Policy Rules and Policy Tools:
The Price of Distributional Stabilization
——*Preliminary*——*

Luis Calderon†
University of Bonn

October 28, 2026

Abstract

This paper examines how much monetary and fiscal policy can stabilize consumption inequality after an oil supply shock, and at what aggregate cost. I model policy instruments jointly with the distribution of U.S. household consumption, income, and wealth in a Bayesian VAR, and compare the outcomes policy can achieve under alternative mandates. The shock raises inflation, depresses real activity, and compresses the consumption distribution. Under single mandates, the policy rate has limited capacity to stabilize oil-induced inflation but substantial capacity to stabilize the consumption tails (the 90/10 spread). A dual mandate over inflation and unemployment mainly stabilizes real activity, with gains partly extending to the untargeted consumption and income spreads. Adding even a small weight on the consumption spread to the dual mandate removes much of the remaining spread loss at a modest aggregate cost, borne mostly through unemployment; stabilizing the spread further comes at a convex cost. Fiscal instruments extend how far consumption inequality can be stabilized, but do not lower its price in achieving it. Finally, stabilizing the tails does not stabilize the whole distribution: a target on all ten consumption deciles relative to aggregate consumption is harder to meet with the policy rate, but benefits more from fiscal policy.

Keywords: *Monetary policy; Fiscal policy; Policy instruments; Heterogeneous agents; Inequality; Bayesian VAR; Oil supply shocks; Inflation*

JEL Classification: *D31, E52, E61, E62, E21, E32*

*I thank Christian Bayer, Benjamin Born, and my many incredible 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

Monetary and fiscal policy are the predominant vehicles for navigating the economy through aggregate shocks. Their primary objective has traditionally been to contain the recessionary and inflationary pressures arising from such shocks. Shocks may also propagate heterogeneously across households, due to differences in their income and wealth composition, and marginal propensities to consume (Auclert, 2019; Bayer, Born, and Luetticke, 2024). This introduces a distributional dimension to the policy problem, which has recently received more attention (Acharya, Challe, and Dogra, 2023; Bhandari et al., 2021; Del Negro et al., 2025; Le Grand, Martin-Baillon, and Ragot, 2026; McKay and Wolf, 2026). This paper focuses on a narrower question and asks: How much can monetary and fiscal policy stabilize consumption inequality, and at what cost to inflation and unemployment stabilization?

I answer this question through a dynamic empirical model (a VAR) that jointly tracks policy and macroeconomic aggregates with a parsimonious representation of the joint distribution of household consumption, income, and wealth (Bayer, Calderon, and Kuhn, 2025). By incorporating the joint distribution in the model, it allows me not only to target consumption inequality directly, but also to trace how policy affects consumption across households with different combinations of income and wealth. An oil supply shock à la Känzig (2021) defines the setting and whose effects policy seeks to alter. Policy instruments comprise the short-rate and two fiscal instruments of taxes and government spending. These instruments will differ in their targeted and untargeted aggregate and distributional effects, which expands the set of outcomes policymakers can achieve. I model policy choices under alternative mandates that trade off these outcomes and construct the corresponding counterfactual policy interventions (Caravello, McKay, and Wolf, 2026).

The paper presents several findings. First, the oil supply shock presents a trade-off: inflation rises, output contracts and unemployment rises, while consumption, measured relative to its mean, falls more near the top of the distribution than at the bottom, as in Calderon (2025). Second, comparing a single mandate over inflation to a single mandate targeting the consumption 90/10 spread, I find conventional monetary policy can better stabilize the consumption spread than oil-induced inflation. Third, the dual mandate's stabilization of real activity carries over to *untargeted* consumption inequality, but also income inequality, consistent with Acharya, Challe, and Dogra (2023). Fourth, the price of distributional stabilization is convex: a small weight on consumption inequality already achieves substantial stabilization at relatively little aggregate cost. Fifth, adding fiscal policy through taxes and government spending achieves greater distributional stabilization, but does not do it at a better price. With or without fiscal instruments, most of this cost falls on unemployment stabilization, as little inflation stabilization is achieved in the first place.

These findings center around a rather reduced description of consumption inequality, focused on stabilization of the consumption tails. This does not, however, necessarily stabilize the entire consumption distribution. I therefore also consider a second objective I call *risk sharing*, which measures differences in consumption responses across all deciles along the consumption distribution relative to aggregate consumption. With the policy rate alone, a triple mandate targeting the consumption 90/10 spread achieves a spread loss of 0.22, but a risk-sharing loss of 0.53.[1] This difference reflects the greater difficulty of stabilizing the entire consumption distribution rather than a single distributional moment. With both monetary and fiscal policy, targeting risk sharing directly achieves the same 90/10 spread loss of 0.22, but reduces risk-sharing loss to 0.31 vs. the 0.53 above, suggesting that fiscal policy extends stabilization beyond the tails to the broader distribution. Altogether, following an oil supply shock, a small weight on consumption inequality already goes a long way in terms of distributional stability; fiscal instruments let policy stabilize the tails further, but not more cheaply; and stabilizing the entire consumption distribution is more difficult, but benefits more from fiscal policy.

Turning to the details of the paper, the analysis is centered around how policy affects consumption inequality and households more generally. A prominent class of models makes the case why we should study the distribution over household outcomes at all; positing ex-ante identical households that become heterogeneous in income and wealth through idiosyncratic income risk, borrowing constraints, and access to different asset markets (Bayer, Born, and Luetticke, 2024). These differences shape how households consume and save, and therefore the very aggregate fluctuations that the *dual* view of stabilizing prices and output seeks to contain (Auclert, 2019). Therefore, the present paper adopts the same distributional view of the recent literature and demands for the analysis not only a representation of the marginal distribution of consumption, but also its joint distribution with income and wealth.

By approaching this empirically however, this invites the challenge of measuring and representing the business cycle fluctuations of the joint distribution in finite form. For this, I represent the distributional component of the model with quarterly functional distributional factors, generated using the methodology of Bayer, Calderon, and Kuhn (2025), which capture in a parsimonious way the requisite high-frequency business-cycle variation of the joint distribution of U.S. household consumption, income, and wealth. Embedding these factors in the VAR lets me recover the distributional consequences of the shock and of policy, and also *target* them.[2] This extends the conventional mandates to ones with a distributional objective, from

[1] The spread loss is measured as the sum of squared impulse responses under the counterfactual path divided by the sum of squared responses under the observed path over a finite horizon.

[2] Because it is linear in the factor responses to first order, it can be targeted in the same way as inflation and unemployment.

a single mandate over consumption inequality to a triple mandate that weighs it alongside inflation and unemployment.

A distributional objective adds something to that view only under two conditions. First, the oil supply shock must have heterogeneous incidence: if it changes every household's consumption in the same proportion, targeting the distribution on top of a dual mandate is redundant (McKay and Wolf, 2026). Second, the instruments must have distinct effects on the targets (acknowledged in Le Grand, Martin-Baillon, and Ragot, 2026). If the oil supply shock redistributes and policy does not, policy cannot correct the shock's distributional incidence. An economy hit by an oil supply shock meets this first condition. Oil supply shocks threaten stagflation (Känzig, 2021), but also affect households unequally (Calderon, 2025; Del Canto et al., 2025; Känzig, 2026). To address the second condition, I use externally identified policy shocks for the short rate, government spending, and taxes (Aruoba and Drechsel, 2024; Auerbach and Gorodnichenko, 2012; Bauer and Swanson, 2023; Romer and Romer, 2010). I show that each lever has distinct aggregate and distributional effects.

With the setting defined, the resulting VAR is large. Furthermore, the system is modeled with 16 quarterly lags to accommodate the slow propagation of oil supply shocks (Hamilton and Herrera, 2004; Kilian and Zhou, 2023). The long lag length is a response to Barnichon et al. (2026), who caution against small VAR systems, but also the fact that they have appealing properties (De Graeve and Westermark, 2025). A Bayesian approach with equation-specific shrinkage makes the system feasible (Chan, 2022) and is known to work well for impulse-response estimation (Baumeister, 2025). The oil supply shock is ordered first (Plagborg-Møller and Wolf, 2021) and scaled to a 10% increase in the oil price on impact.[3]

Estimating the system generates dynamic causal effects of both the oil supply shock and policy shocks. Under the assumptions of linearity and that private behavior depends on policy only through its realized and expected paths, these dynamic causal effects are sufficient to construct counterfactual policy under alternative mandates (Caravello, McKay, and Wolf, 2026). For my analysis, I consider three counterfactual mandates: a single mandate over inflation, a dual mandate that adds unemployment, a triple mandate that adds a consumption inequality term. I also consider for one result a single mandate targeting consumption inequality exclusively to assess monetary policy's distributional stabilization capacity. Each mandate is evaluated across three different sets of policy instruments: **Span A** is the policy rate alone, **Span B** is the policy rate and taxes, and **Span D** is the policy rate with taxes and government spending.[4]

[3] In Calderon (2025), I show that not all oil supply shocks are the same: it is the Känzig (2021) shocks that raise *core* CPI, as opposed to oil supply shocks à la Baumeister and Hamilton (2019). I keep only posterior draws that imply a stable system.

