Oil Supply Shocks and Monetary Policy

——*Work in Progress!*——[*]

Luis Calderon[†]

University of Bonn

June 5, 2026

Abstract

I revisit the role of monetary policy in the transmission of oil supply shocks within a single Bayesian VAR that jointly identifies a Känzig (2021) oil supply news shock and six orthogonal monetary policy shocks. I combine the causal mediation framework of Dufour and Wang (2024), to recover what the Federal Reserve's reaction function loads on, with the sufficient-statistics counterfactual of Caravello, McKay, and Wolf (2024), to evaluate *three* alternative rules: a nominal rate peg, strict CPI-inflation targeting, and an optimal dual-mandate rule. Four findings emerge. (i) Aggregate responses are largely insensitive to the rule at short constraint horizons, with the Fed reacting primarily to inflation, the oil price, and demand-side mediators, and negligibly to the oil-market block. (ii) Policy duration is what discriminates the rules at the macro level: held for only a few quarters the three rules look near-identical, but enforced over the full twenty-quarter horizon the rate peg and dual-mandate converge to one another and deliver materially better output, inflation, and unemployment stability than the realized path. (iii) Distributional responses discriminate even at short horizons: the dual-mandate rule tracks the realized Gini paths, while the rate peg and strict inflation targeting both compress income, wealth, and consumption inequality, with strict-π amplifying the realized compression by roughly twenty per cent for income, two-and-a-half-fold for consumption, and threefold for wealth. (iv) The strict-π compression is driven by redistribution toward the bottom of the distribution rather than top-tail destruction. The cross-section, and the duration over which a rule is held, thus discriminate between rules that headline macro aggregates at short horizons cannot.

Keywords: *Oil shocks; Monetary policy counterfactuals; Sufficient statistics; Income, consumption, and wealth inequality; Bayesian VAR; External instruments*

JEL Classification: *C32, C38, D31, E31, E52, E58, Q43*

---

[*] I thank colleagues at the Institute for Macroeconomics and Econometrics for early-stage discussions. All errors are my own.

[†] Institute for Macroeconomics and Econometrics—University of Bonn, luis.calderon@uni-bonn.de.

1 Introduction

Oil supply shocks tend to be stagflationary: they increase the price of goods and services and depress real activity simultaneously. This was precisely the situation the central bank faced in the 1970s and history seems to be repeating itself. This presents a trade-off for the central bank: raise rates to deter realized or expected inflation, but at the risk of causing a recession or *look through* with the potential risk of inflation spiraling. Previous work has suggested the central bank's role is not innocuous for the transmission of these shocks and may even be responsible for the large macroeconomic outcomes historically observed (Bernanke, Gertler, and Watson, 1997). This view has been challenged in subsequent empirical and theoretical work (see e.g., Kilian and Lewis, 2011, for a review). This paper re-examines the presented evidence, under the same class of VAR models employed, adopting recent insights as to the specification of these models (Kilian and Zhou, 2023; Baumeister, 2025), with naturally more years of data, and exploits new econometric frameworks that provide new insights on the role of the central bank for the transmission of these shocks.

Four questions structure the paper. First, what share of the macroeconomic response to an identified oil supply shock is driven by the systematic component of monetary policy rather than by the shock's direct transmission? Second, what does the central bank reaction function mostly load on along the oil-shock horizon—the real price of oil itself, the inflation pass-through it generates, real activity, or some combination? Third, following the suggestion of Kilian and Lewis (2011), how does the Fed reaction function look like had it not responded to the oil shock specifically, while still responding normally to inflation and real activity? Fourth, against three counterfactual rules—a nominal interest-rate peg, strict inflation targeting, and an optimal dual-mandate rule that minimizes a quadratic loss in inflation and the unemployment gap—could the Fed have improved aggregate or distributional outcomes, and through which channel?

To answer the first three questions, I apply the generalized impulse response (GIR) framework of Dufour and Renault (1998) and the causal mediation interpretation of Dufour and Wang (2024). Together, these papers provide a structured econometric framework to quantify how different channels (mediators) contribute to the dynamic causal effects of an exogenous shock, offering a more nuanced view of the transmission mechanism. In my setting, the framework allows me to quantify the contribution of the ongoing stance of monetary policy to both the aggregate and distributional effects of oil supply shocks by decomposing impulse response functions (IRFs) at each horizon into the channels represented in the VAR.

My focus is the contribution of the Federal Funds Rate (FFR) to the level of the IRFs and whether it amplifies or dampens them. The FFR IRF itself captures the central bank's reaction to the shock. The same framework also allows me to decompose the FFR IRF and identify the variables driving the policy response, thereby recovering the central bank's reaction function. Under policy invariance in the Lucas critique sense, this further allows me to infer the central bank response absent the oil shock, as emphasized by Kilian and Lewis (2011).

The understanding I can obtain thus far is that of the typical Fed response and this is important because it answers the question of the average historical role played by the central bank in response to the history of oil supply shocks thus far. If the characterization of such oil supply shocks is unchanged, as well as the set of structural shocks and monetary regime in the economy, I can expect precisely this typical response; however, the central bank may at times want to deviate, either temporarily or permanently, from this typical response and it is important to understand whether such deviations lead to better economic stability or further contraction. For this, I use the sufficient-statistics counterfactual framework of Caravello, McKay, and Wolf (2024) (CMW) to ask how the aggregate and distributional consequences of an oil supply shock would have changed had the central bank operated under a different monetary policy rule.

Regarding the technicalities of the procedure, the construction of the counterfactuals requires at least as many auxiliary monetary policy innovations as there are independent constraints T in the rule. A nominal interest-rate peg, for example, imposes one constraint per horizon; the constraint being holding the deviation of the federal funds rate to zero from its baseline path. Strict inflation targeting analogously imposes T inflation constraints. The optimal dual-mandate rule is softer (a quadratic loss rather than a hard zero) but still pins down a T-dimensional optimum. In each case the rule violation is a T-vector that the counterfactual must project onto the span of the available monetary shocks. With a single monetary shock the span is one-dimensional and the projection collapses to least-squares onto a single instrument.

Increasing the number of shocks K expands the span, reduces the projection residual, and—provided the shocks are not collinear—makes the projection well-conditioned in the sense of Wang (2026, Sect. 3). The empirical payoff is that one does not have to litigate which monetary shock is "best": by including a battery of identified series, the framework leverages their joint information and is robust to mis-specification of any single proxy. The trade-off is that adding shocks too aggressively introduces multicollinearity and thus singularity issues. As shown below, the shock space in our setup will be free of these concerns.

The number of constraints T is a choice and the convention of choosing T equal to H, the IRF horizon, is perhaps too stringent and unrealistic. Policymakers think about the implementation of policy in the near future; therefore, this paper adopts the view of Wang (2026) and enforces rules over shorter horizons ($T \in \{4, 6, 8\}$ quarters) that are closer to following paths than rules. With the headline $K = 6$ specification the projection is admissible for $T \geq K = 6$; for the shorter $T = 4$ horizon I fall back to a $K = 4$ baseline that retains the standard MP-shock set (Luetticke, 2024; Bauer and Swanson, 2023; Hack, Istrefi, and Meier, 2024) and is exactly identified at $T = K = 4$. The dual-mandate rule is over-determined at every T in the sweep but nonetheless delivers a better fit than $T = 20$. Overdetermined systems are not automatically invalid; the diagnostics in Appendix C indicate the degree to which the Lucas critique bites. After T periods the rule no longer binds and the policy rate adjusts endogenously through the VAR coefficients, using the same systematic reaction function that produced the original IRF.

To enforce these rules in the short run I borrow *six* distinct monetary policy shocks from the literature for the headline $K = 6$ specification: **(i)** the Luetticke (2024) unified narrative series, which chains the conventional and unconventional halves of the post-1969 shock history into a single continuous measure; **(ii)–(v)** the four orthogonal high-frequency factors of Jarociński (2024), identified by sign and zero restrictions on FOMC-window surprises in FFR futures, forward-guidance futures, long-rate futures, and the S&P 500—the short-rate, forward guidance, large-scale-asset-purchase, and residual factors, respectively; and **(vi)** the Hack, Istrefi, and Meier (2024)-orthogonalized Romer and Romer (2004) narrative shock, which retains pre-2008 identification power that the Jarociński HF factors lack. A $K = 4$ baseline consisting of the Lütke unified series, Bauer–Swanson high-frequency surprises, and the Hack, Istrefi, and Meier (2024)-orthogonalized Romer and Romer (2004) and Aruoba and Drechsel (2024) narratives is retained for the $T_c = 4$ constraint horizon, where $K = 6$ would render the peg and strict-π projections under-determined ($T_c < K$). Each shock captures a different slice of the monetary-policy partition, and the joint span at $K = 6$ delivers the lowest pairwise correlations achievable from the union of identification schemes available in the literature (see Section 3.1).

Both methodologies will rely on a VAR, as did the conclusions of Bernanke, Gertler, and Watson (1997) and Kilian and Lewis (2011), incorporating on top new insights as to the specification of these oil VARs (Baumeister and Hamilton, 2019; Kilian and Zhou, 2023) and more generally, structural VARs (Baumeister, 2025). The usual suspects remain, namely, the modeling of global demand and supply of oil, the relevant oil price, and the representation of domestic activity in the U.S.. What will be different here is (1) the choice of inventories variable, which will be more robust to measurement error concerns (Baumeister and Hamilton, 2019) (2) long-lags (De Graeve and Westermark, 2025) to (i) achieve a better representation of the covariances along the oil cycle (ii) adopt the mispecification robustness of LPs (Ludwig, 2024) (iii) without increased uncertainty around the estimates (De Graeve and Westermark, 2025) (3) a distributional block (Bayer, Calderon, and Kuhn, 2025) to further isolate the supply component of the shock, but of course to understand the broad distributional implications and (4) all under a Bayesian approach (Chan, 2022) to model richer environments without violating the regularity conditions of frequentist VARs. In a companion paper, I realize and employ these insights, and additionally show they are important for the implications of oil price shocks.

Because this paper is extensively on counterfactuals, it is important to discuss the extent to which the Lucas critique (Lucas, 1976) applies. If agents form expectations using the prevailing policy rule, changing the *rule* should change the reduced-form dynamics of the economy, and IRFs estimated under one rule cannot be used to evaluate another. The recent sufficient-statistics literature, however, has converged on a common defense. Wang (2026) states it most explicitly. The construction holds the private-sector decision rules and the baseline policy rule fixed and alters only the sequence of policy shocks within a finite window; in his words, "the policy peg is not a new rule, it is a finite string of temporary deviations from the baseline rule" (Wang, 2026, p. 7). The counterfactual is credible when the implied shock sequence is short, moderate in magnitude, and balanced in sign; it is strained otherwise because it would be informative about a regime shift and may trigger belief revision in the sense of Lucas (1976).

Wang (2026) calls these conditions of short, moderate, and balanced deviations as the *policy invariance* restriction, and argues that his approach is appropriate. For our baseline counterfactuals, I will follow Caravello, McKay, and Wolf (2024), who also make the same restriction operational, but at the population level: the alternative rule is admissible if it lies in the linear span of the auxiliary monetary innovations identified in the VAR (Caravello, McKay, and Wolf, 2024). McKay and Wolf (2023b) show that, under linear-Gaussian dynamics, the resulting counterfactual is identified up to a remainder term that vanishes in the limit of small rule perturbations. In my setup, the rate peg, the strict-IT rule, and the dual-mandate rule are enforced through linear combinations of K monetary innovations whose pairwise correlations are low (Table 1); the implied shock sequences are unprecedented in neither magnitude nor sign-persistence, and they pass diagnostic tests proposed by Wang (2026). For now, I implement the counterfactual procedure of Caravello, McKay, and Wolf (2024). Down the line, I would like to implement Wang (2026) for comparison.

Findings. My application embeds this construction in a Bayesian VAR that jointly identifies a Känzig (2021) oil supply news shock and $K = 6$ monetary innovations (the headline specification; the $K = 4$ baseline is retained for the $T_c = 4$ robustness, see Section 3.1). The endogenous block contains eight macroeconomic series (the real oil price, world oil production, OECD oil inventories, world industrial production, real GDP, unemployment, CPI inflation, and the federal funds rate) and eight orthogonal factors that summarise the joint income–wealth–consumption distribution. I evaluate *three* counterfactual policy rules that span the policy space from two extremes to an interior optimum:

(i) a *nominal rate peg* that pins the federal funds rate deviation to zero over the constraint window—the most accommodative limit, accepting unbounded inflation pass-through to hold the rate fixed;

(ii) *strict inflation targeting* that pins CPI inflation deviation to zero over the constraint window—the most aggressive limit, accepting unbounded rate movements to hold inflation at target; and

(iii) an *optimal dual-mandate* rule that minimizes a textbook quadratic loss in inflation and the unemployment gap (equal weights, $\lambda = 1$), sitting between the two extremes and mapping directly onto the standard central-bank loss function.

The first two rules bound the policy space from two sides; the third nests them in a single quadratic loss and lets the Fed trade off inflation against employment rather than committing to one pole.

The macro responses to the oil supply shock under all three counterfactual rules differ only modestly from the realized path *at short constraint horizons*. The federal funds rate, inflation, GDP, and unemployment under the rate peg, strict inflation targeting, and the dual-mandate rule all sit broadly within the 68% credible band of the realized response when the rule is enforced for $T_c \in \{4, 6, 8\}$ quarters, differing mostly in second-order features of the medium-run path. The Dufour and Wang (2024) mediator decomposition of the realized FFR IRF identifies what the Fed is reacting to: a considerable share of the FFR response is mediated by inflation, the oil price, and demand-side variables (GDP and the distribution), with negligible feedback from the upstream oil-market block (oil production, inventories, and world IP).

What discriminates the rules at the macro level is policy duration. The rate peg would do little if held for only four quarters, but when enforced over the full twenty-quarter horizon it keeps the federal funds rate close to its steady state and delivers materially better output, inflation, and unemployment stability than the realized path. The dual-mandate rule, held for the same long window, converges toward the same accommodative profile—in the limit it even lowers rates on impact in exchange for smaller cuts in the medium run—so at $T_c = H = 20$ the two rules essentially coincide and dominate the realized response on macro stability. Strict inflation targeting remains the outlier: its medium-run tightening required to neutralise oil-shock pass-through is quantitatively similar to the realized path, so it offers little macro improvement over the realized response at any constraint horizon. The headline macro aggregates therefore separate the rules only when the policymaker is willing to commit to a long enforcement window, an observation made explicit in Appendix A.

