Forward Guidance and Expectation Formation: Cross-Country Evidence from Survey Data

Christopher S. Sutherland[1]

[1]*Bank of Canada, Ottawa, Ontario, Canada*

January 27, 2019

Abstract

How forward guidance and its attributes influence forecaster expectations is not yet fully understood. Two main challenges explain this gap in the literature. First, measuring forward guidance is difficult. So, I construct a new central bank data set and then combine it with survey data that tracks individual forecasters over time. In doing so, I propose a methodology to categorize forward guidance. I argue that forward guidance is multifaceted and therefore has overlapping *attributes* rather than mutually exclusive *classes*. One attribute, Delphic (predictive) forward guidance, characterizes the vast majority of observed cases. Odyssean (promissory) forward guidance, by contrast, has only been observed on two occasions. Second, identification is difficult. Extending the approach of Altavilla and Giannone (2017), I find that forward guidance causes private-sector forecasters to revise their one-year outlook for the policy rate in the intended direction by about five basis points on average. Further, the commitment that accompanies Odyssean forward guidance may be the only way to amplify its influence. Indeed, neither quantitative easing announcements nor the release of central bank policy rate projections appeared to augment the effectiveness of forward guidance. In fact, another attribute, state-contingent forward guidance, actually appears to dampen effectiveness compared to time-contingent forward guidance.

JEL Classification: D83, E37, E52, E58.
Keywords: forward guidance, central bank communication, expectations, survey data, signalling, applied econometrics.

Acknowledgments

This paper was made possible with the help of excellent research assistance from my colleagues at the Bank of Canada (Liam Lindsay, Francine Saumur), Bank of England, European Central Bank, Norges Bank (Cathrine Fiva Waller), Reserve Bank of Australia (Erin Walker), Reserve Bank of New Zealand (Raewyn Peters), Sveriges Riksbank. I am very grateful for comments and suggestions from my colleagues at the Bank of Canada, Bruno Feunou, Jean-Sébastien Fontaine, Toni Gravelle, Monica Jain, Rodrigo Sekkel, and Jonathan Witmer.

1 Introduction

Numerous event studies have found that forward guidance in a particular country affects market interest rates in that country (e.g. Gürkaynak et al. (2005), Moessner (2013), Swanson (2017), Hubert and Labondance (2018b)). However, no study has attempted to estimate the response of private-sector policy rate forecasts to forward guidance. Some studies have focused on the effect of forward guidance on forecaster expectations (e.g. Campbell et al. (2012), Kool and Thornton (2015), Andrade et al. (2019)), but such studies focus on forecaster disagreement and forecast accuracy rather than the extent to which forecasters revised their outlooks in the intended direction. Fewer papers still have investigated how various features of forward guidance influence expectation formation (Coenen et al. (2017), Jain and Sutherland (2018)). Two main reasons explain these gaps in the literature.

The first is an econometric challenge. Identification is very difficult due to the endogeneity of forward guidance to policy rate forecasts. That is, the barrier to estimating the effects of forward guidance is to disentangle its effects from those of other macroeconomic trends and policy interventions (Besley et al. (2008), Dell'Ariccia et al. (2018)). Accordingly, most papers that study forward guidance typically take a high-frequency event study approach using hourly or minutely market interest rate data. Although indispensable as an identification strategy, this approach has a number of drawbacks. First, it is not able to say much about the persistence of the effects. Second, by focusing on the narrow windows around policy interventions, these studies are unable to account for the fact that many central bank policy actions (policy rate changes, forward guidance changes, asset purchase program changes, etc.) have already become priced by the markets. As such, they measure asset price reactions to the surprise components of policy actions. Third, some results could partially be artifacts of financial markets (e.g. risk sentiment) and not necessarily reflective of pure expectations. Finally, these studies may be unable to distinguish between macroeconomic signals and monetary signals.

The second challenge is a data gap. To provide estimates about the effects of forward guidance more generally would require a large forward guidance data set spanning multiple years, and ideally, countries rather than just isolated episodes as reflected in many of the event studies found in the literature. Yet it is not enough to simply record all the dates of forward guidance episodes by country. We must also consider the valence (dovishness or hawkishness) of forward guidance if we are to measure whether such guidance moved policy rate expectations in the intended direction. Further, we must consider the effects of other central bank policies such as policy rate changes and asset purchase programs if we are to isolate the effects of forward guidance. Such data would afford also us a great deal of heterogeneity, which would help with the identification issues discussed above. It would also be desirable to understand the effects of different attributes of forward guidance. To date, no such data set has been available however. So, to conduct this analysis, I first construct a new central bank data set and then combine it with carefully assembled private-sector survey data.

The central bank data set records the forward guidance, quantitative easing, and policy rate changes of eight inflation-targeting central banks over the past thirty years. I took a narrative approach to forward guidance classification. I gathered and read about 30 years of monetary policy statements across eight central banks and, inter alia, identified all *changes* in forward guidance. These changes were categorized as either dovish (-1), neutral (0), or hawkish ($+1$). I am able to identify about 500 total months when forward guidance was used (totalled across eight central banks). I use a similar approach to record central banks' use of quantitative easing. I also gather *individual* forecasts and their full set of domestic macroeconomic expectations over time. This allows me to control for each individual's macroeconomic outlook and identify the effects of forward guidance, quantitative easing, and policy rate changes respectively.

Further, to improve our understanding of particular forward guidance attributes, I also record metadata for empirical analysis. Following the literature (Campbell et al. (2012), Andrade et al. (2016), Coenen et al. (2017), Moessner et al. (2017), Andrade et al. (2019), Jain and Sutherland (2018)), I record whether the forward guidance was either state-contingent, time-contingent, or qualitative and whether it was either "Delphic" (predictive) or "Odyssean" (promissory). To do so, I propose a methodology to categorize forward guidance. The methodology is intended to standardize how we describe and conceptualize forward guidance and to resolve some lingering ambiguity in the literature about the attributes of forward guidance central banks use. I argue that forward guidance can seldom be sorted into one mutually exclusive *class* but instead that it typically has one or many overlapping *attributes* and that some classifications in the literature are really attribute subsets.

The identification strategy in this paper is multifaceted. First, identifying so many examples of forward guidance and quantitative easing, which is best accomplished by taking a cross-country approach, allows me to overcome the challenge that forward guidance and quantitative easing often happen simultaneously and to isolate numerous cases of each. Second, the frequency of the survey data is monthly so the problem of anticipated forward guidance is, although not eliminated, mitigated substantially. Third, by limiting the analysis to inflation-targeting central banks and periods, I am able to make the identifying assumption that no other time-varying confounder exists outside the framework used to conduct this analysis. In doing so, I am able to adjust for the macroeconomic shocks that could confound the relationship between forward guidance changes and policy rate forecast revisions without introducing simultaneity bias. To do so, I follow Altavilla and Giannone (2017) who use individual revisions to current-year macroeconomic forecasts as a control variable.

The main finding is that the average revision to forecasts of the policy rate in twelve months' time in response to forward guidance is about five basis points in the intended direction. That is, hawkish or dovish forward guidance tends to move interest rate forecasts in the intended direction (up and down respectively) by about five basis points on average. This result is important for a number of reasons. These estimates provide central banks some guidance as to how much they can expect to influence expectations ex ante before implementing forward guidance. They also provide new evidence of the ability of forward guidance to influence survey expectations in the intended manner. The results also provide additional evidence that private forecasters update their expectations in the manner suggested by models of Bayesian learning (Coibion et al. (2018b)). Interestingly, the estimated average forward guidance effect is surprisingly small given that changes in forward guidance frequently provide material updates to the outlook for the path of policy rates. This suggests that perhaps central bank credibility issues, noise (Sims (2003), Woodford (2003)), or forecaster inattention (Mankiw and Reis (2002), Coibion and Gorodnichenko (2012a), Coibion and Gorodnichenko (2015), Coibion et al. (2018a), Coibion et al. (2018b)) may constrain forward guidance expectation formation but not so much so as to render forward guidance ineffective.

By addressing the data and identification challenges discussed above, I am able to make numerous other novel contributions. First, we can now better understand how and where forward guidance works. To begin with, the influence of forward guidance is much weaker during periods at the effective lower bound. This may help explain why I find that the influence of forward guidance over policy rate expectations has been weaker in recent years. It could also help explain why forward guidance has been the most effective in Canada, Australia, New Zealand, and Sweden—countries that started using forward guidance early in their histories. I also show that forward guidance can have spillover effects. For example, Federal Reserve forward guidance has had strong spillover effects to Canadian policy rate expectations.

Second, the only way to amplify forward guidance appears to be to make a commitment. Surprisingly, I find that quantitative easing announcements do not appear to amplify forward guidance. Similarly, when central banks released both forward guidance and policy rate forecasts, no amplification of the influence of forward guidance was detectable. Many papers have advocated for the publication of central bank policy rate forecasts (e.g. Faust and Leeper (2005), Woodford (2005), Svensson (2006), Rudebusch and Williams (2008), Woodford (2013), Svensson (2015)). The efficacy of central bank policy rate forecasts is still an open question, but the results in this paper add to the evidence provided by Jain and Sutherland (2018) that policy rate forecasts do not noticeably influence private-sector forecasters. Interestingly, state-contingent forward guidance actually appears to dampen effectiveness compared to time-contingent forward guidance. One potential explanation for these latter two findings is that the purpose of both state-contingent forward guidance and policy rate projections may be more so to communicate about the reaction function and monetary policy inertia (Coibion and Gorodnichenko (2012b)) than to amplify the influence of forward guidance. As such, studying forecast errors would probably be more instructive about the efficacy of these particular policy tools than analysing forecast revisions.

By contrast, Odyssean forward guidance is extremely rare, but perhaps extremely effective. Indeed, only two instances appear in the sample data that cover eight central banks and thirty years. Accordingly, our ability to study Odyssean forward guidance is limited. Nonetheless, estimates of the effect of Odyssean forward guidance, which, naturally, come with large standard errors, suggest that Odyssean forward guidance greatly amplifies the influence of forward guidance on expectation formation. In fact, it appears to double the effect. This result provides some empirical evidence for the ideas of Krugman et al. (1998) and Eggertsson et al. (2003). These authors suggest that if a central bank can credibly commit to keep the policy rate below the levels suggested by the central bank's reaction function at some point in the future, then the central bank should be able to lower longer-term interest rates due to the expectations hypothesis. Indeed, I find that forward guidance has much larger effects on forecasts of the policy rate at the twelve-month forecast horizon than the three-month forecast horizon.

2 Data

This study combines two types of panel data. First, I use individual private-sector survey forecasts of interest rates, which is the outcome variable of interest, and inflation rates, which is a control variable. Second, I use data on central bank policy signals, which include policy rate decisions, forward guidance, quantitative easing, and for robustness, central bank projections. This paper is concerned with how the central bank policy signals influence the private-sector forecasts of interest rates. The data come from twelve inflation-targeting countries: Australia, Canada, France, Germany, Italy, Netherlands, New Zealand, Norway, Spain, Sweden, United Kingdom, and the United States of America. All data have a monthly frequency as the survey frequency is monthly. The data are constructed such that all central bank policy signals strictly precede all private-sector forecasts in time. The data are depicted in several figures below and summarized in Table 3.

2.1 Private-sector forecast data

The individual private-sector survey forecast data come from Consensus Economics. Each month, private-sector forecasters submit their forecasts for interest rates, inflation, growth, and other macroeconomic data to Consensus Economics. The interest rate forecasts are fixed-horizon forecasts and the other macroeconomic forecasts are fixed-event forecasts. Each forecaster provides his or her forecast for a domestic three-month interest rate (typically the Treasury bill rate) for *both* three months into the future and twelve months into the future. Each forecaster also provides his or her forecast for the inflation rate for *both* the current year and the next year. Similarly, they provide their forecast for the domestic economic growth rate and the unemployment rate for *both* the current year and the next year respectively.

The uncertainty associated with fixed-horizon forecasts is fundamentally different from the uncertainty associated with fixed-event forecasts. The uncertainty associated with each fixed-horizon interest rate forecast is, all else equal, constant as each forecast looks either three or twelve months into the future. Conversely, the uncertainty associated with fixed-event inflation forecasts (i.e. the inflation rate for that year) should, all else equal, monotonically decrease as each month in the calendar year passes. Forecast uncertainty should be highest in January because the forecaster is either anticipating inflation for all twelve months of the current year or all twelve months of the next year. Forecast uncertainty is lower in July, for example, because the forecaster has already observed as many as six months of realized inflation data and therefore must only forecast six more months for the current year instead of twelve as in January; the forecast for next year's inflation is also a less distant prospect in, say, July than it was January.

As the purpose of this study is to determine how individual forecasters *revise* their forecasts of the policy rate in response to a signal provided by a central bank, we must track individual forecasters over time and calculate the changes in their forecasts on a monthly basis. The fundamental challenge of tracking individual forecasts over time relates to naming conventions. Typically, there is no challenge because each forecaster is recorded using the same name each month. The challenge lies in the exceptions to this rule.

Least seriously, the names of the firms are occasionally misspelled in some months, which if not detected, would cause that monthly observation to be treated as separate from its correctly spelled counterparts from other months. Most seriously, the names of firms change over time. This is usually because of mergers, acquisitions, and bankruptcy. Accordingly, we must track individual forecasters across these events. Doing so is a nontrivial task and is likely one reason why most studies aggregate the data across forecasters. To track individual forecasters over time, I worked with the data provider to track firms through mergers and acquisitions by country across time. This allowed me to consolidate many individual forecaster time series that were seemingly separate into fewer, harmonized individual forecaster time series. The consequence of this consolidation is to decrease the total number of firms in the sample data and to increase the number of observations per firm.

I take the first difference of each monthly observation for each forecaster to obtain forecaster revisions. The calculations for interest rates are trivial because we simply take the first difference of a fixed-horizon forecast. The calculations for inflation forecasts are more complex. In eleven out of the twelve months, the calculation is also simply a first difference. However, in January, the reference points for current-year forecasts and next-year forecasts move forward one year. As discussed in Section 3.1, we are only interested in current-year inflation forecasts. Hence, to calculate January forecast revisions, we can take the January forecast of current-year inflation from year t and subtract the December forecast of next-year inflation from year $t-1$. The result represents how the forecaster has revised his or her outlook for current-year inflation in January of that same year.

2.2 Central bank data

This study uses four types of central bank data: policy rate, forward guidance, and quantitative easing data as well as central bank projections data for robustness. These data take the role of both treatment and control variables in the regressions to follow.

First, daily central bank policy rates are obtained from the Bank for International Settlements. These data are matched to the date of each monthly private-sector forecast. I take the first difference of each central bank policy rate and call this change p_{ct}. Intuitively, had a central bank, for example, raised its policy rate by twenty-five basis points, then one might reasonably expect a private- sector forecaster to revise his or her forecast of the future policy rate commensurately. Of course, any such revision would depend on a number of factors, not least how the central bank's change to its policy rate conformed with the forecaster's expectations. So, p_{ct} serves not only as a control, but as a baseline to approximate the transmission of changes in short-term policy rates to changes in expectations of policy rates in the future.

2.3 Forward guidance data

The second type of central bank data collected was forward guidance data. I took a narrative approach to forward guidance classification. In essence, I read about 30 years of monetary policy statements across eight central banks and, inter alia, and identified all *changes* in forward guidance. These changes were categorized as either dovish (-1), neutral (0), or hawkish ($+1$). This narrative approach is very similar in nature to that used in Istrefi (2016) and a follow-up paper Bordo and Istrefi (2018). "Istrefi (2017) collects the perceptions of Fed watchers and other analysts as reflected in the US media and builds a measure of policy preferences (a hawkdove index) of the FOMC. The narrative record in the media is used as a public source and a filter of all relevant information about these policymakers backgrounds, their political interests and supporters and their economic beliefs. These beliefs are expressed in their writings, testimonies and speeches before joining and during their time at the Fed and in their policy actions (votes and dissents). To build the HawkDove measure, about 20,000 articles or reports, from more than 30 newspapers and business reports of Fed watchers, referencing to 130 FOMC members were consulted" (Bordo and Istrefi (2018)).

The first stage of forward guidance classification for this paper was to track down all monetary policy communication from each central bank. The vast majority of this communication takes the form of press releases announcing and explaining a monetary policy decision. In rare cases, especially in the 1990s, some central banks also included explanations of their monetary policy decisions in their accompanying monetary policy reports. The second stage was to read and analyse this communication to identify forward guidance statements. Naturally, it is therefore crucial to establish some very particular definitions of forward guidance. In this study, forward guidance is defined as a statement that provides direct information about the probable stance of monetary policy in the future.[1]

[1] "Forward guidance in monetary policy means providing some information about future policy settings" (Svensson (2015), page 20). "Communication about the likely future course of monetary policy is known as "forward guidance"" (Federal Reserve, 2018). "forward guidance, which means communicating how the ECB expects its policy measures to evolve in the future and what conditions would warrant a change in the policy stance"

The forward guidance periods and classifications in this paper have also been cross-referenced with the periods and classifications in Gürkaynak et al. (2005), Rudebusch and Williams (2008), Svensson (2010), Woodford (2013), Femia et al. (2013), Kool and Thornton (2015), Svensson (2015), Charbonneau and Rennison (2015), Obstfeld et al. (2016), Coenen et al. (2017), Moessner et al. (2017), Swanson (2017), Bhattarai and Neely (2018), Kuttner (2018), Hubert and Labondance (2018b), Andrade et al. (2019). Almost invariably, the periods in this paper are consistent with those in the papers cited above. Whenever the classifications of forward guidance differed from those in the papers cited above (most papers identify dates rather than types or attributes of forward guidance however), I revisited these classifications and occasionally made adjustments.

It is also important to stress both the depth and the breadth of the data. The data cover eight central banks and span 29 years. The first central bank to begin using forward guidance regularly in this data set was the Reserve Bank of New Zealand on December 1, 1996: "The prospects for further loosening in monetary conditions will depend on how quickly inflation pressures recede." Early forms of forward guidance tended to be somewhat subtler than contemporary forward guidance. Despite the prescience the Reserve Bank of New Zealand, it appears that the first instance of forward guidance in the data actually took place six years earlier in Australia by Governor Bernie Fraser on January 23, 1990: "Decisions in respect of any further easing in monetary policy will be made in the light of developments in the economy, especially demand and wage and price levels." The Reserve Bank of New Zealand, the Norges Bank, and the Sveriges Riksbank have used forward guidance the most often—each beginning in the 1990s. The Bank of England and the European Central Bank have used forward guidance the least frequently in the sample period—neither beginning until well after the onset of the financial crisis. The Bank of Canada, Federal Reserve, and Reserve Bank of Australia lie somewhere in between.

The third stage was to categorize each type of forward guidance. Three types of categorizations (and corresponding quantifications) were made. Forward guidance was labelled as either dovish (-1), neutral (0), or hawkish ($+1$). The terms dovish, neutral, and hawkish forward guidance refer to the future path of monetary policy. Hawkish forward guidance would suggest that the next change in the policy rate is more likely to be up than down. Dovish forward guidance would be the opposite. Neutral forward guidance would suggest that the next change in the policy rate is as likely to be up as it is down.

