

# When Less is More: Central Bank Transparency and the Publication of Monetary Policy Committee Voting Records\*

Caitlin T. Ainsley<sup>†</sup>

February 5, 2020

## Abstract

Should monetary policy committees publish voting records? This is a question at the heart of the current debate over the optimal design of central banks. Due to the sweeping transparency reforms of the 1990s, it has been difficult to isolate the causal effect of publishing voting records. I overcome this challenge by analyzing a quasi-natural experiment in Brazil in which a national Freedom of Information law mandated the publication of voting records and incidentally was applied to the central bank. Contrary to conventional wisdom concerning the benefits of transparency for the predictability of monetary policy and inflation, I show publishing voting records corresponded with less accurate private inflation forecasts and reduced markets responsiveness to central bank communication.

white.

Word Count: 9814

---

\*My thanks to Chris Adolph, Jeff Arnold, Cliff Carrubba, Adam Glynn, Julia Gray, Robert Franzese, Dorothy Kronick, and Eric Reinhardt as well as seminar participants at Emory University, the Browne Center at the University of Pennsylvania, and University of Washington.

<sup>†</sup>Assistant Professor, Department of Political Science, University of Washington. 101 Gowen Hall, Box 353530, Seattle, WA, 98195. cainsley@uw.edu

Since the global financial crisis, policymakers have shown renewed interest in the transparency and accountability of the monetary policymaking process. While the “transparency revolution” of the late nineties and early aughts made central banks far less opaque (Blinder 2004), the role non-elected policymakers at central banks played during both the financial crisis and recovery has spurred criticism and calls for greater openness and opportunities for public oversight. Recently, advocates of improved transparency have focused primarily on the publication of voting records, as in many countries central banks remain one of the few public decision-making bodies with secret votes. For example, in recent years EU lawmakers and some central bankers themselves have argued the legitimacy and accountability of the European Central Bank hinges on a more transparent record of individual votes (Curtin 2017). Similar discussions are ongoing in Australia, India, Taiwan, and elsewhere.

However, academics and policymakers increasingly agree that in the case of monetary policymaking, the *optimal* level of transparency may not be the *maximum* level transparency (Van Der Cruijssen, Eijffinger and Hoogduin 2010; Issing 2014). This is an important departure from much of the early work on transparency which focused on composite measures and aggregate effects of increased information concerning the policymaking process and outcomes. While substantial evidence suggests monetary policy and economic performance indicators improve with increased transparency in aggregate, we know surprisingly little about the implications of individual policy changes that would lead us closer to identifying the elusive, *optimally transparent* central bank. In addressing the publication of voting records, this paper contributes to this empirical gap in the literature by examining the implications of one of the most highly debated areas surrounding central banks and monetary policymaking today.

No consensus exists as to whether monetary policy committees (MPCs) should publish voting records from rate-setting meetings. While scholars have warned against such transparency measures that open the monetary policymaking process up to the private interests of shadow principals (Adolph 2013; Hansen, McMahon and Prat 2018), advocates of publishing voting records claim they provide both a mechanism for public accountability and aid in the anchoring of inflation expectations by closing the information asymmetry between the market

and policymakers (Buitter 1999; Van Der Cruijssen, Eijffinger and Hoogduin 2010; Neuenkirch 2012). In the contemporary economic climate, this objective of improving the performance of market expectations cannot be overstated. In a statement central bankers themselves often quote, economist Michael Woodford (2003) says that “not only do expectations about policy matter, but, at least under current conditions, very little else matters.” Thus, while many have demonstrated the ways in which transparency can open central bankers up to pressure from private interests (Adolph 2013; Hansen, McMahon and Prat 2018), it may still be desirable if publishing voting records affords markets valuable information that improves the quality of private forecasts. In this case, the optimally transparent central bank likely moves forward publishing voting records while building in necessary safeguards for the aforementioned potential threats to independence.

This article investigates whether the publication of attributed voting records does improve the performance of private forecasts and anchoring of inflation expectations. While the debates over the transparency of voting records have often implicitly assumed this relationship exists, there is limited empirical evidence of this causal claim in the extant literature. Drawing on a quasi-natural experiment occurring with the 2012 implementation of a Freedom of Information (FOI) law in Brazil, I provide evidence that the publication of voting records both worsened the accuracy of market expectations and limited the central bank’s ability to influence and manage expectations. I examine this relationship in three stages. First, I explore the consequences of publishing voting records for the voting behavior of monetary policy committee (MPC) members. In a descriptive examination of voting records that employs matching methods before and after the transparency intervention, I demonstrate the shift from publishing anonymous vote-splits to attributed voting records has suppressed dissent on Brazil’s MPC. While this result is consistent with seminal findings in the extant literature on central bank transparency (Meade and Stasavage 2008), it highlights a distinct mechanism for this behavior. Since the policy intervention under examination shifted reporting practices from an anonymous vote-split to attributed votes, the declining rate of dissent cannot be explained as it often has as an effort to hide internal disagreements from markets.

Given this finding that minutes containing attributed voting records appear to report higher levels of consensus, I next turn to how this change in reporting practices affects how market expectations respond to the publication of minutes. In light of an extant literature that emphasizes dissenting votes as key predictors of future rate changes, I argue publishing attributed voting records ought to dampen the responsiveness of private actors' expectations over future rates to the publication of minutes. Drawing on daily survey data reporting private expectations over the central bank's primary policy instrument, the Selic interest rate, I demonstrate that the publication of attributed voting records has diminished the central bank's ability to influence market expectations through its announcements following rate-setting meetings.

Finally, I consider the consequences of publishing voting records for the accuracy of inflation expectations. To estimate the effect publishing voting records has had on the accuracy of inflation expectations, I draw on time series of monthly inflation expectations survey data and employ a pair of Bayesian structural time series models to estimate synthetic control series reflecting the accuracy of inflation expectations in the absence of the transparency policy change. In the first model, I construct a synthetic control for the accuracy of inflation expectations from responses in the same survey for expectations over other economic outcomes unrelated to inflation. While this design has the advantage of incorporating in the synthetic control series idiosyncratic changes affecting the formation of forecasts in Brazil during the post-intervention period, it cannot incorporate idiosyncratic changes in inflation forecasting unrelated to the transparency policy. To address this, I estimate a second synthetic control series with surveys of market inflation expectations in Chile, Mexico, Peru, Israel, Poland, and Thailand. In contrast to the first approach, this synthetic control better approximates the performance of inflation expectations and issues related specifically to inflation, but cannot account for idiosyncratic changes occurring in Brazil during the post-intervention period. Fortunately, comparing the observed inflation forecast accuracy to either synthetic control yields the same inference: since the BCB began publishing individually attributed voting records, inflation forecast inaccuracy has increased significantly. Taken together, publishing

voting records appears to have limited the ability of the BCB to shape private expectations over future policy and proven detrimental to the accuracy of inflation expectations.