The distributional responses are more discriminating across the three rules *even at short constraint horizons*. The dual-mandate rule delivers Gini paths nearly indistinguishable from the realized response on all three margins. The rate peg and strict inflation targeting, in contrast, both *compress* inequality further than the realized path: strict-π amplifies the realized medium-run Gini compression on every margin (by roughly twenty per cent for income, two-and-a-half-fold for consumption, and threefold for wealth), and the rate peg compresses the income and consumption Ginis but tracks the realized response on wealth. The mechanism under strict-π is a large redistribution toward the bottom of the distribution rather than top-tail destruction; the top-decile/bottom-half decomposition (Section 5.3.2) makes the mechanism explicit. Two margins thus discriminate between the rules: the duration over which the rule is held, on the macro side; and the cross-section, even at short horizons, on the distributional side. A central bank evaluated on macro aggregates and short enforcement windows alone would not separate the three rules cleanly; one that evaluates on the cross-section, or that is willing to commit to a long enforcement window, can—and the ranking is unambiguous.

Contribution. The paper makes three contributions. First, it provides the first model-free counterfactual answer to the question Kilian and Lewis (2011) pose at the end of their paper. Where they criticize the BGW counterfactual and call for "models that allow policy responses to depend on the underlying causes of oil price shocks," the subsequent literature (Leduc and Sill, 2004; Carlstrom and Fuerst, 2006; Nakov and Pescatori, 2010; Bodenstein, Guerrieri, and Kilian, 2012; Plante, 2014; Gagliardone and Gertler, 2025) has answered the call only by committing to a structural DSGE model. My exercise computes oil-shock counterfactuals using only sufficient statistics for the monetary policy span, in the spirit of Caravello, McKay, and Wolf (2024), and so trades the structural interpretability of the DSGE exercises for a much weaker reliance on modeling assumptions. To my knowledge, while Caravello, McKay, and Wolf (2024) apply the sufficient-statistics framework empirically to monetary policy shocks, and Kilian and Lewis (2011) apply a coefficient-restriction counterfactual empirically to oil shocks, this paper is the first to combine them—model-free CMW counterfactuals on identified oil shocks under multiple alternative rules—and to do so jointly with a household distributional dimension. Second, the multi-shock identification gives the alternative rules meaningful enforcement power, in contrast to the single-MP-shock VARs of the older counterfactual literature (Bernanke, Gertler, and Watson, 1997; Kilian and Lewis, 2011). Third, I show that the CMW linear-superposition formula extends, draw-by-draw, to the underlying distributional factor IRFs and hence to any non-linear function of the joint distribution; the substantive payoff is a quantitative estimate of the distributional cost of the three rules. The third contribution generalizes beyond the specific application: any household-level outcome that admits a finite-dimensional functional representation can be projected through.

Outline. The remainder of the paper is organized as follows. Section 2 situates the paper against the long oil–monetary-policy debate initiated by Bernanke, Gertler, and Watson (1997) and Kilian and Lewis (2011), the recent sufficient-statistics counterfactual literature (McKay and Wolf, 2023a; Caravello, McKay, and Wolf, 2024; Wang, 2026), and the literature on monetary policy and inequality. Section 3 describes the data: the K monetary policy shocks (Section 3.1) and the macroeconomic controls (Section 3.2); the oil supply shock and distributional block are documented in the companion paper. Section 4 sets out the empirical framework in three parts: the multi-shock BVAR with the asymmetric conjugate prior of Chan (2022) (Section 4.1); the CMW counterfactual algebra (Section 4.2); and the three counterfactual rules in closed form (Section 4.3). Section 5 presents the empirical findings in three subsections: the macro counterfactuals (Section 5.1), the channel decomposition of the realized IRFs and the per-rule decomposition under each counterfactual (Section 5.2), and the distributional counterfactuals (Section 5.3). Section 6 concludes.

2 Related literature

The paper overall is relatively underdeveloped with considerable literature to cover. With that, the paper devotes a large part to the literature review to make clear the contributions one could make i.e., what is still missing in my understanding of the oil-monetary policy debate.

Kilian–Lewis (2011). An important empirical precedent is Kilian and Lewis (2011), henceforth KL, who revisit the influential claim of Bernanke, Gertler, and Watson (1997) (BGW) that the Federal Reserve's systematic tightening in response to oil-price shocks was responsible for *all* the negative impact of the oil shock. The claim is based on a proposed interest-rate peg counterfactual in which the central bank does not change the policy rate in response to the oil price shock. KL, together with Hamilton and Herrera (2004), show that the result is fragile across lag lengths, sample windows, and identification of the underlying oil shock. Hamilton and Herrera (2004) additionally show that the BGW peg required policy deviations of unprecedented magnitude, raising issues with the Lucas-critique problem that Bernanke, Gertler, and Watson (2004) also acknowledge. Overall, KL finds (i) no empirical evidence that systematic monetary policy is an important source of fluctuations in the US economy (ii) that in general, oil price shocks have negligible effects on the macroeconomy and (iii) argues that it would be a mistake to respond to oil price shocks rather than its underlying determinants e.g., the oil price responding to supply or demand factors.

DSGE literature. The KL prescription is to design models that take into account the endogeneity of the real price of oil and that allow policy responses to depend on these underlying causes of oil price shocks. This has been almost exclusively answered with DSGE models. Leduc and Sill (2004) attribute roughly 40% of the post-shock output drop to the policy response. Carlstrom and Fuerst (2006) attribute essentially none. Nakov and Pescatori (2010) and Plante (2014) try different rules and rank them by welfare; the latter finds that strict

CPI-inflation targeting performs poorly relative to a core-inflation or output-gap rule once oil-shocks are decomposed by source. The most direct DSGE response to KL is Bodenstein, Guerrieri, and Kilian (2012), who derive shock-specific optimal responses and show that the Fed should ease against oil-driven aggregate-demand shocks but tighten only modestly against oil-supply shocks.

Most recently, Gagliardone and Gertler (2025) estimate a New Keynesian DSGE model with oil as a complementary good for households and a complementary input for firms, search-and-matching unemployment, and real wage rigidity, and match impulse responses to an SVAR identified with the Känzig (2021) oil supply news shock and the Bauer and Swanson (2023) monetary policy shock. Their historical decomposition over 2010–2023 attributes the post-2021 inflation surge mainly to oil shocks combined with accommodative monetary policy. Bernanke and Blanchard (2023)'s reduced-form decomposition of pandemic-era inflation finds that energy and supply-side factors account for the bulk of the headline CPI movement once expectations are accounted for. Their structural counterfactual under a strict inflation-stabilization rule produces sharply larger output and unemployment losses than the estimated baseline, while a near-accommodating rule moderates the real losses at the cost of higher inflation.

Caravello, McKay, and Wolf (2024). The methodological foundation for the model-free exercise performed herein is Caravello, McKay, and Wolf (2024) (CMW), with theoretical underpinnings in McKay and Wolf (2023b) and Barnichon and Mesters (2023). The CMW insight is that the counterfactual response of any endogenous variable to a structural shock under a new policy rule is a linear combination of its responses to that shock and to a set of auxiliary monetary policy innovations identified in the same VAR. As long as the rule is a linear constraint on outcomes and the auxiliary block spans a sufficiently rich policy-shock space, the counterfactual is identified without re-estimating the model and without the Lucas-critique problem that sank the BGW exercise: the rule is enforced through historically observed policy variation rather than through unprecedented shifts.

In practice, what this will mean is agents in period zero, in addition to the oil supply news shock, are hit by a single bundle of K auxiliary monetary policy shocks calibrated to enforce the rule; the economy then evolves under its historical reduced-form dynamics, with the systematic Fed reaction function continuing to operate. The appealing feature of this design is that it does not require the economy to be continuously surprised: agents experience one set of policy innovations at the moment of impact, not a fresh sequence of unanticipated deviations every quarter. The latter is the approach formalized by Sims and Zha (2006), which was actually applied informally in Bernanke and Mishkin (1997).

Under linear-Gaussian dynamics, McKay and Wolf (2023b) prove that the one-shot counterfactual (only date 0 shocks) is identified up to a Lucas-critique residual—the adjustment in agents' decision rules induced by the policy change—that is second-order in the size of the policy perturbation and therefore vanishes for small auxiliary impulses. Repeatedly surprising agents, as in the Sims–Zha sequential design, inflates the cumulative implied shocks and pushes this residual outside the small-perturbation regime where it remains bounded.

Dufour and Wang (2024). The aggregate part of my analysis—decomposing the realized IRFs of an oil supply shock into a direct effect and an indirect effect through the federal funds rate—has a precedent in the same KL paper. KL argue that "the relevant counterfactual" for the BGW question "is not one in which the Federal Reserve holds the interest rate constant in response to an oil price shock [. . .] but a counterfactual in which the Federal Reserve reacts to fluctuations in other macroeconomic state variables (such as inflation and real output) as it normally would with only the direct response to the real price of oil being shut down." The construction is informal: KL reorder the VAR so the FFR depends on the oil price last, set the contemporaneous oil-price coefficient in the FFR equation to zero, and recompute IRFs. Dufour and Wang (2024) formalize the same idea using the short-run/long-run causality machinery of Dufour and Renault (1998): any structural shock admits a decomposition of its IRF into a direct effect on outcome Y that bypasses any specified set of mediator variables $\mathcal{M}$, and an indirect effect that travels through $\mathcal{M}$. KL's exercise is the special case in which $\mathcal{M}$ is the federal funds rate and the shock is oil-supply; my generalization adds two ingredients absent from KL: $\mathcal{M}$ can be any subset of the VAR's mediators, and the framework extends to the multi-shock VAR in which K monetary innovations enter alongside the oil disturbance. I adopt this formalism in Section 5.2, where I also discuss the conditions under which the mediation decomposition admits a counterfactual reading.

Most related literature. Presented here are papers that are empirical in nature and are the most recent empirical antecedents to my paper. Wang (2026) develops a local-projection (LP) counterpart to Sims and Zha (2006) *but also* Caravello, McKay, and Wolf (2024) and applies it, in his first empirical exercise, under counterfactual short-rate pegs of length three, twelve, and twenty-four months, using the Känzig (2021) supply news surprise and the Bauer and Swanson (2023) high-frequency monetary instrument. Broer, Kramer, and Mitman (2025) go further on the distributional side: using German administrative microdata, they ask how the systematic response of monetary policy to the oil supply news shock of Känzig (2021) shapes the earnings and employment distribution, conducting a policy rate non-response counterfactual under both Caravello, McKay, and Wolf (2024) and Sims and Zha (2006).

Closest in methodological spirit to my exercise is Ider et al. (2025), who apply the McKay and Wolf (2023b) sufficient-statistics counterfactual to the European Central Bank and energy prices. They document that ECB monetary policy decisions significantly influence global energy prices and show, under the same Lucas-critique-robust framework I adopt, that this ability roughly halves the tightening required to stabilize euro-area inflation after an energy supply shock. Their optimal-policy analysis under a stylised dual-mandate loss rationalises the ECB's observed "look-through" response to energy shocks as quasi-optimal precisely because the ECB *can* influence energy prices. Their setting is the ECB and they do not consider distributional outcomes; the $K = 6$ counterfactual rules and joint income–consumption–wealth distribution reported in Section 5 are complementary to their aggregate ECB-side evidence.

I differ from Wang (2026) and Broer, Kramer, and Mitman (2025) on three margins. First, the policy span. My auxiliary monetary block stacks K shocks—the Luetticke (2024) unified CMP/UMP narrative, Bauer and Swanson (2023) high-frequency surprises, and the Hack, Istrefi, and Meier (2024)-orthogonalized Aruoba–Drechsel and Romer–Romer narratives—giving the CMW projection more degrees of freedom to enforce the rule. Several papers, including Broer, Kramer, and Mitman (2025), build their span out of one or two monetary instruments, which leads to trouble in enforcing their supposed counterfactuals—a limitation Wang (2026) himself flags when discussing the cost of one-shot designs.

Second, the distributional object: Wang (2026) does not discuss distributional consequences and Broer, Kramer, and Mitman (2025) work with moments from marginal earnings/employment distribution. I use the joint income–consumption–wealth *functional* representation of Bayer, Calderon, and Kuhn (2025), which lets me evaluate counterfactuals for any moment of the joint distribution (Gini coefficients, top 10% shares, bottom 50% shares). Third, the rule set: I evaluate two boundary rules (a nominal rate peg and a strict CPI inflation target) and an optimal interior rule (a dual-mandate loss-minimizing response), rather than just a non-response benchmark.

Functional data in macroeconomics. The distributional data come from Bayer, Calderon, and Kuhn (2025), which fuses microdata at the household level with macroeconomic indicators of higher frequency in a Bayesian state-space framework. Their construction follows Sklar (1959): any joint distribution over $\mathbb{R}^d$ decomposes into its marginal CDFs and a copula encoding the dependence structure, and both objects can be projected onto an orthonormal Legendre basis to obtain a finite-dimensional state vector. The eight leading principal components underlying the functional data captures all the business-cycle variation in the joint distribution and enter my VAR as the distributional block. This functional-VAR strategy follows Chang, Chen, and Schorfheide (2024) and Chang and Schorfheide (2024), Lenza and Savoia (2024), and Bjørnland, Chang, and Cross (2023), but to my knowledge this paper is the first to combine it with a CMW counterfactual exercise in oil context.

Monetary policy and inequality. An established literature documents that monetary contractions raise income and consumption inequality (Kaplan, Moll, and Violante, 2018; Coibion et al., 2017; Auclert, 2019; Bayer, Born, and Luetticke, 2020; Andersen et al., 2021); the mirror-image finding for monetary expansions is mixed (Lenza and Slacalek, 2018). This paper speaks to the same question but with two differences. First, the impulse is an identified oil supply shock rather than a monetary shock, so the relevant object is whether the distributional incidence of the oil shock depends on the *systematic* component of the Fed's response. Second, the rule is varied counterfactually via the sufficient-statistics machinery of Section 4.3, holding the shock fixed. I find that the realized oil shock compresses the income, wealth, and consumption Ginis in the medium run; the dual-mandate rule delivers Gini paths nearly indistinguishable from the realized response; and the rate peg and strict inflation targeting both compress inequality further than the realized path, with strict-π the larger amplifier on all three margins. The top-decile/bottom-half decomposition (Section 5.3.2) attributes the strict-π compression to redistribution toward the bottom of the distribution rather than top-tail destruction.

3 Data

This section presents data unique to this paper and refers the reader to the companion paper on the description of the distributional data I use for the estimation and the oil supply shocks as well. This section describes the monetary policy shocks used to discipline the counterfactuals using the methodology of Caravello, McKay, and Wolf (2024) and describe the macroeconomic aggregates, which have slightly changed from the companion paper to accommodate the rules specified in my counterfactuals i.e., any variable in some proposed rule must be within the variable space of the VAR.

3.1 Monetary policy shocks

The auxiliary monetary block enters the VAR (Section 4.1) as a set of externally identified shocks. I report two specifications: a $K = 6$ headline (used throughout Section 5 and the appendix robustness exercises) and a $K = 4$ baseline retained for when the number of constraints, T_c, is equal to 4.[1]

Headline ($K = 6$). The $K = 6$ specification stacks the following six monetary-policy shocks, chosen to span the 1969Q1–2023Q4 sample with as little overlap in identifying variation as possible:

---

[1] $K = 6$ would render the peg and strict-π projections under-determined ($T_c < K$).