For example, on October 25, 2012, the Sveriges Riksbank stated "It is now more probable that the repo rate will be cut rather than being raised during the winter," and on May 10, 2018 the Reserve Bank of New Zealand stated that "The direction of our next move is equally balanced, up or down." Generally, each term would correspond to an implicitly downward-sloping, flat, or upward-sloping path of future policy rates respectively. For example, the Federal Reserve's use of the statement "higher interest rates will be warranted" would be categorized as hawkish and implies an upward-sloping policy rate path. Helpfully, the Reserve Bank of New Zealand, the Norges Bank, and the Sveriges Riksbank release projections of the policy rate path, so the slope of these paths are actually explicitly stated rather than only implicitly articulated by the central bank. Once all of the published policy rate paths for each central bank were gathered, I also used these projections to assign forward guidance values.[2] Next, each forward guidance

(European Central Bank, 2018)

[2]To be clear, I consider central bank policy rate projections to be a subset of forward guidance. Each time a central bank policy rate projection is released, that period would automatically be considered to be a period with forward guidance. So a policy rate projection implies forward guidance in the data. Fortunately, this is an inconsequential assumption because central banks that release policy rate projections on a regular basis almost invariably also provide verbal forward guidance as well. Of course, the converse is not true: the provision of

statement was also labelled as either Delphic or Odyssean. Finally, forward guidance was also labelled as either qualitative, time-contingent, or state-contingent (more on these classifications below).

To assign values to forward guidance requires a very systematic approach. First, all monetary policy press releases and monetary policy reports for each country issued between 1990 and 2018 were gathered. In many cases, older press releases and monetary policy reports were not available online so it was necessary to collaborate with the central banks in our sample to collect these documents. Second, I read each statement in chronological order, one country at a time. The purpose of doing so was to keep track of each central bank's evolving narrative and to detect shifts in this narrative, which often correspond to changes in forward guidance. Third, all instances of forward guidance in those statements were identified.

Fourth, the categorizations described above were made. Fifth, all significant *changes* in forward guidance were identified, which allowed for the creation of the main variable of interest in this study, f_{ct}—the coding of which is now described. When a central bank significantly shifted the tone of its forward guidance, this change was recorded as either a -1 (a shift to a more dovish stance) or a $+1$ (a shift to a more hawkish stance). When the tone of forward guidance did not change, this was recorded as a 0. If no forward guidance was found at all, this was also recorded as a 0. The matrix below illustrates how numerical values were assigned to changes in forward guidance. For example, if a central bank shifted from no forward guidance (neutral) to forward guidance with a bias towards rate increases (hawkish), then a value of $+1$ would be assigned.

		To:		
		Dovish	Neutral	Hawkish
From:	Dovish	0	+1	N/A
	Neutral	−1	0	+1
	Hawkish	N/A	−1	0

Naturally, the measurement of forward guidance raises two questions. First, why not measure forward guidance using a more precise numerical measure? Indeed, I initially attempted to do so by scoring forward guidance in increments of 0.2 on a scale ranging from -1 to 1 (as in Feroli et al. (2017)). This required assigning the bias and forward guidance into one of eleven categories. Despite the guidance of a carefully created rubric intended to reduce the decisions to a simple algorithm, the assignment of forward guidance to such narrowly defined categories resulted, in the author's opinion, in false precision. Hence, it was abandoned for a more practical, albeit less precise, methodology.

Second, is it really possible to measure forward guidance? Central banks strive to be clear in their communications. They tend to place more scrutiny on statements that are likely to be received as signals. As such, it is typically very straightforward to glean whether a central bank has attempted to send a dovish signal, hawkish signal, or, far more often, no signal at all. Indeed, financial and economic journalists do this on a regular basis and typically manage to reach consensus. Although there have been many instances of ambiguous forward guidance, these are greatly outnumbered by instances of clear forward

verbal forward guidance does not imply the availability of a central bank policy rate projection. Nor is the inverse true: the absence of a central bank policy rate does not imply the absence of verbal forward guidance. It is only the contrapositive that is generally true: the absence of verbal forward guidance generally implies the absence of a central bank policy rate projection.

guidance. Whether one dovish signal is stronger than another is, however, regrettably, hidden within the data.

As discussed above, all instances of forward guidance are also categorized as either Delphic or Odyssean, which are defined as follows. Delphic forward guidance is a statement that provides information about the probable stance of monetary policy in the future. Odyssean forward guidance is a statement that strongly resembles a commitment on the stance of monetary policy in the future.[3] Delphic forward guidance is more predictive and Odyssean forward guidance is more promissory. A clear lesson emerges from this exercise: forward guidance is overwhelmingly Delphic (see Figure 1). Moessner et al. (2017) argue that there is a disconnect between the theoretical forward guidance literature and the applied forward guidance literature. They observe that most theoretical studies of forward guidance assume commitment, that is, Odyssean forward guidance. In practice, however, central banks almost never use Odyssean forward guidance. In fact, the author was able to record only two instances of Odyssean forward guidance: one from the Bank of Canada and one from the Sveriges Riksbank. Each is discussed below.

I also categorized all instances of forward guidance as either state-contingent, time-contingent, or qualitative. Time-contingent forward guidance is defined as a statement that provides information about the probable stance of monetary policy at a specific time in the future. State-contingent forward guidance is defined as a statement that provides information about the central bank's reaction function that is either more specific than or different from the central bank's mandate. Qualitative forward guidance is that which does not fall into either category above. It is much more vague in nature and much more common.

Moessner et al. (2017) (pg. 679) correctly observe that "Qualitative, calendar-based and threshold-based forward guidance are also referred to as open-ended, time-contingent and state-contingent forward guidance, respectively." Indeed, qualitative and open-ended forward guidance appear to be equivalent although I favour the term qualitative forward guidance because it implies a clear counterpart, quantitative forward guidance, which as discussed below, has been used in practice. Conversely, it is unclear what closed-ended forward guidance would entail. Threshold-based forward guidance, by contrast, is a special case of state-contingent forward guidance. It is far more common for a central bank to condition forward guidance on some (occasionally vague) future state than to define a numerical threshold, which is actually quite rare in practice. Similarly, calendar-based forward guidance is a special case of time-contingent forward guidance. Central banks often refer to a specific time in their forward guidance, but that time is seldom as specific as a calendar date or even a calendar month. More often, time-contingent forward guidance refers to something like a quarter or even a half of a year. Occasionally, central banks have referred to specific upcoming meetings, which could appropriately be considered calendar-based forward guidance.

Note that none of these definitions require any kind of commitment. Even some empirical studies of forward guidance suggest that commitment plays a role in these definitions. "When FG is time dependent, the central bank commits to a particular path for a given time horizon" (Kool and Thornton (2015)). Further, the forward guidance literature often suggests that forward guidance is *either* state-contingent or time-contingent, *either* Odyssean or Delphic. Such classifications, although useful, risk creating the perception of distinct instruments with distinct objectives and distinct outcomes. In fact, forward guidance is frequently multidimensional. Bhattarai and Neely (2018) observed that forward guidance can even be both Delphic and Odyssean in the same announcement. Mario Draghi observed at a Peterson Institute for International Economics conference on October 12, 2017, "Our rate guidance has both time- and state-based dimensions since rates cannot rise until we see an improvement in the inflation outlook

[3]The words "strongly resembles a commitment" are used because no central bank would ever truly issue an unconditional commitment on the future state of monetary policy.

sufficient to end net asset purchases." Janet Yellen echoed this point shortly afterwards at the European Central Bank communications conference on November 14, 2017, "I believe every bit of forward guidance—even when it's been calendar-based over the years—the FOMC has used the words *we think such and such will be appropriate in light of the outlook for the economy.*" Perhaps, then, we should begin discussing forward guidance in terms of overlapping *attributes* rather than mutually exclusive *classes*. This subtle shift in the debate would allow for a more nuanced discussion of forward guidance and better recognize its multi-faceted nature. The two examples of Odyssean forward guidance illustrate this point.

On April 21, 2009 the Bank of Canada used the following header for the press release announcing the monetary policy decision of the Governing Council.

> "Bank of Canada lowers overnight rate target by 1/4 percentage point to 1/4 per cent and, conditional on the inflation outlook, commits to hold current policy rate until the end of the second quarter of 2010."

Clearly, the Governing Council made a commitment. Of course, this commitment was not unconditional. This conditionality illustrates that the forward guidance was also state-contingent. The clear reference to a specific date demonstrates that the forward guidance was also time-contingent. Two statements contained in the press release also have distinctly Delphic attributes. "The Bank of Canada today announced that it is lowering its target for the overnight rate by one-quarter of a percentage point to 1/4 per cent, which the Bank judges to be the effective lower bound for that rate." "Today's decision to lower the policy rate by 25 basis points brings the cumulative monetary policy easing to 425 basis points since December 2007. It is the Bank's judgment that this cumulative easing, together with the conditional commitment, is the appropriate policy stance to move the economy back to full production capacity and to achieve the 2 per cent inflation target." These statements suggest that it is unlikely that the policy rate will be lowered further. The use of the word *judgement* illustrates the probabilistic nature of the statements making them much more Delphic than Odyssean. The next example—and the only other example of Odyssean forward guidance I found—reinforces the point that forward guidance is frequently multifaceted.

On October 28, 2014, the Riksbank included the following text in its press release announcing that the Executive Board decided to cut the repo rate by 0.25 percentage points to zero per cent:

> "**Zero repo rate until inflation clearly picks up**
> It is important that inflation rises towards the target of 2 per cent. The repo rate is therefore being cut by 0.25 percentage points to zero per cent. The low repo rate increases demand in the economy, which contributes to higher inflationary pressures. The highly-expansionary monetary policy may also contribute to keeping inflation expectations anchored around 2 per cent by sending a clear signal that monetary policy is focused on inflation approaching the inflation target.
>
> The repo rate needs to remain at this level until inflation has clearly picked up. Slow increases in the repo rate are expected to begin until the middle of 2016 and it should reach 1.75 per cent towards the end of 2017. This is an unusually low repo rate at a time when economic activity is good, resource utilisation is close to its normal level and CPIF inflation is 2 per cent."

The statement "Zero repo rate until inflation clearly picks up" is the phrase that seems to suggest a commitment. Of course, the commitment is less explicit than the one discussed above, but the language is certainly suggestive. The strength of the statement is echoed in the later sentence "The repo rate needs to remain at this level until inflation has clearly picked up." These statements about the path of the repo rate are also stronger than those used in the previous (September) press release, which included a similar sentence, "The repo rate needs to remain low for a long period of time for inflation to rise towards the target," but did not include the stronger statement, "zero repo rate until inflation clearly picks up." The October 2014 forward guidance also has aspects of both time-contingent and state-contingent forward guidance. First, the Executive Board makes it clear that the repo rate will stay at zero until inflation has clearly picked up. In other words, the level of the repo rate is most concerned with the low rates of inflation. Second, the statement includes a clear reference to a date at which we can expect the repo rate to increase. Of course, because the Executive Board releases its projections of the path for the repo rate, this time-contingent forward guidance is not extraordinary but, for the Riksbank, a regular feature. Finally, unlike the two examples below, one feature of this forward guidance that is conspicuously absent is any insistence on behalf of the Executive Board that this forward guidance does not represent a commitment. Of course, the absence of the denial of a commitment is not evidence of a commitment, but as far as forward guidance commitments go, this is about as close as it gets.

Two very famous examples of threshold-based forward guidance also warrant discussion. Because they are such famous examples, they have come to typify state-contingent forward guidance. In fact, they represent the exception rather than the rule. On December 12, 2012 the Federal Open Markets Committee included the following statement in its press release announcing the monetary policy decision:

> To support continued progress toward maximum employment and price stability, the Committee expects that a highly accommodative stance of monetary policy will remain appropriate for a considerable time after the asset purchase program ends and the economic recovery strengthens. In particular, the Committee decided to keep the target range for the federal funds rate at 0 to 1/4 percent and currently anticipates that this exceptionally low range for the federal funds rate will be appropriate at least as long as the unemployment rate remains above 6-1/2 percent, inflation between one and two years ahead is projected to be no more than a half percentage point above the Committees 2 percent longer-run goal, and longer-term inflation expectations continue to be well anchored. The Committee views these thresholds as consistent with its earlier date-based guidance. In determining how long to maintain a highly accommodative stance of monetary policy, the Committee will also consider other information, including additional measures of labor market conditions, indicators of inflation pressures and inflation expectations, and readings on financial developments. When the Committee decides to begin to remove policy accommodation, it will take a balanced approach consistent with its longer-run goals of maximum employment and inflation of 2 percent.

The first sentence is considered qualitative forward guidance due to the open-ended nature of the phrase "considerable time." The second sentence is considered state-contingent forward guidance because of the phrase "at least as long as the unemployment rate remains above 6-1/2 percent." This particular instance of forward guidance has been referred to as threshold-based forward guidance in the literature as well. The third sentence references "earlier date-based guidance." This refers to a phrase that was used in the statements leading up to this FOMC meeting, "In particular, the Committee also decided today to keep the target range for the federal funds rate at 0 to 1/4 percent and currently anticipates that exceptionally low levels for the federal funds rate are likely to be warranted at least through mid-2015." This reference illustrates that the Committee has shifted from time-contingent forward guidance to state-contingent forward guidance.[4] Although it is tempting to consider this episode of forward guidance Odyssean

[4] Indeed, this earlier date-based guidance was classified as time-contingent forward guidance.

because of the use of an explicit threshold, the statement is in fact rather Delphic. The phrase "currently anticipates that" indicates that this forward guidance is far more predictive than promissory. The later sentences also make it clear that the threshold of a 6-1/2 percent unemployment rate is one of many guidelines and that the committee intends to use numerous macroeconomic indicators to reach its future decisions. Janet Yellen later reinforced these inferences two months later in a speech in Washington:

> It deserves emphasis that a 6-1/2 percent unemployment rate and inflation one to two years ahead that is 1/2 percentage point above the Committee's 2 percent objective are thresholds for possible action, not triggers that will necessarily prompt an immediate increase in the FOMC's target rate. In practical terms, it means that the Committee does not expect to raise the federal funds rate as long as unemployment remains above 6-1/2 percent and inflation one to two years ahead is projected to be less than 1/2 percentage point above its 2 percent objective. When one of these thresholds is crossed, action is possible but not assured.

On August 7, 2013 the Bank of England took a similar step:

> In particular, the MPC intends not to raise Bank Rate from its current level of 0.5% at least until the Labour Force Survey headline measure of the unemployment rate has fallen to a threshold of 7%, subject to the conditions below.

The Monetary Policy Committee (MPC) goes on to outline a set of three so-called knockout conditions that would render the forward guidance obsolete along with some qualitative and balance sheet forward guidance.[5] Once again, we see state-contingent forward guidance by use of a threshold. And once again, this forward guidance is clearly more Delphic than Odyssean because of the use of the word "intends" as well as the addition of the knockout conditions. In his first speech as Governor of the Bank of England, Mark Carney made the following remarks about the threshold:

> Furthermore, thinking unemployment will come down faster than we expect isnt enough to believe interest rates will rise soon. As I said earlier, the 7% threshold is a staging post to assess the economy. Nobody should assume that it is a trigger for raising rates.

Charles Bean made it even more clear that this forward guidance was not a commitment at a speech at the Jackson Hole Symposium:

> This guidance is intended primarily to clarify our reaction function and thus make policy more effective, rather than to inject additional stimulus by pre-committing to a time-inconsistent lower for longer policy path in the manner of Woodford (2012). While such a time-inconsistent policy may be desirable in theory, in an individualistic committee like ours, with a regular turnover of members, it is not possible to implement a mechanism that would credibly bind future members in the manner required.

2.4 Quantitative easing data

I also collected quantitative easing data for each central bank over the sample period. The process was very similar to that described above for forward guidance. First, I checked each monetary policy decision press release for statements about quantitative easing. This was done as part of the data entry for

[5] The full text of the forward guidance can be found on page 7 of the MPC's August 2013 Inflation Report.

forward guidance. If a central bank initiated a quantitative easing program, that this was considered a dovish change to quantitative easing. Quantitative easing was labelled as either dovish (-1), neutral (0), or hawkish (+1). If a central bank made no change to its quantitative easing program or if you did not have one, this was labelled with a 0. If a central bank announced that it intended to end its quantitative easing program, this was labelled as a +1.

Changes to quantitative easing were also recorded if a central bank announced either an expansion (-1) or a contraction (for, a reduction in pace) (+1). The dates and scoring were also crossed checked with the literature (Gagnon et al. (2011), Krishnamurthy and Vissing-Jorgensen (2011), Swanson (2017), Altavilla and Giannone (2017)).

Figure 1: **Forward guidance changes and forecasts of the policy rate**

Each figure in the panel charts changes to forward guidance, central bank policy rates, and individual private-sector forecasts of the policy rate at the one-year horizon. Each individual private-sector forecast is colour-coded by the type of forward guidance.

