$$
$$+(1-F(\bar{I}^*)^N)\mathbb{E}\Big[\mu-p(I_{max})-(\frac{\partial p}{\partial y}+\theta\sigma^2)y-\theta\sigma^2 e\Big]=0 \tag{E.7}$$

where,

$$\frac{\partial(1-F(\bar{I}^*)^N)}{\partial y}=\frac{\partial(1-F(\bar{I}^*)^N)}{\partial\bar{I}^*}\frac{\partial\bar{I}^*}{\partial y}=-NF(\bar{I}^*)^{N-1}f(\bar{I}^*)\frac{\partial\bar{I}^*}{\partial y}$$

$\frac{\partial\bar{I}^*}{\partial y}$ can be computed applying the implicit function theorem with respect to the dealer's indifference condition E.6, i.e.

$$\frac{\partial\bar{I}^*}{\partial y}=-\frac{F(\bar{I}^*)^N\Big[\bar{p}-\mu-\gamma\sigma^2(y-\bar{I}^*)\Big]}{(N-1)f(\bar{I}^*)C+F(\bar{I}^*)^N\gamma\sigma^2 y}$$

and from the pricing rule, we have

$$\frac{\partial p}{\partial y}=\frac{1}{2}\gamma\sigma^2\Big[1-\Big(\frac{F(\bar{I}^*)}{F(I_{\max})}\Big)^{N-1}\Big]$$

The expectation of Eq. (E.7) operates on the conditional distribution of $I_{\max}$ given $I_{\max}>\bar{I}^*$. Its CDF and PDF are given as follows

$$Pr(I_{\max}\le x|I_{\max}>\bar{I}^*)=\frac{Pr(I_{\max}\le x,I_{\max}>\bar{I}^*)}{Pr(I_{\max}>\bar{I}^*)}=\frac{\big[F(x)^N-F(\bar{I}^*)^N\big]}{\big[1-F(\bar{I}^*)^N\big]}$$
$$f(I_{\max}\le x|I_{\max}>\bar{I}^*)=\frac{Nf(x)F(x)^{N-1}}{\big[1-F(\bar{I}^*)^N\big]}$$

Put together, we can now solve for y and $\bar{I}^*$ jointly with two equations, i.e. the dealer's indifference condition, Eq. (E.6), and the F.O.C of the investor problem, Eq. (E.7).

E.2.0.1 Voice vs. Electronic Trading

Consider a continuum of investors characterized by their relative bargaining power $\delta\equiv 1-\eta\in[0,1]$. Further assume δ follows some exogenous distribution $H(\cdot)$.

We first consider the investor choice between electronic trading and voice one.

Investor's gain from the voice trading can be written as

$$V_{\text{voice}}(\delta) = \frac{\delta\sigma^2}{2(\gamma+\theta)}\mathbb{E}\big[(\gamma I + \theta L)^2\big]$$

Apparently, investor's gain from voice market is increasing in his bargaining power δ. In the meantime, investor's gain from the electronic auction is given by

$$V_{\text{elec}} = \Big[1 - F(\bar{I}^*)^N\Big] y_{\text{elec}} \mathbb{E}\Big[\mu - p_{\text{elec}} - \frac{1}{2}\theta\sigma^2(y_{\text{elec}} - 2L)\Big]$$

where $\mathbb{E}[p_{\text{elec}}]$ follows the pricing rule Eq. (E.5) and operates under the conditional distribution of $I_{\max}$ given $I_{\max} > \bar{I}^*$.

If $V_{\text{voice}}(1) > V_{\text{elec}}$, there must exist some δ^* such that $V_{\text{voice}}(\delta^*) > V_{\text{elec}}$. Hence, investor's participation rate of electronic trading is $H(\delta^*)$. Otherwise, if $V_{\text{voice}}(1) \leq V_{\text{elec}}$, the participation rate of electronic trading is simply 1.