[4] I find adding taxes alone vs. adding government spending alone (**Span C**) deliver similar stabilization of the *targeted* outcomes, similar to the result of Wolf (2025) for interest-rate paths and transfers. This is shown in

Results show the dynamic causal effects of an oil supply shock present a meaningful trade-off for policymakers: prices increase, GDP contracts and unemployment rises, while consumption, measured relative to its mean, falls more near the top of the distribution than at the bottom. Focusing on the situation where policymakers are limited to using only the policy rate to combat these effects (**Span A**), the policy rate has far greater capacity to stabilize consumption inequality than inflation. A single mandate over inflation removes only 6% of the inflation loss. The reason is the shape of the shock: inflation rises and then reverses within the policy window, so a policymaker restricted to one policy choice at time $t = 0$ cannot offset both phases of inflation. Given the consumption response is not as volatile, the policy rate has greater capacity to stabilize it: A single mandate over consumption inequality removes about 80% of the spread loss, and it also stabilizes unemployment, output, and income inequality, which it does not target. A dual mandate central bank gives up inflation stabilization for unemployment stabilization, removing 68% of the unemployment loss (84% with all instruments), and the consumption-spread loss falls by a third as a spillover effect.

Fiscal instruments would extend the reach of distributional stabilization. Under the triple mandate, the consumption-spread loss falls from 0.22 with the policy rate alone to 0.16 with taxes and 0.11 with all instruments. The extension offered by fiscal policy is more visible when considering a more comprehensive target of consumption stabilization I call *risk sharing*, which measures differences in consumption responses across all deciles along the consumption distribution relative to aggregate consumption. With the policy rate alone, targeting the risk-sharing measure achieves 0.50 loss; 0.39 with taxes and 0.31 with all instruments. It achieves this by giving up some stability at the tails, but still the consumption spread remains substantially stabilized $(0.22, 0.16, 0.11) \to (0.30, 0.26, 0.22)$.

Thus far, we have described the capacity of monetary and fiscal policy to stabilize consumption, but this distributional stabilization comes at a price. Moving from the dual mandate that stabilizes prices and unemployment to the triple mandate that targets the $90/10$ consumption spread removes 67% of the dual mandate's consumption-spread loss with the policy rate alone, 80% with taxes, and 88% with all instruments, at an aggregate cost of 11%, 28%, and 38%. Most of the gain comes at small weights: with a weight of 0.2 on the spread (vs. a weight of 1.0 in the triple mandate), the mandate already removes 59–77% of the dual mandate's spread loss at a cost of 7–14%. Most of the cost is unemployment stabilization given up: inflation accounts for 12% of it with the policy rate alone and 26–35% with fiscal instruments. Around the dual mandate, the trade-off is therefore between the distribution and real activity, not inflation. The reason is because the triple mandate can only give back stabilization the dual mandate achieved, and with the policy rate alone that is unemployment stabilization.

Appendix C.

The tension between distributional stabilization and price stability emphasized by Bhandari et al. (2021), Bilbiie (2025), and Del Negro et al. (2025) becomes operative only once fiscal instruments enter the policy toolkit. They expand the set of attainable outcomes and attain greater inflation stabilization under the dual mandate: its inflation loss is 1.28 with the policy rate alone and 1.02 with all instruments. With greater stabilization achieved, moving to the triple mandate can then mean giving up some of that stabilization; hence, the rise in the inflation cost share ($12\%_A \to 26\%_B \to 35\%_D$). This would not be the only tension, however. By modeling also consumption with income and wealth, I find targeting consumption inequality does not necessarily stabilize income and wealth inequality. Under the triple mandate where we target the consumption 90/10 spread, the untargeted income-spread loss is *above* the dual mandate's, and the wealth spread stays near its realized level or over for most mandates and policy spans considered.

Finally, I ask how much of these gains survive uncertainty about policy transmission. Thus far, the policymaker selects one bundle for *each* posterior draw. To model uncertainty around economic paths, I force the planner to choose one bundle for the *whole* posterior. In this scenario, the inflation gains largely vanish and unemployment stabilization shrinks but survives. The gains in distributional stabilization from moving from a dual to a triple mandate shrinks substantially as well. Remarkably, under the dual mandate, the consumption spread and risk sharing measure under uncertainty *improve* as untargeted outcomes vs. dual mandate with full information, but the gains do not extend at all to inflation or unemployment—they become worse.

The remainder of the paper is organized as follows. Section 2 situates the paper in the literature on optimal policy and inequality. Section 3 describes the data. Section 4 sets out the empirical framework and the counterfactual mandates. Section 5 presents the results, and Section 6 concludes.

2 Literature Review

The results in this paper speak to the literature on policy design in heterogeneous-agent economies. In making policy prescriptions for economic stabilization, the literature has identified five key elements: (1) the distributional incidence of the shock under consideration; (2) the set of instruments available to policy makers in handling the aggregate shock and the extent to which they can adjust them at the relevant time; (3) the distributional incidence of this set of available policy instruments, which may further compromise or insure households; (4) the objective of the planner in structural normative analyses; and (5) the supply block of the model, which governs the aggregate trade-offs associated with policy intervention. In the following, I

first discuss the most closely related papers and then other related literature.

Relevant literature The literature central to this paper asks how household heterogeneity changes the design of stabilization policy. Several prominent papers have contributed to this strand. To isolate the insurance capacity of monetary policy, Acharya, Challe, and Dogra (2023) study optimal monetary policy in an analytical HANK model with CARA utility, access to a riskless real bond and no hard borrowing constraint. They find when idiosyncratic earnings risk is acyclical or countercyclical, consumption risk is countercyclical, giving the planner an additional motive to stabilize the *level* of output. Relative to RANK, the planner therefore places more weight on the *level* of output and less on price stability and the output gap.[5] Bilbiie (2025) also posits an analytical HANK model where inequality enters the policymaker's objective *directly*, generating an additional inequality-stabilization motive and leading the planner to tolerate greater inflation volatility than in RANK.

Bhandari et al. (2021) focus on the *insurance* role of stabilization policy in an incomplete-markets economy. Following a positive markup shock, the optimal HANK response can reverse the standard RANK prescription: because the shock raises dividends and lowers real wages, the planner cuts interest rates to support wages and insure workers, despite the resulting inflation. Following a TFP shock, heterogeneity in nominal asset positions creates a distinct insurance problem between borrowers and lenders; the planner uses surprise inflation and a real rate below the natural rate to lower the ex-post real return on nominal debt. When sufficiently flexible fiscal instruments are available, however, these missing insurance markets can be addressed more directly, shrinking the redistributive role assigned to monetary policy.

A second strand asks how the set and incidence of available policy instruments shape what the planner can implement. Wolf (2025) establishes under general conditions, an output-inflation path implementable through interest-rate policy can also be implemented through time-varying transfers, even though the two policies need not have the same distributional incidence. McKay and Wolf (2026) provide a complementary perspective and ask whether monetary policy alone has enough distributional incidence to justify using it. They find that adding a distributional objective (stabilizing households' consumption shares) changes only modestly the optimal monetary response relative to a dual mandate. This arises because a monetary easing raises consumption *uniformly* across households, while flexible goods prices and a fixed labor share make the supply block essentially RANK-like. Under sticky prices, where markups and the labor share fluctuate, the optimal policy response deviates more from the dual mandate. They find fiscal transfers are more effective than monetary policy at stabiliz-

[5] Related to element (4) above, in their baseline, they allow the planner to intervene and use fiscal instruments at date-0 to abstract from pre-existing wealth inequality; a common approach in the optimal policy literature to avoid the otherwise redistributive inflation bias through surprise rate cuts by the planner (Dávila and Schaab, 2023); but also the focus of the paper is on insurance over consumption risk and not policy-driven redistribution.

ing distributional outcomes since they have progressive incidence. Le Grand, Martin-Baillon, and Ragot (2026) ask whether fiscal policy can replicate the distributional incidence of monetary policy. They show with sticky prices/wages, a rich, flexible set of fiscal instruments determines the extent to which monetary policy should concern itself with redistributive matters.

A related literature studies how household heterogeneity shapes preferences over stabilization policy. Gornemann, Kuester, and Nakajima (2022) show that household preferences over inflation versus unemployment stabilization vary systematically with labor-market exposure and household balance sheets: wealth-rich or retired households tend to prefer inflation stabilization, whereas wealth-poor working households place relatively greater value on unemployment stabilization. Del Negro et al. (2025) emphasize that these distributional trade-offs also depend on the aggregate shock under consideration. In their estimated HANK model, supply shocks disproportionately hurt poorer households through wages and employment while benefiting wealthy households through profits, making aggressive inflation stabilization particularly costly for the consumption of poorer households. Demand shocks instead have relatively larger effects on wealthy households through their exposure to procyclical profits. Bayer et al. (2025) use a two-country HANK model to show how heterogeneity across households can rationalize observed differences in debt policy, providing a political-economy explanation beyond simply attributing them to national preference.