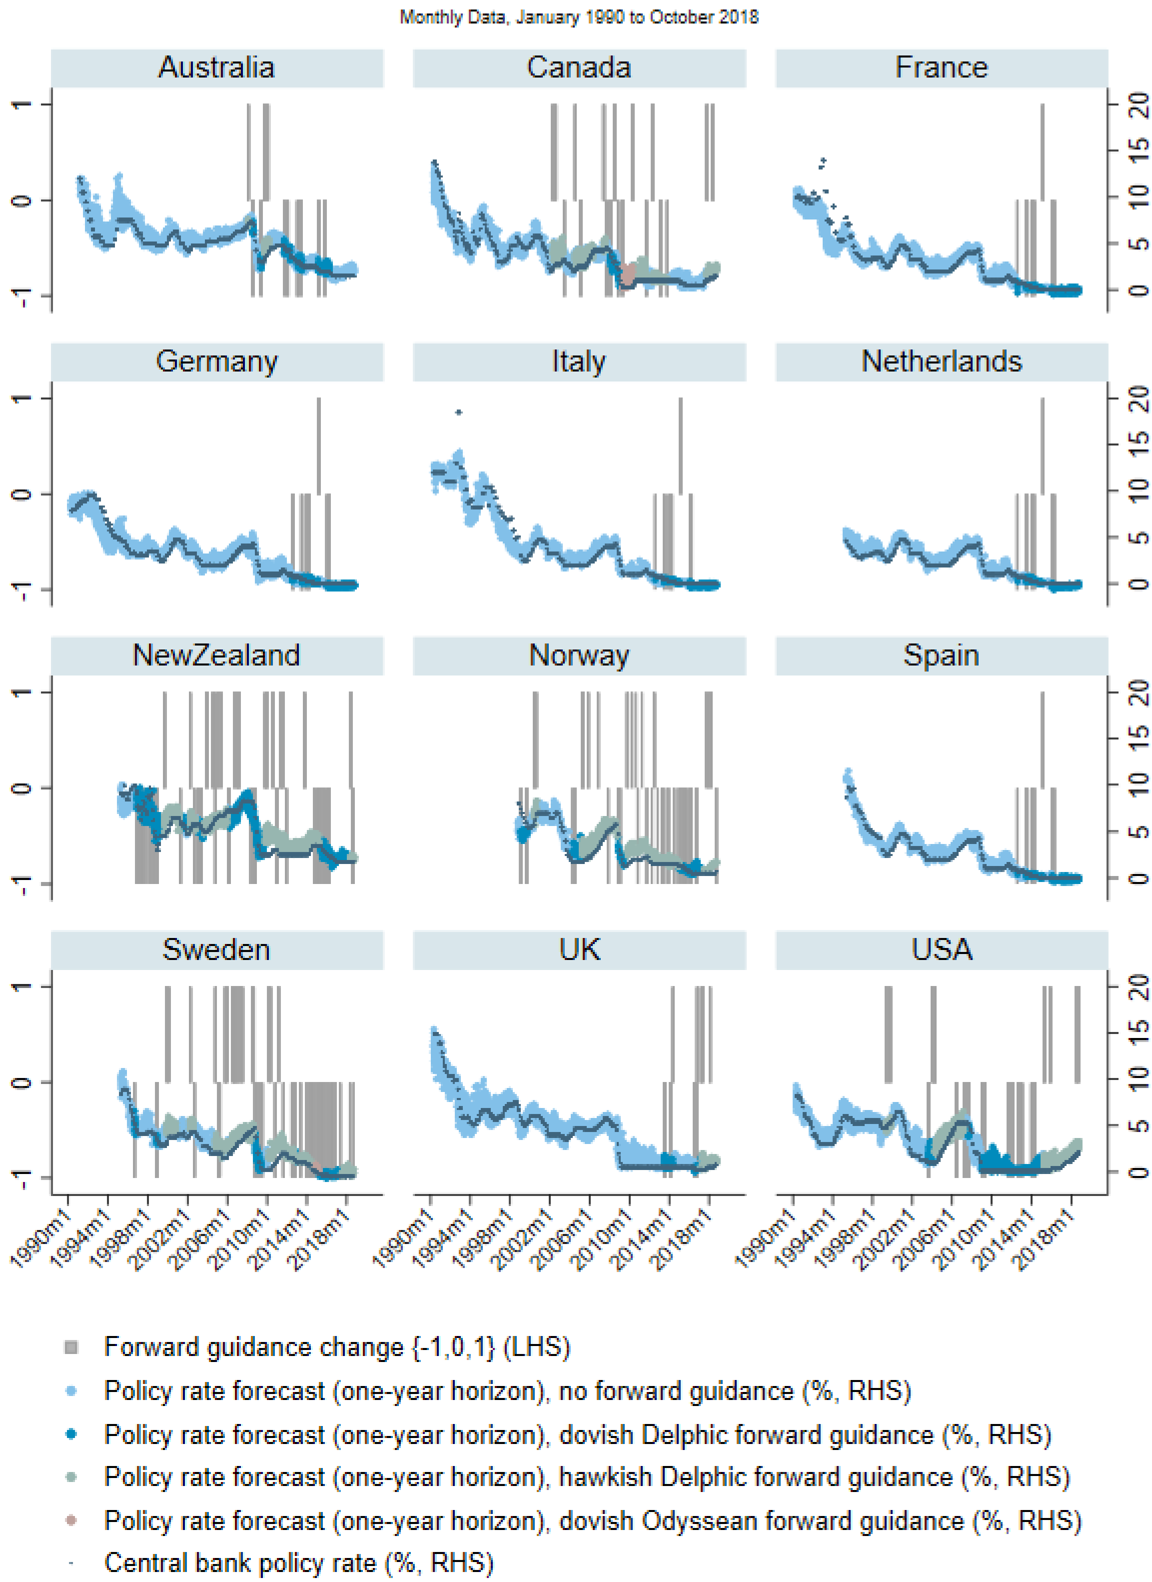

Figure 2: **Quantitative easing changes and forecasts of the policy rate**

Each figure in the panel charts changes to quantitative easing, central bank policy rates, and individual private-sector forecasts of the policy rate at the one-year horizon. Each individual private-sector forecast is colour-coded by the type of forward guidance.

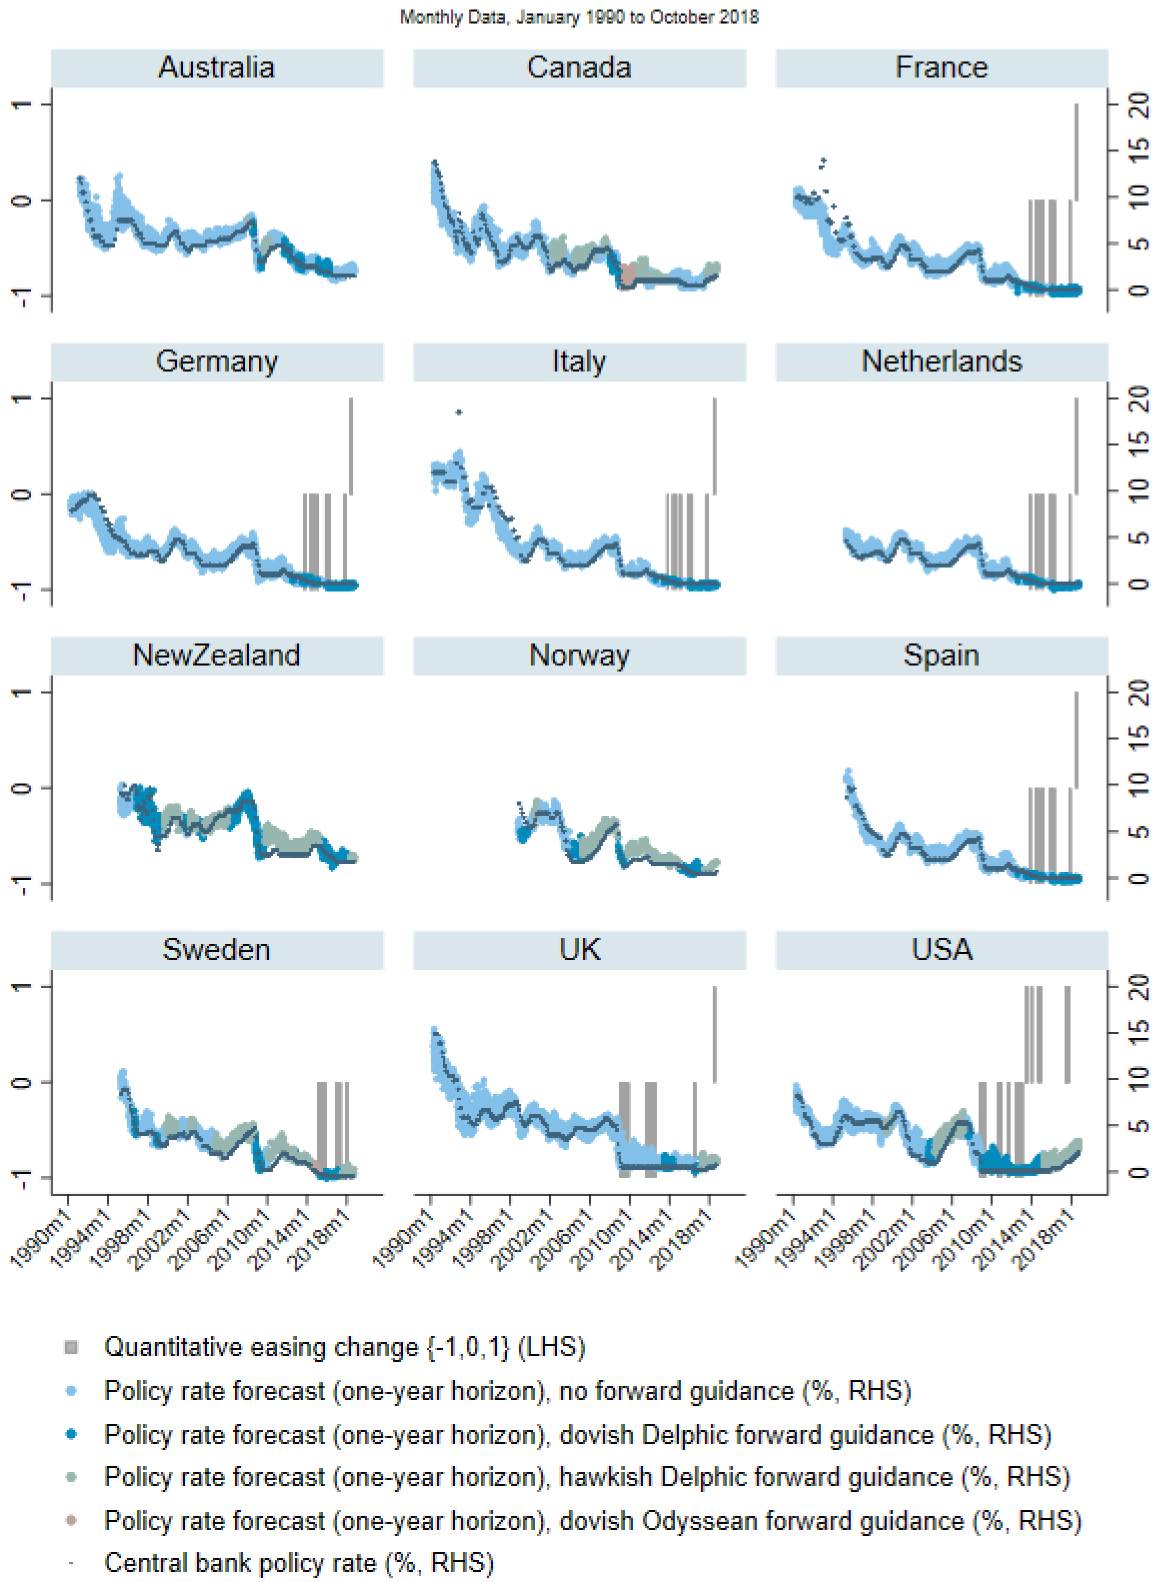

3 Framework

3.1 Causes, alternate potential causes, and confounds

This paper attempts to estimate the causal effect of forward guidance on private-sector interest-rate forecasts. To do so, the identification assumptions will need to take into consideration numerous important factors such as, inter alia, alternate potential causes, confounds, intermediate outcomes, simultaneity, and the interaction between treatment (forward guidance) and outcome (interest rate forecast revisions).

Revisions to private-sector interest rate forecasts could have many alternate causes other than monetary policy signals. To isolate the causal effect of forward guidance, we may need to condition on such alternate causes. However, to make realistic identification assumptions, model parsimony will be crucial. The key strategy for parsimony will be to focus on inflation-targeting central banks only. Private-sector forecasters know that inflation-targeting central banks maintain a near singular focus on the alignment of future inflation with the central bank's inflation target. They also know that central banks often provide monetary policy signals. It therefore seems reasonable to assume that the two primary inputs to private-sector interest rate forecast revisions are monetary policy signals and inflation forecasts.

Of course, there are many other possible causes of these revisions, such as the domestic growth and unemployment outlooks. In an inflation-targeting regime, however, each can also be considered to be an *input* to either monetary policy signals or the inflation outlook. Were we to consider non-inflation targeting central banks, we would have to consider other inputs to central banks' reaction function, such as the exchange rate or the unemployment rate. In that case, placing a more narrow focus on inflation would lead to omitted variable bias. I provide evidence of the validity of this assumption in subsection B.4.

Model parsimony is also critical because the treatment variable, forward guidance (f_{ct}), is necessarily measured with error, which, as argued below, will tend to introduce attenuation bias to β. Let $\tilde{f}_{ct}$ represent the unobservable, true state of forward guidance in a given country at a given time and let ν_{ct} represent the measurement error.

$$f_{ct} = \tilde{f}_{ct} + \nu_{ct} \qquad [1]$$

Now $f_{ct} \in \{-1, 0, 1\}$ and we assume that $\mathbb{E}[\tilde{f}_{ct}] = 0$ as $\mathbb{E}[f_{ct}] = 0$ and $\mathbb{E}[\nu_{ct}] = 0$. $\tilde{f}_{ct}$ should follow a fairly symmetric, continuous distribution between -1 and $+1$, whereas f_{ct} follows a fairly symmetric, discrete, trinomial distribution $\{-1, 0, 1\}$. The nature of these two distributions implies that $Cov(f_{ct}, \nu_{ct}) < 0$ and that $Cov(\tilde{f}_{ct}, \nu_{ct}) = 0$. ν_{ct} is assumed to be uncorrelated with all other regressors x_{ict}, $Cov(x_{ict}, \nu_{ct}) = 0$, and the error term ϵ_{ict}, $Cov(x_{ict}, \nu_{ct}) = 0$. In this case, because $Cov(f_{ct}, \nu_{ct}) \neq 0$, the probability limit of $\beta_{f_{ct}}$ can be characterized as follows (Wooldridge (2010)).

$$plim(\beta_{f_{ct}}) = \beta_{f_{ct}} \left(\frac{\sigma^2_{\omega_{ct}}}{\sigma^2_{\omega_{ct}} + \sigma^2_{\nu_{ct}}} \right) \qquad [2]$$

where ω_{ct} is the error term from the following linear projection and where x_{ict} is some matrix of covariates,

$$f_{ct} = x_{ict} + \omega_{ct} \quad [3]$$

The addition of relatively extraneous covariates will tend to introduce attenuation bias into the measurement of β because, to the extent that such additional regressors have positive correlation with f_{ct} and weak or no positive partial correlation with the dependent variable, $r_{ic,t+h}$, those additional regressors will tend to capture some of the signal value of f_{ct} and further attenuate the partial correlation between $r_{ic,t+h}$ and f_{ct}, thereby biasing β towards zero. Conversely, the omission of useful regressors that have positive correlation with f_{ct} and positive partial correlation with $r_{ic,t+h}$ will tend to introduce positive omitted variables bias. β would be biased upwards as it would inappropriately capture some of the positive effect attributable to the omitted, useful regressor. So we are faced with the trade-off between attenuation bias and omitted variables bias. Hence, we must strike a balance between parsimony and completeness and be very careful deciding which regressors to condition on and which to omit. The following figure helps address this trade-off by addressing the key channels that influence interest rate expectations.[6]

Figure 3: **Changes to forward guidance and interest-rate forecast revisions** π_{ict} represents a revision to a private-sector inflation forecast for the current year from forecaster i in country c at time t. π_{ct} represents a revision to a central-bank inflation forecast for the current year. f_{ct} represents a change to forward guidance. $r_{ic,t+h}$ represents a revision to a private-sector forecast of the policy rate h months in the future. The black arrows imply a causal relationship. The absence of such arrows implies the lack of a causal relationship. The nodes represent observable data.

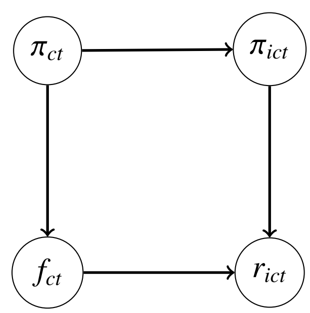

The key component of Figure 3 is that f_{ct} is assumed to have a causal effect on $r_{ic,t+h}$. In other words, we assume that private-sector forecasters adjust their forecasts for future policy rates in direct response to central banks' updates to their forward guidance. Additionally, π_{ict} is also assumed to have a causal influence on $r_{ic,t+h}$. That is, we assume that private-sector forecasters adjust their forecasts for future policy rates in direct response to their own revisions to their own current-year inflation forecasts (see subsection 3.1 for more on this). Stated differently, we assume that the latest inflation data contains information about the future path of inflation. In an inflation-targeting regime, information about the future path of inflation should directly translate into information about the future path of policy rates. So to isolate the causal effect of f_{ct} on $r_{ic,t+12}$, we must condition on the appropriate variables. Based on Figure 3, we can do so using one of three conditioning strategies. We can condition on π_{ict} , π_{ct} , or both.

Conditioning on π_{ict} seems to be the most straightforward and elegant approach. First, because π_{ct} strictly precedes π_{ict} in the data-generating process, then π_{ict} should incorporate π_{ct} and so π_{ct} should contain no

[6]This can be thought of as a more general, nonparametric method, which in the causal inference literature, is known as a directed acyclic graph (Pearl (2009)).

information about $r_{ic,t+h}$ once we have conditioned on π_{ict}. Second, because we consider π_{ict} to be an alternate cause of $r_{ic,t+h}$, it is preferable to condition on that alternate cause. For example, π_{ict} could also incorporate time-varying idiosyncratic expectations that could affect $r_{ic,t+h}$. Indeed, we would expect this to be the case otherwise all individual forecasters would produce identical forecasts. This luxury is afforded to us because π_{ict} comes from the same individual forecaster, i, as does $r_{ic,t+h}$ in the Consensus Economics data and so we have about as much data for π_{ict} as $r_{ic,t+h}$. Nonetheless, later, we will add π_{ct} as a covariate as a robustness check. This will come at the cost of lowering T because the data for π_{ct} are not as expansive as π_{ict}.

We must also consider how quantitative easing contributed to monetary policy signalling. Coenen et al. (2017) emphasize the importance of considering the roles of quantitative easing and forward guidance together. Again, we will consider the change to the signal, as a *change* to central banks' quantitative easing, which we will call q_{ct}. Similarly, $q_{ct} \in \{-1, 0, 1\}$ where -1 represents a dovish signal, 0 represents no signal, and 1 represents a hawkish signal.

Figure 4: **Changes to forward guidance and interest-rate forecast revisions with quantitative easing**
π_{ict} represents a revision to a private-sector inflation forecast for the current year from forecaster i in country c at time t. π_{ct} represents a revision to a central-bank inflation forecast for the current year. f_{ct} represents a change to forward guidance. q_{ct} represents a change to quantitative easing. $r_{ic,t+h}$ represents a revision to a private-sector interest rate forecast h months into the future. The black arrows imply a causal relationship. The absence of such directed edges implies the lack of a causal relationship. The nodes represent observable data.

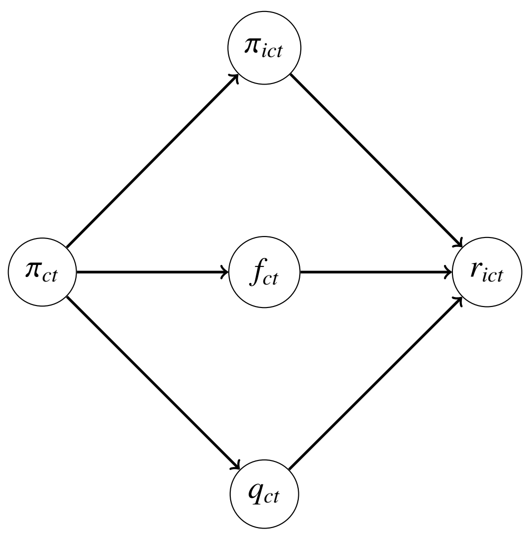

The assumptions above about the main causes of interest rate forecast revisions are well aligned with the literature on the effects of unconventional monetary policy. "These dimensions represent the three aspects of FOMC announcements that had the greatest systematic effect on asset prices over the sample; intuitively, these three dimensions are likely to correspond to changes in the federal funds rate, changes in forward guidance, and changes in LSAPs" (Swanson (2017)). We will also condition on r_{ct}, the change in the central bank policy rate over the last period, to adjust for the more mechanical changes to an individual's forecast of the policy rate, r_{ict}. Of course, Swanson (2017) is a high-frequency event study of, inter alia, market short-term interest rates, whereas this paper considers lower-frequency survey data so we must also condition on π_{ict} as well as these three factors.

