

Identification via Heteroskedasticity when Attention is Endogenous

Alistair Macaulay*

University of Surrey

January 29, 2026

Abstract

This paper demonstrates that rational inattention biases structural vector autoregressions identified via heteroskedasticity. Inattentive agents optimally acquire more information during high-volatility regimes, so shock transmission changes with shock variance, which violates the key identifying assumption of constant structural parameters. I characterize the resulting bias and develop two correction methods: one using attention proxies, another exploiting multiple variance regimes. Applying the first of these methods to unconventional monetary policy at the zero lower bound reveals that standard estimates substantially underestimate the role of expectations in policy transmission, with signalling effects on inflation expectations particularly prominent.

JEL codes: C32, D83, E52

*School of Economics, University of Surrey. Email: a.macaulay@surrey.ac.uk.

1 Introduction

Households, firms, and investors acquire more information about economic objects when they have strong incentives to do so. This is the core premise of models of ‘rational inattention,’ in which agents choose how much information to process about the variables affecting their decision problems, subject to a cost. Following [Sims \(2003\)](#), a large literature has explored how this endogenous information acquisition affects a wide range of macroeconomic dynamics.¹ There is also a growing body of empirical evidence supporting the mechanism.²

In this paper I show that endogenous information acquisition also affects the identification of structural vector autoregressions (VARs) via heteroskedasticity. This approach, pioneered by [Rigobon \(2003\)](#) and [Rigobon and Sack \(2003, 2004\)](#), and reviewed recently in [Lewis \(2025\)](#), is a commonly-used tool for applied macroeconomics. It is often motivated as requiring fewer structural economic assumptions than other identification schemes, as the key assumptions are statistical, not economic. However, once we allow for the possibility of endogenous information acquisition, as is now common in other areas of macroeconomics, this ceases to be true. Rational inattention therefore has econometric implications, not just macroeconomic ones.

The key argument is straightforward. Identification via heteroskedasticity exploits changes over time in the variance-covariance matrix of the structural shocks hitting the economy. These changes can be used to identify the transmission of the shocks, *if that transmission is itself constant across variance regimes*. In other words, the mapping from structural innovations to observable variables must remain constant. However, a core lesson of rational inattention models is that when the variance of a shock rises, decision-makers optimally pay more attention to it ([Sims, 2003, 2010](#)). At the same time, greater attention affects how shocks transmit to decisions and aggregate outcomes. The ‘constant shock transmission’ required for identification via heteroskedasticity therefore fails, unless we impose a strong economic assumption that agent information processing does not vary, or does not affect macroeconomic dynamics.

I make three contributions. Theoretically, I show formally that the presence of rational inattention biases structural VAR estimates identified via heteroskedasticity. Methodologically, I develop two techniques to correct for this bias: one using an external proxy for attention, another using observations of three or more variance regimes. Empirically, I apply this bias-corrected identification to assess the role of expectations and attention in the transmission of unconventional monetary policy shocks.

¹See e.g. [Maćkowiak and Wiederholt \(2009, 2015\)](#); [Pasten and Schoenle \(2016\)](#); [Stevens \(2019\)](#); [Song and Stern \(2024\)](#); [Macaulay \(2021, 2025\)](#), and the review in [Maćkowiak et al. \(2023\)](#).

²e.g. [Mondria and Quintana-Domeque \(2013\)](#); [Cavallo et al. \(2017\)](#); [Coibion et al. \(2018\)](#); [Roth and Wohlfart \(2020\)](#); [Dean and Neligh \(2023\)](#); [Weber et al. \(2025\)](#), among many others.

Acknowledging the effects of rational inattention on identification via heteroskedasticity has benefits beyond bias correction. When attention is endogenous, impulse responses to structural shocks differ depending on the variance regime, as that regime determines information choices. I show that an identification that does not take this into account will recover a non-convex combination of the true regime-specific responses, which may even lie outside the span of the true responses. However, the bias correction methods I propose recover the distinct impulse responses present in each variance regime. By comparing them, we therefore reveal how expectations, and information acquisition, affect shock transmission. This insight is not available from other VAR identification schemes that yield a single impulse response per shock.³

I apply these insights to the study of unconventional monetary policy, extending the canonical implementation of identification via heteroskedasticity in [Wright \(2012\)](#) to the full zero lower bound (ZLB) period in the US (2008-2015). Wright's identifying assumptions are that (i) the variance of monetary policy shocks is greater on days with scheduled Federal Open Markets Committee (FOMC) announcements than on other days; and (ii) the transmission of a unit monetary policy shock is constant over time. Rational inattention calls the second of these into question.

To check if rational inattention is important in this context, I conduct two tests. First, I show that in both the full ZLB period and the 2008-2011 subsample analysed by [Wright \(2012\)](#), the difference between the covariance matrices of reduced-form innovations in the two regimes has an eigenvalue structure inconsistent with the assumption of constant shock transmission, but which *is* predicted by the model with rational inattention. Second, I use daily Google Trends searches for "Federal Reserve" as a proxy for attention to monetary policy. Search volume increases substantially on FOMC announcement days, regardless of the realized monetary policy shocks that occurred (as identified from high-frequency asset price movements by [Swanson, 2021](#)). This suggests that attention to monetary policy is systematically higher in the high-variance regime, precisely as the theory predicts.

Given these findings, I then use the Google Trends series to implement the first of the correction methods I propose in the earlier part of the paper. This involves using the attention proxy to infer the ratio of structural shock variances across regimes. When combined with an assumption that one variable in the system is unaffected by attention on impact, this pins down the rotation among observationally equivalent impact vectors. Intuitively, knowing how much more agents pay attention in the high-variance regime reveals how much of the observed covariance difference is due to changing variances versus changing transmission, allowing us to separate the two.

The second correction method does not require a proxy for attention, and so is more conservative

³This includes recursive identification (e.g. [Christiano et al., 1999](#)), long-run restrictions (e.g. [Blanchard and Quah, 1989](#)), as well as approaches using external instruments (e.g. [Romer and Romer, 2004](#); [Gürkaynak et al., 2005b](#)).

in that dimension. However, it does require three or more observable variance regimes. The insight is that with a single shock of interest, the change in impulse responses between regimes 1 and 2 is collinear with the change between regimes 2 and 3, because both shifts occur due to a one-dimensional change in agent attention to that structural shock. This collinearity provides an extra restriction that can be used to disentangle the true structural parameters of the model. Since the monetary policy application has only two clearly-defined variance regimes, I use the attention-proxy method throughout.

The results imply that the updating of beliefs about the stance of monetary policy, through attentive agents, is of central importance to the transmission of unconventional monetary policy shocks. Uncorrected impulse responses from conventional identification via heteroskedasticity are strongly distorted by endogenous information acquisition, often lying outside the span of the corrected regime-specific responses. These uncorrected estimates imply greater persistence in the effects of unconventional monetary policy on Treasury and corporate bond yields than is observed once the bias correction is complete, and they miss a large short-term swing in breakeven inflation rates.

Specifically, I find that the immediate impact of policy on Treasury yields is similar across regimes, consistent with mechanical portfolio rebalancing that operates independently of beliefs. However, the *persistence* of these effects differs markedly by regime: after 50 business days, the effect on the 10-year yield is approximately half as large in the low-attention regime than in the high-attention regime. Without the updates to expectations facilitated by high attention, the effects of unconventional policy on the yield curve are largely transitory. High-grade corporate bond yields follows similar patterns, indicating that they are indeed close substitutes for long-term Treasuries, as argued by [Krishnamurthy and Vissing-Jorgensen \(2011\)](#) and others.

For riskier investment-grade corporate bonds, however, attention acts as a powerful amplifier. BAA yields respond 25% more strongly on impact in the high-attention regime, consistent with a credit channel where information about monetary policy causes investors to update default expectations. The impulse responses in the two regimes subsequently converge, suggesting that inattentive agents eventually learn the relevant information for this market, but with a delay that dampens the initial transmission.

Finally, the bias correction reveals pronounced effects of unconventional policy on inflation expectations. Theoretical models of the signalling channel of unconventional monetary policy such as [Eggertsson and Woodford \(2003\)](#) predict that inflation expectations should increase strongly after expansionary monetary policy shocks. However, uncorrected impulse responses, such as those in [Wright \(2012\)](#), display only muted effects on market-based expectations. I find that this is driven

by a strong ‘rotation’ in the low-attention regime, in which medium-term breakeven inflation rises but longer-term breakevens fall, plausibly reflecting mechanical or liquidity-driven distortions. In contrast, when attention is high, expectation effects offset this, delivering small positive effects overall. The effect of unconventional policy on inflation expectations is therefore substantially stronger than raw breakevens data would suggest, as the small positive effects on forward breakevens in the high attention regime include a large negative mechanical effect, which is offset by updates to expectations.

Taken together, these findings add new evidence to the view that signalling effects are a dominant driver of unconventional monetary policy transmission. Looking forward, this highlights a trade-off for policymakers engaging in quantitative tightening: effective management of inflation expectations requires high attention, but this same attention amplifies the immediate tightening of financial conditions in corporate credit markets.

Related Literature. This paper principally contributes to the literatures on rational inattention and the identification of structural VARs using heteroskedasticity. In the former, a large literature since [Sims \(2003\)](#) has shown that endogenous information acquisition has important effects across macroeconomics and finance, including in price setting ([Maćkowiak and Wiederholt, 2009](#); [Pasten and Schoenle, 2016](#); [Matějka, 2016](#); [Stevens, 2019](#)), consumption ([Tutino, 2013](#)), business cycles ([Maćkowiak and Wiederholt, 2015](#); [Flynn and Sastry, 2024](#)), monetary policy ([Paciello and Wiederholt, 2014](#)), misallocation ([Gondhi, 2023](#)), trade ([Dasgupta and Mondria, 2018](#)), financial markets ([Mondria and Quintana-Domeque, 2013](#); [Kacperczyk et al., 2016](#)), labor markets ([Acharya and Wee, 2020](#); [Ellison and Macaulay, 2021](#)), and household finance ([Lei, 2019](#); [Macaulay, 2021, 2025](#)). Empirical evidence has been found in experiments ([Dean and Neligh, 2023](#)), surveys ([Roth and Wohlfart, 2020](#); [Link et al., 2023](#)), and in VARs involving expectations ([Geiger and Scharler, 2021](#)). I extend this literature by showing that endogenous attention also affects macroeconometric analysis, specifically structural VARs identified via heteroskedasticity.⁴

This method of identifying structural VARs was popularized by [Rigobon \(2003\)](#) and [Rigobon and Sack \(2003, 2004\)](#). While these consider discrete observable variance regimes, the method has since been extended to other forms of heteroskedasticity, including Markov-switching models ([Lanne and Lütkepohl, 2010](#)), smooth volatility transitions ([Lütkepohl and Netšunajev, 2017](#)), and other parametric volatility processes (e.g. [Normandin and Phaneuf, 2004](#)).⁵ Applications

⁴Perhaps closest in spirit to this paper is [Maćkowiak and Wiederholt \(2024\)](#), who study how rational inattention affects the parameters identified in information-treatment randomized controlled trials.

⁵See [Lütkepohl and Netšunajev \(2015\)](#) for an early survey of possible models, and [Lewis \(2021, 2025\)](#) for more recent developments. [Magnusson and Mavroeidis \(2014\)](#) discuss identification from more general sources of reduced-form instability.

include estimating the spillover effects of sovereign bonds (Rigobon, 2003), cross-border financial market linkages (Ehrmann et al., 2011), monetary policy (Rigobon and Sack, 2003; Wright, 2012; Nakamura and Steinsson, 2018; Schlaak et al., 2023), fiscal multipliers (Lewis, 2021), macro-financial feedback effects (Brunnermeier et al., 2021), and pandemic shocks (Miescu and Rossi, 2021). I show that rational inattention, now a relatively common assumption in macroeconomic theory, can bias structural VARs identified in this way if my correction methods are not applied. Although for clarity I focus on discrete variance regimes and a single shock of interest, the core mechanism extends to more general cases: attention responds to volatility, and this causes the key constant-transmission identifying assumption to fail.

The rational inattention mechanism I study causes the matrix governing the impact effects of structural shocks to vary with the variance regime. Bacchicocchi and Fanelli (2015), Bacchicocchi et al. (2018), and Angelini et al. (2019) also consider variation in this impact matrix, though they allow it to vary in a general way, and thus have to provide extra restrictions to identify the model parameters. As Brenna et al. (2023) point out, these extra restrictions are often difficult to defend. My approach differs in two important ways. First, rather than a general relaxation of the constant-impact matrix assumption, rational inattention provides a concrete economic reason for time-varying transmission, enabling both testing and bias correction. Second, the correction recovers distinct impulse responses under varying levels of attention, which reveals how attention affects the transmission of macroeconomic shocks. This insight is unavailable from approaches that treat impact-matrix variation as a nuisance, or other structural VAR identification schemes.

My application follows from Wright (2012) and Nakamura and Steinsson (2018), who similarly identify the effects of monetary policy shocks by assuming higher variance on FOMC announcement days. More broadly, this exercise contributes to the literature on unconventional monetary policy, which has used a variety of theoretical and empirical approaches (e.g. Gagnon et al., 2011; Gertler and Karadi, 2011; Swanson, 2011; Baumeister and Benati, 2013; Gaballo, 2016; de Groot and Haas, 2023; Ikeda et al., 2024).⁶ Within this, many papers have argued theoretically that expectations play an important role in the transmission of unconventional monetary policy (e.g. Iovino and Sergeyev, 2023; Bhattacharai et al., 2023), and found empirical evidence consistent with this (e.g. Bauer and Rudebusch, 2014; Boneva et al., 2016). I provide direct evidence of expectational effects: the bias correction recovers impulse responses to unconventional monetary policy shocks under high and low attention, and thus under different degrees of expectation updating.

⁶See Dell’Ariccia et al. (2018) and Kuttner (2018) for reviews of this vast literature.

Outline. Section 2 lays out a baseline structural VAR model, and highlights the assumptions necessary for identification via heteroskedasticity to succeed. Section 3 adds rational inattention to this model, and shows how this biases the parameter estimates and impulse responses obtained if the econometrician proceed with identification via heteroskedasticity as usual. Section 4 develops two methods for correcting the bias, and Section 5 implements the first of them to study the transmission of unconventional monetary policy shocks. Section 6 concludes.

2 Identification via Heteroskedasticity: Baseline Framework

2.1 Structural VAR Representation

Consider a structural VAR(p) of the form

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{B}\boldsymbol{\varepsilon}_t, \quad (1)$$

where \mathbf{y}_t is an $n \times 1$ vector of observable variables, $\mathbf{A}(L) = \mathbf{I}_n - \mathbf{A}_1 L - \mathbf{A}_2 L^2 - \cdots - \mathbf{A}_p L^p$ is a matrix polynomial in the lag operator, \mathbf{B} is an $n \times n$ invertible matrix mapping structural shocks to observables, and $\boldsymbol{\varepsilon}_t$ is an $n \times 1$ vector of structural shocks. The structural shocks are such that $\mathbb{E}[\boldsymbol{\varepsilon}_t] = 0$ and $\mathbb{E}[\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'] = \boldsymbol{\Sigma}_\varepsilon$, where $\boldsymbol{\Sigma}_\varepsilon$ is diagonal, implying the structural shocks $\boldsymbol{\varepsilon}_t$ are mutually uncorrelated.

The reduced-form representation is

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{u}_t, \quad (2)$$

where $\mathbf{u}_t = \mathbf{B}\boldsymbol{\varepsilon}_t$ are reduced-form innovations with a covariance matrix given by

$$\boldsymbol{\Omega} = \mathbb{E}[\mathbf{u}_t \mathbf{u}_t'] = \mathbf{B}\boldsymbol{\Sigma}_\varepsilon\mathbf{B}'. \quad (3)$$

The fundamental identification problem is that there are infinitely many combinations of the structural matrices $\mathbf{B}, \boldsymbol{\Sigma}_\varepsilon$ that could rationalize the observed reduced-form covariance matrix $\boldsymbol{\Omega}$. Without further restrictions, the structural parameters are not separately identified.

2.2 Identification via Heteroskedasticity

The heteroskedasticity-based approach, following Rigobon (2003), exploits variation in shock variances across regimes. Suppose there exist $K \geq 2$ variance regimes indexed by $s \in \{1, \dots, K\}$.

In regime s , structural shocks have a covariance matrix given by

$$\Sigma_{\varepsilon}^{(s)} = \text{diag}(\sigma_{1,s}^2, \sigma_{2,s}^2, \dots, \sigma_{n,s}^2), \quad (4)$$

where the variances $\sigma_{j,s}^2$ may differ across regimes. Throughout, I will use a superscript (s) to denote matrices specific to regime s .