Distributional transmission The incidence of policy varies between households because they differ in their marginal propensities to consume, exposure to labor and capital income, and net wealth composition, among other dimensions (Auclert, 2019; Auclert, Rognlie, and Straub, 2024; Kaplan, Moll, and Violante, 2018). These distributional responses can in turn amplify or dampen aggregate responses, the precise objects current mandates are specified over. This speaks to a broader literature studying how household distributions shape the transmission of aggregate shocks and, in turn, influence macroeconomic dynamics (Bayer, Born, and Luetticke, 2024; Luetticke, 2021); see Auclert, Rognlie, and Straub (2025) for an overview. A complementary literature measures the distributional effects of monetary policy in reduced-form (e.g., Andersen et al., 2023; Cloyne, Ferreira, and Surico, 2020; Holm, Paul, and Tischbirek, 2021). More recently, functional representations of distributions over micro-units have been embedded directly in VARs (Bayer, Calderon, and Kuhn, 2025; Chang, Chen, and Schorfheide, 2024; Chang and Schorfheide, 2026; Lenza and Savoia, 2024).

Oil supply shocks and monetary policy This paper operates in the context of an oil supply shock. Oil supply shocks confront the central bank with a genuine trade-off between inflation and activity, so that rules are forced to disagree (Hofmann, Manea, and Mojon, 2025). Their

incidence is also heterogeneous across households. Känzig (2026), Del Canto et al. (2025), and Calderon (2025) document the unequal incidence of oil and inflation shocks, Pallotti et al. (2024) measure who bore the 2021–2023 inflation, and Auclert et al. (2024), Chan, Diz, and Kanngiesser (2024) and Bobasu, Dobrew, and Repele (2025) study the monetary response to an energy shock in heterogeneous-agent models, finding that credit constraints make aggressive tightening more costly and that low-wealth households lose through both labor income and direct price exposure. Del Negro et al. (2025) also contribute to this literature, showing an aggressive inflation response helps poorer households under demand shocks, but harms them under supply shocks. The results below, among other things, can speak to these predictions.

The setting also connects to the longstanding literature on oil shocks and monetary policy. Bernanke, Gertler, and Watson (1997) argued that much of the output cost of oil shocks was the Federal Reserve's own systematic response; Hamilton and Herrera (2004) and Kilian and Lewis (2011) disputed the counterfactual behind that claim, and a recent sufficient-statistics literature has revisited it with identified shocks (Bjørnland, Cross, and Hölz, 2025; Castelnuovo, Mori, and Peersman, 2024; Ider et al., 2026), alongside the structural treatments of Bodenstein, Guerrieri, and Kilian (2012), Nakov and Pescatori (2010) and Gagliardone and Gertler (2026).

Methodology Methodologically, the paper combines recent work on semi-structural monetary policy counterfactuals with advances in high-dimensional Bayesian VARs. Counterfactual outcomes are constructed from estimated impulse responses without specifying a full structural model, while accounting for the Lucas critique (Barnichon and Mesters, 2023; Caravello, McKay, and Wolf, 2026; Chan, 2022; McKay and Wolf, 2023). The VAR system is in the class of long-lag VARs (De Graeve and Westermark, 2025) and regularized following a Bayesian approach (Chan, 2022).

3 Data

The analysis covers several data sources that define the set of variables in the VAR. They can be described as spanning three blocks: a block of policy shocks, a block of macroeconomic aggregates, and a distributional block. Ordered above the three blocks is an oil supply shock measure (Känzig, 2021), which is the non-policy shock bludgeoning the economy. The policy block includes one externally constructed conventional monetary policy shock and two fiscal policy shocks spanning taxes and government spending. The third block of variables is characterized by macroeconomic aggregates, which includes the oil-market block of Känzig (2021) together with some U.S. macroeconomic aggregates from FRED-QD (McCracken and Ng, 2021). These also include the policy aggregates over which the mandates are defined.

The final block concerns the distribution over households' consumption, income, and wealth, whose parsimonious representation is a set of principal components, characterizing the business cycle fluctuations of the distribution. The sample is post-Volcker and pre-pandemic: it runs from 1983Q1 to 2019Q3 at quarterly frequency. Each source is described in turn.

3.1 Oil supply shock

The oil supply shock I use is from Känzig (2021), defined as the common component of WTI crude oil futures price changes in a narrow window around OPEC announcements. These would be price changes of contracts spanning the first-year term structure. This shock is treated as capturing news regarding the future supply of oil by OPEC and not an observed physical supply disruption as in Baumeister and Hamilton (2019). The surprises are aggregated to quarterly frequency by summing within each quarter, and the series enters the VAR ordered first, so that the identified oil shock is ultimately the innovation of the shock equation and the policy levers below are orthogonalized against it contemporaneously. In a different paper (Calderon, 2025), I show that not all oil supply shocks are alike and it is precisely the Känzig (2021) oil shock that is inflationary: core CPI and the other price components respond positively and credibly, real activity contracts, and it has distributional consequences. Because of this, it creates competing targets: a policy that stabilizes one need not stabilize the other.

3.2 Policy instruments

To first order, to identify counterfactual paths, two sets of statistics are sufficient: the dynamic causal effects of the aggregate shock of interest, which trace the baseline response under the prevailing policy rules, and the causal effects of policy interventions, which determine how different policy levers can alter that baseline. The oil supply shock provides the first set. Identifying the second is more challenging because monetary and fiscal authorities respond endogenously to economic conditions when using their policy instruments in pursuit of their mandates. Observed policy changes are therefore not random.

To then recover the causal effects of the different policy levers, I rely on the literature that has thought about this endogeneity problem more carefully (Aruoba and Drechsel, 2024; Auerbach and Gorodnichenko, 2012; Bauer and Swanson, 2023; Hack, Istrefi, and Meier, 2025; Romer and Romer, 2010). Relying on this literature provides me a set of externally identified shocks that have a mapping to a policy lever. They are discussed in turn.

3.2.1 Monetary policy

The central bank, along the business cycle, relies on several instruments to implement specific policy for economic stabilization. The set of instruments will vary between monetary unions, but the monetary-policy literature mainly distinguishes between two classes of instruments: conventional monetary policy and unconventional monetary policy. The paper focuses on the first and selects one externally constructed monetary policy shock that represents *variation* in the policy rate. For the monetary block in the VAR, the objective will be to span this policy lever using this shock, so the planner if she so wishes can deviate from the prevailing policy plans.[6]

Conventional monetary policy These are changes in the current target for the policy rate. This would be the conventional monetary policy instrument.[7] These changes are typically empirically measured from the reaction of federal funds futures around policy announcements, which isolates market revisions over the near future policy rate (Gürkaynak, Sack, and Swanson, 2005; Kuttner, 2001). Although the central bank achieves these discretionary movements through changes in their administered rates and physical operations—interest on reserve balances, overnight reverse repurchase, the discount window, and open market operations—I view all as a means to achieve changes in the short-term rate.

With that, to identify variation in the short-rate, I use one series built from two shocks. Up to 2007Q4, I use the narrative shock of Aruoba and Drechsel (2024), purged of predictable systematic-policy movements, leaving a cleaner measure of discretionary monetary innovations for the counterfactual exercise (Hack, Istrefi, and Meier, 2025). From 2008Q1, I splice it with the high-frequency surprises of Bauer and Swanson (2023), measured from FOMC announcement-window changes in interest-rate futures and orthogonalized to the economic and financial news available before the announcement.[8] Each piece is standardized over its own window.[9]

[6] Other levers may include reserve requirements as a separate bank-liquidity instrument and targeted lending operations (TLTROs), which again affect liquidity by offering favorable loans to banks to increase their lending.

[7] While central banks have long *influenced* short-term interest rates, the Federal Reserve set federal funds rate targets before October 1979, operated a nonborrowed-reserves procedure between 1979 and 1982, and returned to targeting the funds rate thereafter.

[8] The spliced shock generates virtually identical median dynamic causal effects as the Aruoba and Drechsel shock alone, but with greater posterior uncertainty around the estimate.

[9] Because the levers have partially overlapping coverage windows (AD-HIM ends in 2007Q4; the government spending lever begins in 1981Q4; the tax lever ends in 2007), I follow Caravello, McKay, and Wolf (2026) and Känzig (2021) and set observations outside each shock's native coverage window to zero before estimation.

3.2.2 Fiscal policy

Monetary policy redistributes and insures through its own aggregate and distributional incidence (Bhandari et al., 2021; Le Grand, Martin-Baillon, and Ragot, 2026), but may be incapacitated in the presence of shocks that induce significant aggregate and distributional tradeoffs. Both papers proposed flexible and rich sets of fiscal instruments on top of conventional monetary policy, finding fiscal policy can replicate the distributional incidence of monetary policy and thus, insure markets to the extent where monetary policy could not. To examine how much of this role can be achieved with a more limited fiscal toolkit, I incorporate two fiscal policy levers into my framework: government spending and taxes.