3.2 Simultaneous Equations

Of course, private-sector forecasters may not consider just current-year inflation when forming expectations of future policy rates, but they may also consider future inflation. However, if we include future inflation expectations in our control variables we introduce simultaneity bias. Instead, consider the following system of structural equations representing both interest rate forecast revisions and inflation forecast revisions. Here it is useful to recall that monetary policy operates with a lag, that both forward guidance and quantitative easing are themselves forward-looking in nature, and that each variable below represents a change rather than a level.

$$r_{ic,t+h} = f(\pi_{ict}, ..., \pi_{ic,t+h}, ..., \pi_{ic,t+h+j}, f_{ct}, q_{ct}), j \in \mathbb{N} \quad [4]$$

$$\pi_{ic,t+h} = f(r_{ic,t}, ..., r_{ic,t+h-1}, f_{ct}, q_{ct}), j \in \mathbb{N} \quad [5]$$

We can observe the following exclusion restrictions. Revisions to forecasts of the policy rate at some point in the future, h, are a function of revisions to the path of inflation up to and beyond that point. They are also a function of changes to forward guidance and changes to quantitative easing. By contrast, revisions to forecasts of the inflation rate at some point in the future, h, are a function of revisions to the path of policy rates up to but not at or beyond that point. More simply, the policy rate is forward-looking while the inflation rate is backward-looking. Hence, if we only include revisions to forecasts of inflation with forecast horizons that do not cross into the future periods in which monetary policy may be influential, then we avoid the endogeneity of inflation forecast revisions.

The only inflation forecasts available in the data are those of current-year and next-year inflation. Theoretically, next-year inflation could be influenced by the level of the policy rate one year from now. So we should rely on current-year inflation to control for contemporaneous changes to the macroeconomy. This is precisely the approach taken by Altavilla and Giannone (2017) who use the same line of reasoning (see pp 959-960). The authors add that "Assuming that the effect of policy on the real economy is delayed and that policy decisions are not affected by current-quarter variations in long-term bond yields is tantamount to the recursive identification scheme used in structural vector autoregressions to identify standard policy (for recent implementations and critical discussions see (Leeper et al., 1996; Bernanke et al., 2005; Uhlig, 2005; Banbura et al., 2010; Giannone et al., 2015)." For robustness, I include larger sets of macroeconomic control variables in subsection B.4 and find that the results do not change at all across the various configurations. The following sections discuss additional and related identification assumptions in much more detail.

3.3 Identification Assumptions

Swanson (2017) lists three identification challenges that apply to this study. First, it may be difficult to disentangle the effects of forward guidance and quantitative easing. After all, during the financial crisis, changes to forward guidance and quantitative easing happened simultaneously in some cases. However, in our panel data, simultaneous implementation is the exception rather than the rule. The correlation between changes to forward guidance and changes to quantitative easing is rather low at approximately 0.10 (see the table below). This luxury is afforded to us because of the size and heterogeneity of our data set.

Second, financial markets are forward-looking and so will incorporate expectations of policy rate changes, forward guidance changes, and quantitative easing changes before they take place. This is a major challenge for high-frequency event studies that use financial market data. The issue is much less severe for this study because the data are monthly and changes to central bank policy are considerably more difficult to anticipate one month in advance than, say, one day in advance.

Third, central banks can surprise forecasters by inaction because of the role expectations play. In this study, we are primarily interested in the ceteris paribus relationship between changes to forward guidance and revisions to interest rate forecasts. The hypothesis is that, all else equal, when f_{ct} takes the value 1 (hawkish) we expect rate expectations to rise, when it takes the value 0, we expect expectations to hold, and when it takes the value -1, we expect rate expectations to fall. Hence, inaction can be defined by the value 0. If forecasters were expecting a 1, but ultimately observed a 0, they may have revised down their rate expectations. If forecasters were expecting a -1, but ultimately observed a 0, they may have revised up their rate expectations. As long as the effect is fairly symmetric, that is, the expected forward guidance forecast error is approximately equal to zero, then they should not bias our estimates. As the mean of f_{ct} is approximately zero, then this seems to be a reasonable assumption.

For valid causal inference, we will need additional assumptions. To identify β, we will start by assuming strict exogeneity and then consider to what extent we can relax this assumption.

$$\mathbb{E}(u_{it}|f_i, q_i, \pi_i, p_i, \alpha_i) = 0 \qquad \forall \quad i = 1, 2, ..., N \quad t = 1, 2, ..., T \tag{6}$$

We will investigate these assumptions in more detail, consider any potential violations, and discuss how to address them. The assumptions below are outlined in Imai and Kim (2019), which states the assumptions necessary for causal inference when using the fixed effects estimator.[7] These assumptions elucidate the importance of the structure of Equation 4. Were these equations to include expected interest rates, expected inflation rates, and the presence of forward guidance in levels form, then most or all of these assumptions would be implausible. However, the equations relate *revisions* of expectations to interest rates to *revisions* of expectations to inflation rates and *changes* in central banks forward guidance and quantitative easing. With this in mind, we now consider the plausibility a number of important assumptions in turn.

First, and most importantly, I assume that no other unobserved time-varying confounder exists. In other words,

$$r_{ic,t+h} = f(f_{ct}, q_{ct}, \pi_{ict}, p_{ct}, \alpha_i, \alpha_t) \tag{7}$$

$$\epsilon_{it} \perp\!\!\!\perp f_{ct}, q_{ct}, \pi_{ict}, p_{ct}, \alpha_i, \alpha_t \tag{8}$$

Because this study is conducted within the confines of inflation-targeting frameworks, we assume that figure 4 captures the principal causes of $r_{ic,t+h}$: central bank signals and the outlook for inflation. To the extent that a private-sector forecaster understands that inflation is the dominant component of a central

[7]Similar assumptions are also covered in Wooldridge (2005) and Wooldridge (2010), for example, which discuss the necessary conditions for consistent estimation of the average treatment effect using the fixed effects estimator.

bank's reaction function, then whatever else can be considered to influence a given forecaster's inflation outlook can be assumed to be subsumed within π_{ict}.

Second, I assume that past revisions to interest-rate forecasts do not affect current revisions to interest-rate forecasts (past outcomes do not affect current outcomes).

$$\mathbb{E}[(r_{ic,t+h})(r_{ic,t-m+h})] = 0, \qquad m = 1, 2, 3, ... \tag{9}$$

however, in the regression,

$$r_{ic,t+h} = r_{ic,t-1+h}\theta + f_{ct}\beta + q_{ct}\gamma + \pi_{ict}\delta + p_{ct}\phi + \alpha_i\zeta + \alpha_t\eta + \epsilon_{ict} = 0 \tag{10}$$

$\theta \approx 0.17$. However, this assumption can be partially relaxed without jeopardizing causal inference (Imai and Kim (2019)).

Third, I assume that individual private-sector interest rate forecast revisions do not affect future central-bank forward guidance revisions (past outcomes do not affect future treatments). This assumption is less self-evident but is quite plausible. Central banks consider a large amount of data when setting monetary policy and their communication strategies, such as macroeconomic data, financial market data, survey data, geopolitics, etc. It is not plausible that private-sector interest rate forecasts alone are enough to materially influence central banks' forward guidance, let alone an *individual* private-sector forecaster's rate forecast.

Fourth, I assume that past treatments do not affect current outcomes (past revisions to forward guidance do not contemporaneously and directly affect revisions to one-year forecasts for the policy rate). Focusing only on the treatment and the outcome,

$$r_{ic,t+h} = f(f_{ct}) = \mathbb{E}[r_{ic,t+h}|f_{ct}] \tag{11}$$

Here, t corresponds to monthly data. This means the forecasters have between one and thirty days to incorporate forward guidance into their expectations because, in the data, changes in forward guidance strictly precede revisions to private-sector interest rate forecasts. Although it would be reasonable to expect that past forward guidance can affect the future level of interest rate forecasts, it is less plausible to postulate that past forward guidance *changes* should affect future interest rate forecast *revisions*. One exception may be very recent revisions to central bank forward guidance. Forecasters may be somewhat to update their forecasts, particularly if they exhibit either sticky expectations (Mankiw et al. (2003)) or rational inattention (Sims (2003)). So, later, we will relax this assumption somewhat and address the violation of strict exogeneity by conditioning on a lagged treatment (forward guidance) variable.

$$r_{ic,t+h} = +f_{ct}\beta_1 + f_{c,t-1}\beta_2 + q_{ct}\gamma + \pi_{ict}\delta + p_{ct}\phi + \alpha_i\zeta + \alpha_t\eta + \epsilon_{ict} \tag{12}$$

Fifth, I make the *stable unit treatment value assumption* (Rubin (1980)), which comprises two sub-assumptions. First, it assumes that there is no interference between units (Cox (1958)); that is, neither $Y_i(1)$ nor $Y_i(0)$ is affected by what action any other unit took (using the notation from Fisher (1923) that indicates treatment (1) and control (0)). In other words, we assume that individual private-sector forecasters form their forecasts independently of one another, or, without collusion. This assumption is especially plausible given the structure of the data. Each month, each forecaster uses a template provided by the data provider to complete his or her forecasts. Once completed, the forecaster returns the completed template to the data provider before the monthly deadline. Once all submissions have been received and the deadline has passed, the data provider consolidates the forecasts into one spreadsheet and distributes the file to its subscribers weeks later. Absent collusion, the forecasters are incapable of contaminating one another's forecasts.

Of course, we should not discount the possibility of *some* amount of collusion. For example, one forecaster could always have discussions with another, or worse yet, one forecaster could release his or her forecasts to the public before the monthly deadline. I argue that this type of collusion is by far the exception rather than the rule, but for robustness, we will probe this assumption further later in the paper. Second, and more formally, it assumes that there are no hidden versions of treatments; no matter how forecaster i received treatment (f_{ct}), the outcome that would be observed would be $Y_i(1)$ and similarly for treatment 0 (Rubin (2005), page 323). The validity of this assumption is demonstrated in subsection B.3.

3.4 Anticipated and unanticipated forward guidance

One limitation of this study is that we are not able to separately estimate the effects of anticipated forward guidance and unanticipated forward guidance. Numerous studies have estimated the effects of unanticipated monetary policy rate changes on the term structure of interest rates (Kuttner (2001), Gürkaynak et al. (2005), Swanson (2017)). To do so, these studies typically calculate the difference between the realized policy rate decision and the market-implied policy rate expectation immediately before the decision as a measure of surprise.

If we could measure anticipated forward guidance ex ante, we could then measure the effects of unanticipated forward guidance by purging our change in forward guidance variable, f_{ct}, of this anticipated element. Efforts to this end have been undertaken in the studies of financial market rates in which the identification strategy has relied on the high-frequency nature of the data, which allows researchers to focus on narrow windows around policy decisions. Unfortunately, this identification strategy is not available to us in studies of survey expectations.

In this study, we are interested in the effect of a change in forward guidance, f_{ct}, on an individual revision to an interest rate forecast, r_{ict}. In the measuring this effect, we control for that same individual's revision to his or her relevant macroeconomic expectations, π_{ict} (formed at the same time as r_{ict}), because we are operating within the confines of inflation-targeting jurisdictions. π_{ict} is an agglomeration of that individual's private macroeconomic signal and the public macroeconomic signal (Morris and Shin (2002), Amato et al. (2002)). The reaction of r_{ict} to f_{ct} therefore represents the individual's reaction to the central bank's forward guidance signal (Morris and Shin (2005), Svensson (2006), Amato and Shin (2006) Morris and Shin (2007), Hubert (2014), Tang (2013), Lustenberger and Rossi (2018), Morris and Shin (2018)). How much of this central bank signal can individual forecasters reasonably anticipate? Perhaps some but certainly not all. Far more importantly, how much of the anticipated signal, f_{ct}, is embedded in $y_{ic,t-1}$, the interest rate forecast that took place one month earlier? The answer should be almost certainly little-to-none.

If $y_{ic,t-1}$ does not incorporate any anticipated forward guidance for f_{ct}, then the (ceteris paribus) relationship between r_{ict} and f_{ct} reflects an unconfounded measure of individual forecasters' average response to forward guidance. Why *would* $y_{ic,t-1}$ reflect some aspect of anticipated forward guidance in period t? One possibility is that period $t-1$ did not contain a central bank policy rate decision and corresponding press release but that period t does and that the individual anticipates the central bank to modify its forward guidance in period t. To clarify the mechanics, recall that $y_{ic,t-1}$ reflects a policy rate forecast at either a three-month or twelve-month horizon. So, in either case, the $t-1$ forecast pertains to a date that follows any policy rate decision in month t by at least a month. However, to incorporate anticipated forward guidance changes one month in advance would reflect an unusual degree of certainty about precisely how the central bank intends to draft its policy communication.

Further, the nature of f_{ct} makes it far more difficult to forecast than the level of the policy rate at some time in the near future. f_{ct} incorporates a panoply of factors: an agglomeration of private views (those of the Governing Council or the Monetary Policy Committee) some of which may have known preferences (Hansen et al. (2014), Hansen et al. (2017), Bordo and Istrefi (2018)), precise choice of language, both intended and perceived tone, historical communication to that point (Woodford (2013)), sentiment (Hubert and Labondance (2018a)), and narrative (Shiller (2017)). This is not an easy variable to forecast. Admittedly, Governing Council or Monetary Policy Committee speeches—especially by the Governor or Chair—may provide some minor insight into upcoming forward guidance (Ehrmann and Fratzscher (2007), Hansen and McMahon (2016)). Ultimately however, no one can forecast future forward guidance—not even the Governor or Chair—because forward guidance is drafted collectively by a group of people within a committee in the days immediately before the policy rate decision. To assume that monetary policy committees not only systematically signal their policy rate decision but also signal upcoming forward guidance would be to presuppose that decisions on monetary policy are agreed in advance of monetary policy committee meetings (Svensson (2010)).

Although this study does not disentangle the effects of anticipated and unanticipated forward guidance, from the discussion above, we can nonetheless draw two conclusions. First, the estimates of the average forward guidance effect may be attenuated somewhat due to anticipated forward guidance. Second, the magnitude of this attenuation is likely to be small because it is unlikely that much anticipated forward guidance is embedded in individuals' interest-rate forecasts.

4 Estimation

I estimate the following model where all variables are defined as above. Estimation will be done using a two-way fixed effects estimator, which is analogous to a difference-in-differences approach. Specifically, the fixed effects estimator uses both firm and monthly fixed effects. Standard errors are clustered by country (Bertrand et al. (2004), Stock and Watson (2008)). Correlation of forecast revisions is likely to be much greater by country than by individual forecasters. If the standard errors were clustered by individual forecasters, then they would be strongly biased downwards. This bias could be in orders of magnitude in some cases because of the large sample size in this study. One drawback of this approach (Bertrand et al. (2004), Cameron et al. (2008), Cameron et al. (2011)) is that the standard errors will be estimated using only twelve clusters (one for each country). So, for robustness, I re-estimate Equation 13 below for each country and cluster the standard errors by forecaster. Doing so allows for about twenty-five clusters per country on average.

$$r_{ic,t+h} = \beta f_{ct} + \delta \pi_{ict} + \phi p_{ct} + \gamma q_{ct} + \zeta \alpha_i + \boldsymbol{C_t \eta} + \epsilon_{ict} \quad [13]$$

Two terms have been added to the equation above. First, α_i is included to capture time-invariant, firm-level heterogeneity. Naturally, this term is purely expositional as we will estimate the model above using the within transformation. Second, C_t is a matrix of time-varying, firm-invariant control variables. In most cases, this includes month fixed effects, α_t. This allows us to control for the type of global shocks that might cause forecasters to revise their policy rate forecasts, such as the deteriorating macroeconomic outlook due to the global financial crisis or changes to the Federal Reserve's monetary policy. Other types of exogenous time control variables (in addition to the monthly dummy variables) were tested but were ultimately not included as they were seldom or never significant.

C_t also includes a dummy variable indicating periods in which countries' central banks are operating at or near the effective lower bound, e_{ct}. This control variable allows for the possibility that the nature of revisions to policy rate forecasts is fundamentally different at the effective lower bound. For instance, the propensity to revise policy rate forecasts at the effective lower bound may be lower, particularly in response to forward guidance. As such, e_{ct} is included as both the control variable and an interaction variable with f_{ct}. These periods are identified in a number of ways.

First, if the central bank has made no statements about its views on the level of the effective lower bound, then the zero lower bound is considered to be the effective lower bound. If the central bank has made statements about its views on the level of the effective lower bound as many central banks have (the Bank of Canada, the Bank of England, the European Central Bank, the Sveriges Riksbank), then the stated level supersedes the zero lower bound from that point onwards (for example, from the date of the statement onwards, the effective lower bound might be considered to be a negative policy rate). In that case, the dummy variable takes the value one whenever the central bank is very close to or at its stated effective lower bound. Many central banks ultimately later revised down their stated level of the effective lower bound (the European Central Bank, the Sveriges Riksbank). In that case, the definition of the effective lower bound is iteratively revised from the point of the statement onwards.

Equation 13 includes a group of variables that are somewhat analogous to the factors described in Gürkaynak et al. (2005) and Swanson (2017). p_{ct} is analogous to the target factor (Gürkaynak et al. (2005)). f_{ct} is analogous to the path factor or the forward guidance factor (Gürkaynak et al. (2005) and Swanson (2017)). q_{ct} is analogous to the quantitative easing (LSAP) factor (Swanson (2017)). Crucially, the factors in Gürkaynak et al. (2005) and Swanson (2017) are, by construction, orthogonal to one another. Fortunately, the analogous factors used in this study, while not orthogonal to one another, have very low correlation.

ρ			
	p_{ct}	f_{ct}	q_{ct}
p_{ct}	1	0.05	0.01
f_{ct}	0.05	1	0.10
q_{ct}	0.01	0.10	1

5 Results

Below I present the main results from this paper. Before showing summary statistics from panel regressions, I present the main data from the regressions graphically to preview the results. Figure 5 includes scatter plots by country that show private-sector policy rate forecast revisions on the y-axis and private-sector current-year inflation forecast revisions on the x-axis. These forecast revisions are colour-coded to illustrate whether they coincided with either no change in forward guidance in their respective country, a hawkish change in forward guidance, or a dovish change in forward guidance. This allows us to visually consider the effect the changes in forward guidance may have had on private-sector forecast revisions. Indeed, we can see that positive interest rate forecast revisions tend to coincide with hawkish forward guidance and that negative interest rate forecast revisions tend to coincide with dovish forward guidance.

In Figure 6, the data are filtered to show only revisions that occurred immediately following forward guidance changes. The data are also colour-coded. Separate colours are used for hawkish and dovish forward guidance changes as well as Delphic and Odyssean forward guidance changes. In this figure, it does not appear that the two instances of dovish Odyssean forward guidance coincide with large negative interest rate forecast revisions.

Similarly, in Figure 7, the data are filtered to show only revisions that occurred immediately following forward guidance changes. These data are also colour-coded so that separate colours are used for hawkish and dovish forward guidance changes as well as time-contingent and state-contingent forward guidance changes. This figure does suggest that time-contingent forward guidance has coincided with larger interest-rate forecast revisions in the intended direction than those that coincide with state-contingent forward guidance.