Further, we compute the market share of electronic trading λ as follow:

$$\lambda = \frac{\text{Vol}_{\text{elec}}}{\text{Vol}_{\text{voice}} + \text{Vol}_{\text{elec}}} \tag{E.8}$$

where,

$$\begin{aligned}\text{Vol}_{\text{voice}} &= \big(1 - H(\delta^*)\big)\mathbb{E}(y_{\text{voice}}) \\ \text{Vol}_{\text{elec}} &= H(\delta^*)\big[1 - F(\bar{I}^*)^N\big] y_{\text{elec}}\end{aligned}$$

Finally, we consider the welfare implication for dealers with different inventory positions regarding the addition of electronic trading.

Dealer's gain from the "voice trading only" is given by

$$U_{\text{voice}} = \frac{1}{N}\frac{\mathbb{E}(1-\delta)\sigma^2}{2(\gamma+\theta)}\Big[\gamma I_j + \theta L\Big]^2$$

In the meantime, the dealer's gain from "voice + electronic trading" is

$$\begin{aligned}U_{\text{both}} = &H(\delta^*)\Big[max\big(F(I_j)^{N-1} y_{\text{elec}}\big(p_{\text{elec}} - \tilde{p}\big) - C, 0\big)\Big] \\ &+ \big(1 - H(\delta^*)\big)\Big[\frac{1}{N}\frac{\mathbb{E}(1-\delta|\delta > \delta^*)\sigma^2}{2(\gamma+\theta)}\big[\gamma I_j + \theta L\big]^2\Big]\end{aligned}$$

If $I_j < \bar{I}^*$, the dealer would choose not to participate in the electronic auction.

Comparing U_{voice} and U_{both}, we have

$$
\begin{aligned}
U_{\text{both}} =& \big(1 - H(\delta^*)\big)\Big[\frac{1}{N}\frac{\mathbb{E}(1-\delta|\delta>\delta^*)\sigma^2}{2(\gamma+\theta)}\big[\gamma I_j + \theta L\big]^2\Big] \\
&< \frac{1}{N}\frac{\mathbb{E}(1-\delta|\delta>\delta^*)\sigma^2}{2(\gamma+\theta)}\big[\gamma I_j + \theta L\big]^2 < \frac{1}{N}\frac{\mathbb{E}(1-\delta)\sigma^2}{2(\gamma+\theta)}\Big[\gamma I_j + \theta L\Big]^2 = U_{\text{voice}}
\end{aligned}
$$

That is, when $I_j < \bar{I}^*$, the dealer would always prefer the "voice trading only".

If $I_j > \bar{I}^*$, $U_{\text{both}} < U_{\text{voice}}$ may still hold for a relatively small I_j. That being said, for dealers with not large enough inventory position, they would still prefer "voice trading only". This pattern may hold true even if there is no participation cost of electronic auction at all, i.e. $C = 0$ (see the numerical example in section 3.4.5).

BIBLIOGRAPHY

BIBLIOGRAPHY

Alesina, A., and N. Fuchs-Schündeln, 2007, "Goodbye Lenin (or Not?): The Effect of Communism on People's Preferences," *American Economic Review*, 97(4), 1507–1528.

Apel, M., and M. B. Grimaldi, 2014, "How Informative Are Central Bank Minutes?," *Review of Economics*, 65(1), 53–76.

Barber, B. M., X. Huang, and T. Odean, 2016, "Which Factors Matter to Investors? Evidence from Mutual Fund Flows," *The Review of Financial Studies*, 29(10), 2600–2642.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer, 2015, "X-CAPM: An extrapolative capital asset pricing model," *Journal of Financial Economics*, 115(1), 1 – 24.

Baxa, J., R. Horváth, and B. Vašíček, 2013, "Time-varying Monetary-Policy Rules and Financial Stress: Does Financial Instability Matter for Monetary Policy?," *Journal of Financial Stability*, 9(1), 117–138.

Berk, J. B., and R. C. Green, 2004, "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 112(6), 1269–1295.

Berk, J. B., and J. H. van Binsbergen, 2015, "Measuring skill in the mutual fund industry," *Journal of Financial Economics*, 118(1), 1 – 20.

——— , 2016, "Assessing asset pricing models using revealed preference," *Journal of Financial Economics*, 119(1), 1 – 23.

Bernanke, B. S., 2010, "Monetary Policy and the Housing Bubble," Speech at the Annual Meeting of the American Economic Association.

Biais, B., 1993, "Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets," *The Journal of Finance*, 48(1), 157–185.