Government spending The spending lever is the forecast error of Auerbach and Gorodnichenko (2012): the realized growth of real federal government consumption and gross investment in a quarter minus the growth the Survey of Professional Forecasters expected for that quarter one survey earlier. It measures spending that actually took place and that forecasters did not anticipate, so it captures spending as it is implemented rather than news about future spending (Ramey, 2011). What this implies for the counterfactual, which treats policy as news about future paths, is discussed in Section 4.5. The raw surprises are larger when the macroeconomic news is bad, which is stabilizing spending that the lever should not carry. I therefore purge the series of the news forecasters received within the quarter, namely the SPF revisions to the current- and next-quarter forecasts of real GDP, the GDP deflator and unemployment, and use the residual of this regression (estimated over 1981Q4–2019Q4) as the lever. The purged series correlates at 0.96 with the raw surprise. Coverage starts in 1981Q4, and a positive value is more spending than expected.

Taxes The tax lever is the series of exogenous legislated tax changes of Romer and Romer (2010), in percent of GDP, where a positive value is a tax increase. Romer and Romer classify every postwar tax change by its motive and keep those taken to reduce an inherited deficit or to raise long-run growth, making their use unrelated to current or prospective economic conditions. The series ends in 2007, with the last event in 2005. Exogenous motives do not make a change unanticipated. Each change is dated by the quarter in which it takes effect, and many were legislated earlier, so the lever is closer to an implemented tax change than to a time-zero announcement.

3.3 Macro variables

The macro block complements the oil supply shock and the policy block. First, the macro block respects the standard oil-VAR specifications seen in the literature (Baumeister and Hamil-

ton, 2019; Känzig, 2021; Kilian, 2009). The specification is also augmented with policy aggregates: U.S. GDP, unemployment, the price level, and the policy rate. Together, they consist of eight series: 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, the GDP deflator, and the effective federal funds rate. The real oil price, oil production, oil inventories, world industrial production, U.S. real GDP, and the GDP deflator enter in log levels and are scaled by 100; unemployment and the federal funds rate enter in percent. In addition, the producer price index for all commodities (PPI all commodities, in logs) enters as a control. It is ordered directly after the oil supply shock and before the policy levers, so that the policy surprises are orthogonal to same-quarter movements in commodity prices, the standard remedy for the price puzzle (Sims, 1992).

3.4 Distribution of Household Consumption, Income, and Wealth

To represent the distributional component of the model, I use the quarterly functional time-series data on the joint distribution of U.S. household consumption, income, and wealth constructed by Bayer, Calderon, and Kuhn (2025). Their methodology synthesizes microdata at the household-level with macroeconomic indicators of higher-frequency (McCracken and Ng, 2021) in a Bayesian state-space framework to produce a smooth and high-frequency representation of the joint distribution. The quarterly series are therefore estimates from a mixed-frequency state-space model, not quarterly observations from a joint survey. The VAR conditions on their point estimates and treats the factors as observed, as they are consistent estimates. Consumption, income and wealth are measured per household; income is before taxes and wealth is net worth; each enters relative to its cross-sectional mean in inverse hyperbolic sine units, so the distribution is scale-free. The distributional block will ultimately have a parsimonious representation: a set of principal components characterizing the business cycle fluctuations of the joint distribution of consumption, income, and wealth. How the data are constructed is explained in Appendix G. How consumption, income, and wealth are operationalized is described in Appendix H.

Representation in BVAR What will be important to know is how these data are represented in the Bayesian VAR. From Bayer, Calderon, and Kuhn (2025), I know the data generation process of the joint distribution(s) estimated here and thus also know the first $r = 8$ principal components, $\boldsymbol{F}_t = (F_{1t}, \ldots, F_{rt})'$, capture nearly all of the business-cycle variation of the consumption, income, and wealth distribution in the chosen basis representation. It is precisely these eight factors that enter the baseline Bayesian VAR. The treatment of the distributional

data thereafter ensures impulse responses still retain their properties as a distribution.[10] Appendix I describes how these factors are mapped back into distributional statistics, including the consumption spread that enters the counterfactual policy objectives.

4 Methodology

This section begins by describing the dynamic empirical model that combines the various blocks of time-series data described in Section 3 to estimate dynamic causal effects of oil supply shocks á la Känzig and the three policy levers (Section 4.1). The system is large in the number of parameters it estimates relative to the number of observations pinning them down, so I render the estimation feasible using a Bayesian approach (Section 4.2). This is followed by a discussion on identification of the causal effects (Section 4.3). The last section concludes by taking these causal effects as inputs for the counterfactual methodology of Caravello, McKay, and Wolf (2026) (Section 4.5). Assumptions and limitations are explained therein, as well as the policy rules we plan to explore in Section 5, when we discuss the results.

4.1 The multi-shock Bayesian VAR

To evaluate how the economy would respond to alternative policy paths relative to the prevailing rules, I follow one of the proposed frameworks of Caravello, McKay, and Wolf (2026). To first order, these alternative paths are identified using two sets of sufficient-statistics, sufficient under the path-sufficiency and linearity assumptions of Section 4.5: the dynamic causal effects of an aggregate shock of interest, which trace the baseline response under the prevailing policy rule and the causal effects of policy interventions, which capture how the baseline can be altered through policy. I identify both within the same model, falling in the class of multi-shock VARs. To restore regularity conditions from the potentially dense parameterization, I adopt a Bayesian approach for the estimation. The section describes each part in turn.

Model Let $y_t \in \mathbb{R}^{n_y}$ denote the variables in the VAR system. The variables are stacked to have the oil supply shock first, the commodity price control second, the $K = 3$ policy shocks third, the aggregate variables fourth, and the distributional factors last. To identify the two sets of sufficient statistics, we need to estimate the structural impulse-response matrix Θ_h. To

[10] Specifically, the reconstruction from factor estimates to distributional data preserves the properties of the underlying distributional objects across the impulse-response horizon $h \in 0, \ldots, H$: the reconstructed marginal quantile functions $\hat{\Xi}^{-1}_{m,h}$ are weakly monotone in $u \in [0,1]$ for each m; the reconstructed copula density $\hat{c}_h$ is non-negative on the trimmed cube $[\varepsilon, 1-\varepsilon]^d$; and the copula integrates to unity, $\int [0,1]^d \hat{c}_h = 1$. Uniform marginals follow from restrictions on the coefficients. Monotonicity and non-negativity are verified ex post on a G^d tensor grid; deviations from the strict constraints are below 10^{-6} and treated as numerical noise.

estimate Θ_h, I first specify the reduced-form BVAR

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

where c is a constant, $B_1, \ldots, B_p$ the lag coefficient matrices, and v_t the reduced-form innovation, with covariance matrix Σ.

The proposed system has $n_y = 21$ variables and $p = 16$ quarterly lags.[11] This is to accommodate the dynamic interactions between variables following a shock, but also (delayed) responses that activate along the shock's transmission. This consideration is particularly made for oil shocks: Hamilton and Herrera (2004) and Kilian and Zhou (2023) show that larger, longer lag systems are needed to allow for the building and contraction of oil market cycles. Furthermore, long-lag VARs have appealing properties, potentially reducing both the bias and variance of impulse response estimates relative to shorter lag VARs and have better coverage rates as well (De Graeve and Westermark, 2025).[12] Taken together, these considerations motivate the use of a long-lag VAR in the present application.

However, there is no free lunch. With $n_y = 21$ variables and $p = 16$ lags, the system is large relative to the amount of data, violating regularity conditions for OLS. Each equation has $21 \times 16 = 336$ lag coefficients plus an intercept, more than twice the number of quarters in the 1983–2019 sample. Shrinkage is therefore what makes the system estimable: the prior, not the data, pins down most of the distant-lag and cross-variable coefficients. Consequently, the unrestricted equation-by-equation OLS problem is underdetermined: there are more regressors than usable observations, and the coefficient estimates are not uniquely identified by the sample alone. Conceptually, it may therefore benefit the system to embed some prior information from our economic understanding of these variables. I therefore adopt a Bayesian approach, using priors to embed this information; this regularizes the estimation and controls the uncertainty associated with the large number of coefficients.

[11] I follow comments from Baumeister (2025) that lag order selection via AIC or BIC targets forecast loss, which may yield specifications that are suboptimal for structural inference. I simply choose a specification that spans the horizon considered.

[12] One of the empirical applications of De Graeve and Westermark (2025) concerned monetary policy and demonstrated there too that the choice of lag length can materially alter previously established conclusions. Antolín-Díaz and Surico (2025) provides another application of a long-lag VAR.