Figure 5: **Private-sector policy rate and inflation forecast revisions colour-coded by forward guidance change value**

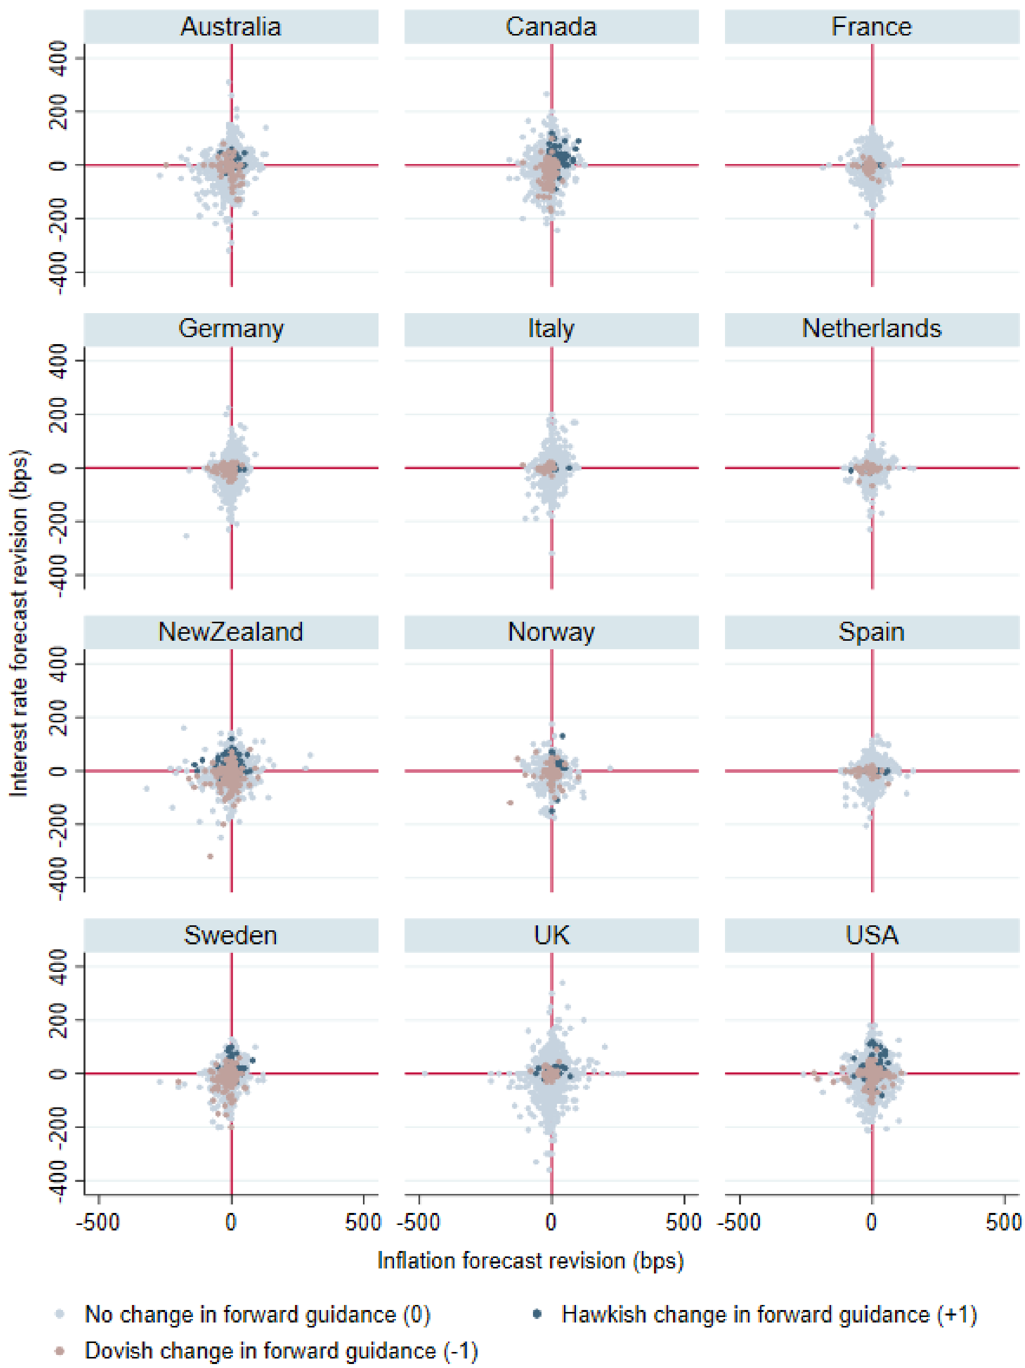

Figure 6: **Private-sector policy rate and inflation forecast revisions colour-coded by type of forward guidance change value**

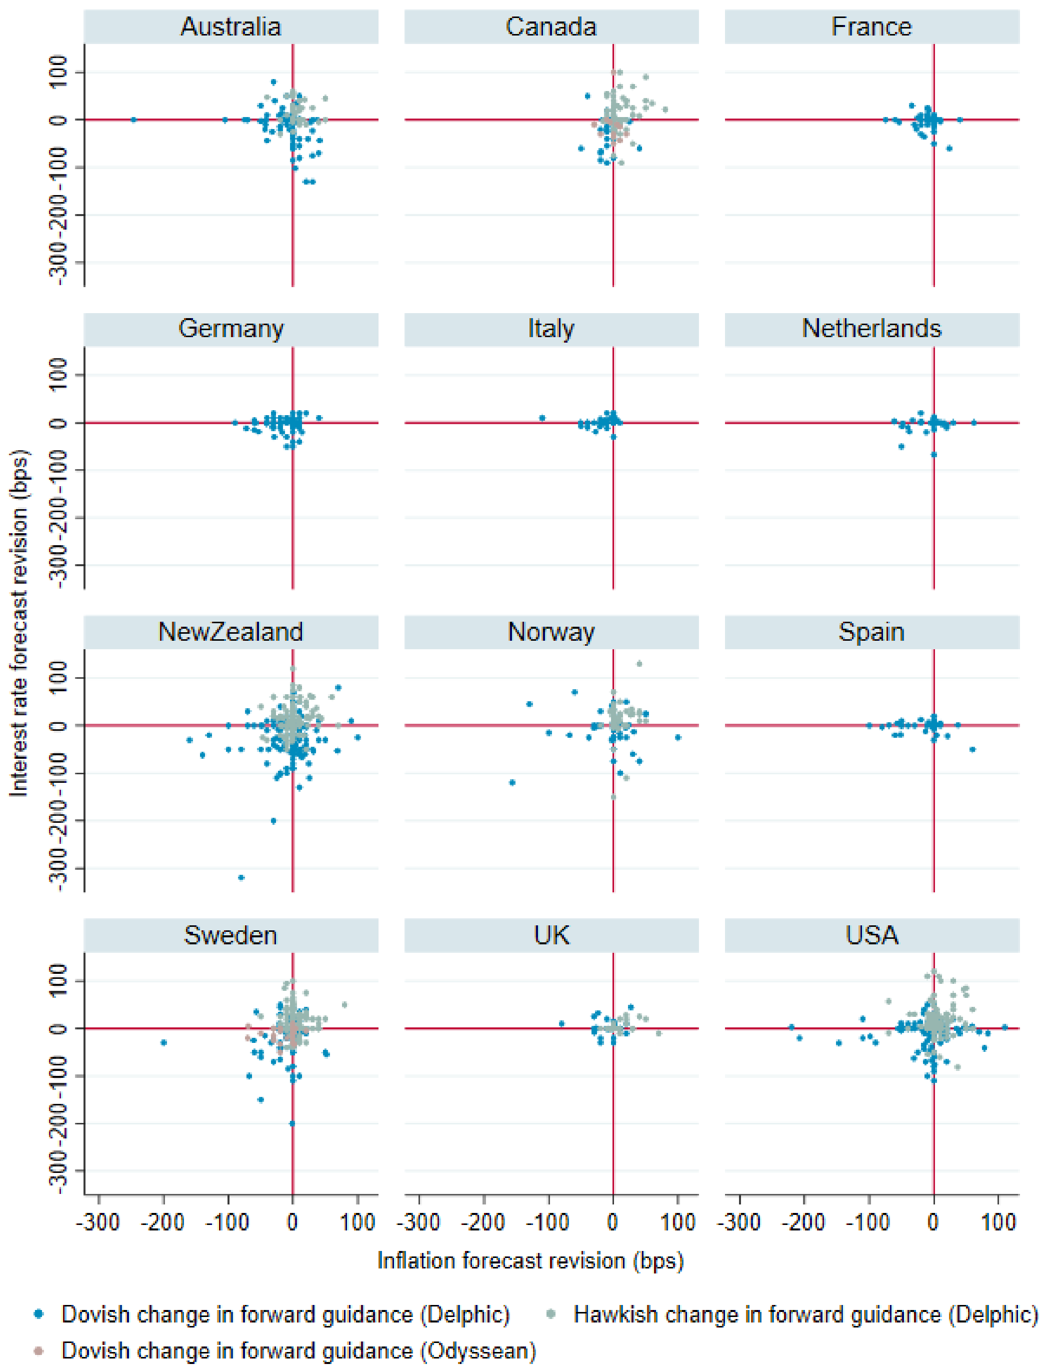

Figure 7: **Private-sector policy rate and inflation forecast revisions colour-coded by type of forward guidance change value**

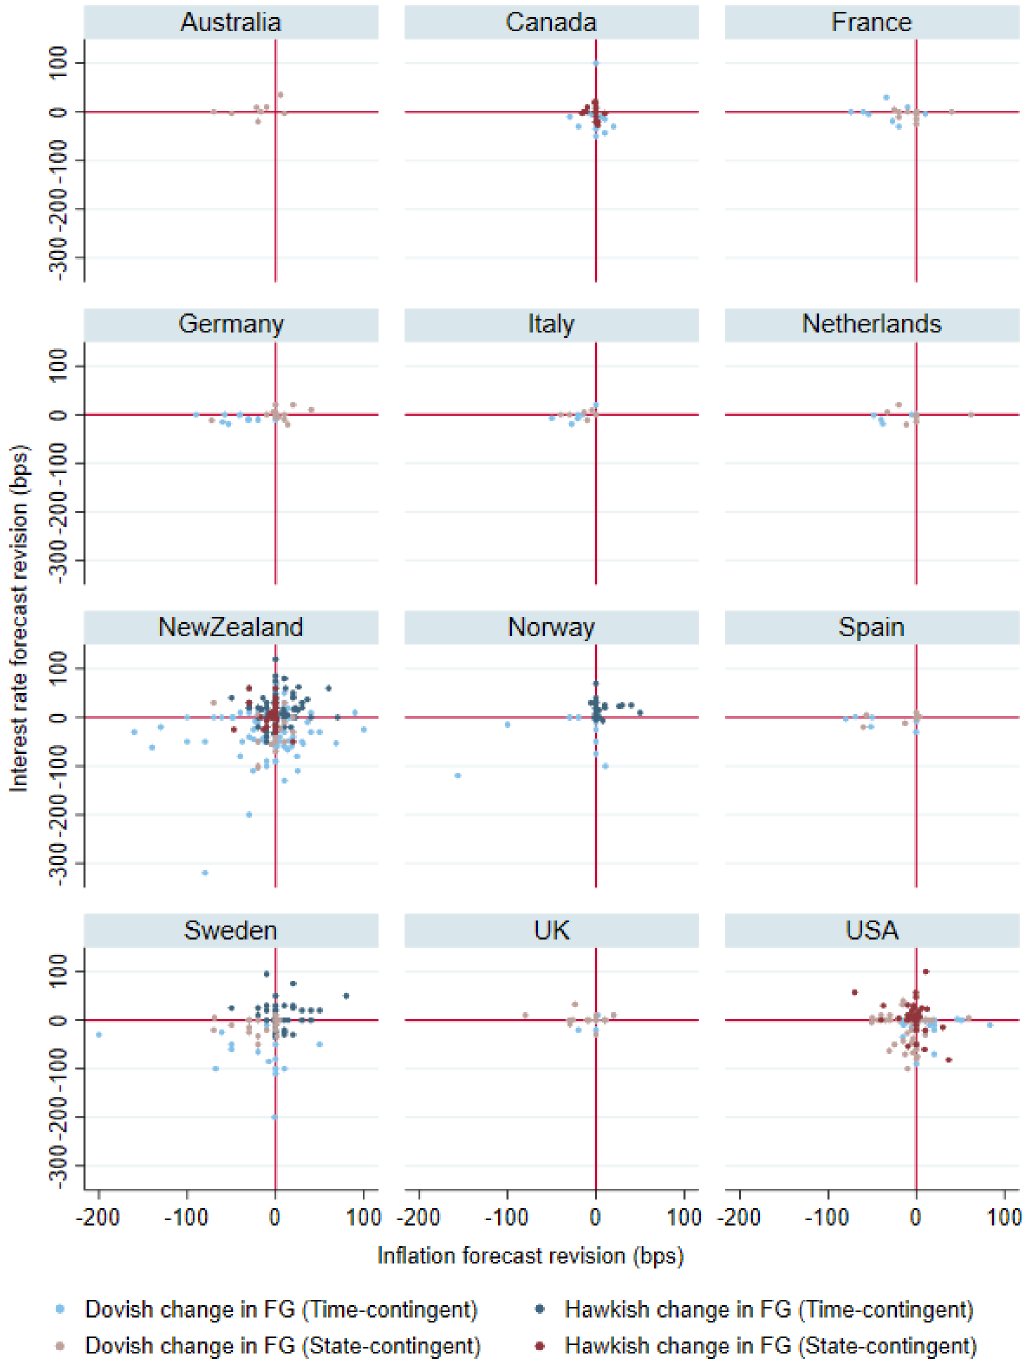

Table 1: **How do changes in forward guidance affect revisions to individuals' forecasts of the policy rate level twelve months from now?**

	[1]	[2]	[3]
(1) Forward guidance $\{-1, 0, 1\}$ change (+1)	4.81***	6.30***	8.62**
	(1.19)	(1.29)	(2.78)
(2) Policy rate (PR) change (+25 bps)	6.66***	6.71***	6.64***
	(0.62)	(0.62)	(0.63)
(3) Private inflation forecast revision (+25 bps)	2.15***	2.17***	2.14***
	(0.35)	(0.35)	(0.35)
(4) Effective lower bound {0, 1}	0.16	-0.09	0.39
	(0.47)	(0.46)	(0.95)
(5) Quantitative easing $\{-1, 0, 1\}$ change (+1)	-1.30	-0.31	-1.11
	(0.96)	(0.88)	(0.82)
(1) x (2)		1.24	
		(0.99)	
(1) x (3)		0.88	
		(1.09)	
(1) x (4)		-4.48***	
		(1.32)	
(1) x |(5)|		-1.51	
		(2.16)	
(1) x Odyssean forward guidance {0, 1}			4.09**
			(1.38)
(1) x Time-contingent forward guidance {0, 1}			0.79
			(2.39)
(1) x State-contingent forward guidance {0, 1}			-5.66**
			(1.91)
(1) x Qualitative forward guidance {0, 1}			-4.67
			(2.80)
Adjusted R^2	0.22	0.22	0.22
N	31174	31174	31174

Standard errors in parentheses

Dependent variable: revision to individuals' forecasts of the policy rate level twelve months from now.

Baseline effects for forward guidance attributes are included in the regressions but omitted here.

Clustered standard errors (at the country level) are shown in parentheses.

This table shows summary statistics from panel regressions with firm and month fixed effects.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.1 Average forward guidance effect, quantitative easing effect, and policy rate target effect

Table 1 shows that the average forward guidance effect on forecaster expectations of the policy rate one year from now (analogous to the path factor or slope factor in other studies) is estimated to be about five basis points. This provides an estimate of how much we can expect a private-sector forecaster to adjust his or her forecast of the policy rate one year from now in response to a hawkish change in forward guidance. This is extremely well aligned with the results in high-frequency event studies of the effect of unconventional monetary policy on market interest rates (Gürkaynak et al. (2005), Swanson (2017)).

The average quantitative easing effect on forecaster expectations of the policy rate one year from now (analogous to the LSAP factor) is estimated to be approximately zero. Quantitative easing does not appear to amplify the influence of forward guidance—a result that differs from those of Coenen et al. (2017). This result also suggests that the signalling channel of quantitative easing may not be as strong as often suggested.

The average policy rate target effect on forecaster expectations of the policy rate one year from now (analogous to the target factor or level factor in other studies) is estimated to be about seven basis points. This provides an estimate of how much we can expect a private-sector forecaster to adjust his or her forecast of the policy rate one year from now in response to a 25 basis point policy rate change.

5.2 The effective lower bound dampens forward guidance

Table 1 shows that the effects of forward guidance on expectations are noticeably dampened during periods at the effective lower bound. Swanson (2017) (Table 4) shows that a similar result holds for the six-month treasury yield. Intuitively, dovish forward guidance, which was much more common during periods at the effective lower bound, has a limited ability to incite forecasters to revise down their policy rate expectations because it may be the case that policy rates cannot foreseeably go lower in many forecasters' eyes. Second, periods of forward guidance at the effective lower bound have coincided with quantitative easing programs on a number of occasions. Hence, quantitative easing programs may have pushed down interest rates and limited the potential stimulative effects of forward guidance at the effective lower bound. Indeed, estimates from panel regressions that include a triple interaction term between f_{ct}, e_{ct}, and a binary dummy variable indicating the presence of a quantitative easing program support this interpretation. Additionally, regressions that constrain the sample to effective lower bound periods only show that forecasters are much less responsive to π_{ict} and much more responsive to p_{ct} at the effective lower bound.

5.3 Odyssean commitments amplify and state-contingent attributes dampen forward guidance

Table 1 shows that, when we interact the change in forward guidance variable with a binary variable indicating Odyssean forward guidance, the effect is roughly doubled. This suggests that the role of commitment in forward guidance is far from trivial and that central banks seeking to cause particularly significant shifts in interest rate expectations should consider adding a commitment. Of course, this result must be interpreted with caution given that this study only recorded two instances of Odyssean forward guidance. As such, the standard error on the Odyssean interaction term is fairly large and the result is only significant at the 5% level. Further, neither of these cases of Odyssean forward guidance were fully anticipated by markets, so some of the additional influence over forecaster rate expectations may be attributable to the unexpected component of the announcements. Nonetheless, these estimates add to the literature by providing perhaps the first empirical evidence to date of the potency of Odyssean forward guidance.

The results also show that state-contingent forward guidance is much less effective than time-contingent forward guidance. The two cases of state-contingent forward guidance used by the FOMC and the Bank of England were two of the most prominent examples of forward guidance and also two of the most complex. Indeed, these cases of forward guidance led to a significant amount of uncertainty among many. Jain and Sutherland (2018) found that time-contingent forward guidance was more effective at reducing private-sector forecaster disagreement than state-contingent forward guidance. These results for state-contingent forward guidance also differ from those of Coenen et al. (2017) who find that state-contingent forward guidance is one of the most effective forms of forward guidance.