The key identifying assumption is:

Assumption 1 (Structural Invariance) *The impact matrix \mathbf{B} is constant across variance regimes:*

$$\mathbf{B}^{(s)} = \mathbf{B} \quad \forall s \in \{1, \dots, K\}. \quad (5)$$

Example: Monetary Policy on FOMC Days. A canonical application is [Wright \(2012\)](#), who identifies monetary policy shocks by comparing FOMC announcement days (high variance for monetary policy shocks) to non-announcement days (low variance). In this context, Assumption 1 implies that the responses to monetary policy shocks are the same on FOMC and non-FOMC days; only the variance of the monetary policy shock is different.⁷

Identification procedure. Under Assumption 1, the reduced-form covariance in regime s is

$$\Omega^{(s)} = \mathbf{B} \Sigma_{\varepsilon}^{(s)} \mathbf{B}'. \quad (6)$$

Identification proceeds by exploiting the differences across regimes. Taking differences between regime s and regime 1, we obtain

$$\Omega^{(s)} - \Omega^{(1)} = \mathbf{B} \left(\Sigma_{\varepsilon}^{(s)} - \Sigma_{\varepsilon}^{(1)} \right) \mathbf{B}'. \quad (7)$$

If the variance changes are concentrated in a single shock (say shock 1), then $\Sigma_{\varepsilon}^{(s)} - \Sigma_{\varepsilon}^{(1)}$ has rank one, and we can write

$$\Omega^{(s)} - \Omega^{(1)} = (\sigma_{1,s}^2 - \sigma_{1,1}^2) \mathbf{b}_1 \mathbf{b}_1', \quad (8)$$

where \mathbf{b}_1 is the first column of \mathbf{B} . The spectral decomposition of $\Omega^{(s)} - \Omega^{(1)}$ can then be used to recover \mathbf{b}_1 up to sign and scale. The sign ambiguity is typically resolved by convention (e.g., a contractionary monetary policy shock raises the policy rate). The scale is pinned down by a

⁷In Wright's notation, "the parameters $\mathbf{A}(L)$, $\boldsymbol{\mu}$ and $\{\mathbf{R}_i\}_{i=1}^p$ are all assumed to be constant," ([Wright, 2012](#), p. F449) where \mathbf{R}_i captures the impact of structural shock i on observables in the VAR.

normalization, such as fixing the variance of the shock in one regime or fixing one element of \mathbf{b}_1 . Wright (2012), for example, uses the normalization that $(\sigma_{1,s}^2 - \sigma_{1,1}^2) = 1$.

The key point of this paper is that when agents engage in optimal information acquisition, Assumption 1 will fail. If the variance of structural shocks changes from one regime to another, this will cause agents to change their information choices, which in turn will affect how they react to shocks, leading to a change in \mathbf{B} . In other words, the response to structural shocks differs systematically across these regimes precisely *because* the variance differs.

3 Rational Inattention and Regime-Dependent Transmission

This section introduces rational inattention into the SVAR framework and shows how this biases estimates identified through heteroskedasticity. For simplicity, I focus here on the case in which the structural shock of interest is i.i.d..

3.1 Environment

Denote the structural shock of interest as $\varepsilon_{1,t}$. In variance regime s , this shock is distributed according to

$$\varepsilon_{1,t} \sim \mathcal{N}(0, \sigma_s^2), \quad (9)$$

where without loss of generality I assume regime $s = K$ has the greatest variance, i.e. $\sigma_K^2 > \sigma_{s \neq K}^2$. In the monetary policy example from Wright (2012) discussed above, $\varepsilon_{1,t}$ is the monetary policy shock, the number of regimes is $K = 2$, and s is equal to 2 (high variance) on FOMC days and 1 (low variance) on non-FOMC days.

I now make two assumptions which are irrelevant in existing treatments of identification via heteroskedasticity, but which are important when introducing information frictions into the model.

Assumption 2 (Transmission Through Expectations) *The effect of shock $\varepsilon_{1,t}$ on the vector of macroeconomic outcomes \mathbf{y}_t depends on the actual shock, and the average perception of the shock across a continuum of agents:*

$$\mathbf{A}(L)\mathbf{y}_t = \boldsymbol{\alpha}_1 \varepsilon_{1,t} + \boldsymbol{\beta}_1 \bar{\mathbb{E}}_t[\varepsilon_{1,t}] + \text{other shocks.} \quad (10)$$

The vector $\boldsymbol{\alpha}_1$ captures transmission channels such as direct asset price effects or mechanical balance-sheet adjustments that operate independently of agent beliefs. In contrast, $\boldsymbol{\beta}_1$ captures transmission channels that depend on agents' perceptions of the shock. In the monetary policy

application, the expectational component may come from consumption, investment, or pricing decisions which require households and firms to form beliefs about the policy stance. I will show below that the average shock perception $\bar{\mathbb{E}}_t[\varepsilon_{1,t}]$ is a deterministic function of the true shock $\varepsilon_{1,t}$, so Assumption 2 will not imply a departure from the structural VAR form in equation (1).

Assumption 3 (Known Variance Regimes) *Agents observe the variance regime s before forming expectations and making decisions.*

Assumption 3 is natural for applications like FOMC announcements, where the schedule is publicly known. This assumption is crucial: if agents did not know the variance regime *ex ante*, attention would not condition on volatility, and my critique would not apply.

3.2 Information Structure

Agents do not observe $\varepsilon_{1,t}$ directly. Instead, I follow the rational inattention literature (reviewed in Maćkowiak et al., 2023) and assume that agents can pay to observe possibly noisy signals about $\varepsilon_{1,t}$, where more precise signals come with a greater cost. Agents observe all other structural shocks directly and without error.⁸

Specifically, each agent i chooses a signal structure, which consists of a signal space \mathcal{Z} and a conditional distribution $\pi(z|\varepsilon_{1,t})$ over signals $z \in \mathcal{Z}$ given the realization of the shock. The agent then takes an action $a_{i,t}$ after observing the signal realization. As is common in the rational inattention literature,⁹ I assume a quadratic objective function, such that agent i wishes to minimize the expected squared deviation of their action from the realized shock:

$$\min_{a_{i,t}} \mathbb{E} \left[(a_{i,t} - \varepsilon_{1,t})^2 \mid z_{i,t}, s \right]. \quad (11)$$

This expectation is formed conditional on the agent's information set, which consists of their idiosyncratic signal $z_{i,t}$ and, via Assumption 3, the regime s . Given this information, the optimal action is the conditional expectation $a_{i,t}^* = \mathbb{E}[\varepsilon_{1,t}|z_{i,t}, s]$, and the minimized loss is the posterior variance $\text{Var}(\varepsilon_{1,t}|z_{i,t}, s)$.

Having found the agent's payoff at the optimal action, conditional on a given signal, I now turn to the choice of signal structure. Following Sims (2003), I impose that the cost of a signal structure

⁸This assumption is made to clearly isolate the role of attention to the shock of interest. If the costs of information are additively separable across variables, as in e.g. Afrouzi and Yang (2021), then relaxing this assumption makes no difference to the results.

⁹See e.g. Maćkowiak and Wiederholt (2009, 2015), who show that this can be motivated as a quadratic approximation to richer objective functions.

is proportional to the mutual information between the shock and the signal, defined as

$$\mathcal{I}(\varepsilon_{1,t}; z) = H(\varepsilon_{1,t}) - \mathbb{E}_z [H(\varepsilon_{1,t}|z)], \quad (12)$$

where $H(\cdot)$ denotes Shannon entropy. Intuitively, mutual information measures the expected reduction in uncertainty about $\varepsilon_{1,t}$ from observing the signal. More precise signals that imply larger expected shrinkage from prior beliefs to posteriors are therefore more costly for the agent to process.

The agent's information-choice problem is therefore to choose a signal structure to minimize the sum of the expected loss from equation (11) and information costs:

$$\min_{\pi(z|\varepsilon_{1,t})} \mathbb{E} [\text{Var}(\varepsilon_{1,t}|z, s)] + \theta \cdot \mathcal{I}(\varepsilon_{1,t}; z), \quad (13)$$

where $\theta > 0$ is the marginal cost of information.

3.3 Optimal Attention Choice

Given a quadratic objective function and Gaussian uncertainty, it can be shown (Sims, 2003) that the optimal signal structure consists of a single noisy signal of the form

$$z_{i,t} = \varepsilon_{1,t} + \eta_{i,t}, \quad \eta_{i,t} \sim \mathcal{N}(0, 1/\kappa_{i,t}), \quad (14)$$

where $\kappa_{i,t} > 0$ is the signal precision.

With this in hand, the posterior expectation of $\varepsilon_{1,t}$ can be written as

$$\mathbb{E}(\varepsilon_{1,t}|z_{i,t}, s) = \tau_{i,t} z_{i,t}, \quad (15)$$

where

$$\tau_{i,t} = \frac{\sigma_s^2}{\sigma_s^2 + 1/\kappa_{i,t}} \quad (16)$$

is the signal-to-noise ratio. Since $\tau_{i,t}$ is monotonically increasing in the amount of information processed $\mathcal{I}(\varepsilon_{1,t}; z)$, I will refer to it as a measure of the agent's attention.

The agent problem (13) can then be reduced to the choice of precision $\kappa_{i,t}$, or equivalently the signal-to-noise ratio $\tau_{i,t}$. Evaluating the posterior variance $\text{Var}(\varepsilon_{1,t}|z_{i,t}, s)$ and the mutual information $\mathcal{I}(\varepsilon_{1,t}; z)$ with the signal structure in (14), we can write the agent's information-choice

problem as¹⁰

$$\min_{\tau_{i,t} \geq 0} \left\{ \sigma_s^2 (1 - \tau_{i,t}) + \theta \cdot \frac{1}{2} \log \left(\frac{1}{1 - \tau_{i,t}} \right) \right\}. \quad (17)$$

Lemma 1 (Optimal Attention) *The optimal signal-to-noise ratio for each agent in period t is:*

$$\tau_s^* = \begin{cases} 1 - \frac{\theta}{2\sigma_s^2} & \text{if } \sigma_s^2 > \frac{\theta}{2} \\ 0 & \text{otherwise} \end{cases}, \quad (18)$$

which is strictly increasing in σ_s^2 for $\sigma_s^2 > \theta/2$.

Proof. The first-order condition for the problem (17) is:

$$-\sigma_s^2 + \frac{\theta}{2} \cdot \frac{1}{1 - \tau_s^*} = 0. \quad (19)$$

Rearranging gives the first case in equation (18). Positivity requires $\sigma_s^2 > \theta/2$. When this condition holds (i.e. at the interior solution for $\tau_{i,t}$) the derivative of τ_s^* with respect to σ_s^2 is

$$\frac{d\tau_s^*}{d\sigma_s^2} = \frac{\theta}{2\sigma_s^4} > 0. \quad (20)$$

■

The key implication of Lemma 1 is:

Corollary 1 *Optimal attention is increasing in the variance of the underlying shock:*

$$\sigma_j^2 > \sigma_k^2 \implies \tau_j^* \geq \tau_k^* \quad (21)$$

for any pair of regimes j, k .

The inequality in signal-to-noise ratios is strict whenever $\sigma_j^2 > \frac{\theta}{2}$.

This is a very standard feature of models of rational inattention. Agents allocate more attention to monitoring monetary policy when monetary policy shocks are more volatile, because large volatility means that agents would make larger expectational errors if they did not process information. The key argument of this paper is that this standard mechanism has consequences for the identification of structural VARs via heteroskedasticity, because it links variance regimes with shock transmission.

¹⁰I specify the Shannon entropy in terms of natural logarithms here for algebraic simplicity. Some rational inattention models instead use base-2 logarithms so that the units of information are bits, but this is simply a rescaling, which makes no difference to the argument presented here.

3.4 Aggregation with Idiosyncratic Signals

With a continuum of agents $i \in [0, 1]$ receiving idiosyncratic signals, the aggregate expectation is

$$\bar{\mathbb{E}}_t[\varepsilon_{1,t}] = \int_0^1 \mathbb{E}_i[\varepsilon_{1,t}|z_{i,t}] di. \quad (22)$$

Combining equations (14) and (15) with the optimal attention from Lemma 1, each agent's conditional expectation is

$$\mathbb{E}_i[\varepsilon_{1,t}|z_{i,t}, s] = \tau_s^* z_{i,t} = \tau_s^*(\varepsilon_{1,t} + \eta_{i,t}). \quad (23)$$

By the law of large numbers, idiosyncratic noise averages out, so that

$$\bar{\mathbb{E}}_t[\varepsilon_{1,t}] = \tau_s^* \varepsilon_{1,t}. \quad (24)$$

The identification failures due to information choice derived below are therefore not caused by noise in expectations, since the noise terms average out across agents. Rather, the failure comes from the structural dependence of attention on the variance of the shock of interest.¹¹

3.5 Regime-Dependent Impact Matrix

Substituting (24) into the expectations-augmented transmission equation (10), we recover the standard structural VAR form in (1), with a regime-dependent impact matrix given by

$$\mathbf{B}^{(s)} = \left(\boldsymbol{\alpha}_1 + \tau_s^* \boldsymbol{\beta}_1, \quad \mathbf{b}_2, \quad \dots \quad \mathbf{b}_n \right). \quad (25)$$

Proposition 1 (Regime-Dependent Transmission) *With known variance regimes and endogenous attention, the structural impact matrix is regime-dependent: whenever at least one regime has $\sigma_s^2 > \theta/2$ and $\boldsymbol{\beta}_1 \neq \mathbf{0}$ we have that $\mathbf{B}^{(s)}$ varies with s .*

This violates Assumption 1 (Structural Invariance), which breaks the usual conditions for identification via heteroskedasticity.

Proof. By Corollary 1, $\tau_K^* > \tau_{K-1}^*$ whenever $\sigma_K^2 > \theta/2$. From (25), $\mathbf{b}_1^{(K)} - \mathbf{b}_1^{(K-1)} = (\tau_K^* - \tau_{K-1}^*) \boldsymbol{\beta}_1 \neq \mathbf{0}$ when $\boldsymbol{\beta}_1 \neq \mathbf{0}$. ■

¹¹Equation (24) also implies that including measures of expectations in the VAR will not address the failure of Assumption 1, because the mapping from $\varepsilon_{1,t}$ to $\bar{\mathbb{E}}_t[\varepsilon_{1,t}]$ still varies with the regime.

3.6 What Heteroskedasticity Identification Recovers

The econometrician observes reduced-form covariances $\Omega^{(s)}$ and (incorrectly) assumes a constant \mathbf{B} . They then compute the difference:

$$\Omega^{(H)} - \Omega^{(L)} = \mathbf{B}^{(H)} \Sigma_{\varepsilon}^{(H)} (\mathbf{B}^{(H)})' - \mathbf{B}^{(L)} \Sigma_{\varepsilon}^{(L)} (\mathbf{B}^{(L)})' \quad (26)$$

$$= \sigma_H^2 (\boldsymbol{\alpha}_1 + \tau_H^* \boldsymbol{\beta}_1) (\boldsymbol{\alpha}_1 + \tau_H^* \boldsymbol{\beta}_1)' - \sigma_L^2 (\boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1) (\boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1)' \quad (27)$$

where H and L denote two regimes such that $\sigma_H^2 > \sigma_L^2$, and we maintain the assumption that only the variance of the first structural shock varies with the regime, all other structural shocks have constant variance and so cancel out in (26).

This is not proportional to $\mathbf{b}_1^{(H)} (\mathbf{b}_1^{(H)})'$ or $\mathbf{b}_1^{(L)} (\mathbf{b}_1^{(L)})'$, outside of the knife-edge case in which $\boldsymbol{\alpha}_1$ and $\boldsymbol{\beta}_1$ are collinear. The identified object is a weighted combination of attention-invariant and attention-sensitive transmission channels, mixed across regimes.

Proposition 2 (Pseudo-parameter recovered when endogenous attention is ignored) *Suppose an econometrician applies the standard spectral procedure used when the impact matrix \mathbf{B} is believed to be constant across regimes: they set \mathbf{b}_1 proportional to the eigenvector associated with the largest (positive) eigenvalue of $\Delta\Omega \equiv \Omega^{(H)} - \Omega^{(L)}$.*

If $(\boldsymbol{\alpha}_1 + \tau_H^ \boldsymbol{\beta}_1)' (\boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1) \neq 0$, then the resulting estimate \hat{b}_1 (up to scale) is given by*

$$\hat{b}_1 \propto (\boldsymbol{\alpha}_1 + \tau_H^* \boldsymbol{\beta}_1) + k \cdot (\boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1), \quad (28)$$

where k is a nonlinear combination of structural parameters $\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1$, variances σ_H^2, σ_L^2 , and attention τ_H^, τ_L^* , given by the largest root of the quadratic*

$$(\sigma_H^2 D) k^2 + (\sigma_H^2 C_H + \sigma_L^2 C_L) k + (\sigma_L^2 D) = 0, \quad (29)$$

where

$$C_s \equiv (\boldsymbol{\alpha}_1 + \tau_s^* \boldsymbol{\beta}_1)' (\boldsymbol{\alpha}_1 + \tau_s^* \boldsymbol{\beta}_1); \quad D = (\boldsymbol{\alpha}_1 + \tau_H^* \boldsymbol{\beta}_1)' (\boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1) \quad (30)$$

Except in knife-edge cases, namely when the regime-specific impact vectors are collinear or orthogonal, \hat{b}_1 is generically different from both $(\boldsymbol{\alpha}_1 + \tau_H^ \boldsymbol{\beta}_1)$ and $(\boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1)$. The resulting impulse responses therefore correspond to a pseudo-parameter rather than the true impact of the shock in either regime.*

Proof. Appendix A.1. ■

That is, an econometrician who ignores endogenous information acquisition will obtain an estimate \hat{b}_1 , which may appear reasonable, but which does not correspond to the true structural impact vector in either regime. This, in turn, means that impulse responses computed using \hat{b}_1 will not reflect the true transmission of the structural shock $\varepsilon_{1,t}$ in either regime.

Corollary 2 (Impulse responses implied by identification ignoring endogenous attention)

Let $\{\Psi_h\}_{h \geq 0}$ denote the reduced-form moving-average matrices implied by the lag polynomial $A(L)$, so that under regime s the structural impulse response of y_{t+h} to a unit $\varepsilon_{1,t}$ shock is

$$IRF_s(h) \equiv \Psi_h(\alpha_1 + \tau_s^* \beta_1), \quad h = 0, 1, 2, \dots \quad (31)$$

If the econometrician identifies \hat{b}_1 as in Proposition 2, then for every horizon h the estimated impulse response is given by

$$\widehat{IRF}(h) \propto \Psi_h \hat{b}_1 = IRF_H(h) + k \cdot IRF_L(h). \quad (32)$$

recalling that k is itself a nonlinear combination of parameters, including the structural parameters determining $IRF_H(h)$ and $IRF_L(h)$.