The implications of these results for governments and central banks currently evaluating their transparency practices in an effort to improve institutional accountability and avoid a recursion of the recent financial crisis in the wake of the financial crisis are threefold. First, given that publishing voting records likely has not been improving a central bank’s ability to manage expectations, this suggests other transparency measures included in aggregate indices must improve the performance of market expectations *more* than existing work has suggested — that is, enough to outweigh the detrimental effects of publishing voting records. Thus, relatively minor reforms to transparency practices stopping short of publishing voting records could create considerable gains. Second, the results also present a cautionary tale for governments and legislatures considering the broad application of FOI laws across institutions and particularly to central banks. Finally, if the current trend of rapid institutional diffusion applies to the publication of voting records, researchers should prepare to examine the implications of heightened uncertainty in the form of inaccurate inflation expectations as well as to consider institutional solutions to remedy this consequence of this particular form of central bank transparency.

## Background

The informational advantages central banks once sought to leverage over the market in order to generate policy surprises are now the subject of considerable debate as policymakers seek to eliminate them. Research on central bank transparency focuses on this goal. Two interrelated objectives underlie motivation for central bank transparency. First, policymakers and researchers widely accept that effective monetary policy hinges on the management of market expectations. As former Governor of the Bank of England Sir Mervyn King put it, central banking ought to be as “boring” as possible in the sense that markets are never (or rarely) surprised (2000). As such, the primary goal of most transparency policies is to

	Definition	Example Policies
<b>Political</b>	explicitly stated policy objectives clarifying MPC members' motives	stating explicit inflation targets
<b>Economic</b>	data that informs monetary policy decisions and MPC models	publishing CB forecasts
<b>Procedural</b>	information about the policy meetings and how decisions are made	publishing meeting minutes and votes
<b>Policy</b>	punctual announcement and explanation of policy decisions following meetings	policy statements and press conferences
<b>Operational</b>	description of monetary policy implementation, including error and unforeseen shocks	announcements of unanticipated shocks

Table 1: Typology of Central Bank Transparency from Geraats (2002)

provide market actors the appropriate information so that their inflation expectations are accurate and consistent. A second though closely related objective of transparency policies is to facilitate the accountability of central banks to the public interest. If modern central banks are to remain independent from political influence and oversight, decision-makers call for public oversight and accountability to fill the gap.

Central banks have implemented many practices and policies to achieve transparency. The seminal works of Fry et al. (2000), Geraats (2002), and Eijffinger and Geraats (2006) provide conceptual frameworks to discuss and measure central bank transparency. From this work emerged the five core categories of central bank transparency: political transparency, economic transparency, procedural transparency, policy transparency, and operational transparency. Table 1 describes these categories and policies that embody them. This framework has been the basis for a large body of cross-national research on the implications of central bank transparency and an organized discussion of the objectives of central bank transparency and evaluation of transparency policies.

The policies noted in Table 1 are increasingly common in central banks. In a sample of 121 countries, Dincer and Eichengreen (2014) demonstrate that between 1998 and 2010 all but 11 implemented policies increasing their central banks' transparency on some dimension. Further,

only one of those 11 exceptions (Uruguay) reduced its transparency during this period.<sup>1</sup> While there were increases on all dimensions of transparency described in Table 1, economic and policy transparency exhibited the greatest gains, due in large part to the growing popularity of explicit inflation targets. In aggregate, countries increased transparency an average of 2.3 points on the 15 point index, from 3.2 to 5.5. Central banks in developed countries remain the most transparent, but 1998–2010 saw meaningful increases in transparency in emerging and developing economies as well.

A substantial literature addresses the effects of transparency generally on the performance of market expectations and of publishing voting records specifically. However, this extant literature often relies on an additive index of transparency that cannot as an explanatory variable evaluate the consequences of discrete policy changes. Of particular interest here is the index’s component of procedural transparency which includes three parts.<sup>2</sup> The first two parts correspond to the presence of an explicit policy rule and timely publication of minutes, respectively, while the third part, of interest here, corresponds to the publication of voting records. However, the publication of voting records is not a binary indicator, as applied research has often treated it. Instead, it can take one of three values: “0” if no voting record is published, “ $\frac{1}{2}$ ” if a non-attributed voting record is published, and “1” if individual voting records are published (or a single central banker determines the outcome).

Estimating the effect of publishing voting records with this indicator requires a pair of related theoretical assumptions I will challenge both theoretically and empirically. First, using this measure to estimate an effect of procedural transparency requires an assumption that there is a constant, linear relationship from publishing no voting record to an unattributed

---

<sup>1</sup>Uruguay ceased publishing inflation reports in English in 2006 but has since reinstated the practice.

<sup>2</sup> The indices that extend the dataset described in Eijffinger and Geraats (2006) employ this variable, including Siklos (2011) and Dincer and Eichengreen (2014), among others. Most researchers have referred to the measure of procedural transparency as index 3, and the sub-sections as 3A-3C, referring to the numbering of the initial survey. I focus here on 3C.

vote-split to individually attributed votes. In other words, this implies we predict that publishing the anonymous vote-split is half as effective at informing market expectations as publishing a complete non-anonymous voting record. A second more implicit assumption necessary for this indicator to be useful for the purposes at hand is that the three levels of procedural transparency either do not affect the voting behavior of central bankers or have a constant effect on voting behavior across levels. In the section that follows, I examine these assumptions and derive a series of testable empirical implications from a more nuanced theory of how central bank transparency in the form of attributed voting records affects the behavior of central bankers themselves and, in turn, market expectations. At best, I argue, complete voting records contain the same information as an anonymous vote-split, but more frequently than not they contain *less information*. This intuition makes the assumptions underlying the use of the aggregate index particularly consequential when we are interested in estimating its effect on some quantity of interest, such as inflation expectations.

In spite of its limitations with respect to understanding the consequences of publishing voting records, the existing literature contains evidence that individual measures of transparency might have differing impacts on market expectations. While the effect of transparency as measured by the entire aggregate index on forecast accuracy is generally positive (Van Der Cruysen and Demertzis 2007; Neuenkirch 2012, 2013), Ehrmann, Eijffinger and Fratzscher (2012) find that transparency has decreasing marginal effects. As the authors speculate, this could suggest that the optimal level of transparency is less than the maximum. One possible explanation for this is that some indicators of transparency, such as publishing voting records, do not boost the performance of market expectations. Another empirical result my argument would explain comes from Neuenkirch (2012), who examines the effect of individual components of the transparency indices. In a sample of ten countries — five of which publish voting records — he cites a positive effect of the categorical voting index on the performance of market expectations. However, this component has the smallest effect of the entire index. While this could suggest publishing voting records in any form has a small effect, it is also consistent with the theoretical argument that while an anonymous vote-split ( $Q3C = \frac{1}{2}$ ) has

a significant positive effect, publishing a complete voting record ( $Q3C = 1$ ) has a negative effect.