5.4 Forward guidance has had the strongest effects in Canada, Australia, New Zealand, and Sweden

In this section, I estimate Equation 13 for each country and provide the results below. I will also test for any spillover effects from the Federal Reserve's monetary policy. To do so, I add an additional variable to Equation 13, $f_{ct}^{US} \in \{-1, 0, 1\}$. This variable simply takes the same value as f_{ct} would if the country were the United States but does so for all countries. That is, this variable simply takes the value of f_{ct} when $c = USA$ regardless of the true value of c. When $c = USA$, $f_{ct}^{US} = 0$ to avoid redundancy with f_{ct}. Here, I also must adjust C_t to avoid collinearity issues with α_t. C_t includes binary dummy variables for each year (less one) and binary dummy variables for each month of the year (less one) in an effort to adjust for the monotonic decrease in uncertainty about current-year inflation (see subsection 2.1).

Overall, Figure 8 shows that forward guidance by the Bank of Canada, Reserve Bank of Australia, Reserve Bank of New Zealand, and Sveriges Riksbank, moved private-sector forecasters' expectations in the intended direction. Conversely, forward guidance does not appear to have been as influential in the euro zone or the United Kingdom. What explains the discrepancy? One simple explanation is that both the Bank of England and the European Central Bank were comparatively late adopters of forward guidance (see Figure 1). Most of their forward guidance has taken place at or near the effective lower bound. As such, these central banks' ability to influence rate expectations with forward guidance would have been dampened (see Table 1 and subsection 5.2) for further explanation. We can also see below in Figure 8 that there are clear spillover effects from Federal Reserve forward guidance to Canadian policy rate expectations. Given that the United States is by far Canada's largest trading partner, this is to be expected.

Figure 8: **Central bank policy decisions' effect on private-sector one-year policy rate forecasts by country**

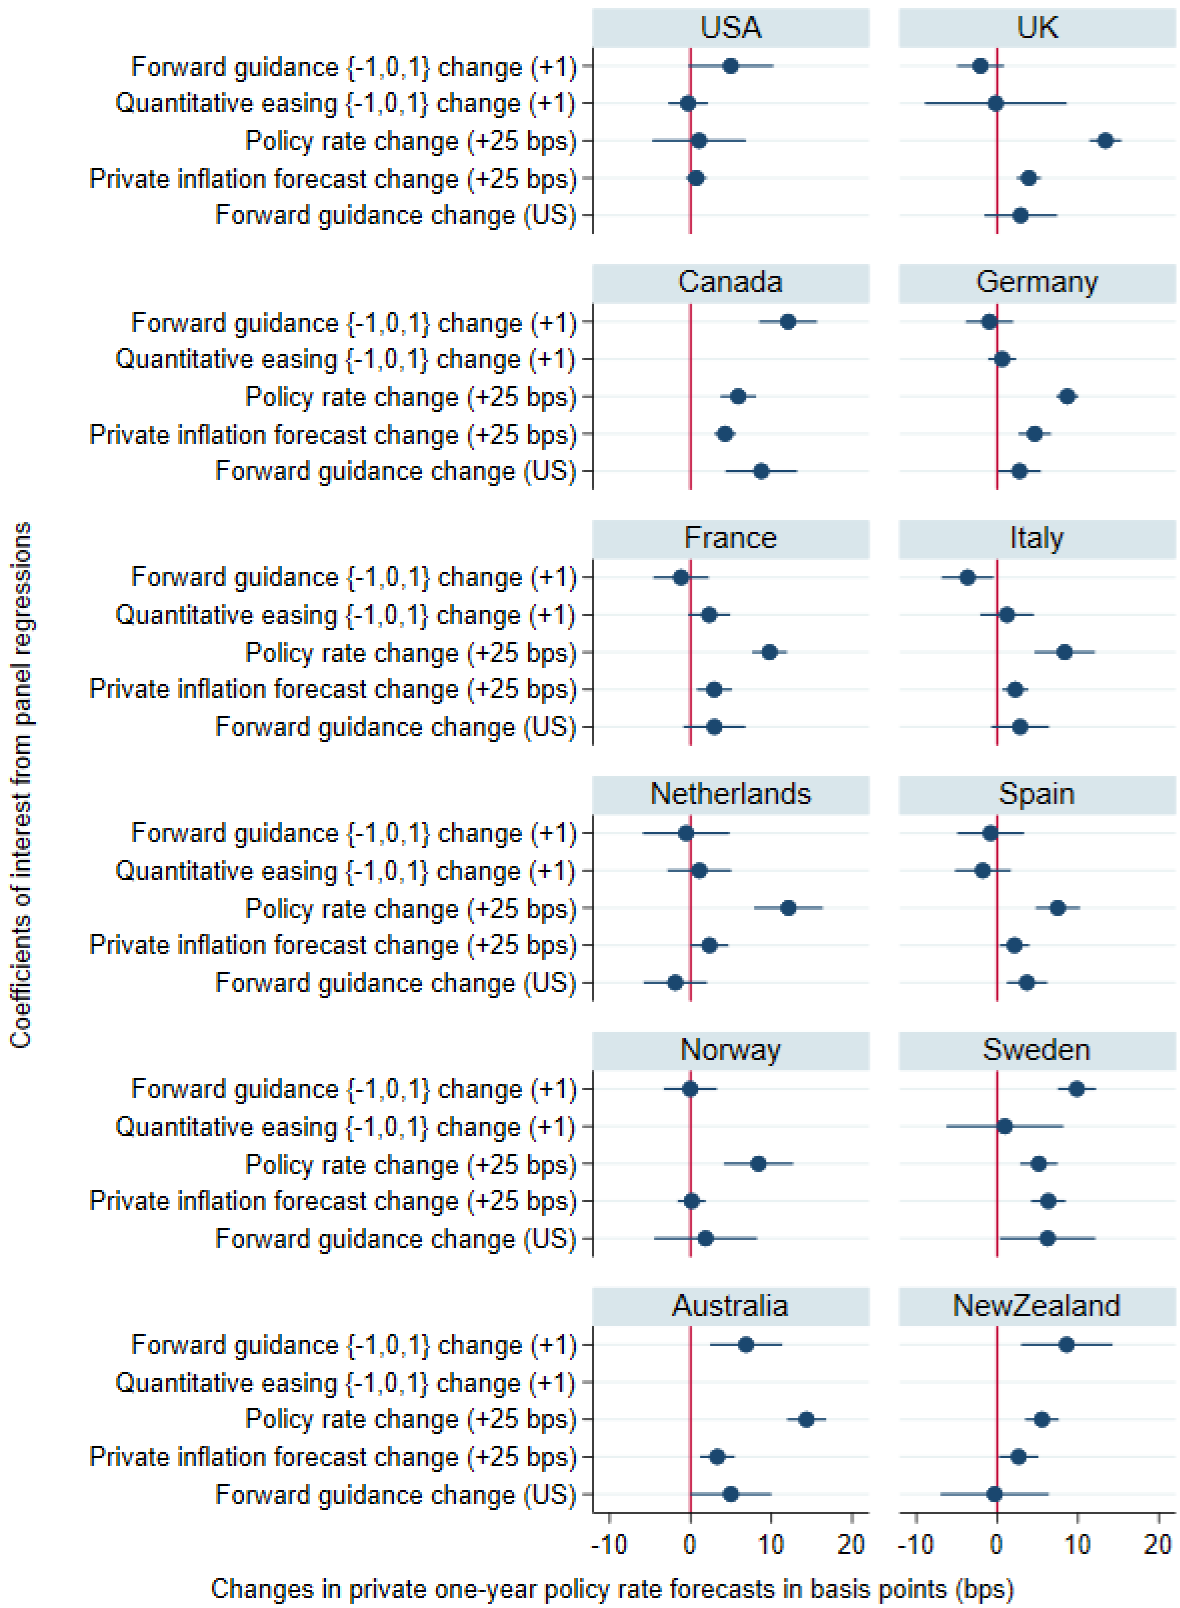

Empirical studies of the Reserve Bank of New Zealand (RBNZ) have found that RBNZ monetary policy surprises (Drew et al. (2007)) and RBNZ interest rate forecasts (Moessner and Nelson (2008), Andersson and Hofmann (2009)) influence the short end of the yield curve. Ferrero and Secchi (2009) document reductions in short-term money market volatility following the introduction of forward guidance. Svensson (2015) shows that the policy rate path projections published by the Reserve Bank of New Zealand have been fairly well aligned with market-implied policy rate paths and that the RBNZ has often appeared to have been successful at influencing market rates. RBNZ forward guidance has also contributed to improved short-term interest rate forecast accuracy and reductions in corresponding forecast disagreement (Kool and Thornton (2015), Jain and Sutherland (2018)).

Using the methodology from Gürkaynak et al. (2005), Brubakk et al. (2017) show that the publication of the policy rate path in Norway influences market-implied policy rate expectations in the manner intended by the Norges Bank. By contrast, Figure 8 shows that this is not the case for private-sector forecasts. This discrepancy should not be attributable to the choice of sample period. Brubakk et al. (2017) restrict the sample data to 2001 onwards as this was when the Norges Bank began inflation targeting. I take the same approach for each country throughout this paper. One possible source of the discrepancy is that Brubakk et al. (2017) investigate the influence of the policy rate path publication, whereas in this paper, I measure and estimate the effects of all forward guidance—both verbal and quantitative.

The estimates for the Sveriges Riksbank suggest that forward guidance in Sweden has been very influential on average. In this case, the results correspond well to those from Brubakk et al. (2017). The Riksbank is among the central banks that have used forward guidance the most. These results may have been stronger had it not been for some challenging instances of forward guidance during the sample period, such as the April 2009 decision and the period from 2011 onwards.

Andersson et al. (2006) study the period before the Riksbank started using forward guidance in 2004. The authors use data from 1996 to 2003 to consider the relative contributions of repo rate changes, inflation reports, speeches, and minutes to changes in the term structure of interest rates. They find that repo rate changes were the dominant driver of short-term rates, which the estimates for Sweden in this paper corroborate. They also find that inflation reports and speeches influence interest rates and conclude that communication is an important tool for the Riksbank.

Svensson (2010) argues that early evidence suggested that the publication of a repo rate path successfully influenced market expectations of the path of monetary policy, but that following the April 2009 policy rate decision, the repo rate path lost some credibility. Svensson (2015) explains that the repo rate paths published by the Riksbank in 2011 and beyond were much higher than that implicitly forecasted by the markets because the Executive Board was attempting to lean against the wind and later to normalize monetary policy. He takes this as a sign that Executive Board forward guidance lacked credibility for a significant period of time.

Nonetheless, Brubakk et al. (2017) show that, from 2001 to 2016, forward guidance in Sweden influenced market-implied policy rate expectations in the manner intended by the Riksbank. Similarly, the estimates provided here show that Riksbank forward guidance moved private-sector policy rate forecasts in the intended direction on average.

The results suggest that, on average, forward guidance has perhaps been the most influential in Canada. The particularly sizable result in Canada could be partially attributable to the success of the Bank of

Canada's so-called conditional commitment. Woodford (2013) showed that, following the commitment made by the Bank of Canada in April 2009, which he describes as unprecedented, market yields along the interest-rate curve fell instantly. He notes that interest-rate contracts with maturities beyond the period mentioned in the conditional commitment (about one year) actually fell by more than those falling within the period covered by the conditional commitment. Woodford (2013) suggests that either rate expectations at longer maturities actually fell more than those at shorter maturities or that uncertainty related to the future path of monetary policy fell thereby lowering term premia. Chang and Feunou (2013) shows that both the implied and realized volatility of interest rate option prices declined following the Bank of Canada's commitment. This tends to favour the uncertainty interpretation, but of course, both aspects could be a factor.

Coenen et al. (2017) analyse forward guidance at the effective lower bound in Canada, Czech Republic, Germany, Italy, Japan, Norway, Sweden, the UK, and the US and document a number of findings. The authors find that certain types of forward guidance attenuate the link between macroeconomic news and bond yields (similar to Swanson and Williams (2014a)) and reduce disagreement about short-term interest rates (as in Jain and Sutherland (2018)), particularly state-contingent and long-horizon, time-contingent forward guidance. More closely related to this study, Hubert and Labondance (2018b) estimate the effect of ECB forward guidance on the term structure of interest rates. Similar to the methodology used in this paper and related to Hubert (2016), the authors take care to control for the information embedded in the Eurosystem macroeconomic projections released by the ECB around the same time that forward guidance is released. They find that Governing Council forward guidance successfully lowered the term structure of interest rates at most maturities.

The results for the United Kingdom are similar to those of the Eurosystem countries. Although private-sector forecasts of the policy rate did not appear to shift down with dovish forward guidance and up with hawkish forward guidance in either the United Kingdom or the Eurosystem countries, this is perhaps to be expected under the circumstances. The Bank of England and European central bank have only utilized forward guidance in recent years during periods when the policy rate was at or near the effective lower bound. As such, it may be unrealistic to think that forward guidance could materially shift private-sector policy rate forecasts. Figure 1 shows that private-sector policy rate forecasts appear to reflect forward guidance issued by the Bank of England and the European Central Bank. That is, they appeared to stay low during periods of dovish forward guidance and perhaps rise marginally during hawkish periods of forward guidance. However, the effects are small as rates remained near the effective lower bound throughout these periods of forward guidance.

The results for the Bank of England and the European Central Bank reported in Figure 8 should not be taken as evidence of ineffective forward guidance, but should instead be interpreted through the lens of the effective lower bound. The prima facie evidence from Figure 1 for the United Kingdom and Eurosystem countries suggests that forward guidance had the intended effects and that we would not really expect those effects to show up in this particular empirical test. Here, the results from Swanson and Williams (2014b) may be more appropriate, which found that United Kingdom government bond yields were substantially constrained by the effective lower bound in 2009 and 2012, but were actually quite responsive to news in 2010 to 2011. The paper also found that German government bond yields were unconstrained by the effective lower bound until 2012. It would be very informative to repeat this analysis for the United Kingdom and the Eurosystem countries covered in this study for the period from 2013 to 2018 when the Bank of England and European central bank have been at or near the effective lower bound and have used forward guidance of numerous types on several occasions.

Figure 8 suggest that FOMC forward guidance moved private-sector forecasters' expectations for the level of the policy rate in twelve months' time in the intended direction by about five basis points on average. FOMC forward guidance began in 2003 and was used frequently thereafter. The FOMC has heavily favoured time-contingent forward guidance, but of course, used state-contingent forward guidance from December 2012 to March 2014. Woodford (2013) observes that the introduction of a specific date into FOMC forward guidance has been found to be particularly influential. This is well aligned with the empirical literature on Federal Reserve forward guidance—both those that use high-frequency market interest rate data and those that use private-sector forecast or expectations.

Svensson (2015) also shows that the policy rate path projections published by the FOMC have been generally well aligned with market-implied policy rate paths, but discusses a discrepancy that began in 2014. From 2014 onwards, markets consistently expected future policy rates to be lower than those forecasted by the FOMC. This gap persisted for a number of years but the two paths have gradually converged over time.

5.5 Forward guidance has influenced forecasts of the policy rate more than bond yields

Empirical studies (mostly event studies) have found that forward guidance influences short-term market interest rates more than the bond yields (Gürkaynak et al. (2005), Moessner and Nelson (2008), Swanson (2017), Gagnon et al. (2011)). Similarly, the effect of forward guidance on forecaster expectations is strongest on the policy rate rather than bond yields and stronger at the twelve-month forecast horizon than at the three-month horizon, which is aligned with these studies. The effect on forecaster expectations of the 10-year bond yield in 12 months' time is estimated to be two basis points, which is lower but comparable to market-based studies. For comparison, Andrade et al. (2016) find that a positive shock to their estimated path factor tends to have a positive impact on the yield curve for horizons greater than two years. A number of factors could explain this result.

> "First, it is hardly meaningful to indicate policy intentions more than a few years ahead since future monetary policy depends on future economic conditions, which become very hard to predict as the forecast horizon increases. Second, the controllability of interest rates declines with maturity since movements in long-term interest rates to a large extent reflect exogenous factors such as global interest rate trends and fluctuating term premia. It is therefore an open empirical issue to determine to what extent monetary policy signaling can affect medium- and long-term interest rates" (Andersson et al. (2006), page 1816).

More recently, Hansen et al. (2018) have argued that central bank communication drives longer-term bond yields by operating through the uncertainty channel. In this paper, forward guidance is defined as communication about the probable path of monetary policy rates, which do not necessarily overlap with central bank communication about macroeconomic uncertainty. This may help explain why the effect of forward guidance on long-term bond yields is so weak. Perhaps it is not forward guidance that affects bond yields but central bank communication about macroeconomic uncertainty and the term premia. It is unsurprising then that Figure 9 below shows that the average quantitative easing effect tends to become larger the further out we move along the term structure, which is consistent with the literature on market interest rates (Krishnamurthy and Vissing-Jorgensen (2011), Swanson (2017)).