(i) **Luetticke (2024) unified series**: a unified narrative monetary policy shock, a single continuous series spanning both the conventional and unconventional (ZLB) regimes. Constructed from using the Romer-Romer econometric setup on shadow rates from Wu and Xia (2015). Available 1969Q1–2015Q4. I refer to them later as Lütke for short.

(ii)–(v) **Jarociński (2024) orthogonal factors**: the four factors u_1 (short rate), u_2 (forward guidance), u_3 (large-scale asset purchases), and u_4 (residual long-rate factor) from Jarociński (2024). The factors are identified by sign and zero restrictions on 30-minute FOMC-window surprises in FFR futures, forward-guidance futures, long-rate futures, and the S&P 500. They are orthogonal to each other by construction; the maximum pairwise correlation among the four is 0.12 (Table 1). Available 1990Q1–2024Q3.

(vi) **Romer–Romer (HIM-cleaned)**: the Romer and Romer (2004) narrative shock with the systematic component of monetary policy removed following Hack, Istrefi, and Meier (2024). Retains pre-2008 identification power for the 1970s and Volcker-era oil-shock episodes that the HF Jarociński factors do not. Available 1969Q1–2007Q4.

As mentioned earlier, for the $T_c = 4$ specification, we can only use systems with $K \leq 4$—otherwise, they are under-determined ($T_c = 4 < K = 6$). The $K = 4$ specification replaces the four Jarociński factors with the more standard pair (Bauer and Swanson, 2023; Hack, Istrefi, and Meier, 2024; Aruoba and Drechsel, 2024):

(ii) **Bauer–Swanson high-frequency surprises**: Bauer and Swanson (2023) FOMC announcement-window surprises in the four-quarter-ahead Eurodollar future, orthogonalized against the Greenbook information set. Available 1988Q1–2023Q4.

(iii) **Aruoba–Drechsel (HIM-cleaned)**: the Aruoba and Drechsel (2024) narrative shock with the systematic component of monetary policy removed following Hack, Istrefi, and Meier (2024). Available 1982Q4–2007Q4.

together with the same Lütke unified series (i) and HIM-cleaned Romer–Romer series listed above.[2]

Why these shocks. The choice of monetary proxies is the central identification decision of the paper, and Table 1 makes the reasoning concrete. The table reports pairwise Pearson correlations across ten candidate shocks; the diagonal is shaded grey (each shock's correlation with itself is unity by construction) and off-diagonal cells with $|\rho| > 0.5$ are highlighted yellow, flagging pairs that move toward collinearity.

---

[2] Because the monetary proxies have non-overlapping or only partially overlapping coverage windows (e.g. HIM-AD ends 2007Q4; HIM-RR ends 2007Q4; Jarociński factors begin 1990Q1), I follow Caravello, McKay, and Wolf (2024) and impute missing values with zero outside each shock's native coverage window before estimation. The same zero-padding convention is used for the K=4 baseline, K=6 headline, and all robustness specifications.

The $K = 6$ headline uses the four orthogonal HF factors u_1, u_2, u_3, u_4 of Jarociński (2024) (rows five through eight of the table). The four factors are orthogonal to each other by construction; the 4×4 block among $u_1, \ldots, u_4$ is essentially diagonal (max $|\rho| = 0.12$). They are moderately correlated with Bauer–Swanson (e.g. BS–$u_1 = 0.51$, BS–$u_2 = 0.57$) which is expected—both factor models extract structured directions from the same FOMC-window HF surprise space—but Bauer–Swanson is no longer in the $K = 6$ set, so these cross-correlations do not appear among the included shocks.

Picking out the $K = 6$ entries from Table 1, the maximum pairwise correlation among the six is 0.63, occurring on the Lütke–HIM-RR pair; the other nine pairs in the $K = 6$ block are at $|\rho| \leq 0.36$, none of them yellow. In the end, the multiple shocks setup give the rule meaningful enforcement power: with $K = 1$ the rule is satisfied only in least-squares sense across $T = 20$ horizons, and the resulting counterfactual may barely deviate from the baseline at any horizon other than the impact response.

Table 1: Pairwise correlations of monetary policy shocks (overlapping observations only).

	BS	Lütke	HIM-R	HIM-A	u_1	u_2	u_3	u_4	JK-MP	JK-CBI
Bauer–Swanson	1.00	0.16	0.06	0.12	0.51	0.57	0.07	-0.10	0.69	0.30
Lütke unified	0.16	1.00	0.63	0.54	0.27	0.22	-0.11	-0.08	0.27	0.14
HIM-Romer	0.06	0.63	1.00	0.59	0.15	0.11	-0.36	0.17	0.05	0.27
HIM-Aruoba	0.12	0.54	0.59	1.00	0.19	0.01	-0.15	-0.11	0.14	0.05
u_1 short rate	0.51	0.27	0.15	0.19	1.00	-0.03	0.12	-0.11	0.72	0.53
u_2 forward guidance	0.57	0.22	0.11	0.01	-0.03	1.00	-0.02	-0.08	0.51	0.01
u_3 LSAP	0.07	-0.11	-0.36	-0.15	0.12	-0.02	1.00	0.07	0.31	-0.05
u_4 residual factor	-0.10	-0.08	0.17	-0.11	-0.11	-0.08	0.07	1.00	-0.41	0.64
JK pure MP	0.69	0.27	0.05	0.14	0.72	0.51	0.31	-0.41	1.00	0.23
JK CB-information	0.30	0.14	0.27	0.05	0.53	0.01	-0.05	0.64	0.23	1.00

Notes. Pearson correlations over the pairwise overlap of observation windows. JK-MP and JK-CBI are not mentioned in the text—they are median monetary-policy and central-bank-information shocks of Jarociński and Karadi (2020), identified by sign-restrictions on the co-movement of rates and equity prices around FOMC announcements (1990–2024, n=145).

3.2 Macro variables

The macro block follows a standard oil-VAR specification (Kilian, 2009; Baumeister and Hamilton, 2019; Känzig, 2021) augmented with the U.S. real side (GDP, unemployment), inflation, and the policy rate. The eight series enter the VAR in the following order: the real WTI crude oil price (deflated by CPI-U), world crude oil production, the OECD-proxy crude oil inventories series of Baumeister–Hamilton, the OECD-plus-six global industrial production index, U.S. real GDP, the U.S. civilian unemployment rate, CPI inflation (quarterly log-difference of CPI-U), and the effective federal funds rate. The five log-level series (real oil price, oil production, oil inventories, world industrial production, U.S. real GDP) enter in log levels and are scaled by 100 for interpretability so that percent-deviation IRFs read off the y-axis directly;

unemployment, inflation, and the federal funds rate enter in levels (percent).

Normalisation. Before turning to the counterfactual, *all* impulse responses are rescaled per posterior draw so that the impact response of the real oil price to the oil shock is $+10\%$. The CMW projection in Eq. (7) below then operates on these scaled IRFs, so the auxiliary m_{aux} is interpreted as the bundle of monetary innovations the Fed would have to deliver to enforce the rule *given* a 10pp oil-price impact.

4 Methodology

4.1 The multi-shock BVAR

Let $y_t \in \mathbb{R}^{n_y}$ stack the variables in the following order: the identified oil supply shock first, the K monetary policy shocks second, the eight macro variables third, and the eight distributional factors last. Placing the oil shock at position one of y_t is the natural recursive ordering: the oil shock is the impulse whose transmission the paper studies, and within-quarter macro and distributional outcomes are assumed not to feed back into the externally measured shock series. The monetary policy shocks occupy positions $2, \ldots, K+1$ so that the CMW auxiliary policy span (Section 4.2) reads off the first K columns of the structural impulse-response matrix Θ_h immediately to the right of the oil-shock column. I estimate the reduced-form BVAR with the same machinery as in the companion paper

$$y_t = c + B_1 y_{t-1} + \cdots + B_p y_{t-p} + u_t, \qquad u_t \sim \mathcal{N}(0, \Sigma), \tag{1}$$

with $p = 16$ lags, an asymmetric Minnesota-style prior with per-equation hyperparameters (Chan, 2022), priors for the long-run (Giannone, Lenza, and Primiceri, 2019), and $D = 5{,}000$ posterior samples that satisfy stability requirements of the VAR after burn-in. Again, the asymmetric structure matters because the five-shock block has very different prior information than the macro and distributional blocks: shock equations are tightly centered on white noise (diagonal $\delta_i = 0$), while macro and distributional equations may adhere to the standard Minnesota persistence prior.

The five shocks are already structurally identified upstream of the VAR: the oil supply shock by Känzig (2021) (with Baumeister and Hamilton (2019)'s sudden oil-production shortfalls as a robustness identification), and the K monetary policy shocks by their respective external identification strategies. Each enters y_t as an externally measured proxy in the first $K+1$ positions, and the recursive structure of B_0 in those rows is a *timing assumption*, not an identifying restriction: within-quarter macro and distributional outcomes are assumed not to feed back contemporaneously into the externally measured shock series (Stock and Watson, 2018; Plagborg-Møller and Wolf, 2021). Stacking the impact matrix B_0 such that the first $K+1$ rows are diagonal yields the structural representation

$$y_t = \sum_{h=0}^{\infty} \Theta_h \, \varepsilon_{t-h}, \tag{2}$$

where $\Theta_h \in \mathbb{R}^{n_y \times n_y}$ is the h-step structural impulse-response coefficient matrix: its (i,j) entry is the response of variable i at horizon h to a unit shock to ε^j at t, for j some structural shock. For each posterior draw I extract two slices of Θ_h that the counterfactual algebra needs:

$$\Psi_i^{\text{oil}}(h) \equiv \Theta_{h,i,1}, \tag{3}$$

$$m_{i,k}^{\text{base}}(h) \equiv \Theta_{h,i,k+1}, \quad k = 1, \dots, K, \tag{4}$$

i.e., the response of the i-th endogenous variable to the oil shock (Ψ_i^{oil}) and to the k-th monetary policy shock ($m_{i,k}^{\text{base}}$) at horizon h.

4.2 CMW counterfactual algebra

Fix a posterior draw and a length of enforcement T_c. For some variable i, write the T-vector of impulse responses to the oil shock as $\Psi_i \in \mathbb{R}^T$ and the $T \times K$ matrix of responses to the K monetary shocks as M_i. A linear policy rule that the central bank follows under the counterfactual is encoded by a collection of $T \times T$ constraint matrices $\{A_i\}_{i=1}^{n_y}$ such that the target is to satisfy

$$\sum_{i=1}^{n_y} A_i y_i^{\text{cnf}} = \mathbf{0}. \tag{5}$$

Counterfactual responses are obtained by adding to the original oil IRF a linear combination of the monetary IRFs:

$$y_i^{\text{cnf}} = \Psi_i + M_i m_{\text{aux}}, \qquad m_{\text{aux}} \in \mathbb{R}^K. \tag{6}$$

Substituting Eq. (6) into Eq. (5) and stacking yields the linear system

$$\underbrace{\Big[\sum_i A_i M_i\Big]}_{\equiv M_1\,(T\times K)} m_{\text{aux}} = \underbrace{-\sum_i A_i \Psi_i}_{\equiv M_2\,(T\times 1)},$$

whose least-squares solution is

$$\boxed{m_{\text{aux}} = -(M_1' M_1)^{-1} M_1' M_2.} \tag{7}$$

This is the closed-form CMW formula: the auxiliary monetary policy shock vector is a linear projection of the rule violation under the oil shock onto the space of monetary IRFs. Once m_{aux} is known, Eq. (6) delivers the counterfactual IRF of every variable in the VAR, including the eight distributional factors.

4.3 Counterfactual policy rules

I evaluate three counterfactual rules. A rate-peg and a strict-inflation-target rule, as well as a closed-form optimal dual-mandate rule that nests both extremes in a single quadratic loss. The first two rules impose hard linear constraints on a *single* target row: the $T \times T$ constraint blocks A_i are zero except for the target variable. With I_T the T-dimensional identity:

- **Zero-rate peg.** $A_{\text{ffr}} = I_T$, all other $A_i = 0$. The rule pins the federal funds rate deviation to zero at every horizon $h \in \{0, 1, \ldots, T-1\}$. The Fed uses the linear combination m_{aux} of monetary innovations to fully neutralise the oil-shock-induced FFR path.

- **Strict inflation targeting.** $A_{\text{infl}} = I_T$, all other $A_i = 0$. The rule pins CPI inflation deviation to zero at every horizon. The Fed uses m_{aux} to fully offset oil-shock pass-through to inflation.

For these two rules the projection is admissible iff $M_1' M_1$ has full column rank, i.e. the IRFs of FFR (resp. inflation) to the four monetary shocks span $\mathbb{R}^K$ in their first T horizons. This is the rank condition emphasized by Wang (2026, Sect. 3) as the identification requirement for any local-projection or VAR-based counterfactual: without enough *independent* monetary instruments, the projection of the rule violation onto the policy span is not unique.

For the third rule is an *optimal* interior rule that does not impose a hard constraint on any row but instead minimizes a standard central-bank quadratic loss in inflation and the unemployment gap:

- **Optimal dual-mandate.** The Fed minimizes $\mathcal{L}(m_{\text{aux}}) = \|\Psi_\pi^{\text{cnf}}\|^2 + \lambda \, \|\Psi_{u_{\text{gap}}}^{\text{cnf}}\|^2$ over m_{aux}, where Ψ_π^{cnf} and $\Psi_{u_{\text{gap}}}^{\text{cnf}}$ are the counterfactual inflation and unemployment-gap IRFs. Substituting Eq. (6) into the loss yields a quadratic in m_{aux} whose first-order condition is the stacked least-squares problem

$$m_{\text{aux}}^{\text{dm}} = -\left(M_{1,\pi}' M_{1,\pi} + \lambda \, M_{1,u}' M_{1,u}\right)^{-1} \left(M_{1,\pi}' M_{2,\pi} + \lambda \, M_{1,u}' M_{2,u}\right), \tag{8}$$

  where $M_{1,\pi}, M_{2,\pi}$ (resp. $M_{1,u}, M_{2,u}$) are the M_1 and M_2 objects of Eq. (7) built from the inflation (resp. unemployment) row of the IRF cube. I set $\lambda = 1$ (equal weights, the textbook benchmark); $\lambda \to 0$ collapses to strict inflation targeting.

The first two bound the space of monetary stances from two sides: a peg is the most accommodative rule that an inflation-targeting central bank could plausibly entertain (it accepts unbounded inflation pass-through to keep the rate fixed), and strict inflation targeting is the most aggressive (it accepts unbounded rate movements to keep inflation at the target). The third sits between them and maps directly to the textbook central-bank loss function: it lets the Fed trade off inflation stabilization against employment stabilization according to λ rather than committing to one pole. Differences between the realized response and the three rules tell me how much of the realized macro and distributional response was driven by the systematic component of monetary policy, and whether a welfare-maximizing central bank with a standard loss function could have done substantially better.