4.2 The Bayesian Setup

Following Chan (2022), I write the VAR in a recursive structural form,

$$\begin{aligned} A_0 y_t &= \widetilde{c} + \widetilde{B}_1 y_{t-1} + \cdots + \widetilde{B}_p y_{t-p} + e_t, \qquad e_t \sim \mathcal{N}(\mathbf{0}_{n_y}, D_e), \\ y_{i,t} &= \sum_{j<i} a_{ij} y_{j,t} + \widetilde{c}_i + \sum_{\ell=1}^{p} \widetilde{b}'_{i\ell} y_{t-\ell} + e_{i,t}, \end{aligned} \tag{2}$$

with $D_e \equiv \operatorname{diag}(\sigma_1^2, \ldots, \sigma_{n_y}^2)$ the diagonal covariance matrix of the structural disturbances. Let the wide tildes indicate parameters from the structural equation. A_0 is unit lower triangular; to match the positive sign in the equation-by-equation representation, its (i, j) entry is $-a_{ij}$ for $j < i$. With this, the structural VAR is a set of n_y independent regressions, so the likelihood factorizes equation by equation. The triangular structure also makes this recursive: equation i contains contemporaneous variables ordered before it, $y_{j,t}$ for $j < i$. The coefficients a_{ij} capture these contemporaneous effects, while $\widetilde{b}'_{i\ell}$ is row i of $\widetilde{B}_\ell$. Define $a_i \equiv (a_{i1}, \ldots, a_{i,i-1})'$ and $\theta_i \equiv \left(a_i', \widetilde{c}_i, \widetilde{b}'_{i1}, \ldots, \widetilde{b}'_{ip}\right)'$. Together, $\{\theta_i, \sigma_i^2\}_{i=1}^{n_y}$ parameterize the equation-by-equation posterior.

Priors Each equation receives a conjugate Normal–Inverse-Gamma prior,

$$\sigma_i^2 \sim \mathcal{IG}(\nu_i, S_i)\,, \qquad \theta_i \mid \sigma_i^2 \sim \mathcal{N}\big(\bar{\theta}_i,\ \sigma_i^2 V_i\big), \tag{3}$$

with shape parameter $\nu_i = 1 + i/2$ and scale $S_i = s_i^2/2$, where s_i^2 is the residual variance from a univariate AR(p) for variable i. Thus, each equation, through S_i, adapts to the residual variability of that series.[13]

The lag-coefficient block of θ_i follows a standard Minnesota structure, shrinking each equation toward an AR(1) in its own first lag with coefficient δ_i, a random walk when $\delta_i = 1$. The prior mean on equation i is δ_i on the first own lag and zero elsewhere, while the corresponding diagonal element of V_i, denoted $V_{i,j\ell}$, is

$$V_{i,j\ell} = \begin{cases} \kappa_{1,i} \,/\, \left(\ell^{2\lambda_3} s_i^2\right) & j = i, \\ \kappa_{2,i} \,/\, \left(\ell^{2\lambda_3} s_j^2\right) & j \neq i. \end{cases} \tag{4}$$

The parameters $\kappa_{1,i}$ and $\kappa_{2,i}$ govern shrinkage on own and cross-variable lags, respectively; λ_3 shrinks distant lags. They describe the system's hyperparameters and are chosen to maximize equation i's closed-form marginal likelihood as in Giannone, Lenza, and Primiceri (2015). For the estimation thereafter, they are then held fixed. Finally, the contemporaneous coefficients a_{ij}

[13] See Chan (2022) Proposition 1 for more details.

are also normally distributed, centered at zero with prior variance σ_i^2/s_j^2; the intercept receives a diffuse prior.

Under the standard conjugate Normal–Inverse-Wishart prior, the variance has the Kronecker form $\Sigma \otimes \Omega$, which forces one shrinkage pattern on every equation. This means the likelihood may be dominated by the prior due to excess shrinkage or vice versa. Adopting the equation-specific prior relaxes this restriction by allowing both the prior mean and prior tightness to vary with i, which matters here because the three blocks (shocks, macro, distribution) of y_t behave very differently. The measured shock series are near white noise by construction, so their equations have $\delta_i = 0$. Macroeconomic aggregates are persistent, and the distributional factors are stationary; for these, δ_i is selected from a grid $\{0, 0.1, \ldots, 1\}$. For the very persistent time-series, for example the price level, this lets the data decide how close to a unit root the equation is, and the stability filter below then keeps only draws inside the unit circle.

Sampling Because the prior is conjugate equation by equation, each (σ_i^2, θ_i) has a closed-form Normal–Inverse-Gamma posterior, and the n_y equations are sampled independently. This ability to handle equations one by one holds given the hyperparameters and the recursive ordering; it is a property of the posterior, and does not mean the structural innovations are economically independent. Draws are therefore i.i.d.: there is no Markov chain and no burn-in. From each draw, I rebuild the structural parameters $(A_0, \widetilde{c}, \{\widetilde{B}_\ell\}_{\ell=1}^p, D_e)$, and keep the draw only if the companion matrix of the VAR is stable i.e., the absolute values of all eigenvalues are within the unit-circle. I report the share of draws the filter rejects. Each posterior draw also delivers the identification directly, since I draw the impact matrix $P \equiv A_0^{-1}\,\mathrm{diag}(\sigma_1, \ldots, \sigma_{n_y})$ directly. The estimation retains $N_{\text{draw}} = 5{,}000$ draws.

4.3 Identification of Causal Effects

The measured shock series are constructed outside the VAR: the oil supply news shock follows Känzig (2021), while the K policy shocks come from their respective external identification strategies. Following the internal-instrument approach (Plagborg-Møller and Wolf, 2021; Stock and Watson, 2018), I place these series, together with the commodity price control, in the first $K + 2$ positions of y_t, so their innovations are included directly in the VAR information set. The oil shock is ordered first to impose that contemporaneous macroeconomic and policy innovations do not feed back into the externally measured shock series. Its impulse responses define the baseline response under the prevailing rule.

Is ordering a problem? Because the counterfactuals use linear combinations of the full policy block, any ordering of these shocks within would still span the same space of responses;

therefore the resulting aggregate and distributional counterfactual paths are invariant to ordering. Formally, within each posterior draw, a reordering of the policy block replaces the K policy columns of the impact matrix P by the same columns times an orthogonal $K \times K$ matrix Q, because both are factorizations of the same conditional covariance. The response matrices become $m_i^{\text{base}} Q$, their column space is unchanged, and the least-squares fit in (10) delivers the same counterfactual. This holds for the full policy block only.

However, the analysis below is over *individual* policy instruments and so analyzing a restricted span, a subset of the policy instruments, may change with the ordering. Comparisons of different policy instrument sets could thus depend on how the policy shocks are orthogonalized i.e., how they are residualized in the Cholesky decomposition, which then may elicit different aggregate and distributional incidence. In other words, different orderings may hand policymakers different levers.

I test this hypothesis in two ways. First, with three shocks, there are only $3! = 6$ Cholesky combinations. Each Cholesky will partial each shock differently, which again has implications about how this precise policy lever affects the set of attainable outcomes. The second way is re-solving the restricted spans in their original, unpartialed self: each lever is a one-standard-deviation innovation in its *own* shock series, with the other policy innovations held at zero, and its effects on the remaining variables follow from the reduced-form covariances. Doing so removes the dependence on the ordering of the policy block. Across both exercises, the targeted losses change little, which could be in part due to the low correlations between the levers. The ordering of the remaining variables does not affect their own responses to the identified oil and policy shocks.

The sufficient-statistics Define standardized structural shocks by $\varepsilon_t \equiv D_e^{-1/2} e_t$, so that $\text{Var}(\varepsilon_t) = I_{n_y}$ and $v_t = P\varepsilon_t$. The corresponding structural moving-average representation is

$$y_t - \mu = \sum_{h=0}^{\infty} \Theta_h \, \varepsilon_{t-h}, \tag{5}$$

where $\mu \equiv (I_{n_y} - B_\Sigma)^{-1} c$ and $\Theta_h \in \mathbb{R}^{n_y \times n_y}$ 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 structural shock $\varepsilon_{j,t}$.

For each posterior draw, I extract two slices of Θ_h that sufficiently characterize the counterfactual:

$$\Psi_i^{\text{oil}}(h) \equiv \Theta_{h,i,1}, \tag{6}$$

$$m_{i,k}^{\text{base}}(h) \equiv \Theta_{h,i,k+2}, \quad k = 1, \ldots, K, \tag{7}$$

i.e., the response of the i-th endogenous variable to the oil shock (Ψ_i^{oil}) and to the k-th policy shock ($m_{i,k}^{\text{base}}$) at horizon h. Impulse responses to the oil shock are rescaled per posterior draw so that the impact response of the real oil price increases by 10%. Responses to the K policy shocks are in units of a one-standard-deviation surprise in the corresponding instrument.

4.4 Are the levers economically credible and statistically distinct?

The policy counterfactuals of Caravello, McKay, and Wolf (2026) rely on combining the dynamic causal effects of identified policy shocks to alter the economy's response to an oil supply shock. For our exercise, these shocks must represent economically meaningful policy interventions. This requires not only identification (via selecting between high-frequency, narrative approaches, etc) but also impulse responses that, conditional on our estimated system, are broadly consistent with established empirical evidence. For example, if a policy rate hike generates an unexplained price puzzle, the policymaker could exploit this questionable response to construct counterfactual policies with possibly implausible implications. It is therefore not sufficient to rely on the identification and interpretation of shocks in the source papers; their estimated effects *within our system* must also be economically credible. As primary evidence, Appendix A, Figure A.1 reports these responses and what the plotted responses corroborate.