Normalizing the estimated impulse response so the impact on variable m at horizon 0 is equal to one (i.e. setting $e'_m \widehat{IRF}(0) = 1$), the estimated IRF is

$$\widehat{IRF}(h) = \frac{IRF_H(h) + k IRF_L(h)}{e'_m (IRF_H(0) + k IRF_L(0))}. \quad (33)$$

which outside of the knife-edge cases discussed in Proposition 2 differs from both $IRF_H(h)$ and $IRF_L(h)$ for generic h .

Proof. Appendix A.2. ■

In this, it is important to recognize that the resulting normalized impulse response obtained from maintaining Assumption 1 is not necessarily a convex combination of the true impulse responses in each regime. The coefficient k need not be within $[0, 1]$, and the denominator in (33) is itself a mixture of structural impulse responses, creating further nonlinearities. This is why allowing for general time-variation in impact matrices has been found to have large effects on IRF estimations in practise (Bacchocchi and Fanelli, 2015; Angelini et al., 2019). The key extra insight I provide is that rational inattention gives a structural reason for why the impact matrix may vary, which thus opens up the possibility of correcting the bias, and gaining economic insights from the variation across regimes.

When is the bias negligible? The bias characterized in Proposition 2 and Corollary 2 arises from the interaction of two forces: agents adjust their attention across variance regimes, and this attention affects shock transmission. If either channel is absent, identification via heteroskedasticity remains valid. This, however, involves strong structural economic assumptions that are unlikely to hold in many contexts of interest to the macroeconometrician.

In general, attention may be constant across regimes for three reasons. First, if attention costs are very large relative to the payoffs of being informed, attention will be 0 in all regimes. In the model above, this occurs if $\theta > 2\sigma_s^2$ for all regimes s . Second, if attention costs are very small, then attention approaches 1 in all regimes. Formally, the difference in attention between two regimes, when both regimes lead to the interior solution for attention, is $\tau_H^* - \tau_L^* = \theta/2(\sigma_L^{-2} - \sigma_H^{-2})$, which approaches 0 as $\theta \rightarrow 0$. The findings of time-varying attention across various contexts and agents suggest that these cases are rare (e.g. [Mondria and Quintana-Domeque, 2013](#); [Song and Stern, 2024](#); [Flynn and Sastry, 2024](#); [Macaulay, 2025](#)).

Third, attention may be constant across regimes if agents do not know that the regime has changed (i.e. if Assumption 3 does not hold). Rational inattention models typically assume that agents know the distribution of shocks they are facing,¹² but if this is not the case then identification via heteroskedasticity would not be subject to the bias characterized above, as long as the econometrician is still able to identify the regimes ex-post. For the monetary policy application in Section 5 this is unlikely to hold, as the high-variance regime consists of scheduled FOMC announcement days. There is also a substantial empirical literature documenting attention responding to volatility in a range of other contexts (e.g. [Andrei and Hasler, 2015](#); [Andrei et al., 2023](#); [Mikosch et al., 2024](#); [Benchimol et al., 2025](#)). Importantly, if the researcher has access to a proxy for attention to the relevant variable(s), the variation of attention across regimes can easily be tested, as I do in Section 5.2 below.

The other situation in which this bias from information acquisition will be negligible is if expectations are themselves a negligible part of shock transmission. Formally, if $\beta_1 = 0$, then equation (28) implies $\hat{b}_1 \propto \alpha_1$, as in standard arguments for identification via heteroskedasticity. However, a vast literature, from [Lucas \(1972\)](#) (and earlier influences) onwards, has shown that expectations are important in a very wide range of macroeconomic and financial dynamics, so such instances where every variable in a VAR is independent of attention are likely to be rare.¹³

¹²See [Ellison and Macaulay \(2021\)](#) for an example where this is not the case.

¹³A further mechanical knife-edge case arises when α_1 and β_1 are collinear, i.e. when $\alpha_1 = \kappa\beta_1$ for some scalar κ . The impact vector then varies in magnitude across regimes but not in direction. The covariance difference matrix remains rank one, and standard heteroskedasticity-based identification recovers the correct direction of the impact vector. This case is unlikely to hold exactly in practice but illustrates that the bias stems from rotation of the impact vector, not merely rescaling.

3.7 Relationship with specification tests

Assumption 1, which I have shown above is affected by rational inattention, is often simply imposed in practical applications. However, in principle it is possible to test it in some contexts. I therefore discuss if and how the bias due to endogenous information acquisition outlined above could be detected in these tests.

Multiple regimes and overidentifying restrictions. When more than two variance regimes are available, Assumption 1 implies that all pairwise differences $\Omega^{(j)} - \Omega^{(k)}$ must be rank one and share a common eigenvector. These restrictions are overidentifying: they imply that the eigenvectors of $\Omega^{(j)} - \Omega^{(k)}$ coincide across all j, k .

Under endogenous attention, these restrictions are generically violated. The reduced-form covariance differences are given by

$$\Omega^{(j)} - \Omega^{(k)} = \sigma_j^2(\alpha_1 + \tau_j^* \beta_1)(\alpha_1 + \tau_j^* \beta_1)' - \sigma_k^2(\alpha_1 + \tau_k^* \beta_1)(\alpha_1 + \tau_k^* \beta_1)', \quad (34)$$

which does not share a common rank-one factorization across j, k unless the vectors $\{\alpha_1 + \tau_s^* \beta_1\}_{s=1}^S$ are all collinear. Thus, even with arbitrarily many regimes, there is in general no regime-invariant impact vector that rationalizes the joint behavior of the reduced-form covariance matrices.

In principle, the overidentifying restrictions that arise when more than two variance regimes are present could be used to test for whether the VAR parameters are constant across regimes. Such tests could detect the specific failure of Assumption 1 generated by endogenous information acquisition (see e.g. [Rigobon, 2003](#), for an example of such a parameter-stability test). However, in practice, these tests often lack power in the sample sizes typical of macroeconomic applications. As [Lewis \(2021, 2022\)](#) highlights, identification via heteroskedasticity is frequently ‘weak,’ meaning that any statistical signal from variance changes is often drowned out by sampling noise. Economically significant shifts in the impact matrix \mathbf{B} due to endogenous attention may still therefore escape rejection in parameter-stability tests.

Rank-one restrictions. If the impact matrix \mathbf{B} is invariant across regimes and only the variance of a single structural shock changes, then the difference in reduced-form covariance matrices $\Delta\Omega_{j,k} \equiv \Omega^{(j)} - \Omega^{(k)}$ is a rank-one matrix. Equivalently, $\Delta\Omega_{j,k}$ has exactly one nonzero eigenvalue in population.

By contrast, under the data-generating process described above, $\Delta\Omega_{j,k}$ is generically of rank *two* whenever $\alpha_1 + \tau_j^* \beta_1$ and $\alpha_1 + \tau_k^* \beta_1$ are linearly independent. Hence, the rank-one condition

required for heteroskedasticity-based identification fails in population, not merely due to sampling error but as a consequence of endogenous information acquisition.

While formal tests for the rank of the covariance difference matrix exist (e.g. [Robin and Smith, 2000](#); [Chen and Fang, 2019](#)), they face similar limitations to the multiple-regime based tests described above. The rational inattention mechanism explored in this section implies a rank-two structure for $\Delta\Omega_{j,k}$, but the second non-zero eigenvalue may still be small relative to the estimation error of the covariance matrices. In this context, a failure to statistically reject the rank-one null in Robin-Smith style tests does not confirm that the impact vector is constant. It may simply reflect that the attention-driven distortion, while sufficient to bias the point estimate, is difficult to distinguish from sampling noise using second moments alone. In the application to monetary policy in Section 5, I do indeed find that $\Delta\Omega_{j,k}$ has an eigenvalue structure consistent with the rational inattention model, but that these eigenvalues are generally small and estimated somewhat imprecisely.

4 Bias correction

The variation in the impact matrix across regimes highlighted above depends structurally on the attention of relevant agents in the economy. In this section I propose two ways of correcting this bias. The first can be used if the researcher has access to a proxy measure of agent attention in each regime, such as internet search data or newspaper article counts. The second exploits variation across more than two variance regimes. Appendix B contains the results of simulations in which the ground truth is known, showing that both methods accurately recover the true impulse responses in each variance regime in finite samples.

4.1 Correction Method 1: using an attention proxy

Geometric preliminaries. The observed difference between two reduced-form covariance matrices across regimes ($\Delta\Omega_{H,L} \equiv \Omega^{(H)} - \Omega^{(L)}$) is defined in equation (34). If Assumption 1 held, this difference matrix would be rank 1, with one non-zero (positive) eigenvalue ([Rigobon, 2003](#)). However, with endogenous attention, Assumption 1 fails and $\Delta\Omega_{H,L}$ is rank 2, with two non-zero eigenvalues. More specifically, since $\Delta\Omega_{H,L}$ is the difference of two rank-1 positive semi-definite matrices, there will be one positive and one negative eigenvalue.

The eigen-decomposition of $\Delta\Omega_{H,L}$ can be written as

$$\Delta\Omega_{H,L} = \lambda_1 u_1 u_1' + \lambda_2 u_2 u_2' \quad (35)$$

where $\lambda_1 > 0 > \lambda_2$ are the eigenvalues, and u_1, u_2 are the corresponding eigenvectors. The true scaled impact vectors, which I denote $w_H \equiv \sigma_H(\alpha_1 + \tau_H^* \beta_1)$ and $w_L \equiv \sigma_k(\alpha_1 + \tau_L^* \beta_1)$, must lie in the plane spanned by u_1 and u_2 .

With this, we can therefore characterize the set of all candidate impact vectors that are consistent with the observed covariance difference. The solution to $w_H w_H' - w_L w_L' = \lambda_1 u_1 u_1' + \lambda_2 u_2 u_2'$ is given by a hyperbolic rotation of the eigenvectors:

$$w_H(\phi) = \sqrt{\lambda_1} \cosh(\phi) u_1 + \sqrt{|\lambda_2|} \sinh(\phi) u_2 \quad (36)$$

$$w_L(\phi) = \sqrt{\lambda_1} \sinh(\phi) u_1 + \sqrt{|\lambda_2|} \cosh(\phi) u_2 \quad (37)$$

where $\phi \in \mathbb{R}$ is an unknown rotation parameter. Identification of the true IRFs reduces to pinning down the scalar ϕ .

Using an attention proxy. From Lemma 1, we can write the ratio of structural shock variances across the two regimes as

$$r \equiv \frac{\sigma_H^2}{\sigma_L^2} = \frac{1 - \tau_L^*}{1 - \tau_H^*}. \quad (38)$$

If we observe the relative levels of (in)attention, we can therefore infer the relative structural shock variances, which are not typically observed in [Rigobon \(2003\)](#)-style exercises.

With this in hand, we can now solve for ϕ and thus the bias-corrected IRFs, if we impose two further restrictions on the structure of transmission.

Assumption 4 (Attention-invariant variable) *The effect of the shock of interest on one of the variables in the system on impact is unaffected by attention: $\beta_{1,k} = 0$ for one specified variable k .*

Assumption 5 (Known attention direction) *The effect of attention on the impact effect of the shock of interest can be signed for one of the variables of the system: we can specify either $\beta_{1,j} > 0$ or $\beta_{1,j} < 0$ for one variable j .*

Assumption 4 requires that there is at least one variable in the system for which the transmission on impact is purely mechanical, and does not depend on expectations, thus $\beta_{1,k} = 0$. The dynamic effects can still depend on attention through the impact effects on other variables, which then feed in to variable k over time through lags of y_t in the VAR. In monetary policy applications, this attention-invariant variable could plausibly be the policy rate, which moves due to the policymaker actions whether or not people pay attention to them. We use this restriction to identify the rotation parameter ϕ in equations (36) and (37), and discuss it further below.

In general, there will be two roots to the resulting equation, resulting in two possible values for ϕ . Assumption 5 is then employed to select between them. If economic theory suggests that a particular variable should respond more strongly on impact to a shock when agents are paying more attention, this provides the necessary restriction. As we are selecting between only two possible roots, we do not need to form a prior on the magnitude of this attention effect. In the running monetary policy example, one could argue that the impact of monetary policy on long-term yields should be larger when agents pay attention to policy, as then their expectations of future interest rates will adjust more strongly.

With these two assumptions, we proceed in the following way. If $\beta_{1,k} = 0$, then we have that the impact of the shock on that variable is invariant across regimes:

$$\frac{w_{H,k}(\phi)}{\sigma_H} = \frac{w_{L,k}(\phi)}{\sigma_L} \implies w_{H,k}(\phi) = \sqrt{r}w_{L,k}(\phi) \quad (39)$$

Substituting equations (36) and (37) into (39) yields a linear equation in $\tanh(\phi)$:

$$\sqrt{\lambda_1}u_{1,k} + \sqrt{|\lambda_2|} \tanh(\phi)u_{2,k} = \sqrt{r} \left(\sqrt{\lambda_1} \tanh(\phi)u_{1,k} + \sqrt{|\lambda_2|}u_{2,k} \right) \quad (40)$$

$$\implies \tanh(\phi^*) = \frac{\sqrt{\lambda_1}u_{1,k} - \sqrt{r|\lambda_2|}u_{2,k}}{\sqrt{r\lambda_1}u_{1,k} - \sqrt{|\lambda_2|}u_{2,k}}. \quad (41)$$

The quadratic nature of the underlying problem yields two potential roots for ϕ , corresponding to the sign ambiguity of the eigenvectors. To ensure a unique and economically meaningful solution, we normalize the recovered vectors such that the impact on a chosen anchor variable has a given sign, and then use Assumption 5 to select between the roots of ϕ^* .

These restrictions, and equation (41), therefore allow us to recover ϕ^* , and thus $\alpha_1 + \tau_H^* \beta_1$ and $\alpha_1 + \tau_L^* \beta_1$. From these we can compute bias-corrected impulse response functions for both regimes.

Sensitivity. If Assumptions 4 and 5 hold exactly and r is measured without error, this method exactly corrects for the bias derived in Proposition 2 in population. Appendix B also shows that in finite samples of the size used in Section 5, the correction recovers the true impulse responses with little error.

In practice, however, both Assumption 4 and the assumption that we can observe a perfect attention proxy may be difficult to satisfy precisely. In simple monetary policy contexts, as mentioned above, the policy rate might plausibly be attention-invariant on impact as required, but outside of this case such a variable may be harder to find. Similarly, available measures of attention

are known to be noisy in many contexts, which may lead to inaccuracy in the calibrated ratio r .

For this reason, in Appendix B I test the sensitivity of the bias correction to violations of Assumption 4 and errors in the calibration of r . I simulate datasets, varying the true impact effect of attention for the variable assumed invariant ($\beta_{1,k}$) or the true variance ratio r . In each simulation, I estimate a VAR and conduct the bias correction using the (now incorrect) baseline assumptions. The results indicate that the correction is robust to moderate errors in $\beta_{1,k}$ and r . When $|\beta_{1,k}|$ is within 10% of the attention effect on other variables in the system, or when r is mismeasured by 10%, the method still delivers more than 90% of the improvement over the uncorrected estimator that would arise under exact assumptions. The method therefore remains valuable even when the identifying assumptions are not exact.

4.2 Correction Method 2: multiple regimes

While Method 1 provides a closed-form solution, it relies on the existence of an attention-invariant variable (Assumption 4), and of a reliable proxy for attention. In applications where rational inattention permeates all variables in the system (i.e., β_1 has no zero elements), or where attention-related data is unavailable, this will not be appropriate. I show here that observing more than two variance regimes allows for an alternative bias correction, which exploits the geometric constraints implied by the attention mechanism.

Specifically, consider the case where we observe three variance regimes, $s \in \{L, M, H\}$. The structural impact vector in any regime s is given by:

$$v_s = \alpha_1 + \tau_s^* \beta_1 \quad (42)$$

Geometrically, this equation implies that the impact vectors $\{v_L, v_M, v_H\}$ must be *collinear*. They lie on a single line in \mathbb{R}^n with intercept α_1 and direction β_1 . This collinearity is a strong over-identifying restriction which allows us to estimate the structural parameters for each regime.

The estimator. First, normalize the parameters specific to the low-variance regime to $\sigma_L^2 = 1$ and $\tau_L = 0$, effectively interpreting α_1 as the baseline transmission and τ_s as the *relative* increase in attention in the medium- and high-variance regimes. We can then define the vector of remaining deep structural parameters as $\Theta = [\alpha'_1, \beta'_1, \tau_M, \tau_H, \sigma_M^2, \sigma_H^2]'$.

The theoretical difference between the covariance matrix in regime s and the baseline regime L is:

$$\Delta\Omega_{s,L}(\Theta) = \sigma_s^2(\alpha_1 + \tau_s \beta_1)(\alpha_1 + \tau_s \beta_1)' - \alpha_1 \alpha'_1 \quad (43)$$

Our aim is to estimate Θ . However, there still remains a scale indeterminacy between τ_s and β_1 : doubling β_1 and halving both τ_M and τ_H does not affect the resulting $\Delta\Omega_{s,L}$. We therefore add one further normalization, that $\tau_H = 1$. If there are n variables in the vector \mathbf{y}_t (i.e. if $\Delta\Omega_{s,L}$ is $n \times n$), this leaves us with $2n + 3$ parameters in Θ to estimate: n elements of α_1 , n elements of β_1 , plus τ_M , σ_M^2 , σ_H^2 .

We estimate Θ by minimizing the distance between these model-implied moments and the data:

$$\hat{\Theta} = \arg \min_{\Theta} \sum_{s \in \{M, H\}} \left\| \text{vec}(\Delta\Omega_{sL}^{obs}) - \text{vec}(\Delta\Omega_{sL}(\Theta)) \right\|_W^2 \quad (44)$$

where W is a positive definite weighting matrix.