The FFR IRF decomposition is a finer-grained counterfactual. The rate-peg rule is the *full-non-response* limit: the Fed does not move the funds rate in response to anything — oil shock, oil price, inflation, real activity, or anything else — for the full T quarters (the FFR is kept at its steady state). Kilian and Lewis (2011) insist that the substantively interesting counterfactual is finer than a full peg:

> "a counterfactual in which the Federal Reserve reacts to fluctuations in other macroeconomic state variables (such as inflation and real output) as it normally would with only the direct response to the real price of oil being shut down."

The Fed equation in the VAR has loadings on lagged macro state variables, including the real oil price itself. KL's counterfactual zeroes out the FFR equation's loading on the oil price specifically, leaving every other loading (inflation, real activity, distribution, lagged FFR, etc.) intact. The Fed continues to react to inflation and output as it normally does—including the inflation and output movements that the oil shock has produced—but its direct rule-of-thumb response to the real oil price is removed.

The Dufour and Wang (2024) channel decomposition of the FFR IRF (Section 5.2) operationalizes exactly this question. The oil-price block of the decomposition isolates the portion of the FFR response that flows through the FFR equation's loading on the oil price. Removing only that block and propagating the remaining (inflation-, activity-, distribution-, own-FFR-driven) path through to any other variable Y gives Kilian and Lewis (2011)'s preferred counterfactual.

The Dufour–Wang shutdown, however, imposes the restriction via some ex-post algebra of the realized IRF; it is in some sense BGW's coefficient-zeroing exercise. The CMW peg implements non-response through a sequence of auxiliary monetary innovations that satisfy the policy-invariance restriction up to a small remainder (McKay and Wolf, 2023b; Wang, 2026). The two methods answer *different* substantive questions (the full peg vs. the shock-specific shutdown). The channel decomposition has no built-in defense against the Lucas critique, so results there are to be interpreted with care. One could do this differently and model a different authority that, as a rule, stabilizes the oil price response, as done in Ider et al. (2025).

5 Results

This section reports the empirical findings in three subsections. Section 5.1 reports the baseline IRFs along with the macro counterfactual IRFs under the three rules across constraint horizons $T_c \in \{4, 8\}$. Appendix 7 and 8 show results for $T_c \in \{6, 20\}$. Section 5.2 decomposes the realized IRFs into mediator channels: first the FFR and oil-price baseline decompositions that recover the Fed's reaction function and the oil-shock pass-through, then the full per-rule decomposition of the four macro outcomes. Section 5.3 reports the distributional counterfactual paths.

Summary of main findings. Four results stand out and motivate the subsections that follow.

(i) On Macro Sensitivity. At first glance, under the Känzig (2021) oil supply news shock and the $K = 6$ baseline, all three counterfactual rules deliver macro paths that are quantitatively close to the realized response on every variable other than the rule's own target (Figure 1). This continues the Kilian and Lewis (2011) verdict *against* Bernanke, Gertler, and Watson (1997)'s "systematic monetary tightening explains most of the macro impact" claim, now with a model-free (Caravello, McKay, and Wolf, 2024) sufficient-statistics counterfactual rather than a coefficient-zeroing exercise. I double-down on Kilian and Lewis (2011) by showing the FFR has little to no role in the propagation of aggregate responses.

(ii) Policy Duration. Policy duration will be an important nuance. At $T_c = 4$, the strict inflation targeting and the dual-mandate generate the largest deviations. Both policies will exchange slightly higher initial rates for greater output and unemployment recovery (medium-run) and less inflation overall. At $T_c = 8$, for dual-mandate, it becomes more salient. The rate-peg would do nothing if held for $T_c = 4$ periods, but convincingly achieves greater output, inflation, and unemployment stability if held for the entire horizon, $T_c = H = 20$. I show that a bank operating under a dual-mandate for 20 periods will converge to something close to a rate-peg.

(iii) Reaction-function. The Dufour and Wang (2024) mediator decomposition of the FFR IRF identifies what the Fed is reacting to: a considerable share of the FFR response is mediated by inflation and real activity, but also the oil price.

(iv) Distributional sensitivity. The headline macro aggregates—FFR, inflation, GDP, unemployment—have similar dynamics across the four rules under $T_c \in \{4, 6, 8\}$, differing mostly in the medium-run path. The three Ginis and distributional shares tell a different story. The dual-mandate

rule delivers Gini paths nearly indistinguishable from the realized response. The rate peg and strict inflation targeting both compress inequality further than the realized path, with strict-π the larger amplifier on every margin (by roughly twenty per cent for income, two-and-a-half-fold for consumption, and threefold for wealth); the mechanism is a redistribution toward the bottom of the distribution rather than top-tail destruction. A central bank evaluated on macro aggregates alone would not separate the four rules cleanly; one evaluated on the cross-section can.

5.1 Macro counterfactuals

Figures 1 and 2 report the counterfactual macro responses to the Känzig (2021) oil supply news shock under two constraint horizons: $T_c = 4$ and $T_c = 8$. For the $T_c = 4$ case, I use the $K = 4$ unified specification referenced above, which provides an exact-fit for the rate-peg and inflation-peg and leaves it free thereafter. For the $T_c = 8$ case, I use the $K = 6$ baseline, where the rule may not be enforced perfectly.[3] Each figure reports the four core macro variables (FFR, CPI inflation, real GDP, unemployment) under the realized policy (solid red) and the three counterfactual rules: rate peg (orange dash), strict π-target (green dot), dual mandate (blue dash-dot). The 68% credible band is shaded only for the realized response. The oil-price column and the additional constraint horizons $T_c \in \{6, 20\}$ are reported in Appendix A (Figures 7, 8, 9).

What the figures show. Three observations stand out across the two figures. *First, the rate peg neutralises the FFR over the constraint window.* By construction the peg rule pins the FFR-deviation to zero over the first T_c horizons. At $T_c = 4$ (Figure 1) the peg path is exactly flat for the first four quarters—the $T_c = K = 4$ system is square and the projection is exact—and drifts back toward the realized response as the systematic Fed reaction function carries on. At $T_c = 8$ (Figure 2) the peg flattens FFR over a wider window, but overall looks no different than the baseline. The rate-peg in our setting is active once it deviates from the baseline—pegging for more than a two years. With $T_c = 20$, the interest-rate is not perfectly pegged, but stays close to zero. In this setting, *it does* achieve greater output, inflation, and unemployment stability and only deviates (lowers) to fight medium-run recessionary movements.

Second, the strict-π counterfactual sits close to the realized path on every variable other than its target. Under strict-π targeting, when enforced perfectly ($T_c \in \{4, 6\}$), we can achieve better medium-run output, but no difference elsewhere. In general, these counterfactuals follow

---

[3] The $K = 6$ baseline cannot be evaluated at $T_c = 4$ for the hard-constraint rules because $T_c \geq K$ is required for M_1 to have full column rank; at $T_c = 4$ the $K = 4$ unified specification is exactly identified ($T_c = K = 4$, square system). The dual-mandate rule does compute at $T_c = 4$ under $K = 6$ (the stacked $2T_c \times K$ system is over-determined) and is reported in the appendix.

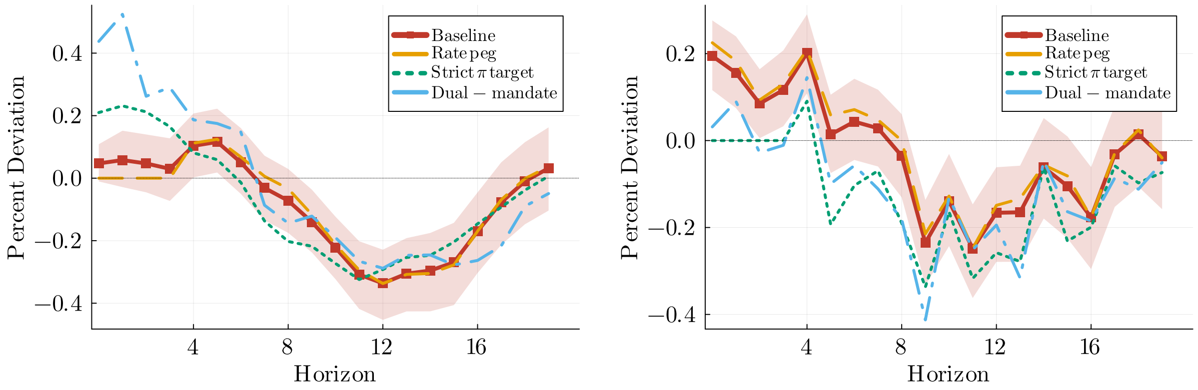

(a) Federal funds rate (b) CPI inflation

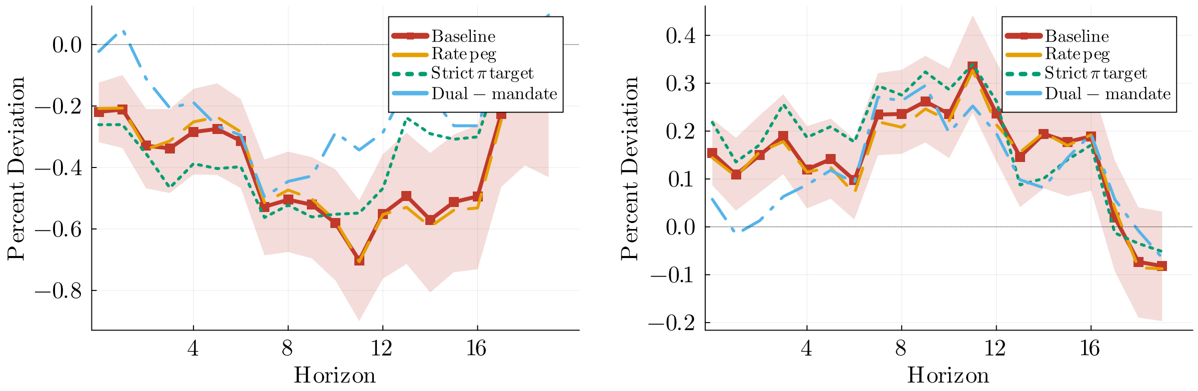

(c) Real GDP (d) Unemployment

Figure 1: Macro counterfactual responses at $T_c = 4$ ($K = 4$ exact-fit case for peg and strict-π). Solid red: baseline response; orange dash: rate peg; green dot: strict π-target; blue dash-dot: dual-mandate. Shaded band: 68% credible interval on the realized response. The peg rule pins FFR to zero exactly over the first four quarters; the strict π rule pins inflation analogously. $H = 20$.

similar FFR paths, most different in the short-run for the $T_c = 4$ case, where the rules are enforced *perfectly*, exchanging higher initial rates for greater output recovery (medium-run) and less inflation overall though not by much.

Third, the dual-mandate rule delivers the empirically largest macro deviation. It minimizes a quadratic loss in both inflation and the unemployment gap, and its m_{aux} path splits the difference between the peg and strict-π extremes. At $T_c = 4$, it exchanges much higher rates near impact to soften the inflation hike, but without the typically associated output drop. The output response is milder relative to baseline and the unemployment response, in the short-run, is improved. At $T_c = 6$, the same pattern emerges. At $T_c = 8$ the dual-mandate path softens the GDP trough even further and the unemployment hump as well, *but* does not hike rates as much as the previous two cases, accepting a bit more inflation. At $T_c = 20$, it even wants to *lower* rates relative to baseline in the short-run, but in exchange for smaller (than baseline) cuts in the medium-run and in some sense, resembles the rate-peg at $T_c = 20$, which achieves the best outcome at the macro level.

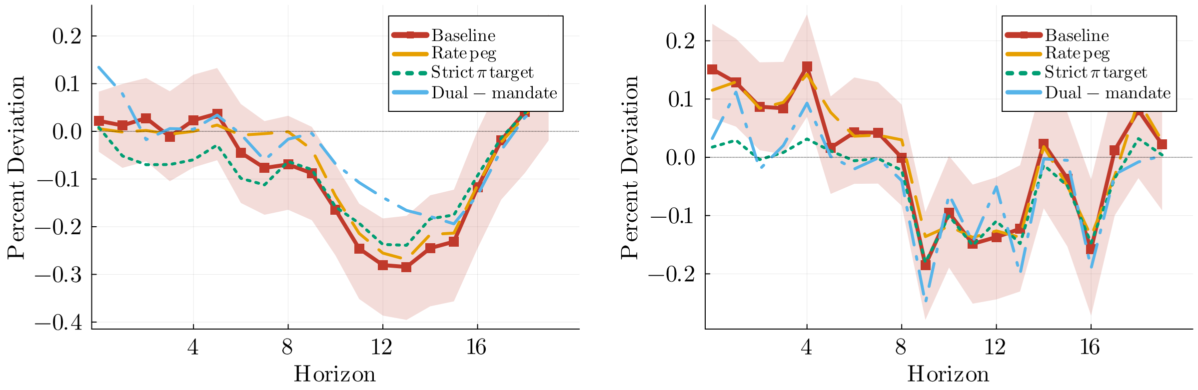

(a) Federal funds rate (b) CPI inflation

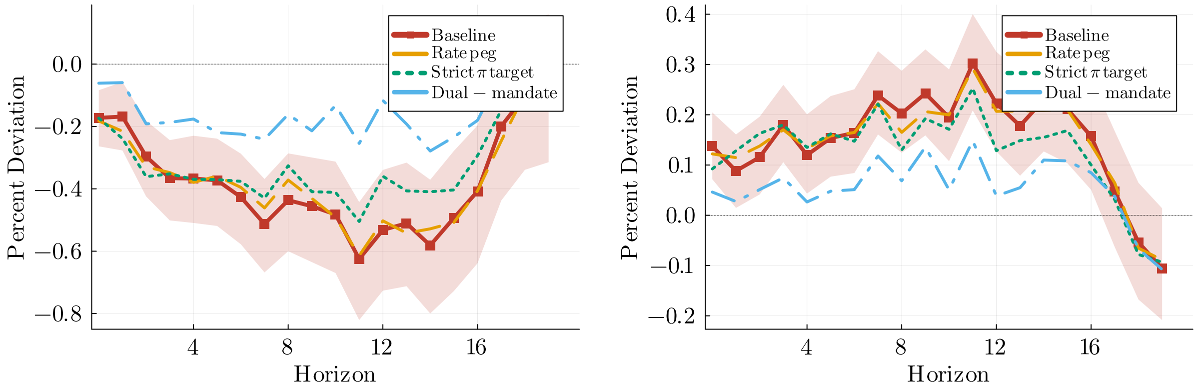

(c) Real GDP (d) Unemployment

Figure 2: Macro counterfactual responses at $T_c = 8$ ($K = 6$ baseline, over-determined system). Lines and bands as in Figure 1.

5.2 Channel decomposition

5.2.1 What is the Fed responding to?

Figure 3 reports the channel decomposition of the realized FFR IRF to the Känzig (2021) oil supply news shock under the in-sample policy. This is the empirical answer to the Kilian and Lewis (2011) question: *what is the Fed actually reacting to when an oil supply shock hits?* Each bar at horizon h stacks the contribution of each mediator block to the FFR response; the sum equals the posterior median IRF (the white markers).