Extant research also contains evidence that aggregate transparency *worsens* the performance of market expectations. Siklos (2013), Naszodi et al. (2016), and Lustenberger and Rossi (2017) all find a negative relationship between aggregate transparency and the accuracy and dispersion of market expectations. The samples these analyses examine are of particular interest given the findings. Siklos (2013) and Naszodi et al. (2016) use samples in which over half of the countries publish individual voting records. While Lustenberger and Rossi (2017) looks at a much larger sample of countries in which fewer publish voting records, the disaggregated sample reveals a pattern consistent with expectations given the issue I have identified with the measure. No region shows a significant improvement in market expectations with increased procedural transparency, though Latin America shows an insignificant improvement and Eastern Europe shows a significant worsening of performance. This is entirely consistent with this paper's argument given the policies that are in place in these regions. Several of the major economies in Eastern Europe began publishing individual voting records during the period Lustenberger and Rossi (2017) examined, 1998 to 2014. By contrast, in Latin America several countries began including anonymous vote-splits in their minutes, but only Brazil published attributed voting records and did so at the very end of the period.

## Theory

While these seemingly competing empirical results are often presented as being at odds with one another, I will argue they are consistent with the following theory that accounts for and differentiates between the first- and second-order consequences of publishing central bank voting records. To understand how publishing voting records affects market expectations, we must examine both the effect such policies have on the informational content of the records released by monetary policy committees as well as how changes in that informational content

affect market expectations. When these two stages are treated individually, the aforementioned problem with the aggregated index for voting records becomes readily apparent.

Advocates of central bank transparency often argue that by providing private market participants with more information about central banks' objectives and the policymaking process, increased transparency will improve the predictability of monetary policy. In aggregate, this claim is not controversial. Specifically with respect to the publication of voting records, there is ample empirical evidence suggesting that dissenting votes cast at MPC rate-setting meetings are predictive of future rate changes (Gerlach-Kristen 2004; Horvath, Smidkova and Zapal 2012; Horváth and Jonášová 2015). Therefore, beginning with the first incremental change in transparency from no voting record to the publication of an anonymous vote split, I expect the informational content of MPC minutes to *weakly* increase and improve the predictability of monetary policy. The reasons I expect this effect to be only *weakly* positive are important and substantively motivated. First, a common concern central bankers themselves express about the publication of voting records is how markets will react — and potentially over-react — to the revelation of a non-unanimous vote and disagreement on the MPC. Because MPCs have historically put a premium on achieving consensus, a policy change to publishing anonymous vote-splits may suppress any dissent that exists on the final vote. In practice, this leads to an informational environment that is observationally equivalent to when no voting record was published at all. However, when vote-splits are recorded but individual votes remain anonymous, dissenting votes that are cast are more likely to reflect true policy positions rather than noisy signals to external government or private principals. While scholars have long-cited the potential capture of central bankers who might be tempted to cast their vote in such a way that curries favor from the government or private industry (Adolph 2013; Hansen, McMahon and Prat 2018), the opportunity for such capture to occur is limited when the votes remain anonymous. Therefore, to the extent central bankers cast dissenting votes when the vote-split is public record, the dissents should be the kind that are informative of actual policy positions on the MPC.

Next, consider how moving from publishing an anonymous vote-split to publishing individually-attributed voting records might affect the voting behavior of monetary policy members through these same two mechanisms. Here, central bankers' incentives to appear as if they agree unanimously on a policy decision should similarly affect voting behavior under the two transparency policies. Publishing votes may deter dissent if central bankers wish for the committee to appear unanimous in their decisions, but this mechanism ought to similarly deter dissent regardless of whether votes are individually attributed or not. However, the second mechanism under consideration — whereby transparency exposes central bankers to external interests and capture — should affect voting behavior differently depending on whether the voting record is anonymous or individually-attributed. Where there exist private interests applying pressure on monetary policymakers' votes, policymakers may vote differently depending on whether their name is attached to their vote. I expect this potential mechanism to further suppress dissent from the majority policy position when votes are individually attributed, weakly reducing the informational content of MPC voting records. Taken together, this logic leads to the following testable implication:

**Hypothesis 1:** *When central banks publish individually attributed voting records from rate-setting meetings, committee members will be less likely to cast dissenting votes than they were when only anonymous vote-splits were published.*

This reflects a theoretical expectation that the identifiability of individual votes affects the voting behavior of central bankers. While this is hardly a novel argument and scholars have found empirical support for similar claims across a variety of contexts (Meade and Stasavage 2008), it is an important link here for our understanding of how publishing voting records may affect the formation of market expectations and why we would observe variation across anonymous and attributed voting records. Critically, if this relationship is in fact borne out in the data, it suggests that we should not assume the information revealed from reporting an anonymous vote-split is effectively contained in an individually-attributed voting record. The information about monetary policymakers' preferences and policy positions that would

have existed if a vote were anonymous cannot be derived from the attributed votes because the votes themselves are cast differently depending on the reporting procedure.

Building on this foundation, I next turn to theorizing how moving from anonymous to attributed votes may affect the formation of market expectations. Central banks and their monetary policy committees increasingly recognize the importance of their communication with private actors in the form of published minutes from their rate-setting meetings. It is in these documents that voting records, whether they are individually attributed or not, are typically reported. As private actors are forming expectations over future monetary policy decisions and inflation outcomes, these minutes provide valuable information both with respect to the final vote as well as glimpse into the decision-making process. Where dissenting votes prove to be predictive of future policy changes, I expect individuals' expectations over such policy changes to be highly responsive to the release of the previous meeting's votes. For example, after observing dissents in favor of a tighter or looser monetary policy, one may update their expectations about future rate decisions in that direction. However, if the shift to publishing an attributed voting record suppresses dissents and thereby reduces the informational content of the voting record, individual expectations about future rate decisions ought to be less responsive to the publication of voting outcomes.

**Hypothesis 2:** *When central banks publish individually attributed voting records from rate-setting meetings, market expectations about future policy will be less responsive to the publication of vote outcomes than they were when only anonymous vote-splits were published.*

Given the importance central bankers assign to managing and influencing market expectations and increased efforts to eliminate policy “surprises” as described by King (2000), this is a more substantively important consequence of publishing attributed votes than it perhaps seems. Minutes from monetary policy meetings currently act as a focal point for communication between central banks and markets. To the extent shifting from anonymous vote-splits to individually attributed voting records marginalizes the perceived informational content of

these minutes among private forecasters, this poses a new challenge for central banks as the coordinate on communication strategies with markets.

Finally, I turn to how the publication of individually-attributed votes ought to affect the accuracy of inflation expectations. On their own, the effects of publishing attributed votes developed in the first two hypotheses may be of little consequence. That is, the presence of fewer dissents does not inherently pose a challenge to the conduct of monetary policy and one may argue there are other avenues such as speeches and technical reports through which the central bank can effectively communicate with markets and shape expectations. However, to the extent these effects reflect broader consequences for the formation and content of inflation expectations, they become critical considerations in the debate over the publication of voting records. Building on the mechanisms outlined in the first two hypotheses, the third and final hypothesis speaks to the consequences of vote transparency for the accuracy of inflation expectations.