Identification conditions and effective degrees of freedom. A crucial feature of the model is that while variance changes are driven by a single structural shock, the resulting covariance difference matrices $\Delta\Omega_{s,L}$ are of rank two (Proposition 2). This structural constraint implies that the elements of $\Delta\Omega_{s,L}$ are not independent; for $n \geq 3$, the matrix contains only $2n - 1$ independent pieces of information rather than the usual $n(n + 1)/2$.

Accounting for this rank constraint, identification requires:

$$\underbrace{2(2n - 1)}_{\text{Effective Moments}} \geq \underbrace{2n + 3}_{\text{Free Parameters}} \quad (45)$$

which implies $n \geq 2.5$. For a two-variable VAR ($n = 2$) this method cannot therefore identify the parameters Θ without further information. For $n > 2$, however, this is not necessary.

Incorporating external information. In principle, the researcher could combine the approach above with proxy measures of attention, as used in Method 1 above. Restrictions on τ_H/τ_M , for example, reduce the dimensionality of the estimation problem, thus making it possible to estimate Θ even with $n = 2$. With $n > 2$, such external information may still be helpful in making the estimates of α_1 and β_1 more precise.

5 Empirical evidence: attention to monetary policy

I now apply the insights developed above to the analysis of unconventional monetary policy shocks at the zero lower bound (ZLB). This is an obvious testing ground for the methods developed above, as [Wright \(2012\)](#) provides a canonical use of identification via heteroskedasticity in precisely

this context. Understanding the transmission channels behind unconventional policies is also of relevance to policymakers in the current period, as they embark on quantitative tightening policies.

I follow Wright (2012) in using daily data on 6 asset prices to estimate the effects of unconventional monetary policy shocks at the zero lower bound, assuming that monetary policy shocks have a greater variance on FOMC announcement days than non-announcement days.¹⁴ I extend the Wright (2012) sample to the full ZLB period (November 3 2008-December 9 2015). The variables are: the 10-year and 2-year zero-coupon Treasury yields, 5-year and 5-to-10-year forward TIPS breakeven inflation rates, and Moody's AAA and BAA corporate bond yields.¹⁵

I start with two tests of the rational inattention model of Section 3: first, I show that the difference between the reduced-form covariance matrices in the high and low variance regimes has an eigenvalue structure that is inconsistent with Assumption 1, but which would be predicted by the rational inattention model. Second, I use Google Trends searches as a proxy for attention to monetary policy, and find that attention is indeed systematically higher on announcement days, even when conditioning on the magnitudes of the ‘news’ realized at each event. This is precisely the prediction of Lemma 1 in Section 3. Finally, I apply the bias-correction method developed in Section 4.1 to the VAR identified via heteroskedasticity, and find substantial differences in the estimated effects of monetary policy, especially on expected inflation.

5.1 Eigenvalue check

Under the standard identifying assumption (Assumption 1), the difference in reduced-form covariance matrices between announcement and non-announcement days $\Delta\Omega_{H,L} \equiv \Omega^{(H)} - \Omega^{(L)}$ is a rank-one positive semi-definite matrix with exactly one positive eigenvalue. In contrast, under the rational inattention model developed in Section 3, this difference matrix is as defined in equation (27). Since this is the difference of *two* rank-one matrices, it generically has rank two, with one positive and one negative eigenvalue.¹⁶

I test this prediction using the data outlined above. Following the specification in Wright (2012), I estimate a VAR(1) on the six variables and compute the covariance matrices of the residuals separately for announcement days ($\Omega^{(H)}$) and non-announcement days ($\Omega^{(L)}$). The one

¹⁴Wright (2012) also includes the days of 5 landmark policy speeches in the high-variance regime. In the absence of an objective criteria for which speeches should count in the later part of the sample (after Wright (2012) was written), I restrict myself to scheduled FOMC announcement days only. A similar identification scheme is explored as a robustness check to the main results in Nakamura and Steinsson (2018).

¹⁵The data are obtained from the Federal Reserve Bank of St. Louis replication archive, accessible at <https://fredaccount.stlouisfed.org/public/datalist/1111>. The sample period contains 56 announcement days. The Wright (2012) subsample ends on September 30 2011, and contains 23 announcement days.

¹⁶The rank is exactly two when $b_1^{(H)}$ and $b_1^{(L)}$ are not collinear, which holds generically when $\beta_1 \neq 0$ and $\tau_H^* \neq \tau_L^*$.

Table 1: Eigenvalues of the Covariance Difference Matrix

	(1) Point Estimate	(2) Bootstrap 95% CI	(3) Point Estimate	(4) Bootstrap 95% CI
λ_1 (smallest)	−0.0027	[−0.0123, −0.0010]	−0.0058	[−0.0294, −0.0029]
λ_2	−0.0013	[−0.0026, −0.0001]	−0.0033	[−0.0060, −0.0005]
λ_3	−0.0001	[−0.0005, 0.0001]	−0.0002	[−0.0013, −0.0000]
λ_4	0.0001	[−0.0001, 0.0007]	0.0000	[−0.0002, −0.0016]
λ_5	0.0009	[0.0001, 0.0036]	0.0018	[−0.0000, 0.0091]
λ_6 (largest)	0.0099	[0.0023, 0.0224]	0.0177	[0.0018, 0.0430]
Sample period	ZLB	ZLB	Wright	Wright

Notes: Eigenvalues of $\Delta\Omega_{H,L} = \Omega^{(H)} - \Omega^{(L)}$ from VAR(1) residuals. Bootstrap confidence intervals based on 2,000 replications. All samples exclude December 1, 2008 (see text).

change I make is that I exclude the data from December 1, 2008, when Ben Bernanke made a speech indicating that the Federal Reserve would begin buying Treasuries. On this day the 5-year TIPS breakeven inflation rate jumped up by 192 basis points, which is 21 times larger than the standard deviation of daily changes over the sample period. This one outlier observation therefore substantially inflates the variance of breakeven inflation on non-announcement dates, giving rise to concerns about weak identification (Lewis, 2022).¹⁷

Table 1 reports the eigenvalues of the covariance difference matrix $\Delta\Omega_{H,L}$. Columns 1 and 2 show the results for the full ZLB sample. The matrix has three negative eigenvalues, with the smallest equal to −0.0027. Under the null hypothesis of constant transmission (Wright’s identifying assumption), all eigenvalues should be non-negative with at most one strictly positive. The presence of negative eigenvalues suggests a violation of this constant-transmission assumption. For comparison, columns 3 and 4 show the results from the restricted subsample used in Wright (2012), which show similar patterns.

To assess statistical significance, I construct bootstrap confidence intervals by resampling announcement and non-announcement days separately with replacement. The 95% confidence interval for the minimum eigenvalue in the full ZLB sample is [−0.012, −0.001], which lies entirely below zero. This rejects the rank-one null hypothesis at the 5% significance level. Moreover, 100% of bootstrap replications produce at least one negative eigenvalue, indicating that this finding is robust to sampling uncertainty.

The two eigenvalues of largest magnitude—one positive (0.010) and one negative (−0.003)—together account for 84% of the total sum of absolute eigenvalues. This concentration is consistent

¹⁷If I include this outlier observation, the results in Table 1 are qualitatively unchanged: the largest eigenvalue in magnitude is 0.0099 and the second largest is −0.0027, with the 95% bootstrapped confidence intervals for both entirely above and below 0 respectively.

with an approximately rank-two structure, as predicted by the rational inattention model. Eigenvalues $\lambda_3 - \lambda_5$ are an order of magnitude smaller and likely reflect estimation noise. λ_2 is less clearly near-zero, possibly suggesting some other failure of Assumption 1 beyond the optimal information choice studied here.

The finding of negative eigenvalues in $\Delta\Omega_{H,L}$ is consistent with the mechanism proposed in this paper. If agents pay more attention to monetary policy on announcement days ($\tau_H^* > \tau_L^*$), then the transmission of policy shocks differs across regimes. The covariance difference matrix is no longer proportional to a single outer product, and the standard heteroskedasticity-based identification procedure recovers a pseudo-parameter rather than the structural impact in either regime. The next subsection provides complementary evidence that attention to monetary policy does indeed spike on FOMC announcement days.

5.2 Attention to monetary policy between regimes

In this monetary policy context, the bias highlighted in Proposition 2 and Corollary 2 arises whenever decision-makers pay more attention to monetary policy on FOMC announcement days than on non-announcement days. In this subsection, I test this directly, using Google Trends search volume as a proxy for attention, as in e.g. [Da et al. \(2011, 2015\)](#).

Specifically, I use the `trendcon` procedure and package from [Eichenauer et al. \(2022\)](#) to construct a daily series of Google searches for the term “Federal Reserve” from January 1 2006 to June 30 2019.¹⁸ Since Google Trends does not provide daily data over this length of sample, the index is constructed by splicing together multiple Google Trends queries, with adjustments to ensure consistency of units across windows (see [Eichenauer et al. \(2022\)](#) for a detailed explanation of the methodology, and [Wang \(2025\)](#) for a similar procedure for searches for the FOMC). This means that the resulting index is not exactly scaled to be within $[0, 100]$ as is typical for stand-alone Google Trends queries. The interpretation, however, is similar, as 98.8% of the observations lie in this range.

Table 2 shows the average index values on announcement and non-announcement days for the ZLB sample that is the main sample of interest, as well as the shorter sample analysed by [Wright \(2012\)](#), and the full sample that includes non-ZLB periods. The patterns are similar across samples,

¹⁸Since this series requires multiple calls to the Google Trends API, it can only cover a single keyword at a time, and cannot sum multiple keyword search intensities. Constructing an index from multiple `trendcon` queries would not produce an interpretable result, as each search term would have its own scale. However, from the monthly data available directly from Google Trends, I find that ‘Federal Reserve’ has more than $10\times$ the search activity of ‘FOMC,’ and more than $20\times$ the activity of ‘Federal Funds Rate’ over the 2006-2025 period. As a robustness check, in Appendix C I re-run the analysis of Table 2 for the search term ‘Monetary Policy,’ which has approximately 15% of the search activity of ‘Federal Reserve.’ All results are robust.

with search activity approximately two-thirds higher on announcement days. Even in the relatively short [Wright \(2012\)](#) subsample, these differences are strongly significant.

Table 2: Attention to the Federal Reserve: FOMC vs. Non-FOMC Days

Sample	Observations		Mean Search Index			<i>t</i> -statistic
	FOMC	Non-FOMC	FOMC	Non-FOMC	Difference	
Wright (2008–2011)	23	1,039	86.106	49.751	36.356	9.011
ZLB (2008–2015)	56	2,537	86.297	50.132	36.164	14.141
Full (2006–2019)	113	4,816	88.975	46.427	42.547	22.882

Notes: Google Trends search index for “Federal Reserve” constructed using the [Eichenauer et al. \(2022\)](#) methodology. FOMC days are scheduled Federal Open Market Committee announcement days. The *t*-statistic tests the null hypothesis that mean attention is equal across FOMC and non-FOMC days.

While indicative, the greater search volume on announcement days is not itself enough to reject the assumption of invariant transmission (Assumption 1). If attention rises only with shock realizations (not variances), then we would see greater attention on announcement days because the average shock magnitude is greater on those days, but there would be no difference in the impact of a unit shock in each regime. To test if this is the case, I therefore estimate

$$\log(A_t) = \alpha + \beta \text{FOMC}_t + \gamma' |S_t| + \delta' X_t + \varepsilon_t, \quad (46)$$

where A_t is the daily search index (trimmed to drop the smallest and largest 1% of observations), FOMC_t is an indicator equal to 1 on FOMC meeting days, and $|S_t|$ are measures of realized monetary policy shocks.¹⁹ Specifically, I use the absolute values of the Federal Funds Rate, Forward Guidance, and LSAP factors from the high-frequency identification in [Swanson \(2021\)](#). X_t includes controls for day-of-week effects, month fixed effects, and two lags of the dependent variable. The trimming of A_t removes the small number of negative values, so using $\log(A_t)$ does not involve any further sample restrictions.

The rational inattention mechanism studied above implies that attention to monetary policy should be higher on announcement days irrespective of the realized shocks. That is, the coefficient $\beta > 0$. If instead the increased searches on announcement days are entirely due to shock realizations, and not the volatility regime, then $\beta = 0$ and all differences in attention between regimes are captured by $\gamma > 0$. Since there are only a small number of announcements in the period analysed by [Wright \(2012\)](#), I estimate this regression only for the ZLB sample (2008–2015) and the full sample (2006–2019).

¹⁹The sample trimming is completed separately for each sample period, so the ZLB-period regressions are run for the days in which A_t is in the middle 98% of the observations for that period.

The results are presented in Table 3. The first column presents the model omitting FOMC_t and $|S_t|$. Column 2 adds FOMC_t , and column 3 adds $|S_t|$. As predicted by the rational inattention model, β is large and significantly greater than 0 across specifications where FOMC_t is included. The γ coefficients on absolute monetary policy shocks are typically close to 0 and not significant, with the exception of LSAP shocks in the ZLB period. Indeed, it is reassuring that this element of γ is positive and significant, as it provides a sense-check on the attention measure: during the ZLB, when asset purchases were a large and unprecedented feature of US monetary policy, we should expect more attention on the Federal Reserve after large LSAP shocks.

A further check that attention is however determined largely by the regime (announcement day or not) rather than the shock realization comes from studying the R^2 of these regressions. Adding the controls for shock size in columns 3 and 6 only adds a minimal amount to the R^2 , an order of magnitude less than is added by the inclusion of FOMC_t in columns 2 and 5. This adds further support to the conclusion that the majority of differences in search volumes between regimes are not driven by the size of the shock realization.

Table 3: Attention to the Federal Reserve

	(1)	(2)	(3)	(4)	(5)	(6)
FOMC day		0.223*** (0.025)	0.163** (0.074)		0.319*** (0.026)	0.271*** (0.052)
Fed funds rate shock (abs.)			-0.217 (0.439)			-0.051 (0.263)
Forward guidance shock (abs.)			0.044 (0.041)			0.091 (0.057)
LSAP shock (abs.)			0.186*** (0.054)			0.023 (0.079)
Observations	2,542	2,542	2,542	4,822	4,822	4,822
R-squared	0.7869	0.7909	0.7913	0.7921	0.7986	0.7987
Sample period	ZLB	ZLB	ZLB	Full	Full	Full

Day-of-week and month fixed effects, and two lags of $\log(\text{attention})$, included in all specifications.

Robust standard errors in parentheses.

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

5.3 Bias-corrected IRFs

I now apply the bias-correction of Section 4.1 to the structural VAR specified as in Wright (2012).²⁰ As described above, this requires two extra assumptions beyond those currently made, and an observation of the relative (in)attention between variance regimes. I choose these as follows.

Assumption 4. The impact effect of monetary policy shocks on the 2-year Treasury yield is unaffected by attention.

Assumption 5. The impact effect of monetary policy shocks on the BAA corporate bond yield is larger in magnitude in the high-variance/high-attention regime.

Attention measure. In the Google Trends series above, the mean of the search activity index is 86.3 in the high-variance regime, and 50.1 in the low-variance regime. Taking an index of 0 to reflect no attention ($\tau_{i,t} = 0$) and 100 to reflect complete attention ($\tau_{i,t} = 1$), and assuming a linear relationship between the search index and $\tau_{i,t}$, we obtain $\frac{1-\tau_L^*}{1-\tau_H^*} = \frac{1-0.501}{1-0.863} = 3.64$.

Discussion. The ideal choice of attention-invariant variable for Assumption 4 in monetary policy settings is the policy rate, which moves mechanically with policy actions. However, since my purpose is to identify the effects of unconventional monetary policy shocks at the zero lower bound, the sample is chosen such that the target federal funds rate is constant at the lower bound (0.25%) in all periods, and so is not included in the VAR. However, prior literature provides a mechanism that suggests the 2-year Treasury yield is a plausible alternative.

Wright (2012) argues that the unconventional monetary policy shocks in the sample operate largely through the behavior of preferred-habitat investors and sophisticated arbitrageurs, as in models such as Vayanos and Vila (2021). This mechanism works through portfolio rebalancing by institutional investors with mandates tied to specific maturities, and arbitrageur responses to resulting price pressure. Since these portfolio adjustments are largely mechanical responses to changes in relative bond supplies they do not require investors to form expectations about the macroeconomic implications of policy.

²⁰Correction method 2 described in Section 4.2 is not applicable here as we only observe two variance regimes. For robustness checks Wright (2012) further divides his high-variance regime into “especially important” days and less important days, but this is based on the realized policy changes on those days, not on the volatility that was known ex-ante to market participants and other economic agents. This extra division would not therefore satisfy Assumption 3 that is required for the model developed in this paper to be applicable.

This mechanism is likely to be especially dominant for shorter maturity assets. Assets with longer maturities may also be affected by the impact of policy announcements on expectations, both of future macroeconomic developments (Boneva et al., 2016; Melosi, 2017) and monetary policy (Bauer and Rudebusch, 2014; Bhattacharai et al., 2023).²¹ For this reason I impose attention invariance on the 2-year Treasury yield rather than the 10-year Treasury yield.

Despite these arguments, Assumption 4 remains rather strong ex-ante. I provide ex-post validation at the end of this subsection, along with a robustness test relaxing the assumption.

For Assumption 5, I consider the impact of monetary policy shocks on corporate bond yields. Part of this transmission is likely to occur through monetary policy's effect on credit spreads: when investors pay closer attention to monetary policy announcements, they update more strongly on what the announcement signals about the Federal Reserve's economic outlook and future policy plans. This would lead them to extract more information about the likely path of the economy, and hence corporate default probabilities, than one who observes policy changes only passively through their effect on Treasury yields. This expectational channel should be particularly prevalent for higher-risk bonds, so I impose the assumption on the BAA yield.