Botsch, M. J., and U. Malmendier, 2016, "Inflation Experiences and Mortgage Choice – Evidence from Residential Mortgages," Working Paper.

Box, G. E., G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, 2015, *Time Series Analysis: Forecasting and Control*. John Wiley & Sons.

Bray, M., 1982, "Learning, Estimation and the Stability of Rational Expectations Equilibria," *Journal of Economic Theory*, 26, 318–339.

Brown, D. P., and Y. Wu, 2016, "Mutual Fund Flows and Cross-Fund Learning within Families," *The Journal of Finance*, 71(1), 383–424.

Bryan, M. F., and S. G. Cecchetti, 1995, "The Seasonality of Consumer Prices," *Federal Reserve Bank of Cleveland Economic Review*, (Q II), 12–23.

Busse, J. A., T. Chordia, L. Jiang, and Y. Tang, 2018, "Mutual Fund Trading Costs," working paper.

Chappell, H. W., T. M. Havrilesky, and R. R. McGregor, 1993, "Partisan Monetary Policies: Presidential Influence Through the Power of Appointment," *Quarterly Journal of Economics*, pp. 185–218.

——— , 1995, "Policymakers, Institutions, and Central Bank Decisions," *Journal of Economics and Business*, 47(2), 113–136.

Chappell, H. W., and R. R. McGregor, 2000, "A Long History of FOMC Voting Behavior," *Southern Economic Journal*, pp. 906–922.

Chappell, H. W., R. R. McGregor, and T. A. Vermilyea, 2005, *Committee Decisions on Monetary Policy: Evidence from Historical Records of the Federal Open Market Committee*. MIT Press.

Chevalier, J., and G. Ellison, 1997, "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, 105(6), 1167–1200.

Chiang, Y.-M., D. Hirshleifer, Y. Qian, and A. E. Sherman, 2011, "Do Investors Learn from Experience? Evidence from Frequent IPO Investors," *Review of Financial Studies*, 24, 1560–1589.

Cho, I.-K., N. Williams, and T. J. Sargent, 2002, "Escaping Nash Inflation," *Review of Economic Studies*, 69(1), 1–40.

Choi, D., B. Kahraman, and A. Mukherjee, 2016, "Learning about Mutual Fund Managers," *The Journal of Finance*, 71(6), 2809–2860.

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

Coibion, O., and Y. Gorodnichenko, 2012, "Why Are Target Interest Rate Changes so Persistent?," *American Economic Journal. Macroeconomics*, 4(4), 126.

De Frutos, M., and C. Manzano, 2002, "Risk aversion, transparency, and market performance," *The Journal of Finance*, 57(2), 959–984.

Desgranges, G., and T. Foucault, 2005, "Reputation-based pricing and price improvements," *Journal of Economics and Business*, 57(6), 493–527.

Dick-Nielsen, J., and M. Rossi, 2016, "The Cost of Immediacy for Corporate Bonds," *Working paper*.

Dyck, A., and L. Pomorski, 2014, "Investor Scale and Performance in Private Equity Investments," working paper.

Elton, E. J., M. J. Gruber, and C. R. Blake, 2001, "A First Look at the Accuracy of the CRSP Mutual Fund Database and a Comparison of the CRSP and Morningstar Mutual Fund Databases," *The Journal of Finance*, 56(6), 2415–2430.

Evans, G. W., and S. Honkapohja, 2001, *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

Fleming, M. J., B. Mizrach, and G. Nguyen, 2017, "The microstructure of a U.S. Treasury ECN: The BrokerTec platform," *Journal of Financial Markets*.

Franzoni, F., and M. C. Schmalz, 2017, "Fund Flows and Market States," *The Review of Financial Studies*, 30(8), 2621–2673.

Frazzini, A., R. Israel, and T. Moskowitz, 2018, "Trading Costs," Working paper.

Friewald, N., and F. Nagler, 2016, "Dealer inventory and the cross-section of corporate bond returns," *Working paper, Available at SSRN 2526291*.

Fuchs-Schündeln, N., and M. Schündeln, 2015, "On the endogeneity of political preferences: Evidence from individual experience with democracy," *Science*, 347(6226), 1145–1148.