For the counterfactual machinery, this is irrelevant. The relevant object there to generate numerically stable counterfactuals is the matrix M_1 of target responses—responses to inflation, unemployment, and our consumption inequality term—in (10). An additional lever expands the policy span if it is linearly independent of the existing levers. The singular values and condition number characterize the rank of M_1 and conditioning of this matrix. Finally, Appendix Table A.1 provides secondary evidence on their *unconditional* correlation, showing that the three shock series are nearly uncorrelated ($|\rho| \leq 0.09$).

4.5 Policy counterfactuals: Assumptions, Limitations, and Estimation

Dynamic causal effects are sufficient to evaluate alternative policy rules (Caravello, McKay, and Wolf, 2026; McKay and Wolf, 2023) and their construction rests on one assumption about private behavior:

Path sufficiency Private behavior depends on policy only through the *realized and expected path* of the policy instrument $\{z_t, \mathbb{E}_t z_{t+1}, \mathbb{E}_t z_{t+2}, \ldots\}$ — not on the policy *rule* generating that path.

Observational equivalence Under path sufficiency, to first order, the equilibrium of an economy where the planner follows an *alternative rule* in response to the oil shock is identical to

the equilibrium under the *historical rule* hit by the same oil shock *plus* a time-zero vector of policy news shocks $m_{\text{aux}} \in \mathbb{R}^K$, where the size of the shocks is chosen to replicate the alternative rule's instrument path. Both produce the same realized and *expected* instrument paths; identical paths imply identical private behavior, hence observational equivalence. This means the counterfactuals can be constructed from estimated impulse responses alone. Now, four clarifications make the result operational.

First, the policy shocks must be news *announced at* $t = 0$. The equivalence requires identical *expected* paths, not only identical realized paths. A sequence of unanticipated surprises dripfed at $t = 1, 2, \ldots$ à la Sims and Zha (2006) could replicate the realized path, but agents at $t = 0$ would still expect the historical rule and would be systematically surprised every period — and thus not a rule — so behavior would differ.

Second, the estimated responses must capture how agents react to anticipated policy *plans*. Announced future paths must be related to short-rate changes, government spending, and/or taxes. Shocks that reflect information about economic fundamentals rather than about policy (e.g., delphic forward guidance) are not: a central bank cannot implement a rule by choosing what to reveal about the economy.[14]

Third, the counterfactuals provide the best linear approximation to the rule conditional on the set of available shocks. In practice, this means a mandate is a weighted least-squares fit of the target responses over $h = 1, \ldots, T_c$, so rules are enforced approximately, not literally each quarter. Enforcing a rule perfectly imposes a dynamic Tinbergen rule: for T horizons, you impose T constraints for *each* target path, which for our application of three targets, requires at least $T \times 3$ linearly independent policy-response paths to accommodate an arbitrary target.[15] With fewer independent paths (which may differ from the number of shocks), exact enforcement is possible only if the desired path lies within their span of course.

Fourth, the Lucas critique is mitigated to first order under these assumptions (Caravello, McKay, and Wolf, 2026). The re-optimization done by agents in response to anticipated future policy paths is embedded in the estimated responses to the identified policy shocks. Hence, under path sufficiency and linearity, matching the expected instrument path reproduces the corresponding first-order equilibrium response. This holds for local perturbations that agents find as credible as the historical ones; the effects of a change in the policy regime itself are not identified.

[14] The fiscal levers are identified policy surprises rather than direct announcements of future policy paths. Nevertheless, their estimated dynamic effects can capture revisions to expected fiscal paths, conditional on the information in our VAR, and thus provide economically meaningful policy interventions. Anticipated and unanticipated fiscal changes may have different effects, as emphasized by Ramey (2011) for government spending and Mertens and Ravn (2012) for taxes. Our counterfactuals should therefore be interpreted as first-order combinations of the estimated fiscal responses, rather than exact implementations of arbitrary announced fiscal plans.

[15] From exploratory estimations, perfect enforcement requires large movements in the instruments, which raises Lucas critique concerns. A similar comment was made in Hamilton and Herrera (2004) about Bernanke, Gertler, and Watson (1997).

Limitations Within each posterior draw, the construction treats each lever's estimated impulse responses as known, easily combined to fit a counterfactual path; estimation uncertainty enters through the distribution across draws (Section 5.5). This relies on three additional assumptions.

First, *certainty equivalence*. First-order decision rules depend only on expected paths, not higher moments. Superposition in (9) is valid because responses to (oil + policy news) add with no interaction terms at this order. This concerns the uncertainty agents face about future shocks and policy.[16] **Second**, *comparable credibility*: agents assign the counterfactual announcement the same credibility as the historical announcements underlying the estimated responses. **Third**, *policy uncertainty*. A true regime change alters both the conditional mean path of the instrument and potentially its future *variance* (a credible peg, for example, may eliminate policy uncertainty). The construction replicates only the expected path, which suffices to first order under certainty equivalence.

The equivalence is thus accordingly challenged if one wishes to consider policy-uncertainty channels, credibility and learning, and nonlinearity (large deviations, the effective lower bound, state dependence) — which is why rules are enforced over a finite window with the historical systematic policy reaction resuming thereafter. Strictly speaking, the historical rule is never "switched off": the bundle m_{aux} is chosen at $t = 0$, and at every horizon the economy evolves under the historical rule plus these shocks. *Resuming* means that no further deviations are imposed and the loss places no weight on horizons after T_c.

Policy counterfactuals: algebra Fix a posterior draw. Let T_c denote the number of quarters for which the planner adheres to the new policy path, and $T \geq T_c$ the horizon of the impulse responses. For each variable i, $\Psi_i^{\text{oil}} \in \mathbb{R}^T$ denotes the oil-shock response over each quarter of the horizon and $m_i^{\text{base}} \in \mathbb{R}^{T \times K}$ the matrix of responses to the K policy shocks, as defined above. A linear policy rule that the planner follows under the counterfactual is encoded by a collection of $C \times T$ target matrices $\{A_i\}_{i=1}^{n_y}$ such that the target is to satisfy

$$\sum_{i=1}^{n_y} A_i y_i^{\text{cnf}} = \mathbf{0}_C, \tag{8}$$

with $C = T_c$ for a rule that pins a single variable, $C = 2T_c$ for the dual mandate, etc. Throughout, $J_c \equiv [\,\mathbf{0} \;\; I_{T_c} \;\; \mathbf{0}\,]$ denotes the $T_c \times T$ matrix that selects horizons $h = 1, \ldots, T_c$ of a response. The bundle is a surprise at time zero, but the loss only counts horizons $h = 1, \ldots, T_c$. The mandate is enforced from one quarter out: the impact-quarter response is not part of the target, and the robustness appendix shows that adding it changes little.

[16] This is different from my uncertainty about the VAR parameters, which the posterior draws carry and Section 5.5 references.

Counterfactual responses y_i^{cnf} are obtained by adding to the original oil IRF a linear combination of the policy IRFs:

$$y_i^{\text{cnf}} = \Psi_i^{\text{oil}} + m_i^{\text{base}} m_{\text{aux}}, \qquad m_{\text{aux}} \in \mathbb{R}^K. \tag{9}$$

Substituting Eq. (9) into Eq. (8) and stacking yields the familiar least-squares solution, provided M_1 has full column rank:

$$\underbrace{\Big[\sum_i A_i m_i^{\text{base}}\Big]}_{\equiv M_1\,(C\times K)} m_{\text{aux}} = -\underbrace{\sum_i A_i \Psi_i^{\text{oil}}}_{\equiv M_2\,(C\times 1)},$$

$$\boxed{m_{\text{aux}} = -(M_1' M_1)^{-1} M_1' M_2.} \tag{10}$$

The closed-form formula gives us the (aux)iliary policy shock vector to be used at $h = 0$, having relied on responses from $h = 1, \ldots, T_c$. It is a linear projection of the rule violation under the oil shock onto the space of policy IRFs. With m_{aux}, Eq. (9) delivers the counterfactual IRF of every variable in the VAR, including the eight distributional factors.

Recall, the counterfactual restrictions are only imposed over a finite horizon. For $h > T_c$, no further restrictions are imposed and the economy evolves under the historical policy rule from the counterfactual state reached at the end of the enforcement window, as implied by the VAR's historical dynamics. All reported losses are computed over the enforcement window $h = 1, \ldots, T_c$.