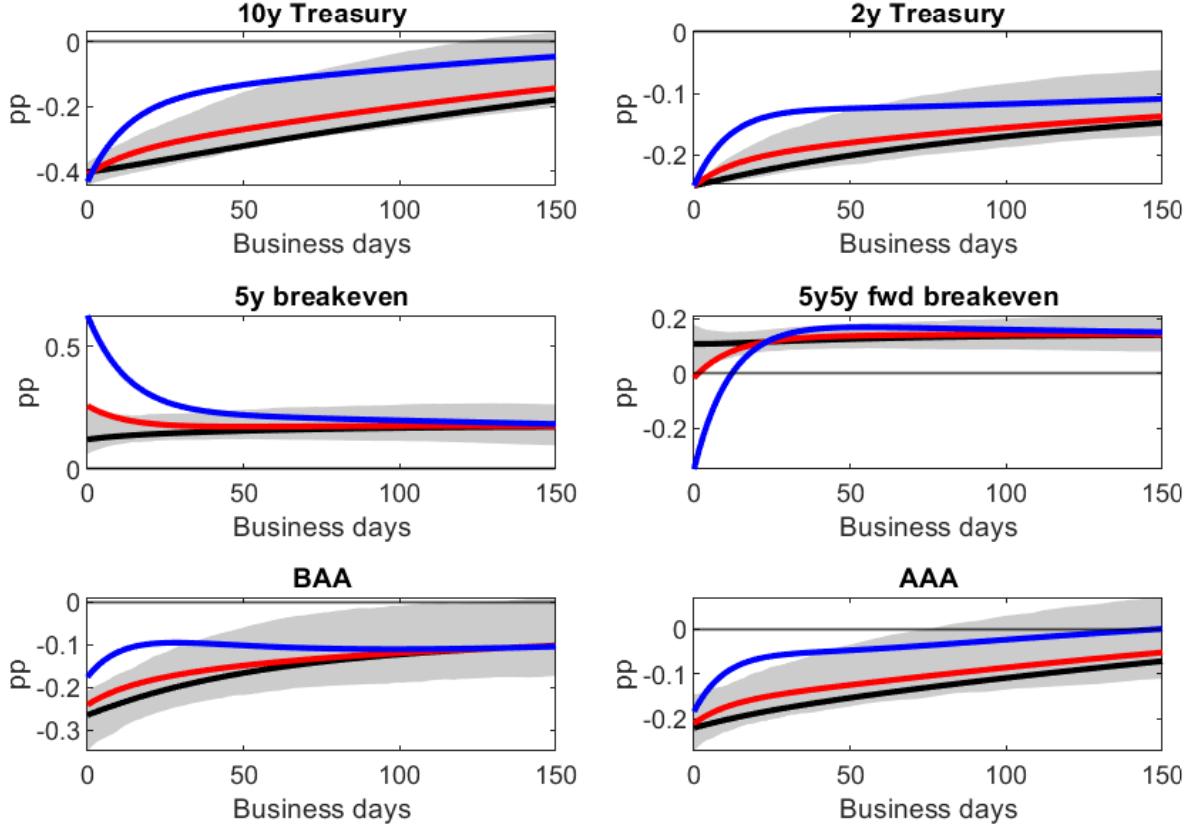
Note that Assumption 5 is only used to select between the two roots of equation (41). As such, there are a range of different possible assumptions that would deliver identical impulse response estimates. For example, Figure 1 below shows that assuming attention amplifies the impact of monetary policy shocks on AAA corporate bond yields, rather than BAA yields, would select exactly the same root for ϕ^* , and thus would imply exactly the same impulse responses. Indeed, while the impact response of the AAA yield is also amplified by attention, the amplification is smaller than in the BAA yield, consistent with the default-risk mechanism outlined here.

Finally, the linear mapping between Google Trends searches and attention is chosen for its simplicity and transparency. However, the core results below are robust to alternative mappings. Appendix C re-computes the same bias-corrected IRFs assuming several other mappings, and in all cases the bias correction substantially alters the results, with the same qualitative changes as those displayed below. Note that in all such mappings, I only make use of the *ratio* between (in)attention across regimes, so scaling by some baseline measure of attention leaves the results unchanged.

Results. Figure 1 shows the results. The un-corrected estimates following the specification in Wright (2012) are shown in black. I normalize the shock such that it delivers a 25 basis point decline in the 2-year Treasury yield on impact.

²¹ Wright (2012) emphasizes the latter channel, noting that “LSAPs could also work in other ways, such as by affecting agents' expectations of the future course of monetary policy”(p. F447).

Figure 1: Impulse response functions with and without the correction for endogenous attention.



Note: Impulse responses to an unconventional monetary policy shock identified via heteroskedasticity, following Wright (2012). The black line shows the uncorrected estimate assuming regime-invariant transmission (Assumption 1), with the shaded region indicating 68% bootstrap confidence intervals, constructed using the Kilian (1998)-adjusted bootstrap. The Kilian bias adjustment is also applied to the point estimates. The red line shows the bias-corrected impulse response for the high-attention regime (FOMC announcement days), and the blue line shows the bias-corrected response for the low-attention regime (non-announcement days). Bias correction follows the method in Section 4.1, imposing that the impact effect on the 2-year Treasury yield is attention-invariant (Assumption 4) and that attention amplifies the impact effect on BAA corporate bond yields (Assumption 5). The responses are normalized such that the 2-year Treasury yield falls by 25 basis points on impact in all cases. Sample: November 2008 to December 2015, excluding December 1, 2008 (see text). The VAR is estimated with one lag on daily data.

The corrected estimates for the high and low variance regimes are shown in red and blue respectively. Since Assumption 4 implies that the 2-year Treasury yield is the same on impact in both regimes, I use the same normalization of the shock size in these corrected IRFs as in the uncorrected estimates.

For all variables, the uncorrected IRFs do not typically lie within the corrected IRFs for the two regimes. This is a consequence of the fact that the uncorrected IRFs are non-convex combinations of the true IRFs (Proposition 2). The bias-corrected IRFs are also not, in general, scaled versions of the uncorrected IRFs or of each other. Even though inattention only affects the impact matrix B in

equation (1), it affects the impact of the shock on each of the variables in the VAR differentially. In the first period after the shock, the differential impacts then imply differential changes in the lags of the outcome vector, thus leading to sometimes very different dynamics across regimes.

The precise way in which impulse responses differ from their un-corrected counterparts, and between regimes, has implications for the overall assessment of unconventional monetary policy shocks and for the role of expectations and attention in their transmission.

Overall shock transmission. The results in the top row of Figure 1 suggest that expansionary unconventional monetary policy shocks have less persistent effects on Treasury yields than would be estimated from a standard identification via heteroskedasticity, particularly in the low-attention regime where the differences are large.

The figure also shows unconventional policy shocks having larger positive effects on 5-year inflation expectations than in the uncorrected estimation, but more muted or even negative effects on longer-term expectations in the initial month after the shock. This initial rotation in breakevens is similar to what [Wright \(2012\)](#) found in his reduced sample, though it is substantially stronger here.

Finally, the transmission to corporate bond yields is also less powerfully expansionary than a standard estimation would suggest. On impact, the difference is largest for BAA yields, but over time these responses converge on each other, while the relative weakening of the AAA yield response persists over time.

These features are important, because they reflect the key transmission channels through which unconventional policies are hypothesized to operate. Forward guidance and quantitative easing directly reduce long-term Treasury rates. Portfolio rebalancing channels imply those changes in government bond yields spill over to other assets, including corporate bonds, making it easier and cheaper for firms to raise financing and expand ([Gagnon et al., 2011](#); [Vayanos and Vila, 2021](#)). Signalling channels imply that expansionary policy increases expected inflation, decreasing real interest rates and thus stimulating consumption and investment ([Bhattarai et al., 2023](#)).²²

The results obtained in Figure 1 imply that, on average, the portfolio balance effect is weaker than the uncorrected estimation would suggest, and the signalling effect is stronger, at least for the short-to-medium term inflation expectations captured by the 5-year breakeven rate. For longer-term inflation expectations, the estimated effects are also greater than the in the uncorrected estimation, but only after the first month after the policy shock. The comparison between the corrected impulse responses in high- and low-variance regimes is informative about the mechanisms behind these results.

²²See [Dell’Ariccia et al. \(2018\)](#) and [Kuttner \(2018\)](#) for reviews of this literature.

Before turning to this, it is worth noting that the deviations from the un-corrected IRFs are much smaller in the high-variance regime. For some questions, the high-variance/high-attention regime IRF may indeed be the relevant object: if the researcher is interested in the effect of policy decisions made on scheduled days then the high-attention regime is the appropriate one to consider. However, for other questions the researcher may be interested in monetary policy in general, such as from speeches that are not highly anticipated before they occur. In this case, the low-attention regime is also relevant. Indeed, this constitutes the vast majority of days in the sample. And in this case the uncorrected IRFs are notably biased for all variables. Moreover, even if transmission on scheduled announcement days is the relevant object for a researcher, comparing corrected IRFs in high- and low-variance periods can help to shed light on the transmission channels behind the estimated effects, which is of interest even to those only considering shocks that arrive in a single regime.

Expectations and attention as drivers of policy transmission. Figure 1 does not only give implications for the understanding of unconventional monetary policies on average. By producing separated impulse responses in conditions of high and low attention, it also allows us to learn about how much of unconventional policy transmission comes through channels associated with the updating of expectations, mediated through possibly limited attention. This also has the extra benefit of aiding the interpretation of the average results.

In general, the large gaps between the corrected IRFs in the high and low variance regimes suggest that a substantial portion of the transmission of unconventional monetary policy comes through expectations, rather than mechanical balance sheet of no-arbitrage effects. The horizon at which attention affects transmission, and whether it amplifies or dampens the shock's effects, vary across variables.

Treasury yields. On impact, unconventional monetary policy shocks have the same impact on the 2-year Treasury yield regardless of the regime. This holds by construction (Assumption 4). However, the impact is also the same across regimes for the 10-year Treasury yield, and this is not imposed by the bias-correction procedure. This is (i) reassuring that placing Assumption 4 on the impact effect on a Treasury yield is appropriate (see Robustness discussion below); and (ii) indicative that the impact of unconventional monetary policy on Treasury yields is largely independent of expectations, consistent with theories of preferred habitat investors rebalancing portfolios (Vayanos and Vila, 2021).

However, after the initial day of the shock, the responses of Treasury yields begin to diverge

between regimes. Specifically, they are substantially less persistent in the low-variance regime, when attention to monetary policy is weaker. 50 business days after the shock, the effects on the 2-year yield are approximately a third smaller in the low-variance than in the high-variance regime. For the 10-year yield, the low-variance regime response is only half as strong as the high-variance regime response at this horizon. This suggests that while the effect of unconventional policy on Treasuries on impact may be mechanical, expectations come to play an important role in the following days. Attention causing greater persistence in Treasury yields is consistent, for example, with unconventional policy shocks in part signalling future policy actions by the Federal Reserve, as explored theoretically by [Bhattarai et al. \(2023\)](#).²³

Corporate bond yields. The AAA corporate bond yields follow similar patterns to Treasury yields: while there is a small amplifying effect of attention on impact, the IRFs for the two regimes diverge over the following days, with low attention associated with substantially less persistent shock effects. This is as would be expected from portfolio rebalancing mechanisms, where high-quality corporate bond yields follow Treasury yields. The fact that attention amplifies the impact effect of the unconventional monetary policy shock is also consistent with the arguments put forward in the discussion of Assumption 5 above, that when investors are more attentive to an expansionary monetary policy shock they revise their default probability expectations more strongly downwards, further reducing corporate bond yields.

BAA corporate bond yields, however, behave quite differently. On impact, attention substantially amplifies the effect of the shock on BAA corporate bond yields. The impact is more than 25% stronger in the high-variance regime, compared with 13% impact amplification for AAA bond yields. This larger amplification from attention is again consistent with the default-expectations mechanism proposed above, as default probabilities are larger and more salient than for AAA-rated bonds. Similarly, it could also reflect a role for expectations in the risk-taking channel of monetary policy ([Bauer et al., 2023](#)): if investors paying attention to monetary policy increase their risk appetite, this would lower credit spreads through the excess bond premium ([Gilchrist and Zakrajšek, 2012](#)). In the low attention regime, investors do not observe the change in monetary policy as precisely, and thus the amplification through risk appetite is weaker.

While this amplification effect initially grows, almost doubling the response one month after the shock, the effects of attention then die away, and the impulse responses in the two regimes converge. This suggests that the attention-sensitive component of this transmission is short-lived, whether it is through default expectations, risk premia, or both. At longer horizons, the effects on

²³Indeed, this is the logic used by [Gürkaynak et al. \(2005a\)](#) and the literature that follows them to separate monetary policy shocks into “Target” and “Path/Forward Guidance” factors.

lower-rated investment-grade bond yields are largely driven by forces independent of monetary policy expectations. If default expectations are the key driver of the initial amplification, this could for example reflect investor learning about default risk at the firm level, which a range of information frictions would render sluggish (Mankiw and Reis, 2002; Angeletos et al., 2020).

Breakevens. The impulse responses for breakeven inflation rates show the most striking differences between high and low variance regimes. The increase in the 5-year breakeven, and the decrease in the 5-10-year forward, are much more pronounced in the low-variance regime. Attention and the response of expectations and perceptions of monetary policy therefore *weakens* these initial responses of market-based inflation expectations.

While this might initially seem strange, it is actually evidence in favor of the intuitive argument given by Wright (2012) to explain the similar short-run rotation in breakevens in his results:

There is a rotation of TIPS break-evens, with five-year break-evens rising and 5–10-year forward break-evens falling. A possible interpretation is that the stronger outlook for demand boosts the short-to-medium-run inflation outlook, but the fact that the LSAPs are overwhelmingly concentrated in nominal (rather than TIPS) securities has an offsetting effect, pushing longer term break-evens lower. Wright (2012), p.F460.

The hypothesis that long-term breakevens are distorted by imperfect arbitrage after large asset purchases by policymakers would suggest a purely mechanical effect in the initial periods after a shock.²⁴ The fact that greater attention offsets this rotation in Figure 1 is exactly what one would expect from such non-expectational mechanisms. The puzzling dynamics die away a month after the shock, again consistent with the driver being a short-run market friction, not a fundamental shift in expectations. At longer horizons breakeven inflation is estimated to increase after the shock.

In fact, the results also imply that the short-run effects of unconventional monetary policy shocks on inflation expectations are substantially greater than a look at raw breakevens data would suggest. The effect on the 5-10-year forward breakeven inflation is close to 0 on impact in the high-variance regime, but this is because attention, and the resulting updating of expectations, is offsetting a very large mechanical effect. This is again consistent with strong signalling effects of unconventional monetary policies, but only when agents are paying attention to the policy announcements.

These insights are important for understanding the zero lower bound period, but also carry relevance for policymakers today. Understanding the transmission channels of unconventional

²⁴See Haubrich et al. (2012) and D'Amico et al. (2018) for discussions of liquidity premia in inflation-linked Treasury markets. Abrahams et al. (2016) in particular find that variation in risk and liquidity premia were especially important for forward breakeven inflation during the financial crisis period.

monetary policy, and how much of that transmission depends on expectations, will be important as central banks embark on quantitative tightening (QT) in the years to come. In particular, the powerful role for expectations and attention found above suggests that communication policies will be of great importance to the effects of QT. The most surprising implication is that if policymakers wish to conduct QT without strong effects on Treasury yields or corporate bond yields, they should not communicate in advance, to aim for low attention to the policy changes. This would, however, come with greater short-run volatility in inflation expectations. There may also be other reasons why communication would be desirable (Haldane et al., 2021): the results here then simply suggest that this communication also may have costs in the markets for government and corporate bonds.

Robustness. While the results presented here align with theoretical developments in the study of unconventional monetary policy, it should be reiterated that they rely on the Google Trends proxy for attention, and its mapping into the model object τ_s^* . Both of these steps are necessarily imperfect: despite its increasingly common use, Google Trends is only a noisy proxy for attention, and the mapping into the model is stylized. In Appendix C I therefore repeat the bias correction in Figure 1 for alternative calibrations of $r = \frac{1-\tau_L^*}{1-\tau_H^*}$, adjusting the Google Trends sample and the mapping into the model. All results are qualitatively consistent with Figure 1.

I also examine the sensitivity of the results to Assumption 4, which imposes that the impact effect of monetary policy on the 2-year Treasury yield is invariant to attention. The key threat to this assumption is that monetary policy shocks affect expectations of future policy or macroeconomic conditions on impact, and that this has an effect on Treasury yields beyond the mechanical portfolio rebalancing discussed above.

While ex-ante this sounds plausible, two pieces of ex-post evidence suggest that the concern is small. First, Figure 1 shows that the impact effect of monetary policy shocks on the 10-year Treasury yield is very similar between regimes, despite the attention invariance only being imposed on the 2-year yield. The corrected high-attention impact effect is within 5% of the low-attention impact, well within the bootstrap confidence intervals for the uncorrected estimates. This is consistent with portfolio balance channels initially dominating in long-term assets, and the expectational effects previously mentioned being minimal. If expectational effects are this small for long-term Treasuries, they are even more likely to be small for shorter maturities.

Second, in Appendix C I relax Assumption 4, and instead choose the rotation ϕ^* to minimize the difference between the effects of attention on impact between the 2-year and 10-year Treasury yields. This approach is motivated by the hypothesis that Treasury markets of different maturities should be affected by similar forces. If expectational effects were large in these markets, this approach

would reveal that, but this is not what I find. On impact, the response of the 2-year Treasury yield is less than 3% larger in the high- vs low-attention regime. Again, the results are consistent with Figure 1, suggesting that the main results derived under Appendix 4 are robust.

6 Conclusion

I have argued that endogenous information acquisition affects macroeconomic practice. Standard models of rational inattention, and a wealth of empirical evidence, predict that economic agents process more information about variables when their volatility rises. Since information processing affects the way those agents react to structural economic shocks, the transmission of those shocks depends on the variance of the process from which they were drawn. This challenges a key assumption required for structural vector autoregressions identified via heteroskedasticity, that the responses to structural shocks are constant, even as their distributions shift.