**Hypothesis 3:** *When central banks publish individually attributed voting records from rate-setting meetings, market inflation expectations ought to be less accurate than they were when only anonymous vote-splits were published.*

## Empirical Analysis

As the preceding review of the extant literature suggests, developing an empirical test of competing predictions with respect to publishing voting records is not as straightforward as it perhaps seems. The institutional shift from near complete secrecy to increasingly transparent banks occurred over a fairly short, 15-20 year period. For the vast majority of central banks, this meant most changes came about in the context of sweeping institutional reforms affecting multiple dimensions of monetary policymaking practices and transparency simultaneously. This pattern of reform implies two things about the nature of the early empirical work that followed. First, isolating the empirical effect of any one aspect became difficult or impossible in most contexts. As scholars we could only examine the effect of transparency writ large —

how higher scores on composite indices of transparency correlate with the accuracy of inflation expectations. Second, though relatedly, the rapid institutional convergence across multiple dimensions of transparency made understanding the singular implications of any given policy reform meaningless. That is, if two transparency measures have always occurred together, they may only affect market performance in tandem, and examining their effects separately is misleading. This seems theoretically possible, for example, in the case of publishing inflation forecasts and publishing minutes from rate-setting meetings. In such a context examining the effects of transparency from large indices provides better guidance to policymakers than examining the effect of one measure in isolation.

However, transparency has continued to evolve. While central banks worldwide have nearly converged on many dimensions of transparency, the publication of voting records remains highly variable across states and levels of development. Some states embrace the policy, including major global economies (i.e., the United States, England, and Japan), and others hesitate to remove what may be the final barrier of institutionalized decision-making secrecy.<sup>3</sup> Thus, the analysis of a cross-national, aggregated index of transparency cannot disentangle the effects of publishing voting records. While the approach this paper employs sacrifices considerable breadth and inevitably introduces concerns of external validity, it does so in order to answer a specific and policy-relevant question facing contemporary governments and central banks today.

## Data

To evaluate the effect of publishing MPC voting records on the performance of market expectations, I focus on the case of the central bank of Brazil, the Banco Central do Brasil (BCB). While relatively few scholars have included the BCB in their cross-national samples of central banks, it is structurally quite similar to those like the Federal Reserve, Bank of England, and Bank of Japan that have historically received more attention. The BCB has

---

<sup>3</sup> Decade long discussions at some central banks over the decision to publish voting records suggests this marked hesitation.

a nine-person monetary policy committee, Comitê de Política Monetária (Copom), that conducts all monetary policy under the leadership of a single governor. The Copom meets eight times a year to discuss and set the Selic interest rate.<sup>4</sup> Since announcing an explicit inflation target in 1999, the Copom has used the Selic interest rate as its primary monetary policy instrument.

While these are attractive features that arguably make this case more generalizable, this paper focuses on Brazil because of a quasi-natural experiment in which Copom did not release its voting records until the national legislature implemented a national Freedom of Information (FOI) law in 2012. Critically, the Access to Information Act (AIA) was broad in its objectives and application. It put in place the foundation for information request procedures applicable across government institutions and included obligations for disclosure and the provision of data in a non-proprietary format. The designers of AIA primarily intended to pressure the national legislature to make all votes open, and the law has no language to suggest it would apply to Copom deliberations. However, it was deemed upon implementation to apply to Copom proceedings. Thus, the minutes of rate-setting meetings began including voting records in April 2012.

The BCB and reporting practices of the Copom were reasonably transparent prior to the implementation of the AIA. In fact, in 2010 aggregate transparency indices put the BCB in the 85th percentile of central banks globally (Dincer and Eichengreen 2014). Copom published detailed minutes that included the number of dissents (without identifying names) seven days following the meeting. It also released immediate statements on policy decisions on the day of meetings as well as quarterly inflation reports. Critical to the research design, none of these institutional practices changed with the implementation of the AIA. While transparency reforms that arise internally at central banks have traditionally come in packages affecting multiple aspects of transparency at once, the AIA only affected the BCB incidentally and therefore did not change any other one aspect of the reporting practices at the BCB. Compliance involved a single change in the final line of Copom's minutes in which the

---

<sup>4</sup>Prior to 2006, the Copom held monthly rate-setting meetings.

announcement of the policy outcome, now listed the names of members next to the policy alternative for which they voted.

While the existence of the AIA intervention motivates my selection of Brazil, focusing on this case also has an advantage in that the BCB resembles central banks in other emerging markets considering the publication of voting records. History shows that in practice, transparency reforms are fairly irreversible. That is, once a central bank has opened up some process of their policymaking process to public scrutiny, it rarely reverts to secrecy. As Issing (2014) noted with respect to debates over transparency at the ECB, “It will be hard, if not impossible, to go back and reduce communication from what has been practiced so far.” Thus, given the main practical purpose of examining the impact of transparency reforms is to inform policy in the states still considering them, the experience in Brazil is a particularly informative case to examine. The BCB is more similar to the central banks in developing countries than are the Federal Reserve, Bank of England, and the central banks of other major economic powers. Brazil is a first mover among a large collection of emerging markets considering such policy adoptions, and thus provides a relevant point of comparison.

## **Vote Transparency and the Casting of Dissenting Votes**

With respect to the first hypothesis, I begin by examining how the publication of attributed voting records affects the voting behavior on the monetary policy committee. Specifically, I am interested in whether dissents are less likely to occur when publication practices shift from anonymous vote-splits to individually attributed voting records. Ideally, we could replicate the design from Meade and Stasavage (2008), comparing individually attributed votes from Copom meetings after the AIA intervention to votes cast by those same individuals members prior to the AIA when committee members did not know their votes would be publicly released. Unfortunately from a research design perspective, the BCB has not followed the approach taken by at Fed to retroactively release voting records from Copom meetings prior to implementation of the AIA. We can, however, gain some intuition over this question by

comparing the occurrence of dissenting votes before and after the AIA-mandated publication of attributed voting records was implemented.

From a purely descriptive standpoint, dissents appear to occur less frequently when voting records are published. While dissenting votes are fairly infrequent both in Brazil and on MPCs more generally, the average number of dissents in the sample of 120 Copom meetings went from 0.58 when votes were anonymous to 0.24 once they began publishing attributed votes ( $p = .049$ ). Similarly, the percentage of these meetings with 1 or more dissenting votes declined from 22.5% to 10.3% ( $p = .071$ ). Of course, the inferences we can draw from such patterns are limited by our ability to establish whether the conditions in which votes are cast at a given meeting prior to the AIA when voting records were anonymous are “all else equal” to those cast at a given meeting in the post-intervention period when voting records are published. That is, an alternative explanation for the relative infrequency of dissents post-AIA that has nothing to do with the change in transparency measures could be that the economic context in which monetary policy was being conducted affected voting behavior. To account for this as much as possible with the observational data at hand, I estimate a series of Bayesian probit regression models that incorporate known predictors of dissenting votes from the extant literature on voting behavior on monetary policy committees. In Table 2 I present the results from a series of models in which the probability of there being one or more dissenting votes at meeting  $t$  is modeled as a function of whether attributed votes are published ( $= 1$ ) as well as macroeconomic indicators for the inflation gap, GDP growth rate, and majority proposed rate change. In addition to including these as controls, I present alongside models estimated from the full sample the results on models estimated on a matched sample constructed by matching on the macroeconomic indicators for the inflation gap and GDP growth.<sup>5</sup>

---

<sup>5</sup>Following the best practices outlined in King and Nielsen (2019), I construct the matched samples using coarsened exact matching implemented in the R package `MatchIt`. See online appendix for descriptive statistics from unmatched and matched samples.