Figure 9: **Central bank policy decisions' effect on private-sector interest rate forecasts over different horizons** (three-month treasury bill yield and 10 -year government bond yield forecasts each at the three-month and twelve-month forecast horizon)

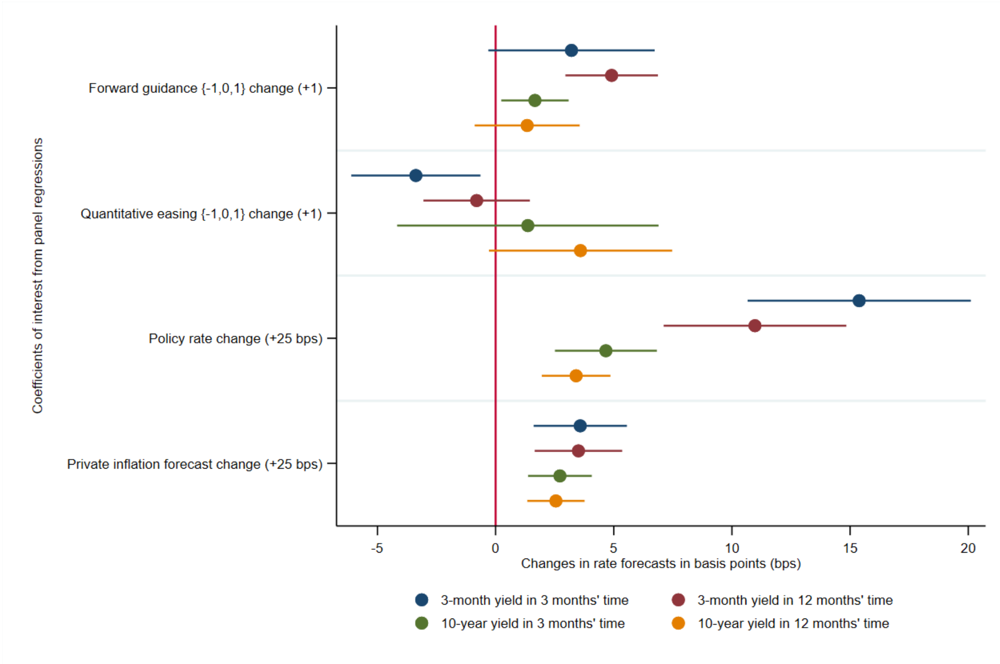

5.6 Forward guidance had the strongest effects in the pre-crisis and crisis periods and the weakest in the post-crisis period

Figure 10: **Central bank policy decisions' effect on private-sector interest rate forecasts in different periods** (Sub-sample analysis for the periods 1990 to 1999, 2000 to 2006, 2007 to 2012, and 2013 to 2018)

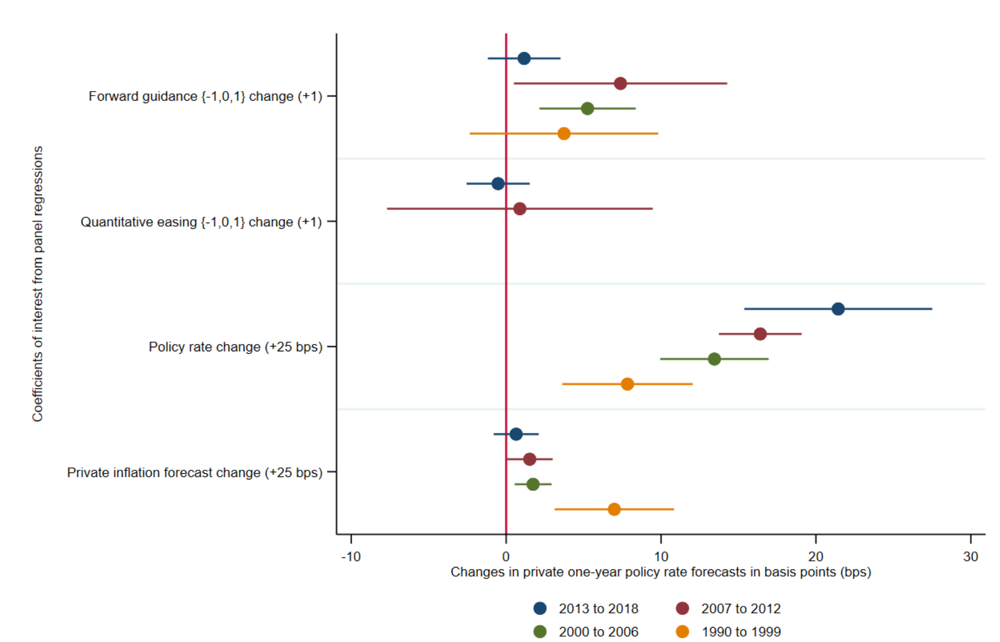

When we segregate our sample data into four periods, we see that forward guidance was particularly influential in both the years leading up to the global financial crisis and the years spanning the crisis. Perhaps most interestingly, forward guidance is estimated to have been much less influential in recent years. Following the analysis for the Bank of England and the European Central Bank above as well as the results from subsection 5.2, this is probably attributable to policy rates at or near the effective lower bound. Similarly, Detmers and Nautz (2012) find that the influence of Reserve Bank of New Zealand policy rate projections waned in the post-crisis period.

It is also not surprising that policy rate changes became more influential over time as central banks approached the effective lower bound and policy rate changes themselves became fewer and farther between. One interesting result is that the relationship between individual forecasters' revisions to their policy rate forecasts and the revisions to their inflation forecasts was strongest in the 1990s. This earlier period in the sample data coincides with a period of higher and more volatile inflation and one in which central banks had not yet established a long track record of successful inflation targeting. That the sensitivity has been much lower in recent years may reflect lower and more stable inflation, anchored inflation expectations, and central banks' reduced needed to respond aggressively to inflation shocks.

5.7 Additional central bank signals

I also performed a robustness check discussed earlier in the paper and tested whether other central bank signals may influence forecaster interest rate revisions. In particular, I included central banks revisions to their inflation and domestic output projections respectively as control variables in the benchmark regression from section 5. When we include these additional regressors, neither is even nearly significant. I then regress each individual private-sector macroeconomic forecasts revision on its corresponding central bank macroeconomic projection revision. That is, I regress private-sector inflation forecast revisions on central bank inflation projection revisions and private-sector domestic output forecast revisions on central bank domestic output projection revisions. Doing so seems to confirm our reasoning from earlier that this information is already embedded in each individual private-sector forecaster's revisions and so we should not control for central bank macroeconomic projection revisions in our regressions.

5.8 Central bank policy rate forecasts do not amplify forward guidance transmission

Numerous studies have argued that it may be desirable for inflation-forecast targeting central banks (Svensson (1997), Woodford (2007)) to publish their policy rate projections on a regular basis (e.g. Faust and Leeper (2005), Woodford (2005), Svensson (2006), Rudebusch and Williams (2008), Woodford (2013), Svensson (2015)). "A published policy rate should affect market expectations of future policy rates and thereby the yield curve and longer market rates that have an impact on economic agents' decision and this way contribute to a more effective implementation of monetary policy" (Svensson (2015)). Jain and Sutherland (2018) found that the publication of central bank policy rate projections did not contribute to lower private-sector forecast disagreement of any kind. Perhaps, however, policy rate projections influence the level of private-sector policy rate forecasts. For example, if a central bank lowers the path of its policy rate projection, do private-sector forecasters lower their policy rate forecasts commensurately?

To test these ideas, I gathered data on central bank policy rate projections. The Reserve Bank of New Zealand began releasing policy rate projections in 1997, the Norges Bank in 2005, the Sveriges Riksbank in 2007, and the FOMC started releasing its so-called dot plot in 2012. I created a binary variable indicating whether, in a given month, a given central bank released a policy rate projection. I then interacted this variable with f_{ct} (a change in forward guidance) in each of the regressions from Figure 9. The hypothesis is that, if the release of a central bank policy rate projection either strengthens or clarifies forward guidance, then the forward guidance should be able to influence expectations more easily. As such, we would expect to see a positive, statistically and economically significant coefficient on the interaction term between f_{ct} and the indicator variable for the release of a policy rate projection. We do not (Table 2). The coefficient is not significantly different from zero. The results were the same when the estimates were produced on a country-by-country basis and when a triple interaction term including the effective lower bound was included. This adds to a growing body of evidence (Jain and Sutherland (2018)) that central bank policy rate projections may not be as useful as commonly argued in the literature.

Interestingly, the estimated coefficients on the interaction term are often negative. In other words, these estimates suggest that releasing central bank projections alongside forward guidance can, on average, actually reduce the influence of forward guidance over the forecasted path of monetary policy to the extent that it is completely ineffective. At best, they have no effect. What might explain this result? Three non-mutually exclusive possibilities seem plausible.

First, forward guidance may be an effective tool for guiding the path of policy rate expectations in the short-term (say, a year or so), but central bank policy rate projections may create some confusion about the terminal policy rate. Indeed, there is significant disagreement about the level of the nominal neutral policy rate (or, as it is often referred to, r*) in the literature. In general, private-sector analysts tend to consider academic estimates of r* to be too high. Perhaps this negative coefficient provides some evidence of this tension.

Second and related, the (admittedly appropriate) amount of uncertainty surrounding central bank projections of the policy rate in the long term (typically about three years into the future), may preclude consensus building. Central banks often embed their mean or median policy rate projections within confidence intervals that expand in size as a positive function of the forecast horizon (Sims (2010)). Graphically, this produces a (usually colourful) cone shape that is frequently referred to as a fan chart. Indeed, this has been the case for the Reserve Bank of New Zealand, Norges Bank, and the Sveriges Riksbank. The FOMC, by contrast, provides longer-term policy rate forecasts for each member, which, naturally, reveals some disagreement within the FOMC (often a lot). This revealed uncertainty may create noise and make it impossible for central banks to influence expected policy rates beyond the horizon at which central banks can be reasonably confident.

Third, central bank commitments (or something resembling a commitment) may be more powerful than a forecast. When the Governing Council or Monetary Policy Committee of a central bank provides forward guidance, to an extent, it puts its credibility on the line. This is probably less true for one of the many forecasts within a monetary policy report or inflation report that a central bank provides on (typically) a quarterly basis.

Table 2: **Does the provision of a central bank policy rate projection influence how changes in forward guidance affect revisions to individuals' forecasts of rates?**

	[1]	[2]	[3]	[4]
(1) Forward guidance $\{-1, 0, 1\}$ change (+1)	3.33*	5.16***	0.90	1.98
	(1.70)	(1.31)	(0.89)	(1.47)
(2) Policy rate (PR) change (+25 bps)	11.39***	6.69***	2.72***	1.13**
	(0.82)	(0.62)	(0.53)	(0.46)
(3) Private inflation forecast revision (+25 bps)	1.46***	2.14***	1.48***	1.56***
	(0.21)	(0.34)	(0.19)	(0.25)
(4) Effective lower bound $\{0, 1\}$	-0.02	0.13	-1.78	-1.13
	(0.39)	(0.49)	(1.22)	(0.86)
(5) Quantitative easing $\{-1, 0, 1\}$ change (+1)	-2.03***	-1.53	-1.33	2.35
	(0.64)	(1.14)	(1.88)	(1.45)
(6) Policy rate projection $\{0, 1\}$	0.61	-2.47	-0.93	-1.82
	(1.36)	(1.83)	(1.17)	(1.50)
(1) x (6)	0.58	-0.53	0.76	-1.87
	(2.68)	(1.93)	(1.22)	(1.56)
Adjusted R^2	0.39	0.22	0.29	0.19
N	30036	30065	29091	29044

[1]: Revision to individuals' forecasts of the policy rate three months from now.

[2]: Revision to individuals' forecasts of the policy rate twelve months from now.

[3]: Revision to individuals' forecasts of the 10-year bond yield three months from now.

[4]: Revision to individuals' forecasts of the 10-year bond yield twelve months from now.

Clustered standard errors (at the country level) are shown in parentheses.

This table shows summary statistics from panel regressions with forecaster and month fixed effects.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6 Conclusion

The main finding is that the average revision to forecasts of the policy rate in twelve months' time in response to forward guidance is about five basis points in the intended direction. That is, hawkish or dovish forward guidance tends to move interest rate forecasts in the intended direction (up and down respectively) by about five basis points on average. This result is important for a number of reasons. These estimates provide central banks some guidance as to how much they can expect to influence expectations ex ante before implementing forward guidance. They also provide new evidence of the ability of forward guidance to influence survey expectations in the intended manner. The results also provide additional evidence that private forecasters update their expectations in the manner suggested by models of Bayesian learning (Coibion et al. (2018b)). Interestingly, the estimated average forward guidance effect is surprisingly small given that changes in forward guidance frequently provide material updates to the outlook for the path of policy rates. This suggests that perhaps central bank credibility issues, noise (Sims (2003), Woodford (2003)), or forecaster inattention (Mankiw and Reis (2002), Coibion and Gorodnichenko (2012a), Coibion and Gorodnichenko (2015), Coibion et al. (2018a), Coibion et al. (2018b)) may constrain forward guidance expectation formation but not so much so as to render forward guidance ineffective.

By addressing the data and identification challenges discussed above, I am able to make numerous other novel contributions. First, we can now better understand how and where forward guidance works. To begin with, the influence of forward guidance is much weaker during periods at the effective lower bound. This may help explain why I find that the influence of forward guidance over policy rate expectations has been weaker in recent years. It could also help explain why forward guidance has been the most effective in Canada, Australia, New Zealand, and Sweden—countries that started using forward guidance early in their histories. I also show that forward guidance can have spillover effects. For example, Federal Reserve forward guidance has had strong spillover effects to Canadian policy rate expectations.

Second, the only way to amplify forward guidance appears to be to make a commitment. Surprisingly, I find that quantitative easing announcements do not appear to amplify forward guidance. Similarly, when central banks released both forward guidance and policy rate forecasts, no amplification of the influence of forward guidance was detectable. Many papers have advocated for the publication of central bank policy rate forecasts (e.g. Faust and Leeper (2005), Woodford (2005), Svensson (2006), Rudebusch and Williams (2008), Woodford (2013), Svensson (2015)). The efficacy of central bank policy rate forecasts is still an open question, but the results in this paper add to the evidence provided by Jain and Sutherland (2018) that policy rate forecasts do not noticeably influence private-sector forecasters. Interestingly, state-contingent forward guidance actually appears to dampen effectiveness compared to time-contingent forward guidance. One potential explanation for these latter two findings is that the purpose of both state-contingent forward guidance and policy rate projections may be more so to communicate about the reaction function and monetary policy inertia (Coibion and Gorodnichenko (2012b)) than to amplify the influence of forward guidance. As such, studying forecast errors would probably be more instructive about the efficacy of these particular policy tools than analysing forecast revisions.

By contrast, Odyssean forward guidance is extremely rare, but perhaps extremely effective. Indeed, only two instances appear in the sample data that cover eight central banks and thirty years. Accordingly, our ability to study Odyssean forward guidance is limited. Nonetheless, estimates of the effect of Odyssean forward guidance, which, naturally, come with large standard errors, suggest that Odyssean forward guidance greatly amplifies the influence of forward guidance on expectation formation. In fact, it appears to double the effect. This result provides some empirical evidence for the ideas of Krugman et al.

(1998) and Eggertsson et al. (2003). These authors suggest that if a central bank can credibly commit to keep the policy rate below the levels suggested by the central bank's reaction function at some point in the future, then the central bank should be able to lower longer-term interest rates due to the expectations hypothesis. Indeed, I find that forward guidance has much larger effects on forecasts of the policy rate at the twelve-month forecast horizon than the three-month forecast horizon.

References

Altavilla, C. and D. Giannone (2017): "The Effectiveness of Non-Standard Monetary Policy Measures: Evidence from Survey Data," *Journal of Applied Econometrics*, 32, 952–964.

Amato, J. D., S. Morris, and H. S. Shin (2002): "Communication and monetary policy," *Oxford Review of Economic Policy*, 18, 495–503.

Amato, J. D. and H. S. Shin (2006): "Imperfect common knowledge and the information value of prices," *Economic Theory*, 27, 213–241.

Andersson, M., H. Dillén, and P. Sellin (2006): "Monetary policy signaling and movements in the term structure of interest rates," *Journal of Monetary Economics*, 53, 1815–1855.

Andersson, M. and B. Hofmann (2009): "Gauging the effectiveness of central bank forward guidance," *European Central Bank Working Paper*.

Andrade, P., F. Ferroni, et al. (2016): "Delphic and Odyssean monetary policy shocks: Evidence from the euro-area," *Discussion Papers in Economics, University of Surrey*, 12, 16.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019): "Forward guidance and heterogeneous beliefs," *American Economic Journal: Macroeconomics*.

Arellano, M. and S. Bond (1991): "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations," *The Review of Economic Studies*, 58, 277–297.

Arellano, M. and O. Bover (1995): "Another look at the instrumental variable estimation of error-components models," *Journal of Econometrics*, 68, 29–51.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How much should we trust differences-in-differences estimates?" *The Quarterly Journal of Economics*, 119, 249–275.

Besley, T., N. Meads, and P. Surico (2008): "Insiders versus outsiders in monetary policymaking," *American Economic Review*, 98, 218–23.

Bhattarai, S. and C. J. Neely (2018): "An Analysis of the Literature on International Unconventional Monetary Policy," *Federal Reserve Bank of St. Louis Working Paper Series*.

Bordo, M. D. and K. Istrefi (2018): "Perceived FOMC: The Making of Hawks, Doves and Swingers," *NBER Working Paper*, 24650.

Brubakk, L., S. ter Ellen, and H. Xu (2017): "Forward guidance through interest rate projections: does it work?" *Norges Bank working paper*.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 90, 414–427.

——— (2011): "Robust inference with multiway clustering," *Journal of Business & Economic Statistics*, 29, 238–249.

Campbell, J. R., C. L. Evans, J. D. Fisher, A. Justiniano, C. W. Calomiris, and M. Woodford (2012): "Macroeconomic effects of federal reserve forward guidance [with comments and discussion]," *Brookings Papers on Economic Activity*, 1–80.

Chang, B. Y. and B. Feunou (2013): "Measuring uncertainty in monetary policy using implied volatility and realized volatility," *Bank of Canada Working Paper*.

Charbonneau, K. and L. Rennison (2015): "Forward guidance at the effective lower bound: International experience," *Bank of Canada Staff Discussion Paper*.

COENEN, G., M. EHRMANN, G. GABALLO, P. HOFFMANN, A. NAKOV, S. NARDELLI, E. PERSSON, AND G. STRASSER (2017): "Communication of monetary policy in unconventional times," *ECB working paper*.

COIBION, O. AND Y. GORODNICHENKO (2012a): "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy*, 120, 116–159.

——— (2012b): "Why are target interest rate changes so persistent?" *American Economic Journal: Macroeconomics*, 4, 126–62.

——— (2015): "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105, 2644–78.

COIBION, O., Y. GORODNICHENKO, AND R. KAMDAR (2018a): "The formation of expectations, inflation, and the phillips curve," *Journal of Economic Literature*, 56, 1447–91.

COIBION, O., Y. GORODNICHENKO, AND S. KUMAR (2018b): "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, 108, 2671–2713.

COX, D. R. (1958): *Planning of Experiments*, Wiley.

DELL'ARICCIA, G., P. RABANAL, AND D. SANDRI (2018): "Unconventional Monetary Policies in the Euro Area, Japan, and the United Kingdom," *The Journal of Economic Perspectives*, 32, 147–172.

DETMERS, G.-A. AND D. NAUTZ (2012): "The information content of central bank interest rate projections: Evidence from New Zealand," *Economic Record*, 88, 323–329.

DREW, A., Ö. KARAGEDIKLI, ET AL. (2007): "Some Benefits of Monetary-Policy Transparency in New Zealand," *Czech Journal of Economics and Finance (Finance a uver)*, 57, 521–539.

EGGERTSSON, G. B. ET AL. (2003): "Zero bound on interest rates and optimal monetary policy," *Brookings papers on economic activity*, 2003, 139–233.

EHRMANN, M. AND M. FRATZSCHER (2007): "Communication by central bank committee members: different strategies, same effectiveness?" *Journal of Money, Credit and Banking*, 39, 509–541.

FAUST, J. AND E. M. LEEPER (2005): "Forecasts and inflation reports: An evaluation," in *A conference on Inflation Targeting: Implementation, Communication and Effectiveness, Sveriges Riksbank, Stockholm, June*, 10–12.

FEMIA, K., S. FRIEDMAN, AND B. SACK (2013): "The effects of policy guidance on perceptions of the fed's reaction function," *Federal Reserve Bank of New York Staff Reports*, 652.

FEROLI, M., D. GREENLAW, P. HOOPER, F. S. MISHKIN, AND A. SUFI (2017): "Language after liftoff: Fed communication away from the zero lower bound," *Research in Economics*, 71, 452–490.

FERRERO, G. AND A. SECCHI (2009): "The Announcement of Future Policy Intentions," *Banca d'Italia Working Papers*.

GAGNON, J., M. RASKIN, J. REMACHE, B. SACK, ET AL. (2011): "The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases," *International Journal of Central Banking*, 7, 3–43.

GÜRKAYNAK, R. S., B. SACK, AND E. T. SWANSON (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*.

HANSEN, S. AND M. MCMAHON (2016): "Shocking language: Understanding the macroeconomic effects of central bank communication," *Journal of International Economics*, 99, S114–S133.

Hansen, S., M. McMahon, and A. Prat (2017): "Transparency and deliberation within the FOMC: a computational linguistics approach," *The Quarterly Journal of Economics*, 133, 801–870.