First note the baseline FFR IRF in Figure 3. It hovers around zero in the short run, then decreases strongly after two years. This near-zero short-run response is the net effect of several mediator channels pulling in opposite directions, which is unobserved otherwise with just an IRF. In the companion paper this is covered, but as a reminder, for some horizon h, a single bar in Figure 3 aggregates the contribution of one mediator (or channel) summed over several mediation-lag pathways ($n = 0, \ldots, 16$), following the truncated Dufour and Wang (2024) decomposition. As an example, the Oil Price bar at horizon $h = 8$ sums the eight pathways 'oil shock → oil price at $t = n$ → FFR at $t = 8$' for $n \in 0, 1, 2, ..., 8$ — i.e., the oil shock activating

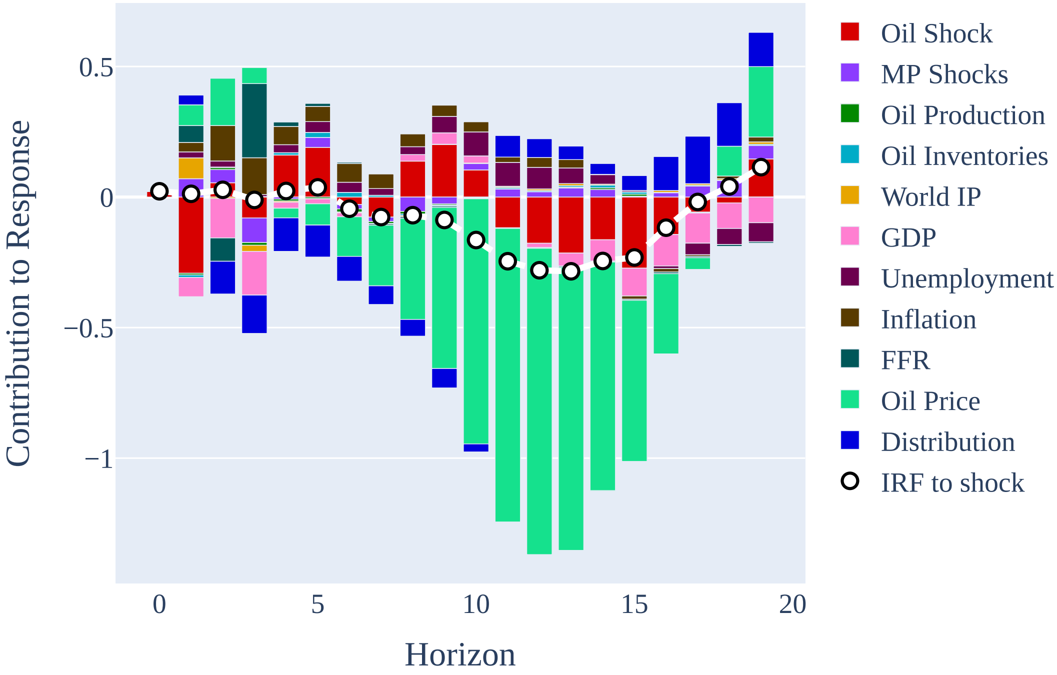

Figure 3: Channel decomposition of the realized FFR IRF to the Känzig (2021) oil supply news shock under the in-sample policy, $K = 6$ baseline. Each bar at horizon h stacks the contribution of each mediator block; the white-circle markers trace the posterior median IRF. The MP Shocks block is zero by construction in the baseline; in counterfactual panels (Figure 4) it carries the contribution of the auxiliary monetary innovations m_{aux}.

the oil-price mediator at impact and at each of the next eight quarters, with the resulting price movement then propagating through the VAR's autoregressive structure from $t = n$ to the response horizon h.

With that, what does the decomposition show? The decomposition shows the bank indeed faces a trade-off between prices and output in the presence of an oil supply shock. Inflation (brown) and the oil price (light teal) carry positive contributions, meaning it induces the bank to want to increase rates, while the demand side (GDP-pink, the distribution-dark blue) induce the bank to decrease rates. Both sets of channels cancel each other out, with slight dominance for price stability. The inflation channel itself would remain positive throughout (the Fed always wants to lean against inflation pass-through), but its magnitude is small relative to the oil-price channel; over the medium horizon the central bank "looks through" headline inflation and responds primarily to the oil-price-induced output contraction (as explicitly shown in the decompositions later). This is the textbook cost-channel pattern (Bernanke and Blanchard, 2023; Gagliardone and Gertler, 2025): short-run hawkish, medium-run dovish, long-run neutral.

Upon an oil supply shock, output and unemployment as mediators would play a secondary role in the Fed's reaction function relative to the oil-price channel. The realized reaction function is also not purely a textbook two-target Taylor rule but also responds to distributional shifts induced by the oil shock. The distributional contribution turns mildly negative through the medium horizon (along with the oil-price channel) and recovers at long horizons. The oil-block channels — oil production, oil inventories, and world industrial production — carry virtually zero contribution throughout the horizon. The interpretation is straightforward: the Fed responds to the oil shock primarily through the relative-price and real-activity transmission, not through the global supply-side dynamics of the oil market itself. The Fed's reaction function is essentially silent on the upstream oil-market block conditional on the downstream macroeconomic state.

Mediation and counterfactual: two views on the same object. The Dufour and Wang (2024) decomposition is identified from the reduced-form VAR alone and is purely descriptive of the realized IRF under the in-sample policy. Reading the indirect effect as a counterfactual—"what would the IRF look like if the channel $\mathcal{M}$ were shut off"—requires the policy-invariance assumption underlying the CMW machinery (McKay and Wolf, 2023b; Caravello, McKay, and Wolf, 2024; Wang, 2026). Under that assumption, when the FFR-channel is shut off through the indirect-effect computation, the resulting paths coincide—in the linear-Gaussian limit, and up to the McKay and Wolf (2023b) small-perturbation residual—with the IRFs that would obtain under a nominal interest-rate peg implemented through the auxiliary policy shocks. The mediation decomposition and the CMW counterfactual are two views on the same object: the former through a channel-attribution lens that does not need to solve for the implied policy shocks, the latter through an alternative-rule lens that does. Their agreement is itself a robustness check on policy invariance, and their *disagreement* is a measure of the Lucas-critique residual. The honest statement of what policy invariance buys is the one Kilian and Lewis (2011) make in their paper: it is not a *defense* of the Lucas critique, it is a *waiver*, and the waiver is credible only insofar as the contemplated rules can be implemented through monetary innovations that look like ones the data have already shown me.

The FFR decomposition can also be used (with caution) to implement an "implicit peg via mediation" counterfactual described in Section 4.3: removing the *oil-price block* of the FFR response (light teal) and propagating the remaining FFR path (inflation-, activity-, distribution-, own-FFR-driven) through to any other variable's IRF. This delivers Kilian and Lewis (2011, Section 3.2.2)'s preferred counterfactual: "the Federal Reserve reacts to fluctuations in other macroeconomic state variables as it normally would with only the direct response to the real price of oil being shut down." In this world, the FFR would definitely increase, leaning and eventually falling into the wind.

5.2.2 Channel decomposition under each rule

Figure 4 reports the full channel decomposition of the four key macro outcomes (GDP, inflation, unemployment, FFR) under the realized policy and the three counterfactual rules, all at the $K = 6$ baseline with $T_c = 6$. A lot can be learned from these decompositions. I will make two points on the machinery, particularly how the implementation of each rule is exemplified through the lens of the decomposition. I close with some discussion on the most prominent channels driving the observed effects.

Rules do what they were set out to do. This speaks to the functionality of the two machinery (Dufour and Wang, 2024; Caravello, McKay, and Wolf, 2024). Under Panel (a) Real GDP, the inflation channel that *is* observed under *Baseline* and *Peg* is *completely gone* under strict inflation targeting and the dual-mandate—precisely what the rules were set out to do. This is also observed when looking at CPI inflation for Panel (b) CPI inflation. For Panel (c) Unemployment, the unemployment and inflation channel are indeed completely suppressed for the dual-mandate and for Panel (d) Federal funds rate, the central bank has no feedback contribution through its own lag dynamics.

MP shocks shows the degree of intervention. The size and sign of MP shocks channel (purple) under each counterfactual give the contribution of the auxiliary monetary innovations m_{aux} to the outcome's response at each horizon. A larger bar means the Fed is intervening more aggressively to enforce the rule; the MP Shocks bar is zero by construction in the baseline column. From all the rules, the dual-mandate relies on more intervention.

The oil price, the distribution, and the FFR. The oil price has considerable impact across all aggregates shown. Without even zooming in on each plot, you can observe that the oil price channel dominates and drives all responses, especially in the medium-run. The effect is strongest for Real GDP and weakest, ironically, for CPI inflation. The distribution seems to amplify inflation, output, and unemployment, with the largest contribution to inflation dynamics. The FFR plays little to no role in the propagation of aggregate responses, confirming and strengthening results by Kilian and Lewis (2011).

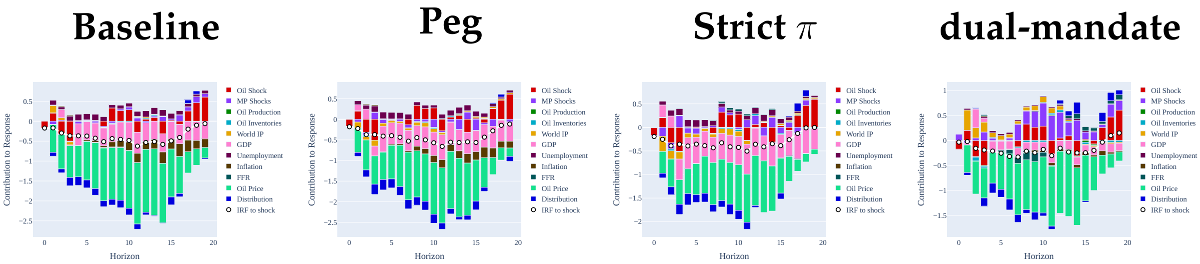

(a) Real GDP

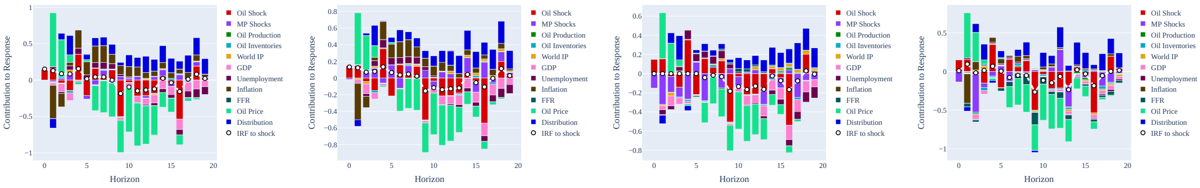

(b) CPI inflation

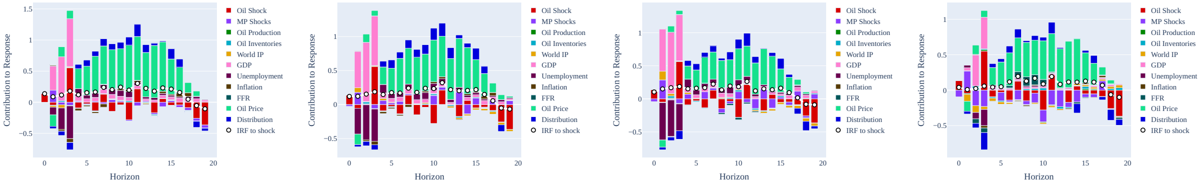

(c) Unemployment

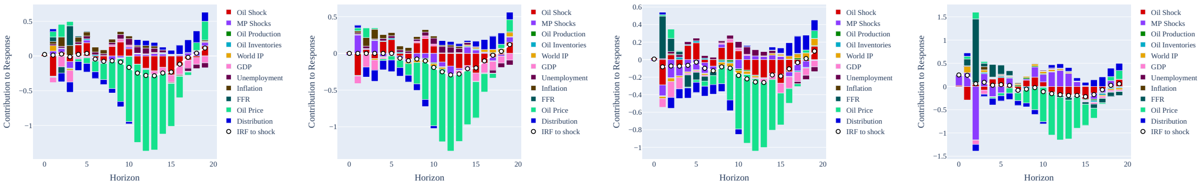

(d) Federal funds rate

Figure 4: Channel decomposition of the four macro outcomes under each policy rule, $K = 6$ specification at $T_c = 6$. Columns correspond to the realized baseline and the three counterfactual rules. The MP Shocks block (zero in the baseline column by construction) gives the share of each outcome's response mediated by the auxiliary monetary innovations m_{aux} under the rule.

5.3 Distributional counterfactuals

This section turns to the distributional responses. I report Gini coefficients, top 10% shares, and bottom 50% shares for income, wealth, and consumption. The main-text figures fix the constraint horizon at $T_c = 6$; the same panels at $T_c = 20$ are reported in Appendix B.

5.3.1 Gini coefficients

Figure 5 reports the three Ginis under the four monetary rules at $T_c = 6$ (the constraint binds for six quarters; the rules then lift and the system evolves under the realized policy). Two features structure the figure: an early-horizon bunching of all four lines while the constraint binds, and a medium-run fan-out in which strict inflation targeting stands apart from the other three rules.

The realized path is itself compressionary. A point worth surfacing before discussing the rules: under realized policy the oil shock generates a small positive Gini bump in the first year on all three dimensions, then turns negative through the medium horizon, troughing in the second-to-fourth year and recovering toward zero. The oil shock therefore *compresses* all three Ginis in the medium run, on average.

Income Gini. All four lines bunch tightly through the constraint window ($h \leq 6$): a small positive deviation of $+0.1$ to $+0.25$ percentage points in the first year, then crossing zero. The rules fan out from $h \approx 8$ onward. The rate peg tracks the realized response in the short-run, but deviates in the medium-run. The dual-mandate rule sits slightly above the realized response in the medium horizon (i.e. marginally less compression). The strict-π target is the standout: it deepens the realized trough by roughly twenty per cent, bottoming at about -0.50 at $h = 14$ against the realized trough of about -0.35, and recovers more slowly. The strict-π rule therefore *amplifies the compressionary effect* of the oil shock on the income Gini.

Wealth Gini. The realized wealth Gini path is muted – a small positive bump ($+0.1$) over the first year, dipping to about -0.1 at $h \approx 10$, recovering. The rate peg and dual-mandate rules sit close to the realized path throughout. Strict inflation targeting, again, is the standout: it pulls the wealth Gini further down to a trough of about -0.30 at $h \approx 14$, a roughly three-fold amplification of the realized compression.