Gallagher, J., 2014, "Learning About an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States," *American Economic Journal: Applied Economics*, 6(3), 206–233.

Goldstein, I., H. Jiang, and D. T. Ng, 2017, "Investor flows and fragility in corporate bond funds," *Journal of Financial Economics*, 126(3), 592 – 613.

Gospodinov, N., and B. Wei, 2015, "A Note on Extracting Inflation Expectations from Market Prices of TIPS and Inflation Derivatives," working paper, Federal Reserve Bank of Atlanta.

Greenwich Associates, 2015, "The Continuing Corporate Bond Evolution," White Paper.

Greenwood, R., and S. Nagel, 2009, "Inexperienced Investors and Bubbles," *Journal of Financial Economics*, 93, 239–258.

Grossman, S. J., and M. H. Miller, 1988, "Liquidity and market structure," *the Journal of Finance*, 43(3), 617–633.

Hansen, S., and M. McMahon, 2016a, "First Impressions Matter: Signalling as a Source of Policy Dynamics," *Review of Economic Studies*, p. forthcoming.

——— , 2016b, "Shocking language: Understanding the macroeconomic effects of central bank communication," *Journal of International Economics*, 99, Supplement 1, S114 – S133.

Harris, L., 2015, "Transaction Costs, Trade Throughs, and Riskless Principal Trading in Corporate Bond Markets," *Working paper*.

Harvey, C. R., and Y. Liu, 2017, "Decreasing Returns to Scale, Fund Flows, and Performance," working paper.

Hasbrouck, J., 2009, "Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data," *The Journal of Finance*, 64(3), 1445–1477.

Hendershott, T., and A. Madhavan, 2015, "Click or Call? Auction versus Search in the Over-the-Counter Market," *The Journal of Finance*, 70(1), 419–447.

Hendershott, T., and H. Mendelson, 2000, "Crossing networks and dealer markets: competition and performance," *The Journal of Finance*, 55(5), 2071–2115.

Huang, R. D., and H. R. Stoll, 1996, "Dealer versus auction markets: A paired comparison of execution costs on {NASDAQ} and the {NYSE}," *Journal of Financial Economics*, 41(3), 313 – 357.

Kahn, L., 2010, "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," *Labor Economics*, 17, 303–3016.

Kaustia, M., and S. Knüpfer, 2008, "Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions," *Journal of Finance*, 63, 2679–2702.

Koudijs, P., and H.-J. Voth, 2016, "Leverage and Beliefs: Personal Experience and Risk-Taking in Margin Lending," *The American Economic Review*, 106(11), 3367–3400.

Laudenbach, C., U. Malmendier, and A. Niessen-Ruenzi, 2018, "The Long-lasting Effects of Experiencing Communism on Financial Risk-Taking," Working Paper.

Lichter, A., M. Löffler, and S. Siegloch, 2016, "The Long-Term Costs of Government Surveillance: Insights from Stasi Spying in East Germany," Working Paper.

Lindsey, D. E., 2003, "A Modern History of FOMC Communication: 1975-2002," working paper, Federal Reserve Board.

Lucca, D. O., and F. Trebbi, 2011, "Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements," Working Paper 15367, National Bureau of Economic Research.

Malmendier, U., and S. Nagel, 2011, "Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?," *Quarterly Journal of Economics*, 126, 373–416.

——— , 2016, "Learning from Inflation Experiences," *Quarterly Journal of Economics*, 131(1), 53–87.

Malmendier, U., and L. S. Shen, 2017, "Scarred Consumption," Working Paper.

Malmendier, U., and A. Steiny, 2017, "Buy or Rent? The Role of Lifetime Experiences of Macroeconomic Shocks within and across Countries," Working Paper.

Malmendier, U., and G. Tate, 2005, "CEO Overconfidence and Corporate Investment," *Journal of Finance*, 60(6), 2661–2700.

Malmendier, U., G. Tate, and J. Yan, 2011, "Overconfidence and Early-life Experiences: The Effect of Managerial Traits on Corporate Financial Policies," *Journal of Finance*, 66(5), 1687–1733.

Marcet, A., and J. Sargent, Thomas, 1989, "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

Markets Committee, 2016, "Electronic trading in fixed income markets," *Bank for International Settlements, Markets Committee Publications*, (7).