I show that ignoring this mechanism biases impulse responses from heteroskedasticity-based identification. However, I also show that the bias can be corrected: one method uses external attention proxies, another exploits variation across three or more variance regimes.

Beyond bias correction, these methods also yield new economic insights. Because attention reacts to the variance of the structural shocks, the true impulse responses to those shocks differ between variance regimes. The correction methods recover regime-specific impulse responses, and comparing them reveals how the updating of expectations contributes to the transmission of the structural shocks of interest.

I apply these methods to the transmission of unconventional monetary policy at the zero lower bound in the US, following and extending Wright (2012). I find that the corrected impulse responses show less persistence in the transmission to Treasury and corporate bond yields, and very different short-run responses of breakeven inflation rates, than those obtained from a standard identification that does not correct for endogenous information. Comparing responses across regimes reveals a powerful role for signalling channels of unconventional monetary policy, especially for long-term inflation expectations.

The bias mechanism I discuss applies in principle to any heteroskedasticity-based identification, including Markov-switching and smooth-transition specifications. Developing practical correction methods for these more complicated settings is a promising avenue for future research. More broadly, the finding that a common econometric technique embeds implicit assumptions about information processing suggests that rational inattention may have unexplored implications for other areas of macroeconomics.

References

Abrahams, M., Adrian, T., Crump, R. K., Moench, E., and Yu, R. (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics*, 84:182–200.

Acharya, S. and Wee, S. L. (2020). Rational inattention in hiring decisions. *American Economic Journal: Macroeconomics*, 12(1):1–40.

Afrouzi, H. and Yang, C. (2021). Dynamic Rational Inattention and the Phillips Curve. *CESifo Working Paper*.

Andrei, D., Friedman, H., and Ozel, N. B. (2023). Economic uncertainty and investor attention. *Journal of Financial Economics*, 149(2):179–217.

Andrei, D. and Hasler, M. (2015). Investor Attention and Stock Market Volatility. *The Review of Financial Studies*, 28(1):33–72.

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2020). Imperfect Macroeconomic Expectations: Evidence and Theory. *NBER Macroeconomics Annual*, 35(1):2–86.

Angelini, G., Bacchicocchi, E., Caggiano, G., and Fanelli, L. (2019). Uncertainty across volatility regimes. *Journal of Applied Econometrics*, 34(3):437–455.

Bacchicocchi, E., Castelnuovo, E., and Fanelli, L. (2018). GIMME A BREAK! IDENTIFICATION AND ESTIMATION OF THE MACROECONOMIC EFFECTS OF MONETARY POLICY SHOCKS IN THE UNITED STATES. *Macroeconomic Dynamics*, 22(6):1613–1651.

Bacchicocchi, E. and Fanelli, L. (2015). Identification in Structural Vector Autoregressive Models with Structural Changes, with an Application to US Monetary Policy. *Oxford Bulletin of Economics and Statistics*, 77(6):761–779.

Bauer, M. D., Bernanke, B. S., and Milstein, E. (2023). Risk Appetite and the Risk-Taking Channel of Monetary Policy. *Journal of Economic Perspectives*, 37(1):77–100.

Bauer, M. D. and Rudebusch, G. D. (2014). The Signaling Channel for Federal Reserve Bond Purchases. *International Journal of Central Banking*, 10(3).

Baumeister, C. and Benati, L. (2013). Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound. *International Journal of Central Banking*, 9(2).

Benchimol, J., Bounader, L., and Dotta, M. (2025). Estimating Behavioral Inattention. *Journal of Economic Behavior & Organization*, 236:107068.

Bhattarai, S., Eggertsson, G. B., and Gafarov, B. (2023). Time Consistency and Duration of Government Debt: A Model of Quantitative Easing. *The Review of Economic Studies*, 90(4):1759–1799.

Blanchard, O. J. and Quah, D. (1989). The Dynamic Effects of Aggregate Demand and Supply Disturbances. *American Economic Review*, 79(4).

Boneva, L., Cloyne, J., Weale, M., and Wieladek, T. (2016). The effect of unconventional monetary policy on inflation expectations: Evidence from firms in the United Kingdom. *International Journal of Central Banking*, 12(3).

Brenna, F., De Graeve, F., and Wouters, R. (2023). Macro- nancial feedbacks through time. *Working Paper*.

Brunnermeier, M., Palia, D., Sastry, K. A., and Sims, C. A. (2021). Feedbacks: Financial Markets and Economic Activity. *American Economic Review*, 111(6):1845–1879.

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35.

Chen, Q. and Fang, Z. (2019). Improved inference on the rank of a matrix. *Quantitative Economics*, 10(4):1787–1824.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, pages 65–148. Elsevier.

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). How Do Firms Form Their Expectations? New Survey Evidence. *American Economic Review*, 108(9):2671–2713.

Da, Z., Engelberg, J., and Gao, P. (2011). In Search of Attention. *The Journal of Finance*, 66(5):1461–1499.

Da, Z., Engelberg, J., and Gao, P. (2015). The Sum of All FEARS Investor Sentiment and Asset Prices. *The Review of Financial Studies*, 28(1):1–32.

D’Amico, S., Kim, D. H., and Wei, M. (2018). Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices. *Journal of Financial and Quantitative Analysis*, 53(1):395–436.

Dasgupta, K. and Mondria, J. (2018). Inattentive importers. *Journal of International Economics*, 112:150–165.

de Groot, O. and Haas, A. (2023). The signalling channel of negative interest rates. *Journal of Monetary Economics*, 138:87–103.

Dean, M. and Neligh, N. (2023). Experimental Tests of Rational Inattention. *Journal of Political Economy*, 131(12):3415–3461.

Dell’Ariccia, G., Rabanal, P., and Sandri, D. (2018). Unconventional Monetary Policies in the Euro Area, Japan, and the United Kingdom. *Journal of Economic Perspectives*, 32(4):147–172.

Eggertsson, G. B. and Woodford, M. (2003). The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity*, 1.

Ehrmann, M., Fratzscher, M., and Rigobon, R. (2011). Stocks, bonds, money markets and exchange rates: Measuring international financial transmission. *Journal of Applied Econometrics*, 26(6):948–974.

Eichenauer, V. Z., Indergand, R., Martínez, I. Z., and Sax, C. (2022). Obtaining consistent time series from Google Trends. *Economic Inquiry*, 60(2):694–705.

Ellison, M. and Macaulay, A. (2021). A Rational Inattention Unemployment Trap. *Journal of Economic Dynamics and Control*, 131(104226).

Flynn, J. P. and Sastry, K. A. (2024). Attention Cycles. *NBER Working Paper*.

Gaballo, G. (2016). Rational inattention to news: The perils of forward guidance. *American Economic Journal: Macroeconomics*, 8(1):42–97.

Gagnon, J., Raskin, M., Remache, J., and Sack, B. (2011). The Financial Market Effects of the Federal Reserve’s Large-Scale Asset Purchases. *International Journal of Central Banking*, 7(1).

Geiger, M. and Scharler, J. (2021). How Do People Interpret Macroeconomic Shocks? Evidence from U.S. Survey Data. *Journal of Money, Credit and Banking*, 53(4):813–843.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34.

Gilchrist, S. and Zakrajšek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, 102(4):1692–1720.

Gondhi, N. (2023). Rational inattention, misallocation, and the aggregate economy. *Journal of Monetary Economics*, 136:50–75.

Gürkaynak, R. S., Sack, B., and Swanson, E. (2005a). Do Actions Speak Louder Than Words? The

Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking*.

Gürkaynak, R. S., Sack, B., and Swanson, E. (2005b). The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models. *American Economic Review*, 95(1):425–436.

Haldane, A., Macaulay, A., and McMahon, M. (2021). The Three E’s of Central-Bank Communication with the Public. In Pasten, E. and Reis, R., editors, *Independence, Credibility, and Communication of Central Banking*, pages 279–342. Central Bank of Chile.

Haubrich, J., Pennacchi, G., and Ritchken, P. (2012). Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps. *The Review of Financial Studies*, 25(5):1588–1629.

Ikeda, D., Li, S., Mavroeidis, S., and Zanetti, F. (2024). Testing the Effectiveness of Unconventional Monetary Policy in Japan and the United States. *American Economic Journal: Macroeconomics*, 16(2):250–286.

Iovino, L. and Sergeyev, D. (2023). Central Bank Balance Sheet Policies without Rational Expectations. *Review of Economic Studies*, 90(6):3119–3152.

Kacperczyk, M., Van Nieuwerburgh, S., and Veldkamp, L. (2016). A Rational Theory of Mutual Funds’ Attention Allocation. *Econometrica*, 84(2):571–626.

Kilian, L. (1998). Small-sample Confidence Intervals for Impulse Response Functions. *The Review of Economics and Statistics*, 80(2):218–230.

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

Kuttner, K. N. (2018). Outside the Box: Unconventional Monetary Policy in the Great Recession and Beyond. *Journal of Economic Perspectives*, 32(4):121–146.

Lanne, M. and Lütkepohl, H. (2010). Structural Vector Autoregressions With Nonnormal Residuals. *Journal of Business & Economic Statistics*, 28(1):159–168.

Lei, X. (2019). Information and Inequality. *Journal of Economic Theory*, 184:104937.

Lewis, D. J. (2021). Identifying Shocks via Time-Varying Volatility. *The Review of Economic Studies*, 88(6):3086–3124.

Lewis, D. J. (2022). Robust Inference in Models Identified via Heteroskedasticity. *The Review of Economics and Statistics*, 104(3):510–524.

Lewis, D. J. (2025). Identification Based on Higher Moments in Macroeconometrics. *Annual Review of Economics*, 17(Volume 17, 2025):665–693.

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information Frictions Among Firms and Households. *Journal of Monetary Economics*, 135:99–115.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–124.

Lütkepohl, H. and Netšunajev, A. (2015). Structural Vector Autoregressions with Heteroskedasticity - A Comparison of Different Volatility Models. *CESifo Working Paper*, 5308.

Lütkepohl, H. and Netšunajev, A. (2017). Structural vector autoregressions with smooth transition in variances. *Journal of Economic Dynamics and Control*, 84:43–57.

Macaulay, A. (2021). The Attention Trap: Rational Inattention, Inequality, and Fiscal Policy. *European Economic Review*, 135:103716.

Macaulay, A. (2025). Cyclical Attention to Saving. *American Economic Journal: Macroeconomics*, 17(4):124–167.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational Inattention: A Review. *Journal of Economic Literature*, 61(1):226–273.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.

Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.

Maćkowiak, B. and Wiederholt, M. (2024). Rational Inattention during an RCT. *American Economic Review: Insights*, DP18875.

Magnusson, L. M. and Mavroeidis, S. (2014). Identification Using Stability Restrictions. *Econometrica*, 82(5):1799–1851.

Mankiw, N. G. and Reis, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Matějka, F. (2016). Rationally Inattentive Seller: Sales and Discrete Pricing. *The Review of Economic Studies*, 83(3):1125–1155.

Melosi, L. (2017). Signalling Effects of Monetary Policy. *Review of Economic Studies*, 84(2):853–884.

Miescu, M. and Rossi, R. (2021). COVID-19-induced shocks and uncertainty. *European Economic Review*, 139:103893.

Mikosch, H., Roth, C., Sarferaz, S., and Wohlfart, J. (2024). Uncertainty and Information Acquisition: Evidence from Firms and Households. *American Economic Journal: Macroeconomics*, 16(2):375–405.

Mondria, J. and Quintana-Domeque, C. (2013). Financial Contagion and Attention Allocation. *The Economic Journal*, 123(568):429–454.

Nakamura, E. and Steinsson, J. (2018). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *The Quarterly Journal of Economics*, 133(3).

Normandin, M. and Phaneuf, L. (2004). Monetary policy shocks:: Testing identification conditions under time-varying conditional volatility. *Journal of Monetary Economics*, 51(6):1217–1243.

Paciello, L. and Wiederholt, M. (2014). Exogenous information, endogenous information, and optimal monetary policy. *Review of Economic Studies*, 81(1):356–388.

Pasten, E. and Schoenle, R. (2016). Rational inattention, multi-product firms and the neutrality of money. *Journal of Monetary Economics*, 80:1–16.

Rigobon, R. (2003). Identification Through Heteroskedasticity. *The Review of Economics and Statistics*, 85(4):777–792.

Rigobon, R. and Sack, B. (2003). Measuring The Reaction of Monetary Policy to the Stock Market. *The Quarterly Journal of Economics*, 118(2):639–669.

Rigobon, R. and Sack, B. (2004). The impact of monetary policy on asset prices. *Journal of Monetary Economics*, 51(8):1553–1575.

Robin, J.-M. and Smith, R. J. (2000). TESTS OF RANK. *Econometric Theory*, 16(2):151–175.

Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, 94(4):1055–1084.

Roth, C. and Wohlfart, J. (2020). How do expectations about the Macroeconomy affect personal expectations and Behavior? *Review of Economics and Statistics*, 102(4):731–748.

Schlaak, T., Rieth, M., and Podstawska, M. (2023). Monetary policy, external instruments, and heteroskedasticity. *Quantitative Economics*, 14(1):161–200.

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

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

Song, W. and Stern, S. (2024). Firm Inattention and the Efficacy of Monetary Policy: A Text-Based Approach. *Review of Economic Studies*.

Stevens, L. (2019). Coarse Pricing Policies. *The Review of Economic Studies*, 124(2):420–453.

Swanson, E. T. (2011). Let's Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2. *Brookings Papers on Economic Activity*.

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118:32–53.

Tutino, A. (2013). Rationally inattentive consumption choices. *Review of Economic Dynamics*, 16(3):421–439.

Vayanos, D. and Vila, J.-L. (2021). A Preferred-Habitat Model of the Term Structure of Interest Rates. *Econometrica*, 89(1):77–112.

Wang, X. (2025). Time-Varying Market Sensitivity to Macro News and Monetary Policy: The Role of Attention. *Working Paper*.

Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B., Kumar, S., Gorodnichenko, Y., Georgarakos, D., Coibion, O., Kenny, G., and Ponce, J. (2025). Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings. *Econometrica*, 93(1):229–264.

Wright, J. H. (2012). What does Monetary Policy do to Long-term Interest Rates at the Zero Lower Bound? *Economic Journal*, 122(564):F447–F466.

A Proofs

A.1 Proof of Proposition 2

Let

$$v_H \equiv \boldsymbol{\alpha}_1 + \tau_H^* \boldsymbol{\beta}_1, \quad v_L \equiv \boldsymbol{\alpha}_1 + \tau_L^* \boldsymbol{\beta}_1,$$

and define

$$\Delta\Omega = \sigma_H^2 v_H v_H' - \sigma_L^2 v_L v_L'.$$

Step 1: Reduction to the span of $\{v_H, v_L\}$.

For any vector x such that $x'v_H = x'v_L = 0$, we have

$$\Delta\Omega x = \sigma_H^2 v_H(v'_H x) - \sigma_L^2 v_L(v'_L x) = 0.$$

Hence, any eigenvector of $\Delta\Omega$ associated with a nonzero eigenvalue must lie in the two-dimensional subspace $\text{span}\{v_H, v_L\}$. It follows that any such eigenvector can be written as

$$x = v_H + kv_L$$

for some scalar $k \in \mathbb{R}$.

Step 2: Characterization of eigenvectors.

Let $x = v_H + kv_L$. Then

$$\Delta\Omega x = \sigma_H^2 v_H(v'_H x) - \sigma_L^2 v_L(v'_L x) = \sigma_H^2 v_H(C_H + kD) - \sigma_L^2 v_L(D + kC_L),$$

where

$$C_s \equiv v'_s v_s, \quad D \equiv v'_H v_L.$$

If x is an eigenvector with eigenvalue μ , then $\Delta\Omega x = \mu x$, which implies

$$\begin{aligned} \sigma_H^2(C_H + kD) &= \mu, \\ -\sigma_L^2(D + kC_L) &= \mu k. \end{aligned}$$

Eliminating μ yields

$$-\sigma_L^2(D + kC_L) = k\sigma_H^2(C_H + kD),$$

which rearranges to the quadratic equation (29). If $D \neq 0$, this quadratic has two real roots, each corresponding to an eigenvector direction of $\Delta\Omega$ in $\text{span}\{v_H, v_L\}$.

Step 3: Identification of the “naïve” estimator.

The standard “naïve” spectral procedure sets \hat{b}_1 proportional to the eigenvector associated with the largest (positive) eigenvalue of $\Delta\Omega$. Hence,

$$\hat{b}_1 \propto v_H + kv_L,$$

where k is the root of the above quadratic corresponding to the largest eigenvalue.

Step 4: Generic distinctness from regime-specific impacts.

If v_H and v_L are collinear, then $\Delta\Omega$ is rank one and all eigenvectors associated with nonzero eigenvalues are proportional to both v_H and v_L . If $v_H \perp v_L$, then $D = 0$ and the unique eigenvector associated with the positive eigenvalue is proportional to v_H .

Outside of these knife-edge cases, the eigenvector $v_H + kv_L$ is not proportional to either v_H or v_L . Therefore, the uncorrected estimator \hat{b}_1 generically differs from both regime-specific impact vectors, implying that the resulting impulse responses correspond to a pseudo-parameter rather than the true structural impact in either regime.

A.2 Proof of Corollary 2

By regime invariance of $A(L)$, the MA matrices Ψ_h do not depend on s . For each regime s , the IRF is $IRF_s(h) = \Psi_h v_s$. The uncorrected estimator satisfies $\hat{b}_1 \propto v_H + kv_L$, hence

$$\widehat{IRF}(h) \propto \Psi_h \hat{b}_1 = \Psi_h(v_H + kv_L) = IRF_H(h) + kIRF_L(h),$$

which is the first claim. The normalized expression follows by dividing both sides by the impact normalization $e'_m(IRF_H(0) + kIRF_L(0))$. The final statement follows because (outside the knife-edge cases) the vectors $IRF_H(h)$ and $IRF_L(h)$ are not collinear for generic h , so their linear combination is not proportional to either one.