	(1)		(2)		(3)	
Published Votes	-0.63 (.007)	-0.81 (.008)	-0.62 (.007)	-1.29 (.018)	-1.03 (.004)	-1.21 (.017)
Inflation Gap	0.09 (.002)	-0.08 (.002)	0.09 (.002)	-0.07 (.003)	0.15 (.002)	-0.01 (.003)
GDP Growth			0.01 (.001)	0.01 (.001)	0.01 (.001)	0.01 (.001)
Rate Change					-0.66 (.007)	-0.51 (.008)
Intercept	-0.83 (.003)	-0.69 (.003)	-0.85 (.004)	-0.77 (.004)	-1.03 (.004)	-0.93 (.005)
Full vs. CEM	Full	CEM	Full	CEM	Full	CEM
Number of Obs.	120	109	120	96	120	96

Table 2: Probability of Dissenting Vote

Results from Bayesian probit regression models for a dependent variable corresponding to whether a dissenting vote is cast in each meeting of Copom. Standard errors provided in parenthesis. For each model specification, results are presented for models estimated on the full sample of votes as well as on a matched sample, in which coarsened exact matching was conducted with covariates for macroeconomic indicators of the inflation gap and GDP growth.

Although we cannot draw substantive inferences from the magnitude of these non-transformed coefficients on their own, the sign and significance of each across specifications comports with expectations and is consistent with the first hypothesis. Focusing on the results of the third model estimated on the matched sample, the predicted probability of a dissenting vote being cast decreases by 16.7% — with a 95% credible interval (-28.5%, -5.5%) — when attributed votes are published. While it is difficult to draw causal inferences given the limitations of the observational data at hand, the results appear to reflect a similar pattern of behavior as recognized in other contexts in which the publication of attributed voting records reduces the likelihood of a policymaker casting a dissenting vote (i.e., Meade and Stasavage (2008)). This finding is particularly relevant for the analysis that follows insofar as it implies that to the extent dissenting votes are informative, the informational content of Copom’s minutes may have declined following the implementation of the AIA and publication of voting records.

## Vote Transparency and the Updating of Policy Expectations

Given that the publication of voting records appears to reduce the probability with which individual central bankers cast dissenting votes, I next turn to the implications of this transparency measure for the formation of market expectations as described in the second and third hypotheses. To examine potential effects of the AIA-mandated publication of attributed votes on private forecasts and expectations, I draw on data from an ongoing survey by the BCB's Investor Relations and Special Studies Department (Gerin). Beginning with BCB's adoption of the inflation targeting regime in 1999, Gerin started fielding the Market Expectations Survey (MES) in an effort to more carefully and systematically monitor market expectations over primary macroeconomic indicators. What began as a relatively small survey of 50 financial institutions and consultancies, the MES now draws regular forecasts from over 100 Brazilian economists, research firms, and the like.<sup>6</sup> Approved forecasters log in to the access-controlled online system at any time and provide updates of their short-, medium-, and long-term forecasts over a variety of price indices, exchange rates, the Selic target rate, and balance of payments variables. While individual forecasts are not identifiable, the MES provides daily updates of the sample statistics for the aggregated forecasts.

Recall the intuition of the second hypothesis: if dissents are informative of future policy changes and the publication of voting records suppresses dissenting votes, then I expect market forecasts for these policy rates to become less responsive to the release of minutes once they began to incorporate individually-attributed votes. Examining this relationship empirically calls for a measure of how much forecasters are updating their expectations about the Selic rate during the period between monetary policy meetings. To construct this indicator, I draw on the reported average daily forecast for the Selic rate in the current year. With this daily time series, I take for each day ( $t$ ) the absolute difference between the reported forecast average the day before ( $t - 1$ ) and the reported forecast average the day after ( $t + 1$ ). This measure of expectations "updating" serves as the dependent variable for

---

<sup>6</sup>Marques (2012) provides a thorough discussion of the origin and evolution of the MES.

the proceeding analyses evaluating the second hypothesis. Then, I code for this daily time series both how many days since the previous rate-setting meeting each occurred and how many (complete) weeks since the previous rate-setting meeting each occurred.

The second hypothesis proposes an interaction effect between the time lapsed since the release of voting records and vote transparency. While shifting to publish attributed votes means monetary policy expectations should update less during the period immediately after their release, there should be no discernible effect of vote transparency on the amount of updating for the remainder of the period. Thus, I first estimate a straightforward linear interaction model where the amount of updating is modeled as a function of the weeks (or days) since the previous meeting interacted with a binary treatment variable indicating whether the observation occurs before or after the BCB began publishing attributed voting records. However, following the best practices outlined by Hainmueller, Mummolo and Xu (2019), I reject this linear specification in favor of presenting marginal effect estimates from a kernel estimator as there is evidence to suggest the marginal effect of vote transparency is non-linear.<sup>7</sup> The marginal effect of transparency on the updating of monetary policy expectations is depicted graphically in Figure 1.

The results are largely consistent with expectations given by the second hypothesis. Regardless of whether one focuses on the first week or first several days following rate-setting meetings and the release of voting records, the marginal effect of transparency is negative. Thus, publishing attributed voting records appears to lessen the degree to which markets respond to the announcement of rate decisions and committee votes. Further, it is important to highlight that this does not imply forecasters are updating less for the duration of the inter-meeting period. As expected, the negative marginal effect observed in the period immediately following the publication of minutes dissipates, with no estimated difference in updating due to the shift in transparency for much of the period after the first week. Somewhat surprisingly, however, the marginal effect of transparency becomes negative again in the period just prior

---

<sup>7</sup>All models are estimated with the R package `Interflex` accompanying Hainmueller, Mummolo and Xu (2019).