Hansen, S., M. McMahon, and C. V. Rivera (2014): "Preferences or private assessments on a monetary policy committee?" *Journal of Monetary Economics*, 67, 16–32.

Hansen, S., M. McMahon, and M. Tong (2018): "The Long-Run Information Effect of Central Bank Communication," *CEPR Discussion Paper. Centre for Economic Policy Research, London. Huang, Y* .

Hubert, P. (2014): "FOMC forecasts as a focal point for private expectations," *Journal of Money, Credit and Banking*, 46, 1381–1420.

——— (2016): "Qualitative and quantitative central bank communication and inflation expectations," *The BE Journal of Macroeconomics*, 17.

Hubert, P. and F. Labondance (2018a): "Central Bank Sentiment," *Working Paper*.

——— (2018b): "The Effect of ECB Forward Guidance on Policy Expectations," *International Journal of Central Banking*, 14, 193–222.

Imai, K. and I. S. Kim (2019): "When Should We Use Linear Fixed Effects Regression Models for Causal Inference with Longitudinal Data?" *American Journal of Political Science*.

Istrefi, K. (2016): "Quantifying Perceived Policy Preferences of the FOMC: 1960-2015," *Working Paper*.

Jain, M. and C. S. Sutherland (2018): "How Do Central Bank Projections and Forward Guidance Influence Private-Sector Forecasts?" *Bank of Canada Discussion Paper*, SWP 2018-2.

Kool, C. and D. Thornton (2015): "How Effective Is Central Bank Forward Guidance?" *Federal Reserve Bank of St. Louis Review*, 97, 303–22.

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

Krugman, P. R., K. M. Dominquez, and K. Rogoff (1998): "It's baaack: Japan's slump and the return of the liquidity trap," *Brookings Papers on Economic Activity*, 1998, 137–205.

Kuttner, K. N. (2001): "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market," *Journal of Monetary Economics*, 47, 523–544.

——— (2018): "Outside the Box: Unconventional Monetary Policy in the Great Recession and Beyond," *The Journal of Economic Perspectives*, 32, 121–146.

Lustenberger, T. and E. Rossi (2018): "The Social Value of Information: A Test of a Beauty and Non-Beauty Contest," *SNB Working Paper*.

Mankiw, N. G. and R. Reis (2002): "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

Mankiw, N. G., R. Reis, and J. Wolfers (2003): "Disagreement about inflation expectations," *NBER Macroeconomics Annual*, 18, 209–248.

Moessner, R. (2013): "Effects of explicit FOMC policy rate guidance on interest rate expectations," *Economics Letters*, 121, 170–173.

Moessner, R., D.-J. Jansen, and J. de Haan (2017): "Communication about future policy rates in theory and practice: A survey," *Journal of Economic Surveys*, 31, 678–711.

Moessner, R. and W. R. Nelson (2008): "Central Bank Policy Rate Guidance and Financial Market Functioning," *International Journal of Central Banking*.

Morris, S. and H. S. Shin (2002): "Social value of public information," *The American Economic Review*, 92, 1521–1534.

——— (2005): "Central bank transparency and the signal value of prices," *Brookings Papers on Economic Activity*, 2005, 1–66.

——— (2007): "Optimal communication," *Journal of the European Economic Association*, 5, 594–602.

——— (2018): "Central bank forward guidance and the signal value of market prices," in *AEA Papers and Proceedings*, vol. 108, 572–77.

Obstfeld, M., K. Clinton, O. Kamenik, D. Laxton, Y. Ustyugova, and H. Wang (2016): "How to Improve Inflation Targeting in Canada," *IMF Working Paper*.

Pearl, J. (2009): *Causality*, Cambridge University Press.

Roodman, D. (2006): "How to do xtabond2: An introduction to difference and system GMM in Stata," *The Stata Journal*.

Rubin, D. B. (1980): "Randomization analysis of experimental data: The Fisher randomization test comment," *Journal of the American Statistical Association*, 75, 591–593.

——— (2005): "Causal inference using potential outcomes: Design, modeling, decisions," *Journal of the American Statistical Association*, 100, 322–331.

Rudebusch, G. D. and J. C. Williams (2008): "Revealing the secrets of the temple: The value of publishing central bank interest rate projections," in *Asset Prices and Monetary Policy*, University of Chicago Press, 247–289.

Shiller, R. J. (2017): "Narrative economics," *American Economic Review*, 107, 967–1004.

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

——— (2010): "Rational inattention and monetary economics," *Handbook of Monetary Economics*, 3, 155–181.

Stock, J. H. and M. W. Watson (2008): "Heteroskedasticity-robust standard errors for fixed effects panel data regression," *Econometrica*, 76, 155–174.

Svensson, L. E. (1997): "Inflation forecast targeting: Implementing and monitoring inflation targets," *European economic review*, 41, 1111–1146.

——— (2006): "Social value of public information: Comment: Morris and shin (2002) is actually pro-transparency, not con," *American Economic Review*, 96, 448–452.

——— (2010): "Policy expectations and policy evaluations: The role of transparency and communication," *Sveriges Riksbank Economic Review*, 1, 43–78.

——— (2015): "Forward Guidance," *International Journal of Central Banking*, 11, 19–64.

Swanson, E. T. (2017): "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets," *NBER Working Paper No. 23311*.

Swanson, E. T. and J. C. Williams (2014a): "Measuring the effect of the zero lower bound on medium- and longer-term interest rates," *American Economic Review*, 104, 3154–85.

——— (2014b): "Measuring the effect of the zero lower bound on yields and exchange rates in the UK and Germany," *Journal of International Economics*, 92, S2–S21.

Tang, J. (2013): "Uncertainty and the signaling channel of monetary policy," *working paper*.

Woodford, M. (2003): "Knowledge, Information and Expectations in Modern Macroeconomics: Essays in Honor of Edmund Phelps, chapter Imperfect common knowledge and the effects of monetary policy," *P. Aghion, R. Frydman, J. Stiglitz & M. Woodford, Princeton: Princeton University Press.[Links]*, 25–58.

——— (2005): "Central bank communication and policy effectiveness," *NBER Working Paper No. 11898*.

——— (2007): "The case for forecast targeting as a monetary policy strategy," *Journal of Economic Perspectives*, 21, 3–24.

——— (2013): "Forward guidance by inflation-targeting central banks," *Working paper*.

Wooldridge, J. M. (2005): "Fixed-effects and related estimators for correlated random-coefficient and treatment-effect panel data models," *Review of Economics and Statistics*, 87, 385–390.

——— (2010): *Econometric analysis of cross section and panel data*, MIT press.

A Appendix: Summary Statistics

Table 3: **Summary statistics: central bank policy changes and private-sector forecast revisions**

	Mean	SD	Count	Min	Max
Forward guidance $\{-1, 0, 1\}$ change	-0.02	0.24	30003	-1	1
Quantitative easing $\{-1, 0, 1\}$ change	-0.01	0.15	30003	-1	1
Effective lower bound $\{0, 1\}$	0.18	0.38	30003	0	1
Policy rate change (+25 basis points (bps))	-0.09	0.86	30003	-8	10
Private policy rate forecast revision, 3-month horizon (bps)	-2.07	26.71	30003	-300	290
Private policy rate forecast revision, 12-month horizon (bps)	-1.84	30.84	30003	-330	300
Private bond yield forecast revision, 3-month horizon (bps)	-1.84	27.10	30003	-259	219
Private bond yield forecast revision, 12-month horizon (bps)	-1.68	27.54	30003	-267	255
Private inflation forecast revision (+25 bps)	-0.02	0.83	30003	-13	12
Private domestic output forecast revision (+25 bps)	-0.05	1.01	30003	-16	8

Notes: +25 bps: variable multiplied by four so regression coefficients can be interpreted as 25 basis point shifts.

B Appendix: Robustness

B.1 Robustness: Placebo Interventions

In this section, we follow the lead of Bertrand et al. (2004) who generate placebo interventions (state-level laws in their case) and replace f_{ct} in Equation 13 with a placebo forward guidance intervention to test whether any bias exists that would give us a false positive for the forward guidance effect observed in our results. Each placebo is randomly drawn from $\{-1, 0, 1\}$. The random draws are structured so that the proportions of each member of the set are the same in the placebo data as the actual forward guidance data. Approximately 95% of the data take the value 0, approximately 3% take the value -1, and approximately 2% take the value $+1$. The figure below shows 100 estimates of the average forward guidance effect and their corresponding 95% confidence intervals, which confirm the absence of bias.

Figure 11: **Results from simulations using randomly generated placebo variables for a change in forward guidance**

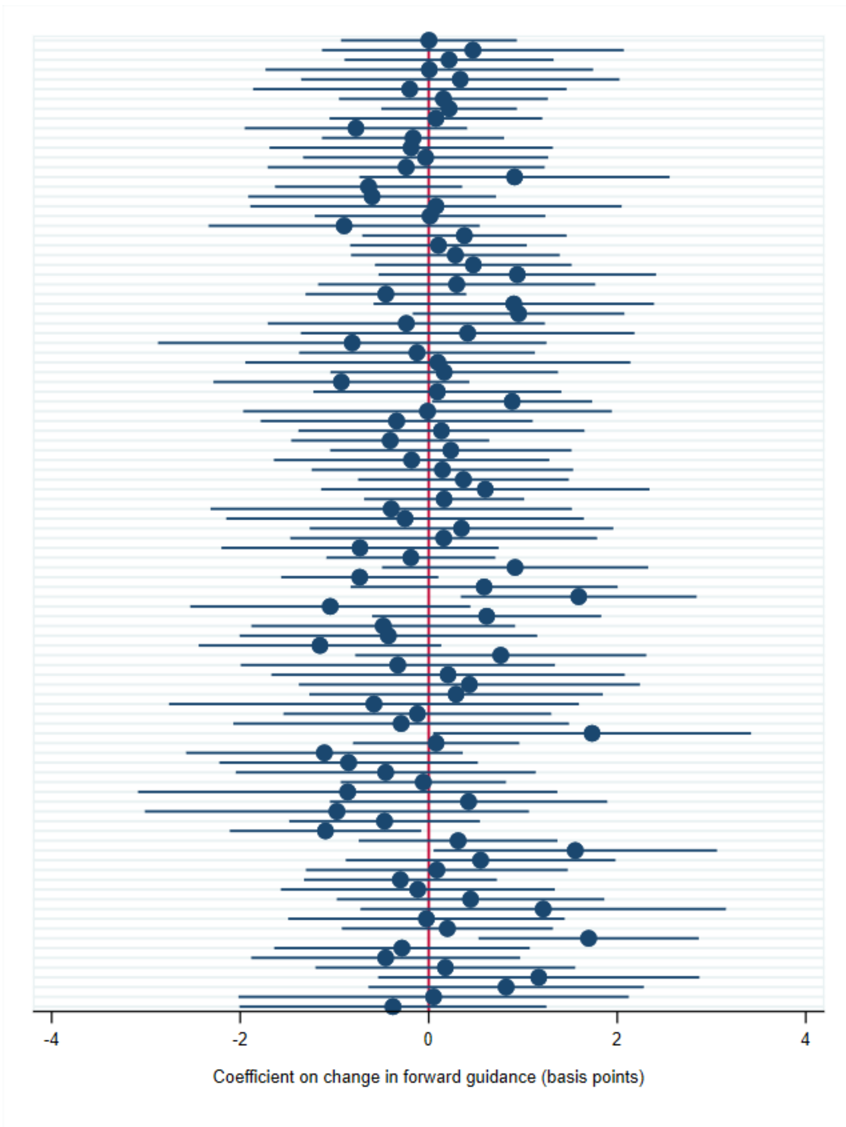

B.2 Robustness: Generalized Method of Moments

In this section, we consider the possibility that forward guidance is not strictly exogenous. To obtain valid estimates under this assumption, we use the Arellano-Bond (Arellano and Bond (1991)) and Arellano-Bover/BlundellBond (Arellano and Bover 1995; Blundell and Bond 1998) dynamic panel estimators. These estimators make a number of assumptions (Roodman (2006)). First, they assume small T, large N panel data sets. In this paper, we have a large T, large N data set. Second, they assume a linear functional relationship. In this paper, we do the same. Third, they assume that the model includes a lagged dependent variable. Fourth, they assume that the independent variables are not strictly exogenous. Thus far, we have assumed something close to strict exogeneity, but we relax that assumption in this section in line with the last two assumptions above. Fifth, they assume time-invariant individual heterogeneity. Sixth, they assume that this individual heterogeneity is orthogonal to the matrix of instruments. Seventh, they assume within-unit heteroscedasticity and autocorrelation.

To estimate this model, we use equation Equation 13; however, we make a couple of adjustments to limit the number of instruments. First, we use yearly time dummies instead of monthly time dummies. Second, we omit the calendar controls, which turn out to be insignificant in this analysis anyhow. Third, we omit the effective lower bound control variable, which also turns out to be insignificant.

To use this estimator, we must take a stance on the exogeneity of each regressor in order to construct an instrument matrix, which is constructed using vectors of lagged regressors. Here, forward guidance and quantitative easing are no longer treated as strictly exogenous. Of course, the lagged dependent variable is considered to be endogenous. All other regressors are treated as strictly exogenous, which is more consistent with the approach taken thus far.

All models are estimated using one-step system generalized method of moments (GMM). I use five lags of the instruments that are treated as not strictly exogenous. I collapse the instruments which results in a matrix of about 50 instruments. The instrument matrix uses first differences and the results are robust to the use of orthogonal deviations as well (Arellano and Bover (1995)). I use standard errors that are robust to heteroskedasticity and arbitrary patterns of autocorrelation within individuals and cluster on country as all countries are used in this panel.

The table below shows the estimates. Overall, the estimates are very similar to those obtained using a standard fixed effects estimator. Hence, when forward guidance is no longer treated as strictly exogenous, the results are substantially unchanged.

Table 4: **GMM Estimation: how do changes in forward guidance affect revisions to individuals' forecasts of the policy rate level twelve months from now?**

	[1]
Forward guidance $\{-1, 0, 1\}$ change (+1)	7.47***
	(1.77)
First lag of forward guidance change	5.21***
	(1.64)
Second lag of forward guidance change	3.12**
	(1.40)
Quantitative easing $\{-1, 0, 1\}$ change (+1)	-0.81
	(1.01)
Policy rate (PR) change (+25 bps)	2.22*
	(1.09)
Private inflation forecast revision (+25 bps)	1.98***
	(0.46)
Private domestic output forecast revision (+25 bps)	2.51***
	(0.33)
Constant	-3.76
	(24.66)
N	32590

Standard errors in parentheses

Dependent variable: revision to individuals' forecasts of the policy rate level twelve months from now.

Clustered standard errors (at the country level) are shown in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3 Robustness: Stable Unit Treatment Value Assumption (SUTVA)

Recall that in subsection 3.3 I made the *stable unit treatment value assumption* (Rubin (1980)), which comprises two sub-assumptions. First, it assumes that there is no interference between units (Cox (1958)); that is, neither $Y_i(1)$ nor $Y_i(0)$ is affected by what action any other unit took (using the notation from Fisher (1923) that indicates treatment (1) and control (0)). In other words, we assume that individual private-sector forecasters form their forecasts independently of one another, or, without collusion. This assumption is especially plausible given the structure of the data. Each month, each forecaster uses a template provided by the data provider to complete his or her forecasts. Once completed, the forecaster returns the completed template to the data provider before the monthly deadline. Once all submissions have been received and the deadline has passed, the data provider consolidates the forecasts into one spreadsheet and distributes the file to its subscribers weeks later. Absent collusion, the forecasters are incapable of contaminating one another's forecasts. Of course, we should not discount the possibility of *some* amount of collusion. For example, one forecaster could always have discussions with another, or worse yet, one forecaster could release his or her forecasts to the public before the monthly deadline. I argue that this type of collusion is by far the exception rather than the rule, but for robustness, we will probe this assumption further later in the paper. Second, and more formally, it assumes that there are no hidden versions of treatments; no matter how forecaster i received treatment (f_{ct}), the outcome that would be observed would be $Y_i(1)$ and similarly for treatment 0" (Rubin (2005), page 323). We can test the validity of this assumption using the robustness check below. The regression below re-estimates Equation 13 after collapsing all individual forecasters data to mean values. That is, I collapse the data by taking the mean of each forecast revision within each country and each month. The results are very similar to those from section 5.

Table 5: **How do changes in forward guidance affect revisions to individuals' forecasts of the policy rate level twelve months from now?**

	[1]
(1) Forward guidance $\{-1, 0, 1\}$ change (+1)	4.26***
	(1.14)
(2) Policy rate (PR) change (+25 bps)	6.08***
	(0.69)
(3) Private inflation forecast revision (+25 bps)	2.38
	(1.41)
(4) Effective lower bound $\{0, 1\}$	0.57
	(0.56)
(5) Quantitative easing $\{-1, 0, 1\}$ change (+1)	-1.25
	(0.98)
Adjusted R^2	0.55
N	3154

The regression is performed on **mean** monthly forecast revisions.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.4 Robustness: Additional Macroeconomic Controls

In this section, I consider whether changing the macroeconomic control variables in Equation 13 influence the estimates of the effect of a change in forward guidance. The estimates of β are not at all sensitive to the omission or addition of these controls. Column [1] does not include any macroeconomic controls and hence shows the effect of omitting π_{ict} from Equation 13. Column [2] includes both π_{ict}, each individual forecaster's revision to his or her current-year inflation forecast, as well as the revision for next-year inflation. Column [3] is the same as column [2] except that it adds each individual forecaster's revision to both his or her current-year and next-year domestic growth forecast. To follow the controls used in Altavilla and Giannone (2017), we would add revisions to current-year unemployment forecasts as well, but these data are only available for eight of the twelve countries. As such, the sample group would not be comparable to the estimates below and so these estimates are omitted, but nonetheless, the results are very similar to those reported below.

Table 6: **How do changes in forward guidance affect revisions to individuals' forecasts of the policy rate level twelve months from now? (various macroeconomic controls configurations)**

	[1]	[2]	[3]
Forward guidance $\{-1, 0, 1\}$ change (+1)	5.04***	4.80***	4.71***
	(1.20)	(1.17)	(0.98)
Policy rate (PR) change (+25 bps)	6.73***	6.58***	6.40***
	(0.62)	(0.62)	(0.72)
Quantitative easing $\{-1, 0, 1\}$ change (+1)	-1.28	-1.25	-1.72
	(0.92)	(0.97)	(0.99)
Effective lower bound $\{0, 1\}$	0.28	0.16	-0.20
	(0.43)	(0.44)	(0.42)
Current-year inflation revision (+25 bps)		1.48***	1.41***
		(0.33)	(0.33)
Next-year inflation revision (+25 bps)		1.90***	1.76***
		(0.25)	(0.28)
Current-year GDP revision (+25 bps)			2.96***
			(0.64)
Next-year GDP revision (+25 bps)			1.73***
			(0.24)
Adjusted R^2	0.22	0.23	0.24
N	31176	30513	30461

Dependent variable: revision to individuals' forecasts of the policy rate level twelve months from now.

Clustered standard errors (at the country level) are shown in parentheses.

This table shows summary statistics from panel regressions with forecaster and month fixed effects.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$