B Bias correction simulations

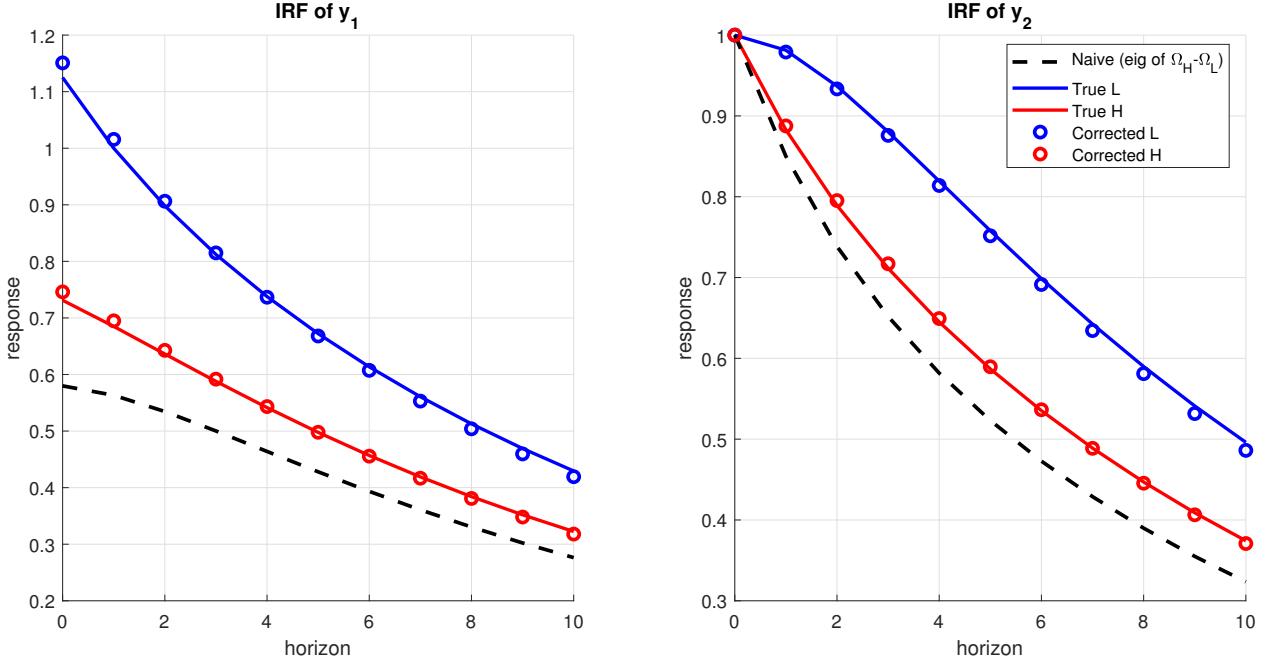
Correction Method 1. Figure 2 displays the results from a VAR(1) in two variables, with two variance regimes and true parameters given by:

$$\mathbf{A}_1 = \begin{pmatrix} 0.8, 0.1 \\ 0.25, 0.7 \end{pmatrix}, \alpha_1 = \begin{pmatrix} 1 \\ 0.5 \end{pmatrix}, \beta_1 = \begin{pmatrix} -0.7 \\ 0 \end{pmatrix}, \mathbf{b}_2 = \begin{pmatrix} 0.2 \\ 0.7 \end{pmatrix}, (\sigma_H^2, \sigma_L^2) = (4, 1), \theta = 0.75. \quad (47)$$

The solid red and blue lines show the true impulse responses to the shock of interest in the high and low-variance regimes respectively. The solid black lines show the impulse responses estimated on a 2,500 observation simulated sample if the researcher does not correct for the bias derived in Proposition 2, and if the high-variance regime occurs with probability 0.15. The red and blue circles show the impulse responses estimated on the same simulated data, using the bias correction method

in Section 4.1, with Assumptions 4 and 5 chosen correctly, and r measured exactly. Since y_2 is the variable invariant to attention on impact, I normalize the shock so that the impact is 1 on this variable in all cases. The correction is able to very closely match the true underlying responses in each regime, while the uncorrected estimates are far from either true response. The length of simulation sample is chosen to match that in the application of Section 5.3.

Figure 2: Impulse response functions in simulated data, using correction method 1.



Note: Impulse responses to structural shock 1 in simulated data from a VAR(1) with two variables and two variance regimes. True parameters are given in equation (47). The solid red and blue lines show the true impulse responses in the high-variance and low-variance regimes, respectively. The dashed black line shows the estimate obtained by applying standard heteroskedasticity-based identification to a simulated sample of 2,500 observations, with the high-variance regime occurring with probability 0.15. Red and blue circles show bias-corrected estimates using the method in Section 4.1, with correct specification of Assumptions 4 and 5 and exact measurement of the inattention ratio r . Responses are normalized such that the impact on y_2 equals one.

Figure 2 shows the results in the ideal situation, in which Assumptions 4 and 5 hold exactly, and the attention proxy (and so the calibration of r) is perfect. Figures 3 and 4 explore the sensitivity of the method to deviations in the accuracy of Assumption 4 and the calibration of r .

First, consider the sensitivity to Assumption 4. I maintain the parameters listed in (47), except I vary the second element of β_1 between -0.35 and 0.35. At each of 21 equally-spaced points in this range, I draw 1000 random samples of 2,500 observations each. I then compute standard and corrected IRFs, with the correction assuming that the second element of β_1 is 0, as in Figure 2. I find the maximum absolute error between each IRF and the true IRF at which they are aiming. Figure 3a then plots the median of these maximum-absolute-errors across simulations. The median maximum-absolute-error for the corrected IRFs is minimized when $\beta_1(2) = 0$, but it is still substantially below

the equivalent errors from the uncorrected estimation for a wide range around that. Indeed, for small deviations from the strict attention-invariance in Assumption 4, the errors in the corrected IRFs do not increase substantially: if the impact effect of attention on y_2 is 30% as strong as on y_1 , then the median maximum-absolute-error increases by 8% (high-attention regime) and 14% (low attention regime) respectively.

Another way to show this is to plot the improvement in estimates resulting from the correction. Figure 4a plots the results from the same simulations as Figure 3a, but instead of the median maximum-absolute-error, it shows the percentage improvement in that error over the uncorrected case. When Assumption 4 is exactly true, the correction removes more than three quarters of the error from the uncorrected estimates (more than 88% in the low-variance regime where uncorrected errors are larger). However, even when y_2 is not exactly attention-invariant, the improvements remain substantial.

Figures 3b and 4b repeat the exercise, but for robustness to mismeasurement of r . In this case the true r is varied between 1.5 and 3.5. Again, at each grid point 1000 simulated samples are drawn, and in each sample the standard and corrected IRFs are computed, with the correction (wrongly) assuming $r = 2$ as in (47). Figure 3b plots the median maximum-absolute-error across simulations, and Figure 4b plots the percentage improvement in that error over the uncorrected estimates. As with Assumption 4, small deviations from the assumptions used in the correction make the errors larger, but not substantially, until the true r goes a long way from the measured value. This is especially an issue if the true r gets too small, as then the dynamics become quite similar in the two regimes.

Correction Method 2. Figure 5 displays the results from a VAR(1) in three variables, with three variance regimes and true parameters given by:

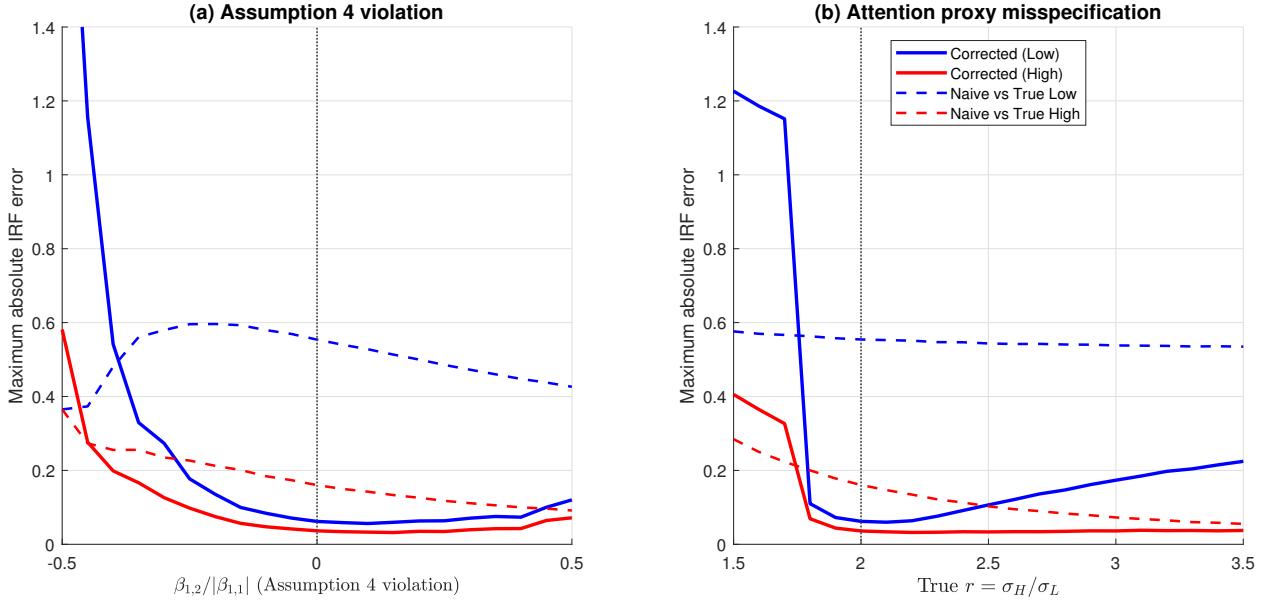
$$\mathbf{A}_1 = \begin{pmatrix} 0.6, 0.3, 0.05 \\ 0.05, 0.5, 0.1 \\ 0, 0.15, 0.55 \end{pmatrix}, \alpha_1 = \begin{pmatrix} 0.4 \\ -0.1 \\ 0.2 \end{pmatrix}, \beta_1 = \begin{pmatrix} 0.6 \\ 0.5 \\ -0.3 \end{pmatrix}, (\mathbf{b}_2, \mathbf{b}_3) = \begin{pmatrix} 0.2, -0.1 \\ 0.3, 0.15 \\ -0.1, 0.25 \end{pmatrix},$$

$$(\sigma_H^2, \sigma_M^2, \sigma_L^2) = (8, 1.2, 0.6), \theta = 1, (p_H, p_M, p_L) = (0.7, 0.2, 0.1). \quad (48)$$

where p_s is the probability of regime s occurring. As with the simulation above, the regimes are drawn i.i.d., and the variances of the other structural shocks not being studied are normalized to 1. The IRFs plotted below normalize the structural shock so that it delivers a unit increase in the first variable in the VAR in all cases.

The solid colored lines show the true impulse responses to the shock of interest in each regime.

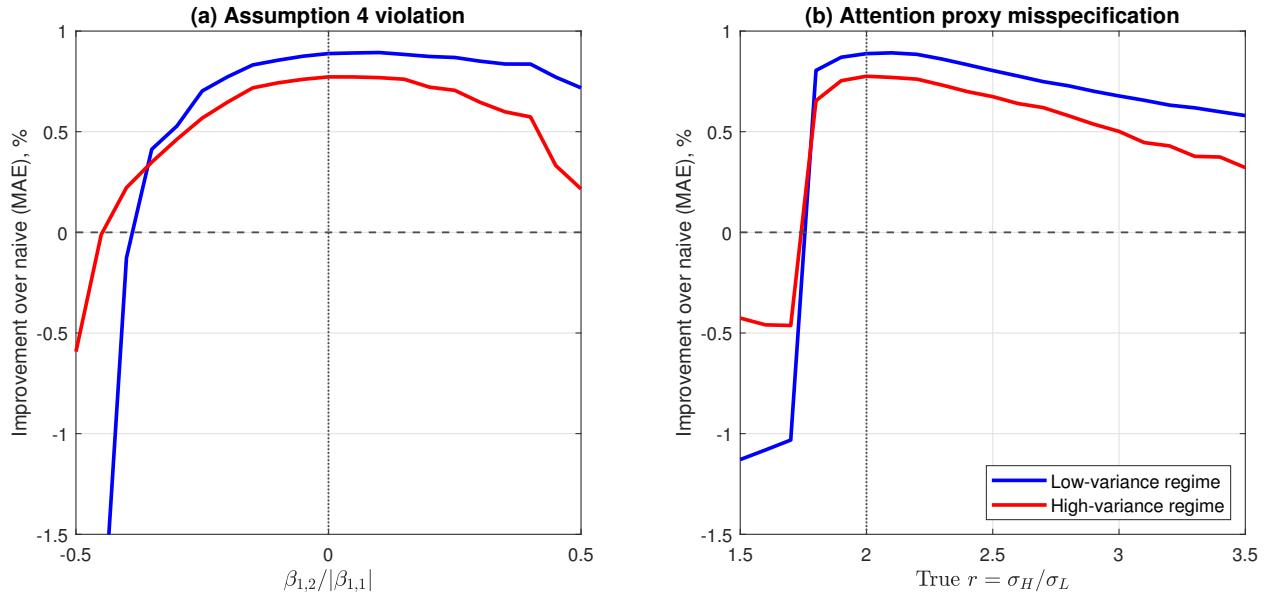
Figure 3: Sensitivity of correction method 1 to Assumption 4 and r measurement: median maximum-absolute-error across simulations.



Note: Sensitivity of correction method 1 to violations of Assumption 4 (panel a) and misspecification of the inattention ratio r (panel b). In panel (a), simulated samples are drawn from the system with parameters described in (47), except for $\beta_1(2)$, which is varied between -0.35 and 0.35. In panel (b), $\beta_1(2)$ is fixed at 0, but the true r is varied between 1.5 and 3.5. At each set of parameters, 1000 random samples are drawn, and both standard and corrected impulse responses are computed. The maximum absolute error between each impulse response and the corresponding true impulse response is computed for each sample over horizons 0-10. The plots show the medians of these maximum-absolute-errors across the simulations. Solid lines show errors for the bias-corrected estimates; dashed lines show errors for uncorrected estimates that ignore endogenous attention.

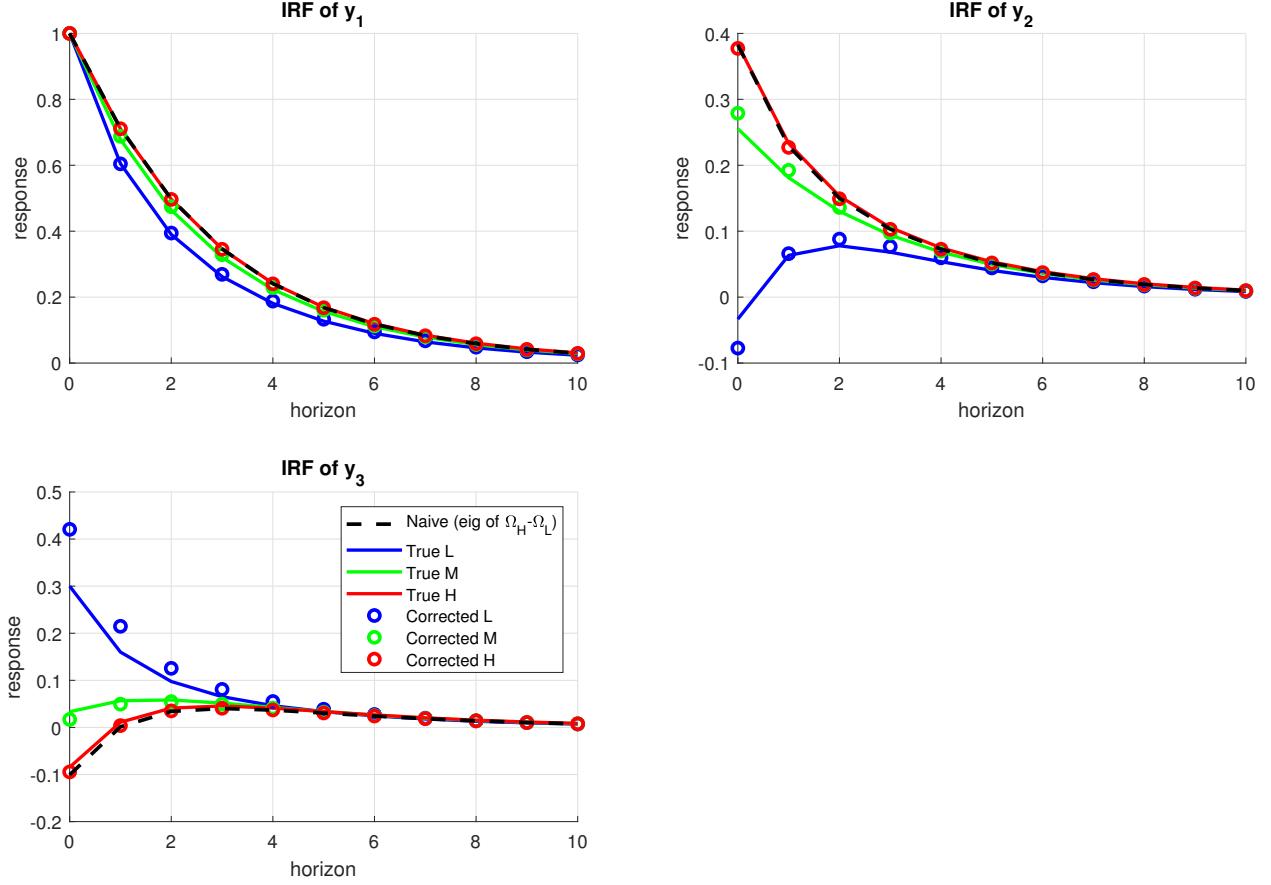
The dashed black lines show the impulse responses estimated on a 2,500 observation simulated sample if the researcher does not correct for the bias derived in Proposition 2. The colored circles show the impulse responses estimated using the bias correction method in Section 4.2. While the results are not as close to the true impulse responses as those from method 1 (Figure 2), especially for the lowest attention regime, this method requires fewer assumptions on the role of attention in each variable's transmission, and still delivers impulse responses close to the true responses in each regime.