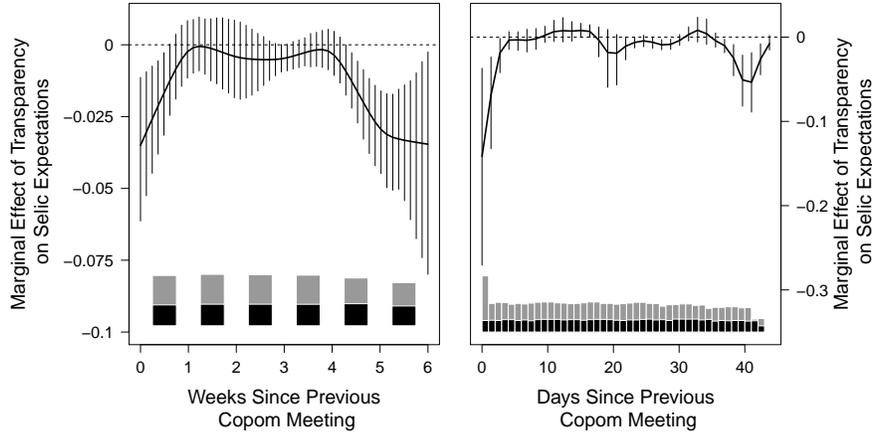


Figure 1: Updating Policy Expectations

Estimates from a kernel estimator of the marginal effect of publishing attributed voting records on the amount forecasts of the Selic rate are updated (y-axis) throughout the inter-meeting period (x-axis). Updating is measured as the absolute change in the average Selic rate forecast. Negative (positive) marginal effects of transparency indicate the forecast average changes less (more) in the post-intervention period when individually attributed votes are published. Vertical bars correspond to 95% confidence intervals estimated with bootstrapped standard errors.

to the next meeting. While there are fewer observations on which to compare for observations over a month from the previous meeting, this negative marginal effect during the pre-meeting period appears both substantively and statistically significant.

## Vote Transparency and the Accuracy of Inflation Expectations

Turning to the third and final hypothesis, I examine the effect publishing attributed voting records has had on the observed accuracy of inflation forecasts in Brazil. To measure the accuracy of inflation expectations, I draw on a second survey item from the Gerin survey of market expectations which reflects the mean forecast of inflation.<sup>8</sup> While Brazil’s MES platform allows forecasters to update their forecasts daily as well as submit inflation forecasts for a variety of time horizons ranging from the current year to over two years in the future, I choose to focus on monthly observations of forecasts of IPCA in one year to facilitate comparisons

---

<sup>8</sup>Though the MES allows survey participants to submit forecasts for a variety of price and inflation indices, I focus here only on results for the Extended National Consumer Price Index (IPCA) forecasts as this is Brazil’s official inflation measure.

with market expectations survey data in other countries. To construct a dependent variable which captures the accuracy of these inflation expectations, I take the absolute difference between the average inflation forecast and the realized inflation rate one year later. From 2005 to 2017, the average inflation forecast error for the sample of 150 monthly observations was 1.965 with a standard deviation of 1.048.

**Estimation of Synthetic Controls.** To examine the effect of publishing attributed voting records on the observed accuracy of inflation forecasts in Brazil, I estimate a pair of synthetic controls which serve as approximations of the unobserved counterfactual of inflation forecast accuracy during the post-intervention period had the BCB continued to publish only anonymous votes from Copom meetings. This estimation strategy has three steps.<sup>9</sup> First, I fit a BSTS model to the “pre-intervention” series of forecast accuracy prior to the shift in reporting practices. As I will discuss further, I fit two separate models that incorporate different theoretically motivated sets of control series. For each model, I then use these weighted control series to estimate a synthetic control to approximate the post-intervention counterfactual. Finally, I compare the observed post-intervention series and estimated counterfactual series, attributing differences in these two series to the shift in the transparency of the voting record from Copom meetings mandated by the AIA.

With the estimation of the first synthetic control, I focus on identifying a comparable indicator of market expectations within Brazil that will incorporate in estimates of the synthetic control domestic variability in the post-intervention period. To accomplish this, I construct from forecasts of industrial production contained in the same survey of Brazilian forecasters a similar indicator of forecast accuracy as used for inflation. While the resulting industrial production control series should similarly reflect economic and political variability in Brazil that is likely to affect the accuracy of market expectations, we should not necessarily expect changes in central bank transparency to affect expectations over industrial production.

---

<sup>9</sup>For a more thorough discussion, see Scott and Varian (2014) Abadie, Diamond and Hainmueller (2015), and Brodersen et al. (2015).

This assumption is critical to the viability of the design, as this approach aims to isolate the effect of the transparency policy shift from other idiosyncratic changes in the Brazilian economy affecting survey expectations more broadly. However, this choice of control series begs the question how likely errors in industrial production forecasts are subject to the same idiosyncratic shocks unrelated to transparency that specifically affect the formation of inflation expectations.

To address the possibility that inflation forecast errors are subject to a set of factors unique to inflation, I construct a second synthetic control for the post-intervention period that draws on six additional series of private inflation expectations from countries unaffected by Brazil’s shift in reporting practices. While this set of controls will not incorporate in the synthetic control idiosyncratic changes coming from within the Brazilian economy that affect inflation expectations, it does provide potential comparison cases of inflation forecasts specifically. Taken together, the estimation of these two synthetic controls seeks to balance an inherent trade-off in identifying appropriate cases to approximate the counterfactual. I draw on a set of six control countries, including three from Latin America (Chile, Mexico, and Peru) and three from outside the region (Israel, Poland, and Thailand). In each country, central banks announce either unattributed vote splits (as the BCB did prior to the implementation of the AIA) or only the monetary policy decision.<sup>10</sup> While these series may appear to be more or less reasonable candidates for approximating the counterfactual case in Brazil, a critical advantage of this approach is that each acts as only a *potential* control group for the estimation of the treatment effect in Brazil. As I will discuss when I turn to estimation, the synthetic control is estimated as a weighted average forecast of the six potential control series. Critically, the model calibrates on the pre-intervention series which of the potential control cases are most informative and predictive of the series in Brazil and then weights those cases accordingly in constructing the synthetic control.

---

<sup>10</sup>Central banks in Chile and Poland made voting records available with multi-year lags for part of the period, but the lag ensures this information was only available *after* markets made their inflation forecasts.

Before turning to estimation and the results, there are several design and methodological assumptions underlying this approach worth noting. First, a key assumption is that market expectations in the control series are independent of the intervention being examined in Brazil. For example, we want to assume the accuracy of inflation expectations in Poland were not affected by the change to publishing voting records in Brazil. For Israel, Poland, and Thailand, this seems to be a fairly reasonable assumption: the economies are only weakly integrated with Brazil's in terms of trade and their exchange rates are not highly correlated. This assumption is perhaps less compelling for the set of Latin American countries. That is, it seems plausible heightened inflation uncertainty in Brazil could spill over into more inflation uncertainty in Mexico, Chile, and/or Peru. However, while this undoubtedly violates the assumption that the control series are independent of the intervention, it does so in a way that further biases against finding evidence of the predicted causal effect in Brazil.<sup>11</sup> Eliminating these series because of a potential downward bias on the results due to spillovers would come at the cost of losing untreated control markets, which provide considerable information in the construction of the synthetic control.