Consumption Gini. The realized consumption Gini follows the income pattern at smaller amplitude: a brief positive bump, then a negative dip troughing around -0.20 at $h \approx 10$, recovery toward zero by $h = 18$. The dual-mandate again tracks the realized path. Strict

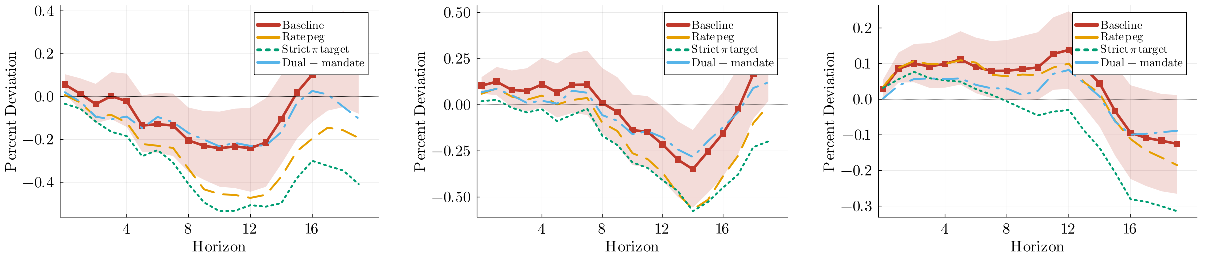

(a) Consumption Gini (b) Income Gini (c) Wealth Gini

Figure 5: Counterfactual Gini responses to the Känzig (2021) oil supply news shock, $K = 6$ baseline, $T_c = 6$. Lines: realized (solid red), rate peg (dashed orange), strict π-target (dotted teal), dual-mandate (dash-dot light blue). Bands shown only for the realized response. Vertical axis: percent deviation of the Gini coefficient from steady state.

inflation targeting and the rate-peg amplifies the trough to about -0.5 at $h \approx 11$, recovering thereafter. The amplification factor under strict-π is the largest of the three Ginis: roughly two-and-a-half-fold relative to the realized trough.

Cross-rule comparison. The summary statement is sharper than at the macro level. One of the three rules – dual-mandate – delivers near-indistinguishable Gini paths on all three margins. The strict-π rule amplifies the medium-run compression on all three margins, with the largest amplification for consumption and the smallest for income. The rate-peg equally compresses, but follows the realized path for wealth.

The trade-off the strict-π rule purchases at the macro level shows up in the distribution as a deeper and more persistent Gini compression. Whether this is welfare-improving or welfare-reducing depends on the social welfare function: a planner that values inequality compression mechanically prefers strict π, but the compression may be from declines at the top, not by gains for the bottom – a point I visit in the next section.

5.3.2 Top decile and bottom half shares

To understand the mechanism behind the Gini results, I decompose each Gini into its top 10% and bottom 50% components. Figure 6 reports the top 10% (row 1) and bottom 50% (row 2) shares for income, consumption, and wealth. The patterns clarify why all three Ginis amplify under strict inflation targeting.

Decomposition of the Gini movement. The three Ginis may differ in which margin does the heavier work. First, I discuss the strict inflation targeting rule (dotted-green line). On consumption, we observe a negative change in the top-10% share and a larger positive change for the bottom-50%. On income, the top-10% movement is slightly more pronounced than the bottom-50% movement, which explains the small difference observed in the ginis. On wealth,

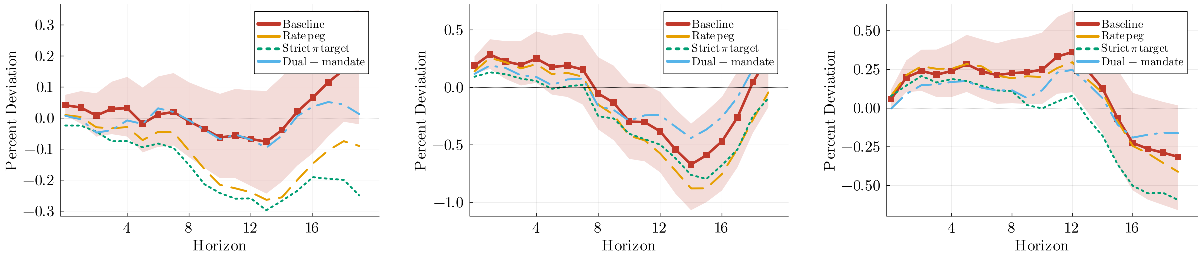

(a) Consumption: top 10% (b) Income: top 10% (c) Wealth: top 10%

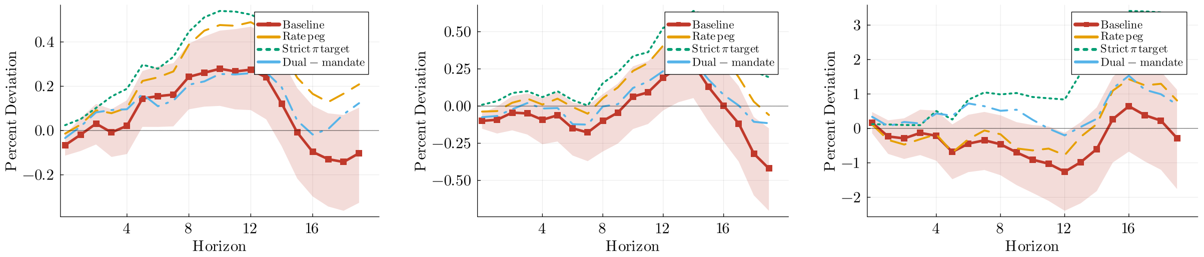

(d) Consumption: bottom 50% (e) Income: bottom 50% (f) Wealth: bottom 50%

Figure 6: Counterfactual top 10% (row 1) and bottom 50% (row 2) shares for income, consumption, and wealth, $K = 6$ baseline, Känzig (2021) oil supply news shock, $T_c = 6$. Lines as in Figure 5.

the bottom-50% movement dominates by a large margin across the IRF horizon, which is why the wealth Gini under strict-π shows the largest amplification of the three. Overall, the rate-peg stays close to baseline, except for consumption, while the strict-inflation target compresses inequality through large redistribution to the bottom. Dual-mandate follows the realized path.

6 Conclusion

I have shown how the CMW sufficient-statistics framework can be married to a multi-shock distributional VAR to recover the counterfactual income, wealth, and consumption distributions that would have prevailed under three alternative monetary rules—a nominal rate peg, strict inflation targeting, and an optimal dual-mandate rule—in response to an identified Känzig (2021) oil supply news shock. The key methodological contribution is the recognition that the linear superposition that delivers the CMW macro counterfactual extends, draw-by-draw, to the underlying factor IRFs and hence to any non-linear function of the joint distribution. The key empirical contributions are quantitative estimates of the macro and distributional cost of each rule across a sweep of constraint horizons $T_c \in \{4, 6, 8, 20\}$. The dual-mandate rule delivers Gini paths nearly indistinguishable from the realized response. The rate peg and strict inflation targeting both compress inequality further than the realized path, with strict-π the larger amplifier on every margin (by roughly twenty per cent for income, two-and-a-half-fold for consumption, and threefold for wealth). The top-decile/bottom-half decomposition attributes the strict-π compression to redistribution toward the bottom of the distribution rather than top-tail destruction.

Two margins discriminate between the rules. The first is policy duration, on the macro side. Held briefly, the three rules deliver broadly similar inflation, GDP, and unemployment paths and look near-identical to the realized response. Enforced over the full twenty-quarter horizon, the rate peg and the dual-mandate rule converge to one another and deliver materially better macro stability than the realized path; strict-π does not, because its medium-run tightening is quantitatively close to what the realized policy already delivered. The macro case for an accommodative rule is therefore conditional on the central bank being willing to commit to a long enforcement window—a feature made explicit in Appendix A. The second margin is the cross-section, which discriminates even at short horizons: the three rules' distributional footprints separate strict-π and the rate peg from the dual-mandate in a quantitatively non-negligible way already at $T_c = 6$. A central bank that incorporates distributional moments into its loss function, or that is willing to credibly commit to a long-duration rule, would rank the rules differently than one that does not. I see this as a strong case for incorporating distributional moments directly into central-bank loss functions, along the lines proposed by Bhandari et al. (2021) and others; the CMW framework offers a tractable empirical complement to that theoretical literature.

Future work. Four extensions are immediate. First, discriminating between *types* of supply shocks matters: an oil-supply news shock, an oil-demand shock, and a broader non-oil supply disturbance can trigger very different Fed reaction functions and distributional responses, and the paper's machinery is well-suited to that shock-by-shock comparison. Second, the Wang (2026) local-projection counterpart to the CMW machinery offers a complementary identification of the auxiliary monetary block, with different small-sample and misspecification properties; running the same rule set under his counterfactual machinery would provide an informative robustness check on the BVAR-based results reported here. Third, the paper has stipulated the central bank as the sole policy authority responding to the oil shock; an oil-side authority—OPEC quota adjustments, strategic-reserve releases, or energy taxes—is a complementary lever whose counterfactual could be enforced through the same sufficient-statistics machinery, identifying the rule through historically observed oil-supply or oil-inventory innovations rather than through monetary innovations. The cross-authority comparison would clarify whether inequality stabilisation is best pursued on the monetary or the oil-policy margin. Fourth, the same exercise can be run with monetary or fiscal impulses in place of the oil shock; the distributional counterfactuals under alternative fiscal rules would be a useful complement to the recent literature on fiscal multipliers and inequality.

References

[1] Asger Lau Andersen et al. “Monetary policy and inequality”. In: (2021).

[2] S Boragan Aruoba and Thomas Drechsel. *Identifying monetary policy shocks: A natural language approach*. Tech. rep. National Bureau of Economic Research, 2024.

[3] Adrien Auclert. “Monetary policy and the redistribution channel”. In: *American Economic Review* 109.6 (2019), pp. 2333–2367.

[4] Regis Barnichon and Geert Mesters. “A sufficient statistics approach for macro policy”. In: *American Economic Review* 113.11 (2023), pp. 2809–2845.

[5] Michael D Bauer and Eric T Swanson. “A reassessment of monetary policy surprises and high-frequency identification”. In: *NBER Macroeconomics Annual* 37.1 (2023), pp. 87–155.

[6] Christiane Baumeister. *Discussion of” Local Projections or VARs? A Primer for Macroeconomists” by José Luis Montiel Olea, Mikkel Plagborg-Møller, Eric Qian, and Christian K*. Tech. rep. Wolf, Working paper, Prepared for the NBER Macroeconomics Annual, 2025.

[7] Christiane Baumeister and James D. Hamilton. “Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks”. In: *American Economic Review* 109.5 (2019), pp. 1873–1910.

[8] Christian Bayer, Benjamin Born, and Ralph Luetticke. “Shocks, frictions, and inequality in US business cycles”. In: (2020).

[9] Christian Bayer, Luis Calderon, and Moritz Kuhn. *Distributional Dynamics*. Tech. rep. 2025.

[10] Ben S. Bernanke and Olivier J. Blanchard. *What Caused the U.S. Pandemic-Era Inflation?* NBER Working Paper 31417. National Bureau of Economic Research, 2023.

[11] Ben S Bernanke, Mark Gertler, and Mark Watson. “Systematic monetary policy and the effects of oil price shocks”. In: *Brookings papers on economic activity* 1997.1 (1997), pp. 91–157.

[12] Ben S Bernanke, Mark Gertler, and Mark W Watson. “Oil shocks and aggregate macroeconomic behavior: the role of monetary policy: a reply”. In: *Journal of Money, Credit, and Banking* 36.2 (2004), pp. 287–291.

[13] Ben S Bernanke and Frederic S Mishkin. "Inflation targeting: a new framework for monetary policy?" In: *Journal of Economic Perspectives* 11.2 (1997), pp. 97–116.

[14] Anmol Bhandari et al. "Inequality, Business Cycles, and Monetary-Fiscal Policy". In: *Econometrica* 89.6 (2021), pp. 2559–2599.

[15] Hilde C Bjørnland, Yoosoon Chang, and Jamie Cross. "Oil and the stock market revisited: A mixed functional var approach". In: *Available at SSRN* (2023).

[16] Martin Bodenstein, Luca Guerrieri, and Lutz Kilian. "Monetary policy responses to oil price fluctuations". In: *IMF Economic Review* 60.4 (2012), pp. 470–504.

[17] Tobias Broer, John Kramer, and Kurt Mitman. "The Distributional Effects of Oil Shocks". In: *IMF Economic Review* 73.3 (2025), pp. 851–889.

[18] Tomás E Caravello, Alisdair McKay, and Christian K Wolf. *Evaluating policy counterfactuals: A var-plus approach*. Tech. rep. National Bureau of Economic Research, 2024.

[19] Charles T Carlstrom and Timothy S Fuerst. "Oil prices, monetary policy, and counterfactual experiments". In: *Journal of Money, Credit and Banking* 38.7 (2006), pp. 1945–1958.

[20] Joshua C C Chan. "Asymmetric conjugate priors for large Bayesian VARs". In: *Quantitative Economics* 13.3 (2022), pp. 1145–1169.

[21] Minsu Chang, Xiaohong Chen, and Frank Schorfheide. "Heterogeneity and Aggregate Fluctuations". In: *Journal of Political Economy* 132.12 (2024), pp. 4021–4067.

[22] Minsu Chang and Frank Schorfheide. *On the Effects of Monetary Policy Shocks on Income and Consumption Heterogeneity*. Working Paper 32166. National Bureau of Economic Research, 2024.

[23] Olivier Coibion et al. "Innocent Bystanders? Monetary policy and inequality". In: *Journal of Monetary Economics* 88 (2017), pp. 70–89.

[24] Ferre De Graeve and Andreas Westermark. *How long is long enough? Long-lag VARs*. Tech. rep. Sveriges Riksbank Working Paper Series, 2025.

[25] Jean-Marie Dufour and Eric Renault. "Short run and long run causality in time series: theory". In: *Econometrica* (1998), pp. 1099–1125.

[26] Jean-Marie Dufour and Endong Wang. *Causal mechanism and mediation analysis for macroeconomics dynamics: a bridge of Granger and Sims causality*. Tech. rep. 2024.

[27] Luca Gagliardone and Mark Gertler. "Oil Prices, Monetary Policy and Inflation Surges". In: *American Economic Journal: Macroeconomics* (2025). Forthcoming. Previous version: NBER Working Paper 31263, 2023.

[28] Domenico Giannone, Michele Lenza, and Giorgio E Primiceri. "Priors for the long run". In: *Journal of the American Statistical Association* 114.526 (2019), pp. 565–580.

[29] Lukas Hack, Klodiana Istrefi, and Matthias Meier. *The Systematic Origins of Monetary Policy Shocks*. Working paper. `https://cepr.org/publications/dp19063`. CEPR Discussion Paper No. 19063 / CRC TR 224 Discussion Paper No. 557, University of Bonn, 2024.

[30] James D Hamilton and Ana Maria Herrera. "Comment: oil shocks and aggregate macroeconomic behavior: the role of monetary policy". In: *Journal of Money, credit and Banking* (2004), pp. 265–286.

[31] Gökhan Ider et al. *Friend, Not Foe? Monetary Policy and Energy Prices*. DIW Discussion Paper. May 2025. DIW Berlin, 2025.

[32] Marek Jarociński. "Estimating the Fed's unconventional policy shocks". In: *Journal of Monetary Economics* (2024).