McLemore, P., 2018, "Do Mutual Funds Have Decreasing Returns to Scale? Evidence from Fund Mergers," *Journal of Financial and Quantitative Analysis*, p. forthcoming.

Meade, E. E., and D. Stasavage, 2008, "Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve," *Economic Journal*, 118(528), 695–717.

Mehra, Y. P., and B. Sawhney, 2010, "Inflation Measure, Taylor Rules, and the Greenspan-Bernanke Years," *FRB Richmond Economic Quarterly*, 96(2), 123–151.

Mishkin, F. S., 2010, "Monetary Policy Flexibility, Risk Management, and Financial Disruptions," *Journal of Asian Economics*, 21(3), 242–246.

Oreopoulos, P., T. von Wachter, and A. Heisz, 2012, "The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates," *American Economic Journal: Applied Economics*, (4), 1–29.

Orphanides, A., 2001, "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, 91(4), 964–9999.

———, 2003, "Historical Monetary Policy Analysis and the Taylor Rule," *Journal of Monetary Economics*, 50(5), 983–1022.

Pastor, L., R. F. Stambaugh, and L. A. Taylor, 2015, "Scale and skill in active management," *Journal of Financial Economics*, 116(1), 23–45.

———, 2017, "Fund Tradeoffs," working paper.

Primiceri, G. E., 2006, "Why Inflation Rose and Fell: Policy-Makers' Beliefs and US Postwar Stabilization Policy," *Quarterly Journal of Economics*, 121(3), 867–901.

Randall, O., 2015, "Pricing and Liquidity in Over-The-Counter Markets," *Working Paper*.

Reis, R., 2013, "Central Bank Design," *Journal of Economic Perspectives*, 27(4), 17–43.

Romer, C. D., and D. H. Romer, 2004, "Choosing the Federal Reserve Chair: Lessons from History," *The Journal of Economic Perspectives*, 18(1), 129–162.

———, 2008, "The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?," *The American Economic Review*, pp. 230–235.

Romer, D. H., 2010, "A New Data Set on Monetary Policy: The Economic Forecasts of Individual Members of the FOMC," *Journal of Money, Credit and Banking*, 42(5), 951–957.

Rotemberg, J. J., and M. Woodford, 1999, "Interest Rate Rules in an Estimated Sticky Price Model," in *Monetary Policy Rules*, ed. by J. B. Taylor. University of Chicago Press.

Sargent, T. J., 1993, *Bounded Rationality in Macroeconomics*. Clarendon Press, Oxford, UK.

———, 1999, *The Conquest of American Inflation*. Princeton University Press, Princeton, NJ.

Seppi, D. J., 1990, "Equilibrium block trading and asymmetric information," *The Journal of Finance*, 45(1), 73–94.

Shiller, R. J., 2005, *Irrational Exuberance*. Princeton University Press, Princeton, NJ.

Sibert, A., 2006, "Central Banking by Committee," *International Finance*, 9(2), 145–168.

Stambaugh, R. F., 1999, "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.

Strahilevitz, M., T. Odean, and B. Barber, 2011, "Once Burned, Twice Shy: How Naïve Learning, Counterfactuals, and Regret Affect the Repurchase of Stocks Previously Sold," *Journal of Marketing Research*, 48, 102–120.

Taylor, J. B., 1993, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Thornton, D. L., and D. C. Wheelock, 2014, "Making Sense of Dissents: A History of FOMC Dissents," *Federal Reserve Bank of St. Louis Review*, 96, 213–227.

Vissing-Jorgensen, A., 2003, "Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions," in *NBER Macroeconomics Annual*.

Williams, R., 2006, "Generalized Ordered Logit/Partial Proportional Odds Models for Ordinal Dependent Variables," *Stata Journal*, 6(1), 58–82.

Yin, X., 2005, "A comparison of centralized and fragmented markets with costly search," *The Journal of Finance*, 60(3), 1567–1590.

Zhu, H., 2014, "Do Dark Pools Harm Price Discovery?," *Review of Financial Studies*, 27(3), 747–789.

Zhu, M., 2018, "Informative fund size, managerial skill, and investor rationality," *Journal of Financial Economics*, 130(1), 114 – 134.