Second, it is worth recognizing how this approach relates to a perhaps more traditional difference-in-differences design and the appropriateness of this choice for this application. In practice, the intuition is quite similar. However, while difference-in-differences designs typically rely on a static regression model, the BSTS model employed here allows for the incorporation of the time series dynamics in this data. Relatedly, a static regression requires we assume observations are iid and the differences are constant over time, both of which are often false. Violating this assumption is of particular consequence in this application because fitting a static model to serially correlated data would provide overoptimistic point estimates and underestimate uncertainty.

---

<sup>11</sup>To the extent expectations over industrial production are similarly affected by the change in transparency policy, this same bias against finding results ought to exist for the first synthetic control constructed from the accuracy of expectations over industrial production in Brazil.

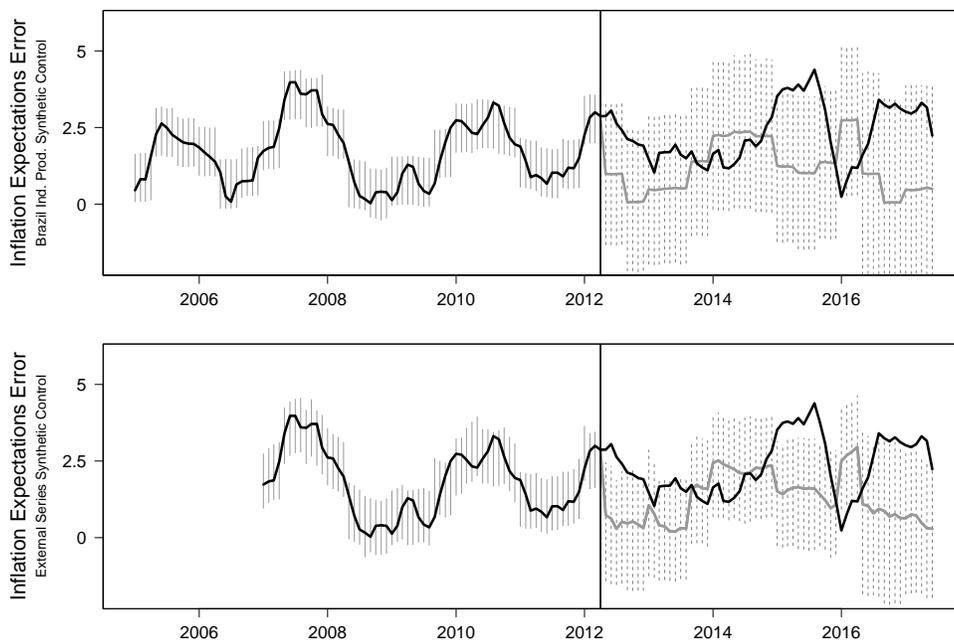


Figure 2: Accuracy of Inflation Expectations

Each panel plots the observed accuracy of inflation forecasts (solid black line) before and after national legislation mandated the monetary policy committee publish attributed voting records from rate-setting meetings. In the top (bottom) panel, the grey line in the post-intervention period corresponds to the posterior density of a synthetic control estimated with a control series of industrial production forecast accuracy in Brazil (inflation forecast accuracy in a sample of six countries).

**Results.** Reserving a more complete specification of the model for the supplementary appendix, Figure 2 shows the estimated synthetic control series plotted alongside the observed accuracy of inflation expectations in Brazil.<sup>12</sup> In each panel, the vertical reference line corresponds to Copom’s shift in April 2012 from publishing anonymous to attributed votes

<sup>12</sup> All posterior distributions are estimated for 10,000 simulations with the R packages **CausalImpact** and **bsts** presented in Brodersen et al. (2015) and Scott and Varian (2014), respectively. Because market expectations in Brazil during the post-intervention period are likely to be highly volatile due to both the nature of the data and the theoretical predictions at hand, all posterior distributions are estimated with a standard error on the state variable of 0.1. This significantly expands the posterior interval and biases the results against finding conclusive evidence.

from their rate-setting meetings. To the left of this vertical line, the black line traces the time series of the dependent variable — the accuracy of one-year inflation forecasts as measured by the absolute distance — when voting records published by the BCB remained anonymous. The vertical bars surrounding this line are 95% credible intervals corresponding to the first stage model fitting of the various control series to the pre-intervention dependent variable. That they adhere closely to the observed data indicates the BSTS model achieves a reasonably good fit to the pre-intervention data. The continuation of the black line into the post-intervention period corresponds to the observed inflation expectations when the transparency policy shifted to include attributed votes. Finally, in each panel of Figure 2, the synthetic control is depicted by the grey line and shaded region, which capture the posterior mean and associated distribution. While there are some minor differences in the two synthetic control series, they are largely comparable to one another and the inferences to be drawn in each case are substantively identical.

Though the credible intervals presented for the posterior density of the synthetic control are rather wide given the choice of a diffuse prior (as explained in the previous section), it is instructive to compare the average observed forecast error for the post-intervention period to the average predicted forecast errors in the post-intervention period had votes continued to be published anonymously. In the post-intervention period, the average inflation expectation in Brazil for a one-year forecast horizon errs by about 2.3 points. In the synthetic control series estimated with the forecast errors on industrial production, the estimated average error is 1.2, with a standard error of 0.41. Thus, while in Figure 2 the observed forecast errors are contained in the credible interval of the synthetic control series at almost any given point, the average difference is statistically significant at all conventional levels of confidence. The same holds for a comparison with the synthetic control constructed from external series of inflation forecast errors. While the average predicted forecast error of 1 is slightly higher for this synthetic control series, the difference between it and the observed error of 2.3 is substantial.

While in absolute terms one point may seem trivial, it represents a meaningful worsening of market expectations in this context. One approach to evaluating the size of this effect in this context is to consider the historic performance of market expectations in Brazil. The estimated effect of increasing the inaccuracy of inflation expectations about one point amounts to a 1.25 standard deviation increase from the pre-intervention series. Further, the estimated effect of publishing voting records significantly reduces the accuracy of expectations in Brazil relative to other countries. In the pre-intervention period, Brazil was second only to Mexico in terms of the average accuracy of forecasters' inflation expectations. However, with the 1.25 point increase, inflation expectations in Brazil outperform only one other country in the sample (Chile), and do so by a fairly narrow margin of one-tenth of a point.

## Discussion

Should central banks publish the voting records from the rate-setting meetings of their MPCs? This is the question at the forefront of the current debate over the optimal design of monetary institutions. The results presented here suggest, if the objective of publishing votes is to better inform markets and improve private inflation forecasts, they should not. Leveraging a natural experiment that occurred with the implementation of the Freedom of Information law in Brazil, I demonstrate that publishing the attributed voting record from MPC meetings corresponds to less accurate and more disperse inflation expectations. Given the primacy of central banks' need to manage and coordinate market expectations, particularly in emerging markets, governments and central banks ought to resist potentially misguided demands for publishing voting records and pursue alternative mechanisms to address concerns of individual accountability at central banks in the post crisis era.