[33] Marek Jarociński and Peter Karadi. "Deconstructing monetary policy surprises—the role of information shocks". In: *American Economic Journal: Macroeconomics* 12.2 (2020), pp. 1–43.

[34] Diego R Känzig. "The macroeconomic effects of oil supply news: Evidence from OPEC announcements". In: *American Economic Review* 111.4 (2021), pp. 1092–1125.

[35] Greg Kaplan, Benjamin Moll, and Giovanni L Violante. "Monetary policy according to HANK". In: *American Economic Review* 108.3 (2018), pp. 697–743.

[36] Lutz Kilian. "Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market". In: *American economic review* 99.3 (2009), pp. 1053–1069.

[37] Lutz Kilian and Logan T Lewis. "Does the Fed respond to oil price shocks?" In: *The Economic Journal* 121.555 (2011), pp. 1047–1072.

[38] Lutz Kilian and Xiaoqing Zhou. "The econometrics of oil market VAR models". In: (2023).

[39] Sylvain Leduc and Keith Sill. “A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns”. In: *Journal of Monetary Economics* 51.4 (2004), pp. 781–808.

[40] Michele Lenza and Ettore Savoia. *Do we need firm data to understand macroeconomic dynamics?* Working Paper Series 438. Sveriges Riksbank (Central Bank of Sweden), 2024.

[41] Michele Lenza and Jiri Slacalek. *How does monetary policy affect income and wealth inequality? Evidence from quantitative easing in the euro area*. Working paper 2190. Published as Lenza and Slacalek (2024), *Journal of Applied Econometrics*. European Central Bank, 2018.

[42] Robert E. Lucas. “Econometric policy evaluation: a critique”. In: *The Phillips Curve and Labor Markets*. Ed. by Karl Brunner and Allan H. Meltzer. Vol. 1. Carnegie–Rochester Conference Series on Public Policy. North-Holland, 1976, pp. 19–46.

[43] Julian F. Ludwig. *Local Projections Are VAR Predictions of Increasing Order*. Tech. rep. SSRN Working Paper 4882149, 2024.

[44] Ralph Luetticke. *A Unified Narrative of Conventional and Unconventional Monetary Policy Shocks*. Working paper. Replication archive: `https://github.com/ralphluet/unified_narrative_mp_shocks`. University College London, 2024.

[45] Alisdair McKay and Christian K Wolf. “Monetary policy and inequality”. In: *Journal of Economic Perspectives* 37.1 (2023), pp. 121–144.

[46] Alisdair McKay and Christian K Wolf. “What Can Time-Series Regressions Tell Us About Policy Counterfactuals?” In: *Econometrica* 91.5 (2023), pp. 1695–1725.

[47] Anton Nakov and Andrea Pescatori. “Oil and the Great Moderation”. In: *The Economic Journal* 120.543 (2010), pp. 131–156.

[48] Mikkel Plagborg-Møller and Christian K Wolf. “Local projections and VARs estimate the same impulse responses”. In: *Econometrica* 89.2 (2021), pp. 955–980.

[49] Michael Plante. “How should monetary policy respond to changes in the relative price of oil? Considering supply and demand shocks”. In: *Journal of Economic Dynamics and Control* 44 (2014), pp. 1–19.

[50] Christina D Romer and David H Romer. “A new measure of monetary shocks: Derivation and implications”. In: *American economic review* 94.4 (2004), pp. 1055–1084.

[51] Christopher A Sims and Tao Zha. “Does monetary policy generate recessions?” In: *Macroeconomic Dynamics* 10.2 (2006), pp. 231–272.

[52] Abe Sklar. “Vol. 8 of Fonctions de Répartition à n Dimensions et Leurs Marges, 229–231”. In: *Paris: Publications de l'Institut de statistique de l'Université de Paris* (1959).

[53] James H Stock and Mark W Watson. “Identification and estimation of dynamic causal effects in macroeconomics using external instruments”. In: *The Economic Journal* 128.610 (2018), pp. 917–948.

[54] Endong Wang. *Local projections identify the same policy counterfactuals as empirical and structural models*. Working paper. `https://arxiv.org/abs/2409.09577`. University of Mannheim, 2026.

A Additional macro counterfactual robustness

This appendix reports the macro counterfactual responses at the two remaining constraint horizons in the robustness sweep ($T_c = 6$ in Figure 7, $T_c = 20$ in Figure 8) and the real oil-price column for all four constraint horizons (Figure 9). The oil-price column is collected separately because monetary policy has at most a small effect on the global oil market at quarterly frequency (Kilian, 2009; Baumeister and Hamilton, 2019), so the counterfactual deviations on the oil-price are mechanically small and do not discriminate between rules.

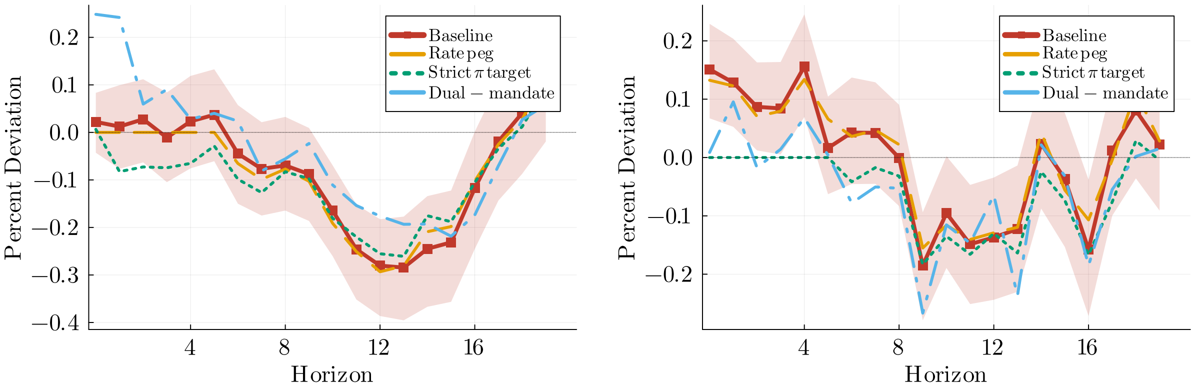

(a) Federal funds rate

(b) CPI inflation

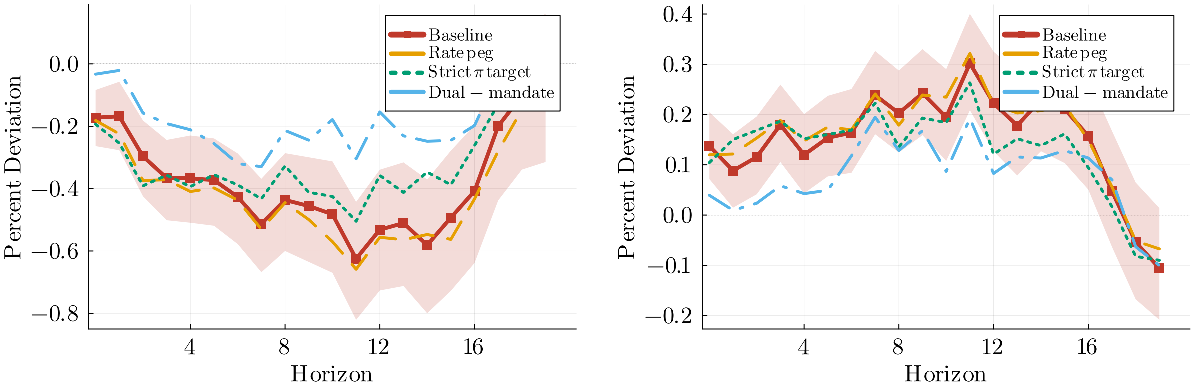

(c) Real GDP

(d) Unemployment

Figure 7: Macro counterfactual responses at $T_c = 6$ ($K = 6$ baseline, exact-fit case where $T_c = K$). Lines and bands as in Figure 1. The square-system case for the $K = 6$ specification: peg and strict-π projections are exact over the first six quarters.

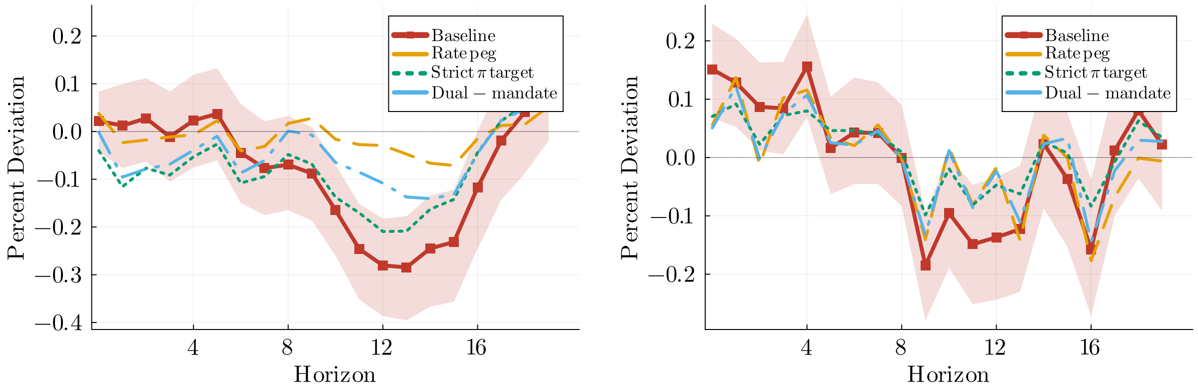

(a) Federal funds rate (b) CPI inflation

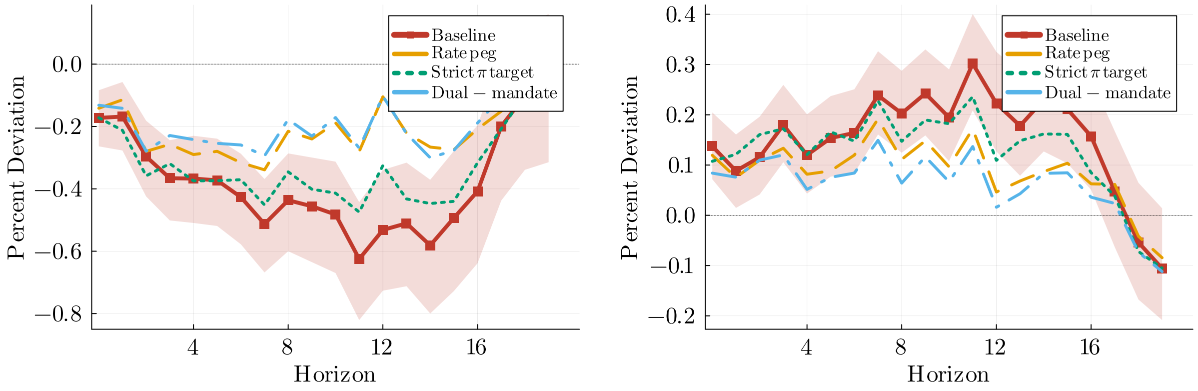

(c) Real GDP (d) Unemployment

Figure 8: Macro counterfactual responses at $T_c = 20$ ($K = 6$ baseline, full-horizon enforcement). Lines and bands as in Figure 1. This is the largest-T_c case in the robustness sweep; the rule is enforced over the entire plotted horizon. The conclusions are qualitatively identical to Figure 2.

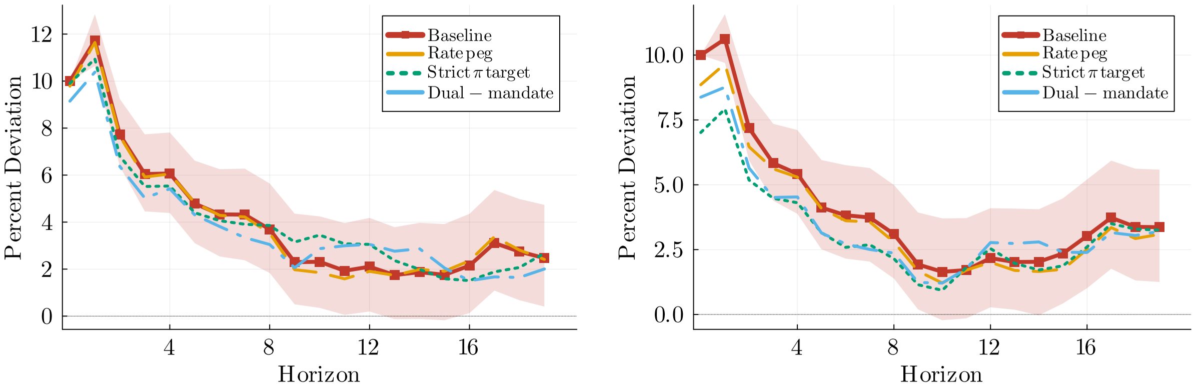

(a) $T_c = 4$ ($K = 6$ unified)

(b) $T_c = 6$ ($K = 6$ baseline)

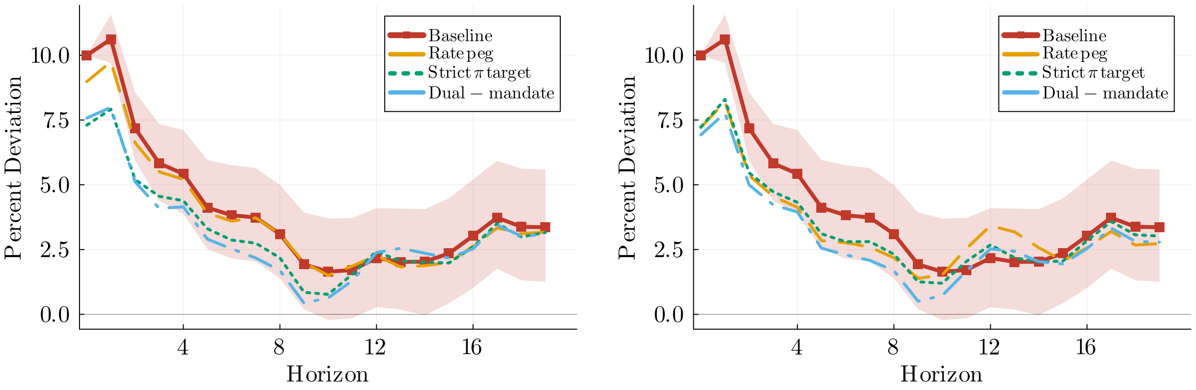

(c) $T_c = 8$ ($K = 6$ baseline)

(d) $T_c = 20$ ($K = 6$ baseline)

Figure 9: Real-oil-price counterfactual responses by constraint horizon T_c. Lines and bands as in Figure 1. Counterfactual deviations on the real oil price are small under all rules at every T_c, consistent with the standard finding that domestic monetary policy has at most a modest effect on the global oil market at quarterly frequency (Kilian, 2009; Baumeister and Hamilton, 2019).

B Distributional counterfactuals at $T_c = 20$

This appendix replicates the distributional figures of Section 5.3 at the long constraint horizon $T_c = 20$, in which the rule binds for the entire response window. The main-text figures report $T_c = 6$ as the headline; the $T_c = 20$ panels here serve as a long-horizon robustness check and amplify the same qualitative ordering across rules.