Counterfactual policy rules I consider a nested set of three policy mandates. The most general specification augments the standard dual mandate with a distributional objective:

$$\mathcal{L}(m_{\text{aux}}) = \|J_c y_\pi^{\text{cnf}}\|^2 + \omega_u \|J_c y_u^{\text{cnf}}\|^2 + \omega_g \|J_c g^{\text{cnf}}\|^2, \tag{11}$$

where y_u^{cnf} and y_π^{cnf} are the counterfactual paths of unemployment and inflation, and $g^{\text{cnf}} \in \mathbb{R}^T$ denotes the counterfactual path of the distributional statistic. J_c selects the enforcement horizons. The three benchmark mandates are nested within (11): the *single mandate* sets $\omega_u = \omega_g = 0$ and stabilizes inflation alone; the *dual mandate* sets $\omega_u = 1$ and $\omega_g = 0$; and the *triple mandate* additionally places weight $\omega_g > 0$ on distributional stabilization. Regarding the distributional statistic, its counterfactual path is linear in the bundle, $g^{\text{cnf}} = \Psi_g^{\text{oil}} + m_g^{\text{base}} m_{\text{aux}}$, exactly like inflation and unemployment, with Ψ_g^{oil} and m_g^{base} obtained by applying a linear map to the factor responses.[17]

[17] Nonlinear statistics such as the Gini lack this property.

For the results, I target two consumption inequality terms. First, the response of the 90/10 consumption spread to the shock, defined as $g = 100 \times [\text{asinh}(Q_{90}/\mu) - \text{asinh}(Q_{10}/\mu)]$, with Q the quantile function of consumption and μ its mean. This is a difference of moments focused on the tails of the distribution. Second, I construct a measure that focuses more on the entire distribution and is designed to be a *descriptive* representation of a risk-sharing target: it is the log consumption responses of the deciles net of the aggregate consumption response (in percentage points). When it is equal to zero, this means every decile moves in accordance with the average. Of course, this does not measure how efficiently households actually share risk, but serves more as a comprehensive measure over the distribution of households as opposed to just the tails.

Both measures move away from stabilizing levels or targeting a preferred level of inequality and directly focus on stabilizing households. All targets enter in points (percentage points for inflation, unemployment, and the risk-sharing measure), and I set the weights to one.[18]

5 Can policy stabilize an oil supply shock?

Oil shocks have been documented to have widespread macroeconomic effects (Baumeister and Hamilton, 2019; Hamilton, 1983; Känzig, 2021; Kilian, 2009): they raise the general price level and depress real activity, confronting policymakers with a trade-off between stabilizing prices and economic activity. They have also been documented to affect households disproportionately (Calderon, 2025; Del Negro et al., 2025; Känzig, 2026), inviting a third dimension to the policy problem. This section thus asks whether policy can, through different mandates and policy instruments, stabilize the aggregate and distributional effects of an oil supply shock.

Section 5.1 begins by documenting how our modeled macroeconomy responds to an oil supply shock à la Känzig (2021), to assess whether such trade-offs are present in the dynamic causal effects. For this, I first report the responses of the conventional policy targets: inflation and unemployment. Alongside, I report the effects on the marginal distributions of consumption, income, and wealth to understand the shock's effects across the distribution. I then show how these same outcomes counterfactually evolve when authorities select between three nested mandates: a single mandate of inflation targeting, a dual mandate that additionally stabilizes unemployment, and a triple mandate that further stabilizes a measure of consumption inequality. Within our setup, when authorities choose to stabilize the consumption distribution, they stabilize the consumption 90/10 spread, which focuses on inequality between the tails; Section 5.4 compares this term with one replaced by a risk-sharing target over all ten consumption deciles, as described in Section 4.

[18] See Appendix I for more details.

In Section 5.2, policymakers are constrained to using only *conventional* monetary policy: the policy rate and no fiscal policy (Span A), to stabilize targeted outcomes. Section 5.3 then asks what it costs to stabilize the consumption distribution: the aggregate stabilization *given up* as we incrementally move away from a dual mandate that does not weigh distributional concerns to a triple mandate that weighs consumption inequality, and how this price changes as fiscal levers enter the policy toolkit, be it legislated tax changes (Span B) or taxes and government spending together (Span D); government spending alone (Span C) is in Appendix C. Section 5.4 asks *whose* consumption policy stabilizes, grouping households by their income and wealth together. Section 5.5 asks policymakers to choose *one* common bundle m_{aux}, the vector of time-zero lever surprises, applied to every posterior draw.

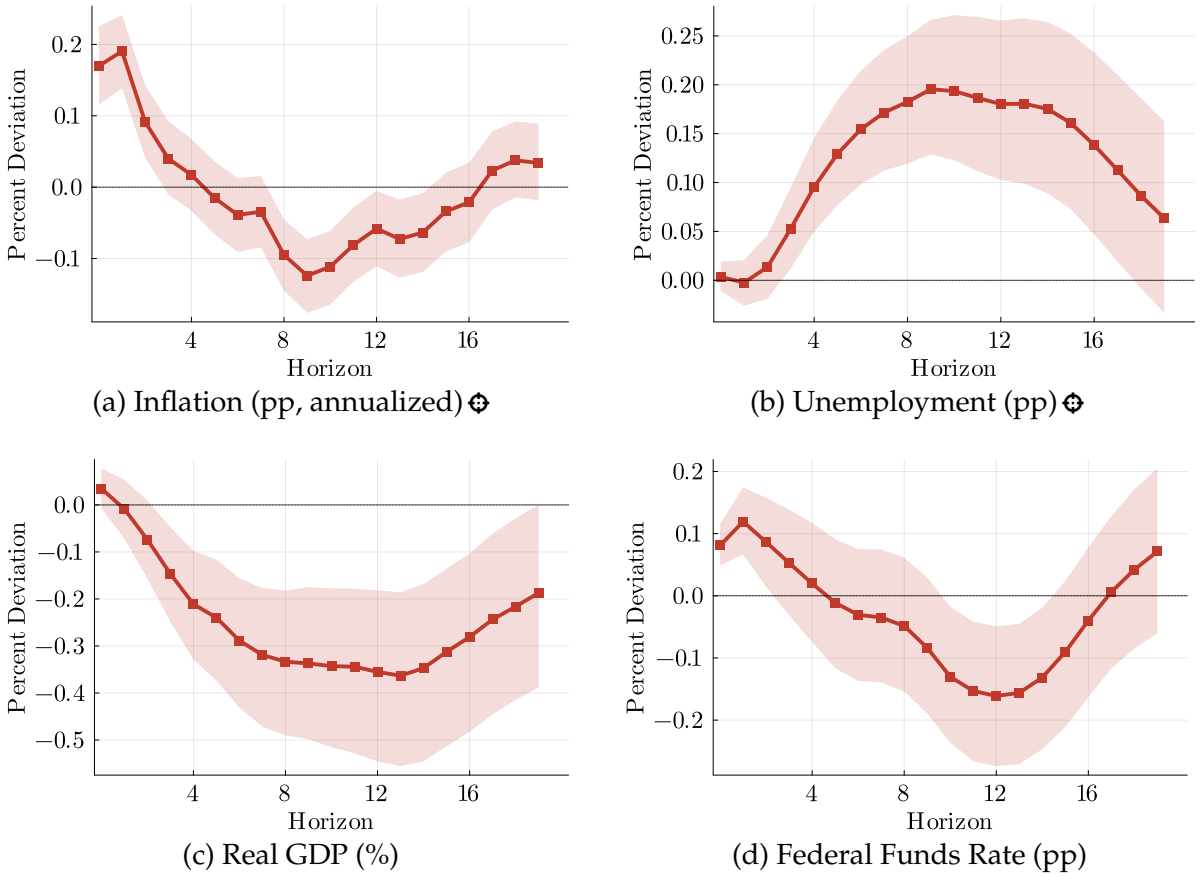
Summary of main findings Four results emerge. First, oil-induced inflation is difficult to stabilize. The inflation-only mandate removes just 6% of the inflation loss with the policy rate, rising to 34% with all instruments, while every mandate that also targets unemployment increases inflation loss. By contrast, the dual mandate removes 68–84% of the unemployment loss. This asymmetry partly reflects the timing of the responses: inflation rises and then falls within a short horizon, whereas real activity responds more persistently and largely without reversing sign.

Second, the dual mandate partially stabilizes consumption inequality, with consumption-spread losses of 0.67–0.94 across policy spans. It also removes 27–45% of income-spread loss but leaves wealth inequality largely unchanged. Third, explicitly targeting consumption inequality removes a further 67–88% of the dual mandate's consumption-spread loss, at an additional aggregate cost of 11–38%, borne mostly through unemployment. Finally, consumption stabilization does not automatically extend to income or wealth inequality: under the triple mandate, income-spread losses remain at 0.69–0.82, higher than under the dual mandate, while wealth-spread losses remain near their realized levels. Thus, the effects of policy on consumption inequality cannot be inferred from its effects on income or wealth inequality alone.