Given the global proliferation of FOI laws like that which mandated the publication of voting records in Brazil, central banks may face little option whether to publish voting records from their MPC rate-setting meetings. While legislatures themselves should exercise caution in the broad application of these laws across institutions and policy domains, it is

also incumbent on central banks to structure their communication strategies such that these voting records do not interfere with market expectations as seen in Brazil. One strategy that would seem to address some concerns about the effect of such transparency measures would be to lengthen the delay in publication of individual voting records. For example, the votes by the MPC at the Central Bank of Iceland are only published in the Annual Report, released in March of the following year. Statements on interest rate decisions as well as the detailed minutes released more immediately following each meeting (and critically before the next meeting) remain anonymous. Adopting a publication schedule such as this seems like an appropriate balance between complying with FOI law that promotes individual accountability without introducing costly noise into markets' formation of inflation expectations.

Beyond the policy implications, the results shed light on several results in the extant empirical literature on central bank transparency. First, the findings here suggest existing estimates of the positive effects of transparency on the performance of market expectations that draw on additive, composite indices provide lower bounds. While this is true for the work that aggregate overall dimensions of transparency, it is especially true for research that focuses specifically on procedural transparency. In practice, this suggests going from no voting record to an unattributed vote-split has an even greater positive effect on accuracy of and the degree of consensus on inflation forecasts than previously expected. This is highly consistent with the theoretical literature on the informational content of central bank voting records (i.e., ?). Given the ease of implementation in the context of already published minutes, central banks that publish no voting information should strongly consider adopting this policy reform.

Moving forward, research on central bank transparency would be well-served to focus more attention on the evaluation of actual policies implemented to achieve the objectives of transparency. While quality research has been conducted on the publication of minutes, inflation targets, and central bank forecasts, policies such as the optimal timing of publications and provisions for accountability should be evaluated. Transparency is a means, not an end; our progress towards an understanding of the "optimal" level of transparency hinges on a

clear definition of the appropriate objectives and empirical observation of the trade-offs these policies involve.

## References

- Abadie, Alberto, Alexis Diamond and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science* 59:495–510.
- Adolph, Christopher. 2013. *Bankers, Bureaucrats, and Central Bank Politics: The Myth of Neutrality*. Cambridge: Cambridge University Press.
- Blinder, Alan S. 2004. *The Quiet Revolution*. Yale University Press.
- Brodersen, Kay H., Fabian Gallusser, Jim Koehler, Nicolas Remy and Steven L. Scott. 2015. "Inferring causal impact using bayesian structural time-series models." *Annals of Applied Statistics* 9(1):247–274.
- Buiter, Willem H. 1999. "Alice in Euroland." *JCMS: Journal of Common Market Studies* 37(2):181–209.
- Curtin, Deirdre. 2017. "'Accountable independence' of the European central bank: Seeing the logics of transparency." *European Law Journal* 23(1-2):28–44.
- Dincer, N. Nergiz and Barry Eichengreen. 2014. "Central Bank Transparency and Independence: Updates and New Measures." *International Journal of Central Banking* 10(1):189–253.
- Ehrmann, Michael, Sylvester Eijffinger and Marcel Fratzscher. 2012. "The Role of Central Bank Transparency for Guiding Private Sector Forecasts." *Scandinavian Journal of Economics* 114(3):1018–1052.
- Eijffinger, S. C W and Petra M. Geraats. 2006. "How transparent are central banks?" *European Journal of Political Economy* 22(1):1–21.
- Fry, Maxwell, DeAnne Julius, Sandra Roger, Lavan Mahadeva and Gabriel Sterne. 2000. Key issues in the choice of monetary policy framework. In *Monetary Frameworks in a Global Context*, ed. Lavan Mahadeva and Gabriel Sterne. Routledge.
- Geraats, Petra M. 2002. "Central Bank Transparency." *The Economic Journal* 112:532–565.
- Gerlach-Kristen, Petra. 2004. "Is the MPC's Voting Record Informative about Future UK Monetary Policy?" *Scandinavian Journal of Economics* 106(2):299–313.
- Hainmueller, Jens, Jonathan Mummolo and Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis* 27:163–192.
- Hansen, Stephen, Michael McMahon and Andrea Prat. 2018. "Transparency and Deliberation within the FOMC: a Computational Linguistics Approach." *Quarterly Journal of Economics* 133(2):801–870.

- Horvath, Roman, Katerina Smidkova and Jan Zapal. 2012. “Is the U.S. Fed Voting Record Informative about Future Monetary Policy?” *Czech Journal of Economics and Finance* 62(6):478–484.
- Horváth, Roman and Júlia Jonášová. 2015. “Central banks’ voting records, the financial crisis and future monetary policy.” *European Journal of Political Economy* 38:229–243.
- Issing, Otmar. 2014. “Communication and transparency – The example of the ECB.” *Journal of Economic Dynamics and Control* 49:70–73.
- King, Gary and Richard Nielsen. 2019. “Why Propensity Scores Should Not Be Used for Matching.” *Political Analysis* 27(4):435–454.
- King, Mervyn. 2000. Monetary Policy: Theory in Practice. In *Joint Luncheon of the American Economic Association and the American Finance Association*.
- Lustenberger, Thomas and Enzo Rossi. 2017. “Does Central Bank Transparency and Communication Affect Financial and Macroeconomic Forecasts? Legal Issues Does Central Bank Transparency and Communication Affect Financial and Macroeconomic Forecasts?”
- Marques, André Barbosa Coutinho. 2012. “Central Bank of Brazil’s market expectations system: a tool for monetary policy.” *IFC Bulletin* 36:304–324.
- Meade, Ellen E. and David Stasavage. 2008. “Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve.” *Economic Journal* 118:695–717.
- Naszodi, Anna, Csaba Csavas, Szilard Erhart and Daniel Felcser. 2016. “Which aspects of central bank transparency matter? A comprehensive analysis of the effect of transparency on survey forecasts.” *International Journal of Central Banking* 12(4):147–192.
- Neuenkirch, Matthias. 2012. “Managing financial market expectations: The role of central bank transparency and central bank communication.” *European Journal of Political Economy* 28:1–13.
- Neuenkirch, Matthias. 2013. “Central bank transparency and financial market expectations: The case of emerging markets.” *Economic Systems* 37(4):598–609.
- Scott, Steven L. and Hal R. Varian. 2014. “Predicting the present with Bayesian structural time series.” *International Journal of Mathematical Modelling and Numerical Optimisation* 5:4–23.
- Siklos, Pierre L. 2011. “Central bank transparency: another look.” *Applied Economics Letters* 18:929–933.
- Siklos, Pierre L. 2013. “Sources of disagreement in inflation forecasts: An international empirical investigation.” *Journal of International Economics* 90:218–231.
- Van Der Cruysen, Carin A B, Sylvester C W Eijffinger and Lex H Hoogduin. 2010. “Optimal central bank transparency.” *Journal of International Money and Finance* 29:1482–1507.

Van Der Cruysen, Carin and Maria Demertzis. 2007. "The impact of central bank transparency on inflation expectations." *European Journal of Political Economy* 23:51–66.

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