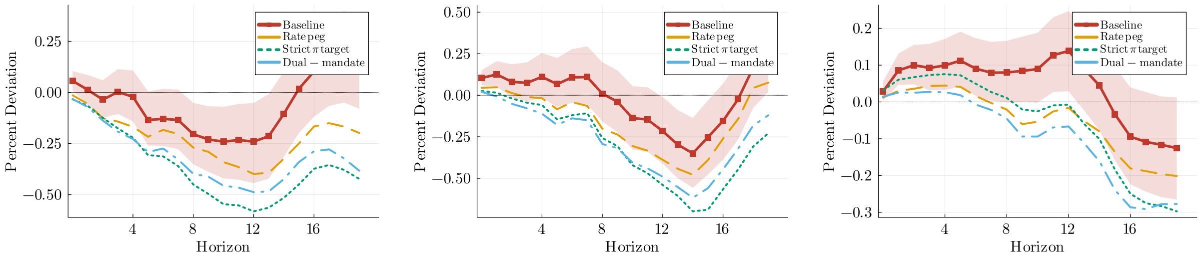

(a) Consumption Gini (b) Income Gini (c) Wealth Gini

Figure 10: Counterfactual Gini responses at $T_c = 20$. Känzig (2021) oil supply news shock, $K = 6$ baseline. Lines as in Figure 5.

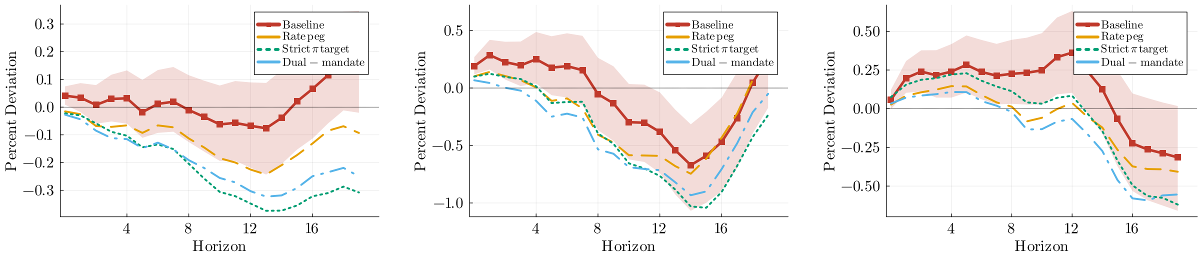

(a) Consumption: top 10% (b) Income: top 10% (c) Wealth: top 10%

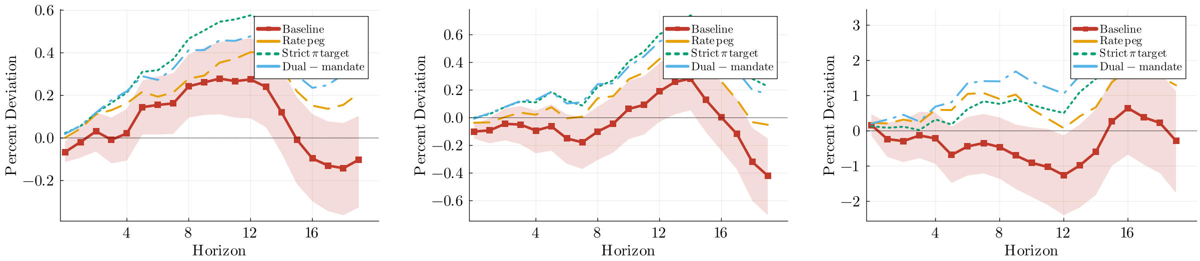

(d) Consumption: bottom 50% (e) Income: bottom 50% (f) Wealth: bottom 50%

Figure 11: Top 10% (row 1) and bottom 50% (row 2) shares at $T_c = 20$, $K = 6$ baseline, Känzig (2021) oil supply news shock. Lines as in Figure 5.

C T_c-robustness diagnostics and the Lucas-critique residual

Table 2 reports the CMW diagnostics for every (specification, T_c, rule) combination in the robustness sweep, separately for the BH and Kanzig oil shocks. Each row reports three quantities at the posterior median (with the 16th–84th percentile band in brackets): the L2 norm of the auxiliary monetary impulse vector m_{aux}, the condition number of $M_1'M_1$ (the policy-span Gram matrix), and the rule-fit ratio $\|M_1 m_{\text{aux}} + M_2\|/\|M_2\|$ (zero = rule exactly enforced, one = no enforcement). All entries use the same 5% per-rule trim on ill-conditioned draws described in Section 4.3; the $K = 6$ entries at $T_c = 4$ for the hard-constraint rules show `NaN` because the projection is under-determined ($T_c < K$) and all draws are masked.

Reading the table. The diagnostics behave as the algebra predicts. As T_c falls toward K, the rule-fit ratio drops monotonically; at $T_c = K$ (unified $T_c = 4$, K6 $T_c = 6$) the projection is exact for the hard-constraint rules and the ratio is essentially zero. For $T_c > K$ the projection is over-determined and the ratio rises with $T_c - K$ until it stabilizes around the values reported at $T_c = 8$. The condition number cond$(M_1'M_1)$ is largest at $T_c = K$ (where the square system has condition-squared sensitivity to the underlying M_1) and falls as additional rows over-determine the projection.

The Lucas-critique residual. The economically interesting quantity is $\|m_{\text{aux}}\|_2$, because McKay and Wolf (2023b) prove that the Lucas-critique residual—the part of agents' decision rules that would actually adjust under a policy change but that the sufficient-statistics counterfactual cannot capture—is $O(\|m_{\text{aux}}\|^2)$. Smaller $\|m_{\text{aux}}\|$ therefore means a smaller Lucas residual and a more credible policy-invariance restriction.

Three patterns in Table 2 bear on how hard the Lucas critique bites in this exercise.

(i) Kanzig is comfortably inside the small-perturbation regime. Under the Kanzig shock, $\|m_{\text{aux}}\|_2$ medians range from 0.14 to 0.85 across all specifications and rules, with all upper-tail (84th-percentile) values below 3 except for the two $K = 6$ $T_c = 6$ square-system entries. Squaring the medians to bound the McKay–Wolf residual gives Lucas-residual upper bounds under 0.04 for the $K = 4$ unified spec and under 0.7 for the largest $K = 6$ cases. These are small relative to the size of the counterfactual deviations themselves (FFR deviations of roughly 0.4 percentage points at peak), so the Lucas residual is quantitatively negligible for the headline Kanzig results.

(ii) The BH shock sits on the edge of the regime. BH $\|m_{\text{aux}}\|_2$ medians are roughly four to five times larger than Kanzig's (e.g. 1.18 for the unified $T_c = 4$ peg, 2.16 for the $K = 6$ $T_c = 6$ peg) with upper-tail values reaching 7.5. Squaring these gives Lucas-residual upper bounds of 1.4 to 30, no longer negligible relative to the counterfactual deviations. The BH results reported alongside the headline should therefore be read as *quantitatively* more sensitive to the Lucas critique than the Kanzig headline; the qualitative ordering across rules is preserved, but the magnitudes carry a larger interpretation tax. This is the empirical reason for adopting the Kanzig HF identification as the headline shock and reporting BH as robustness rather than the other way around.

(iii) Smaller T_c does not always mean smaller $\|m_{aux}\|$. The natural intuition is that pinning the rule over a shorter window requires a smaller impulse. The table shows that this is true comparing $T_c = 20$ to $T_c = 8$ within an over-determined spec, but is reversed in the exact-fit case ($T_c = K$). At $T_c = K$ the system is square and the projection must absorb *all* of the rule violation M_2 through m_{aux}, with no over-determination to soften the load; the result is a larger $\|m_{\text{aux}}\|$ and a worse cond$(M_1'M_1)$ than the over-determined neighbours. Under the BH shock, $\|m_{\text{aux}}\|_2$ for the $K = 6$ $T_c = 6$ peg is 2.16 vs. 1.08 at $T_c = 8$ – a factor of two reduction by moving one quarter further over-determined. The empirically Lucas-safest specification under the BH

shock is therefore $T_c = 8$ with $K = 6$, not the exact-fit $T_c = 6$.

Implication for the headline results. The two main-text specifications in Section 5.1 ($T_c = 4$ $K = 4$ and $T_c = 8$ $K = 6$) sit at posterior-median $\|m_{\text{aux}}\|_2$ values of 0.19 and 0.14 under Kanzig, respectively, with corresponding Lucas-residual upper bounds of 0.04 and 0.02. These are well inside the small-perturbation regime in which McKay and Wolf (2023b)'s identification result is tight, and the policy-invariance restriction is empirically defensible in the sense of Wang (2026)'s $\Delta_{\mathcal{Q}}$ persistence test. The Lucas critique bites quantitatively when the exercise is pushed to large-T_c $K = 6$ specifications or to the BH shock, but neither configuration is what the headline reports.

Table 2: CMW diagnostics across T_c and rule, by oil shock. Median [16th, 84th percentile] across posterior draws (5% top-cond trim).

			$\|m_{\text{aux}}\|_2$		cond$(M_1'M_1)$		rule-fit ratio	
Variant	T_c	Rule	med	[16,84]	med	[16,84]	med	[16,84]
Panel A. BH supply shock								
unified	4	peg	1.18	$[0.42, 4.07]$	$3.3\cdot 10^4$	$[3.4\cdot 10^3, 9.3\cdot 10^5]$	0.00	$[0.00, 0.00]$
unified	4	strict π	1.54	$[0.59, 5.39]$	2285	$[219, 7.2\cdot 10^4]$	0.00	$[0.00, 0.00]$
unified	4	dual-mandate	0.59	$[0.29, 1.25]$	510	$[83, 5938]$	0.70	$[0.52, 0.85]$
unified	6	peg	0.48	$[0.20, 1.18]$	7193	$[1143, 9.3\cdot 10^4]$	0.37	$[0.18, 0.61]$
unified	6	strict π	0.59	$[0.29, 1.32]$	510	$[84, 6636]$	0.53	$[0.30, 0.75]$
unified	6	dual-mandate	0.40	$[0.19, 0.75]$	341	$[58, 3987]$	0.80	$[0.67, 0.91]$
unified	8	peg	0.49	$[0.20, 1.00]$	3969	$[732, 5.2\cdot 10^4]$	0.48	$[0.29, 0.69]$
unified	8	strict π	0.45	$[0.21, 0.89]$	358	$[64, 4485]$	0.67	$[0.47, 0.83]$
unified	8	dual-mandate	0.32	$[0.15, 0.64]$	278	$[49, 3297]$	0.86	$[0.76, 0.94]$
K6	4	peg	NaN		NaN		NaN	
K6	4	strict π	NaN		NaN		NaN	
K6	4	dual-mandate	1.22	$[0.65, 2.42]$	2207	$[318, 3.0\cdot 10^4]$	0.36	$[0.19, 0.55]$
K6	6	peg	2.16	$[0.94, 7.45]$	$1.5\cdot 10^5$	$[1.4\cdot 10^4, 5.8\cdot 10^6]$	0.00	$[0.00, 0.00]$
K6	6	strict π	1.91	$[0.83, 6.73]$	7833	$[699, 2.4\cdot 10^5]$	0.00	$[0.00, 0.00]$
K6	6	dual-mandate	0.72	$[0.40, 1.20]$	726	$[129, 1.0\cdot 10^4]$	0.60	$[0.44, 0.74]$
K6	8	peg	1.08	$[0.54, 2.11]$	$1.8\cdot 10^4$	$[3048, 2.8\cdot 10^5]$	0.27	$[0.13, 0.45]$
K6	8	strict π	1.05	$[0.58, 1.99]$	1298	$[222, 2.0\cdot 10^4]$	0.44	$[0.22, 0.66]$
K6	8	dual-mandate	0.68	$[0.38, 1.14]$	534	$[97, 7427]$	0.71	$[0.58, 0.83]$
Panel B. Kanzig oil supply news shock								
unified	4	peg	0.19	$[0.07, 0.67]$	$1.3\cdot 10^4$	$[1301, 3.7\cdot 10^5]$	0.00	$[0.00, 0.00]$
unified	4	strict π	0.40	$[0.17, 1.32]$	1379	$[130, 4.5\cdot 10^4]$	0.00	$[0.00, 0.00]$
unified	4	dual-mandate	0.23	$[0.12, 0.42]$	240	$[44, 3230]$	0.50	$[0.33, 0.67]$
unified	6	peg	0.15	$[0.07, 0.36]$	2641	$[416, 3.1\cdot 10^4]$	0.30	$[0.14, 0.53]$
unified	6	strict π	0.24	$[0.12, 0.49]$	274	$[48, 4075]$	0.46	$[0.23, 0.67]$
unified	6	dual-mandate	0.20	$[0.10, 0.37]$	160	$[29, 1957]$	0.65	$[0.50, 0.79]$
unified	8	peg	0.14	$[0.06, 0.30]$	1633	$[276, 2.0\cdot 10^4]$	0.46	$[0.28, 0.68]$
unified	8	strict π	0.18	$[0.09, 0.34]$	202	$[36, 2267]$	0.58	$[0.39, 0.77]$
unified	8	dual-mandate	0.19	$[0.10, 0.35]$	128	$[24, 1587]$	0.73	$[0.59, 0.86]$
K6	4	peg	NaN		NaN		NaN	
K6	4	strict π	NaN		NaN		NaN	
K6	4	dual-mandate	0.50	$[0.27, 0.93]$	1400	$[220, 2.2\cdot 10^4]$	0.33	$[0.17, 0.52]$
K6	6	peg	0.68	$[0.27, 2.30]$	$1.2\cdot 10^5$	$[1.2\cdot 10^4, 3.2\cdot 10^6]$	0.00	$[0.00, 0.00]$
K6	6	strict π	0.85	$[0.36, 2.87]$	5255	$[463, 1.6\cdot 10^5]$	0.00	$[0.00, 0.00]$
K6	6	dual-mandate	0.33	$[0.20, 0.56]$	459	$[92, 7540]$	0.57	$[0.41, 0.73]$
K6	8	peg	0.39	$[0.19, 0.77]$	$1.4\cdot 10^4$	$[2429, 2.0\cdot 10^5]$	0.26	$[0.12, 0.47]$
K6	8	strict π	0.39	$[0.21, 0.72]$	946	$[155, 1.6\cdot 10^4]$	0.41	$[0.22, 0.62]$
K6	8	dual-mandate	0.33	$[0.19, 0.51]$	359	$[75, 6267]$	0.64	$[0.50, 0.79]$

Notes. All entries pool 2000 posterior draws after the per-rule 5% top-cond trim. $K = 4$ unified spec uses Lütke + BS + HIM-RR + HIM-AD; $K = 6$ spec uses Lütke + 4 Jarociński factors + HIM-RR. `NaN` = under-determined system ($T_c < K$ for hard-constraint rules); all draws masked. The square-system case ($T_c = K$) has rule-fit ≈ 0 for peg and strict π by construction.