Figure 4: Sensitivity of correction method 1 to Assumption 4 and r measurement: percentage improvement in median maximum-absolute-error across simulations.



Note: Sensitivity of correction method 1 to violations of Assumption 4 (panel a) and misspecification of the inattention ratio r (panel b). In each panel, simulations are drawn for varying parameters as in Figure 3 (see associated note for details). The median maximum-absolute-error across simulations is computed for standard and bias-corrected impulse responses as in Figure 3. The lines plotted here show the percentage improvement in this median maximum-absolute-error in the bias-corrected impulse responses, relative to the uncorrected impulse responses. Specifically, they show $(\text{medmax}(\text{corrected}) - \text{medmax}(\text{uncorrected})) / \text{medmax}(\text{uncorrected})$. Positive numbers reflect a reduction in median maximum-absolute-errors due to the bias correction, negative numbers reflect an increase.

Figure 5: Impulse response functions in simulated data, using correction method 2.



Note: Impulse responses to structural shock 1 in simulated data from a VAR(1) with three variables and three variance regimes. True parameters are given in equation (48). Solid red, green, and blue lines show the true impulse responses in the high-, medium-, and low-variance regimes, respectively. The dashed black line shows the estimate obtained by applying standard heteroskedasticity-based identification to a simulated sample of 2,500 observations, with regime probabilities $(p_H, p_M, p_L) = (0.7, 0.2, 0.1)$. Colored circles show bias-corrected estimates using the method in Section 4.2. Responses are normalized such that the impact on y_1 equals one.

C Section 5 robustness

Table 4 repeats Table 2, using an alternative Google Trends index reflecting searches for the term “monetary policy.” Although quantitative magnitudes vary, the key result that search volume is higher on announcement days is robust.

Table 4: Attention to monetary policy: FOMC vs. Non-FOMC Days

Sample	Observations		Mean Search Index			t-statistic
	FOMC	Non-FOMC	FOMC	Non-FOMC	Difference	
Wright (2008–2011)	23	1,039	53.461	38.223	15.238	3.056
ZLB (2008–2015)	56	2,537	50.342	37.251	13.090	4.290
Full (2006–2019)	113	4,816	49.525	35.421	14.104	6.444

Notes: Google Trends search index for “Monetary Policy” constructed using the Eichenauer et al. (2022) methodology. FOMC days are scheduled Federal Open Market Committee announcement days. The *t*-statistic tests the null hypothesis that mean attention is equal across FOMC and non-FOMC days.

Figures 6 and 7 show the bias-corrected IRFs, as in Figure 1, for alternative mappings between Google Trends data and τ_L^*, τ_H^* . In Figure 6, I maintain the linear mapping used for Figure 1, but calculate average Google Trends volume in each regime after trimming the sample to exclude the largest and smallest 1% of observations, which yields $r = 2.04$. In Figure 7, I assume a linear mapping between $\tau_{i,t}$ and *log* Google Trends volume, which yields $r = \frac{\log(100) - \log(50.1)}{\log(100) - \log(86.3)} = 4.69$. In both cases, the core qualitative results of Figure 1 remain: correcting for the rational-inattention bias substantially alters the estimated effects of monetary policy shocks, and in particular implies that the low-variance regime features substantially less persistent transmission to Treasury yields, and an initial sharp rotation in breakevens.

Finally, Figure 8 shows the bias-corrected IRFs from an alternative procedure in which ϕ^* is chosen to minimize the gap between impact effects of attention on the two Treasury yields, as an alternative to the strict invariance of the 2-year yield in Assumption 4.

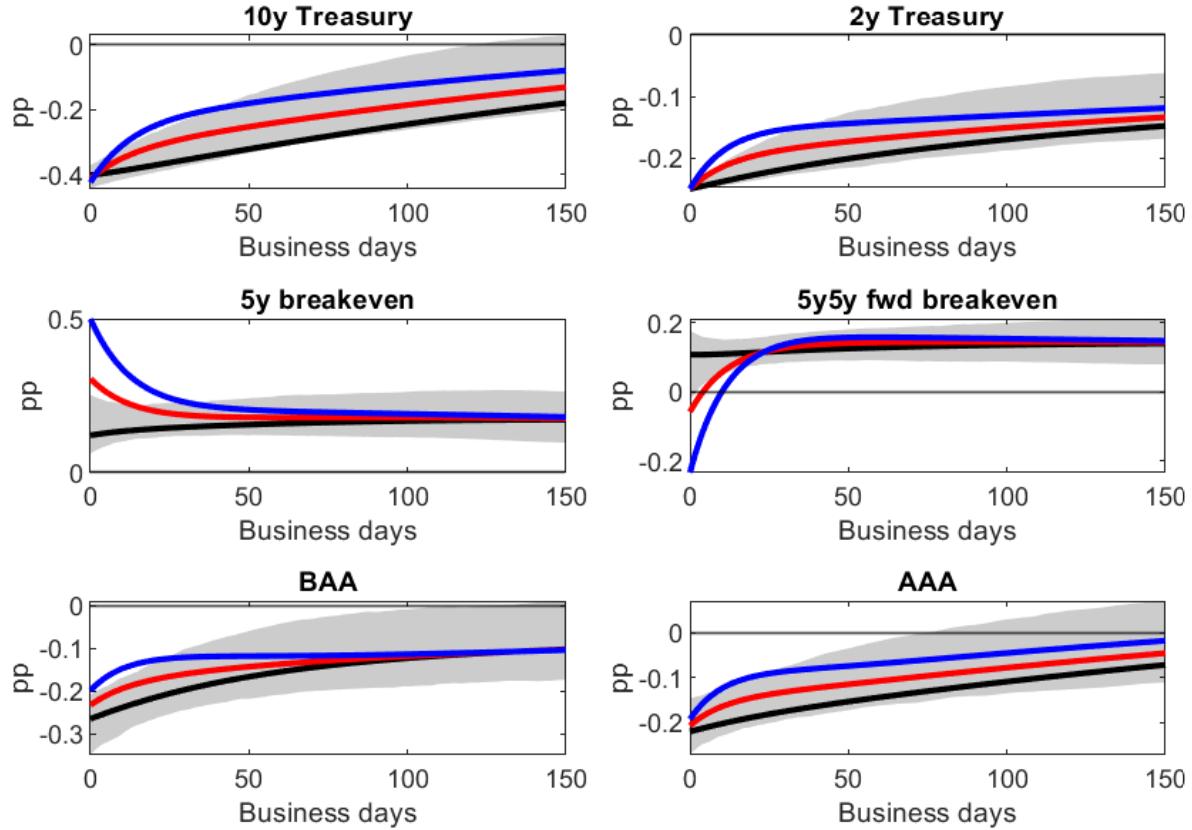
Specifically, let $v_H(\phi)$ and $v_L(\phi)$ denote the regime-dependent impact vectors for a candidate rotation ϕ . I select the rotation ϕ^* that minimizes the transmission gap across regimes for both the 2-year and 10-year Treasury yields. That is, I replace equation (41) with:

$$\phi^* = \arg \min_{\phi} \sum_{k \in \{2y, 10y\}} \left(\frac{v_{H,k}(\phi)}{v_{L,k}(\phi)} - 1 \right)^2 \quad (49)$$

Using squared fractional differences ensures that the identification is not disproportionately driven by the maturity with the largest absolute response. This therefore seeks the structural model

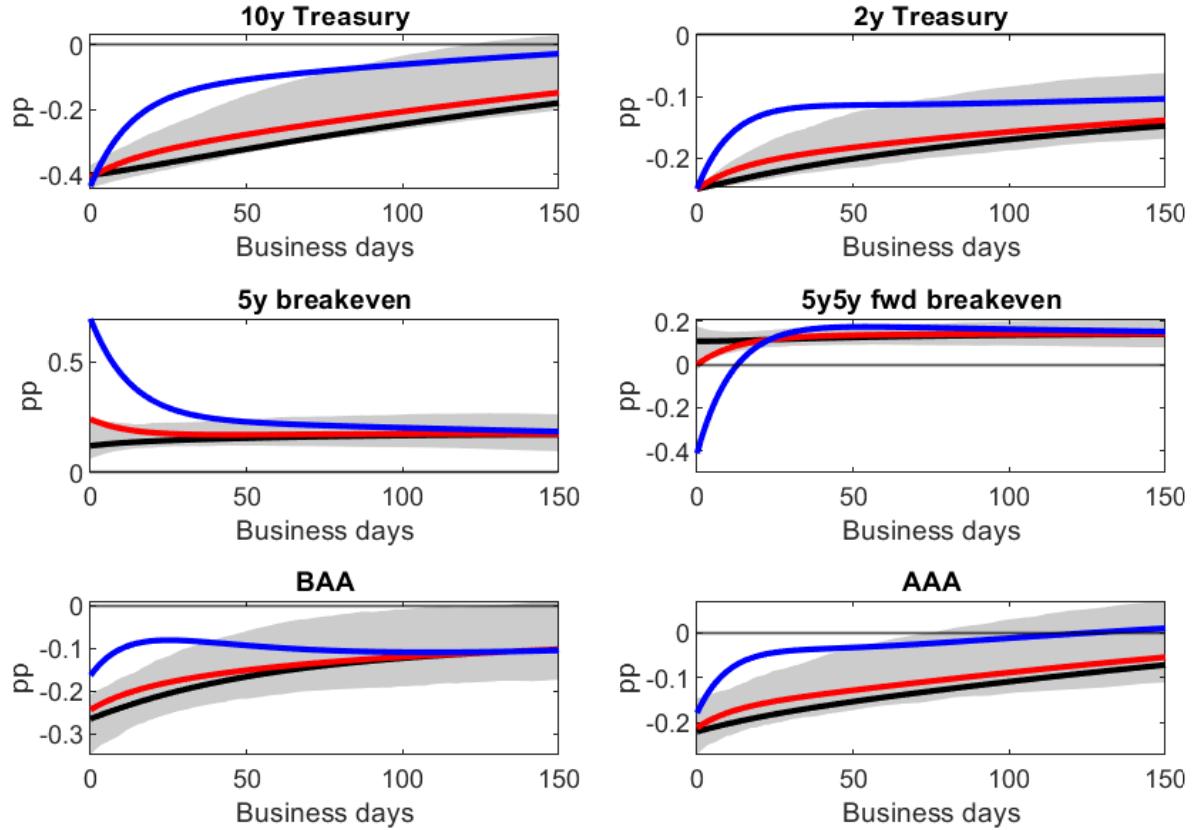
most consistent with the hypothesis that the Treasury curve is driven on impact by equally sized attention effects across maturities, without forcing any single maturity to be perfectly invariant. To make the extent of departure from Assumption 4 clear, I normalize the corrected IRFs for both regimes by the same scaling factor, chosen such that the impact is a 25bp fall in the 2-year Treasury yield on average across the regimes, rather than normalizing each impulse response individually. The results are qualitatively identical to Figure 1, and also quantitatively extremely similar.

Figure 6: Impulse response functions with and without the correction for endogenous attention, with trimmed Google Trends sample as the attention measure.



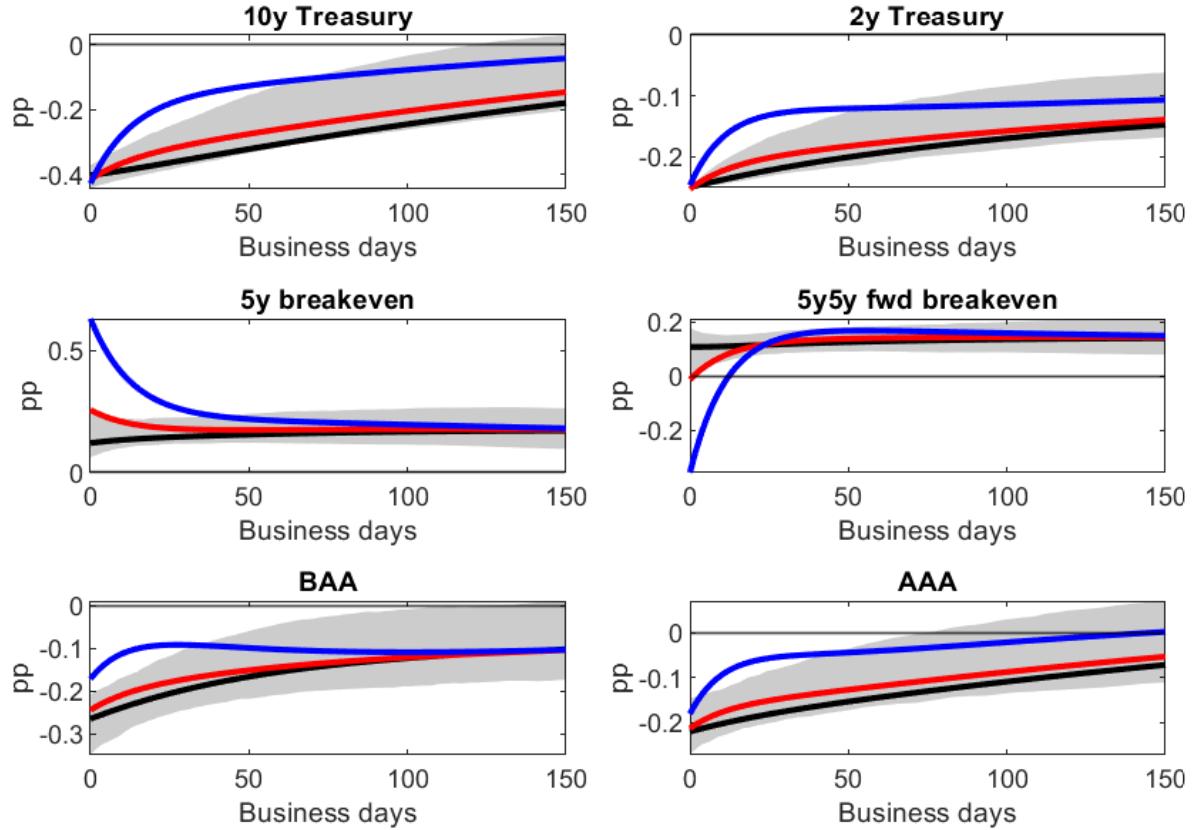
Note: Impulse responses to an unconventional monetary policy shock identified via heteroskedasticity, following Wright (2012). The black line shows the estimate assuming regime-invariant transmission (Assumption 1), with the shaded region indicating 68% bootstrap confidence intervals, constructed using the Kilian (1998)-adjusted bootstrap. The Kilian bias adjustment is also applied to the point estimates. The red line shows the bias-corrected impulse response for the high-attention regime (FOMC announcement days), and the blue line shows the bias-corrected response for the low-attention regime (non-announcement days). Bias correction follows the method in Section 4.1, imposing that the impact effect on the 2-year Treasury yield is attention-invariant (Assumption 4) and that attention amplifies the impact effect on BAA corporate bond yields (Assumption 5). The inattention ratio r is measured using the Google Trends data described in Section 5.2, after trimming the top and bottom 1% of the observations. The responses are normalized such that the 2-year Treasury yield falls by 25 basis points on impact in all cases. Sample: November 2008 to December 2015, excluding December 1, 2008 (see text). The VAR is estimated with one lag on daily data.

Figure 7: Impulse response functions with and without the correction for endogenous attention, with log Google Trends sample as the attention measure.



Note: Impulse responses to an unconventional monetary policy shock identified via heteroskedasticity, following Wright (2012). The black line shows the estimate assuming regime-invariant transmission (Assumption 1), with the shaded region indicating 68% bootstrap confidence intervals, constructed using the Kilian (1998)-adjusted bootstrap. The Kilian bias adjustment is also applied to the point estimates. The red line shows the bias-corrected impulse response for the high-attention regime (FOMC announcement days), and the blue line shows the bias-corrected response for the low-attention regime (non-announcement days). Bias correction follows the method in Section 4.1, imposing that the impact effect on the 2-year Treasury yield is attention-invariant (Assumption 4) and that attention amplifies the impact effect on BAA corporate bond yields (Assumption 5). The inattention ratio r is measured using the log of the Google Trends series described in Section 5.2. The responses are normalized such that the 2-year Treasury yield falls by 25 basis points on impact in all cases. Sample: November 2008 to December 2015, excluding December 1, 2008 (see text). The VAR is estimated with one lag on daily data.

Figure 8: Impulse response functions with and without the correction for endogenous attention, with minimum distance criterion for selecting ϕ^* .



Note: Impulse responses to an unconventional monetary policy shock identified via heteroskedasticity, following Wright (2012). The black line shows the estimate assuming regime-invariant transmission (Assumption 1), with the shaded region indicating 68% bootstrap confidence intervals, constructed using the Kilian (1998)-adjusted bootstrap. The Kilian bias adjustment is also applied to the point estimates. The red line shows the bias-corrected impulse response for the high-attention regime (FOMC announcement days), and the blue line shows the bias-corrected response for the low-attention regime (non-announcement days). Bias correction follows the method for Figure 1, except that Assumption 4 is replaced with equation (49). The responses are normalized such that the 2-year Treasury yield falls by 25 basis points on impact in the uncorrected case, and on average across the two corrected cases. Sample: November 2008 to December 2015, excluding December 1, 2008 (see text). The VAR is estimated with one lag on daily data.