5.1 What does an oil supply shock do to the macroeconomy and the distribution?

Figure 1 reports the realized aggregate responses, that is, the outcomes under the historical monetary and fiscal rules: the posterior median and its 68% credible band. The figures show realized responses only; the counterfactual mandates are summarized by the loss tables below. Panel titles with a crosshair are targeted variables across the different mandates. Figure 2 shows the corresponding realized paths for the 90/10 spread of consumption, income and wealth. Only the consumption spread is targeted. Because the spread is linear in the VAR's

Figure 1: Aggregate Responses to an Oil Supply Shock

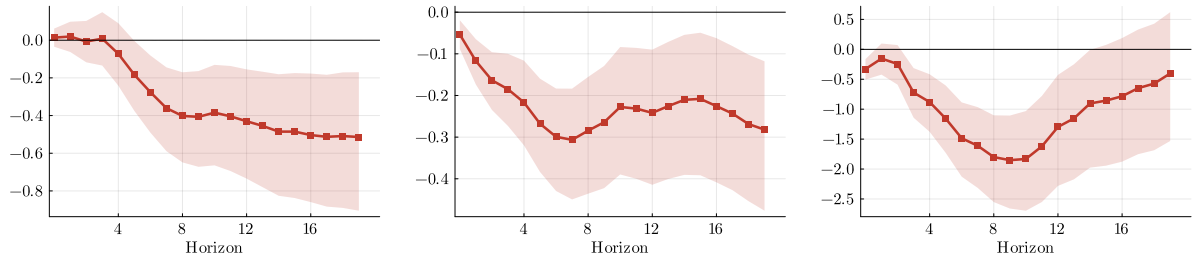


(a) Inflation (pp, annualized) ⌖ (b) Unemployment (pp) ⌖

(c) Real GDP (%) (d) Federal Funds Rate (pp)

Notes: Realized responses to the oil supply shock, i.e. under the historical monetary and fiscal rules: posterior median and 68% band. Horizontal axis is time horizon after the shock in quarters. ⌖ marks a variable that some mandate targets in Section 5.2–Section 5.5.

Figure 2: Distributional 90/10 Responses to an Oil Supply Shock



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

Notes: Realized responses to the oil supply shock, i.e. under the historical monetary and fiscal rules: posterior median and 68% band. Horizontal axis is time horizon after the shock in quarters. The 90/10 spread of each variable in asinh units: $100 \times [\text{asinh}(Q_{90}/\mu) - \text{asinh}(Q_{10}/\mu)]$, with Q the quantile function and μ the mean of each variable. Only the consumption spread is targeted ⌖.

impulse responses to first order, holding the steady-state distribution used in the reconstruction fixed, the triple mandate targets it in the same way as the other mandates target inflation (Appendix I). The figures show posterior median paths; the statistics reported in the tables are medians over draw by draw ratios.

Baseline and counterfactual responses Figure 1 shows an oil supply shock is stagflationary: inflation jumps on impact and fades within a year, followed by disinflation; unemployment builds over two years and stays elevated; and real GDP contracts with the same persistence. The realized funds rate rises near impact and is then cut at the height of the contraction two-three years later, reaching 16 basis points below its pre-shock level, i.e., the historical rule leans against inflation first and against the recession later. Figure 2 shows the shock *compresses* the consumption distribution with a delay: the 90/10 spread falls by about half a point over the first three years, over the same horizons in which oil prices and inflation fall and the funds rate is cut; consumption falls more for the rich than for the poor, consistent with Calderon (2025). Income and wealth inequality compress as well with similar persistence, but respond on impact.

5.2 What can conventional monetary policy stabilize?

This section proposes the first set of counterfactuals, asking policymakers to stabilize the effects of an oil supply shock under different mandates, while limited to using only conventional monetary policy (**Span A**). Table 1 reports the median of ratios, where the ratio is the sum of squared counterfactual responses over T_c for some posterior draw relative to the realized response of the same draw. This quantifies the degree of stabilization achieved over the enforcement window T_c: values below one indicate greater stabilization relative to the realized path, zero indicating full stabilization.

Under the single mandate of inflation targeting (Single_π), the inflation loss falls only to 0.94, and unemployment, GDP and the consumption spread as untargeted outcomes are all *less* stable than under the historical rule. The reason is not that the rate is ineffective. A rate increase lowers the price level (Figure A.1), so in principle the policymaker can lower inflation if she wishes. Two issues are the policymaker must consider the inflation dynamics over the entire window (the enforcement window T_c) and is limited to selecting one bundle at time $t = 0$, as opposed to perhaps an additional bundle at a later horizon.

Why does this matter here? Inflation changes sign within the enforcement window; thus, a rate hike that offsets the initial rise in inflation *deepens* the later *disinflation* we observe, which matters for a price stability mandate. Had the policy rate's inflation response lined up with the shock's inflation response over the same window, inflation-only central bank would be able to stabilize; the sign change makes that alignment poor. Furthermore, the rule here chooses *one* time-zero surprise to the rate, not an arbitrary path of rates, so an inflation-only central bank held to the same accord cannot stabilize both the inflation and disinflation; a sequence of rate announcements at several dates might, but would be a different policy span.[19]

[19] From Figure 1, inflation does not change sign in the first year. When the counterfactuals are then estimated for

Under the dual mandate, the central bank eases about 15 basis points more than the realized rule in the first quarters, which cuts the unemployment loss to 0.32 and the GDP loss to 0.46 but raises the inflation loss to 1.28; there is no divine coincidence. Furthermore, the consumption distribution under the dual mandate stabilizes relative to the realized response, but to a lesser degree, with a consumption-spread loss of 0.67. I also explore the income and wealth spread as untargeted outcomes. The income-spread loss falls to 0.73. The wealth spread is slightly less stable (1.10). In this sense, the dual mandate primarily achieves unemployment stabilization. This is because the policy rate's effects on real activity—on unemployment (the targeted outcome) and on output (an untargeted outcome)—are similarly persistent to the oil shock responses. They do not reverse sign over the policy horizon, giving the rate greater capacity to stabilize real activity. The same intervention also reduces the income- and consumption-spread losses.

The triple mandate augments the dual mandate with the consumption spread. The consumption spread loss falls to 0.22: this removes 78% of the realized spread loss and, relative to the dual mandate, about two thirds of the dual mandate's spread loss (a removed share of 0.67 in Table 2). Conventional monetary policy thus has considerable capacity to stabilize consumption inequality. The cost is mechanically what the dual mandate stabilized in real activity: unemployment loss rises from 0.32 to 0.56 and GDP loss from 0.46 to 0.56, while inflation loss falls slightly ($1.28 \to 1.17$) because the triple mandate does not ease the interest rate as much as the dual.

The substantial improvement in consumption stabilization under the triple mandate illustrates the point of McKay and Wolf (2026): distributional objectives can materially affect policy when the available instruments have differential incidence across households. The policy rate has such incidence for consumption. Indeed, targeting the consumption spread alone (Single$_g$) reduces its loss only marginally further, from 0.22 to 0.21, suggesting that aggregate objectives impose little additional constraint on consumption stabilization. These gains, however, do not extend to other dimensions of inequality. Relative to the dual mandate, the triple mandate raises income-spread loss from 0.73 to 0.82, while wealth-spread loss remains near its realized level (1.07).

Overall, four main results can be inferred. First, with the policy rate lever alone, inflation stabilization is limited under every mandate over the enforcement horizon T_c considered: the inflation-only mandate removes 6% of the inflation loss and raises the other losses. Second, the dual mandate's stabilization of real activity carries over to untargeted consumption inequality, but also income inequality. The first part is what Acharya, Challe, and Dogra (2023) would predict, but to a lesser degree: when income risk is countercyclical or acyclical, stabi-

$T_c = 4$, a policymaker who is targeting inflation can hike the interest rate further and can clear most of the inflation on impact (about half), but activity would worsen three-fold.

Table 1: Losses by mandate, span A: policy rate

	Aggregates			90/10 spreads		
Mandate	Inflation ⌖	Unemployment ⌖	GDP	Consumption ⌖	Income	Wealth
Single$_\pi$	0.94	1.08	1.12	1.33	1.02	0.98
Dual	1.28	0.32	0.46	0.67	0.73	1.10
Triple	1.17	0.56	0.56	0.22	0.82	1.07
Single$_g$	1.17	0.60	0.59	0.21	0.83	1.06

Notes: Each entry is the median over ratios. The ratio is the $\sum_{h=1}^{12} y_h^2$ of the counterfactual response relative to the realized one within the same draw. A value of 1 means no change and below 1 means more stable. Targeted moments are denoted by ⌖. Single$_\pi$: single mandate on Inflation; Single$_g$: single mandate on the Consumption spread. The spreads are 90/10 spreads in asinh units.

lizing activity stabilizes incomes and thus, consumption. Third, consumption inequality can be stabilized by an explicit target, at a cost of unemployment, and without further stabilizing income or wealth inequality. Fourth, the wealth spread is unchanged across different policy objectives. Moving from the dual to the triple mandate, the trade-off in Span A is between consumption-spread and unemployment stabilization: the spread loss falls from 0.67 to 0.22 while the unemployment loss rises from 0.32 to 0.56, and the inflation loss changes little (1.28 to 1.17).

5.3 The price of distributional stabilization: what does fiscal policy buy?

The dual mandate is nested in the triple mandate: $\omega_g = 0$ places zero weight on the consumption spread term, while $\omega_g = 1$ is the triple mandate considered so far. Yet, ω_g need not be one. This section asks: as we move incrementally from the dual toward the triple mandate, increasing ω_g, what is the aggregate cost of that additional unit of distributional consideration, and how does this price change with the set of available instruments?

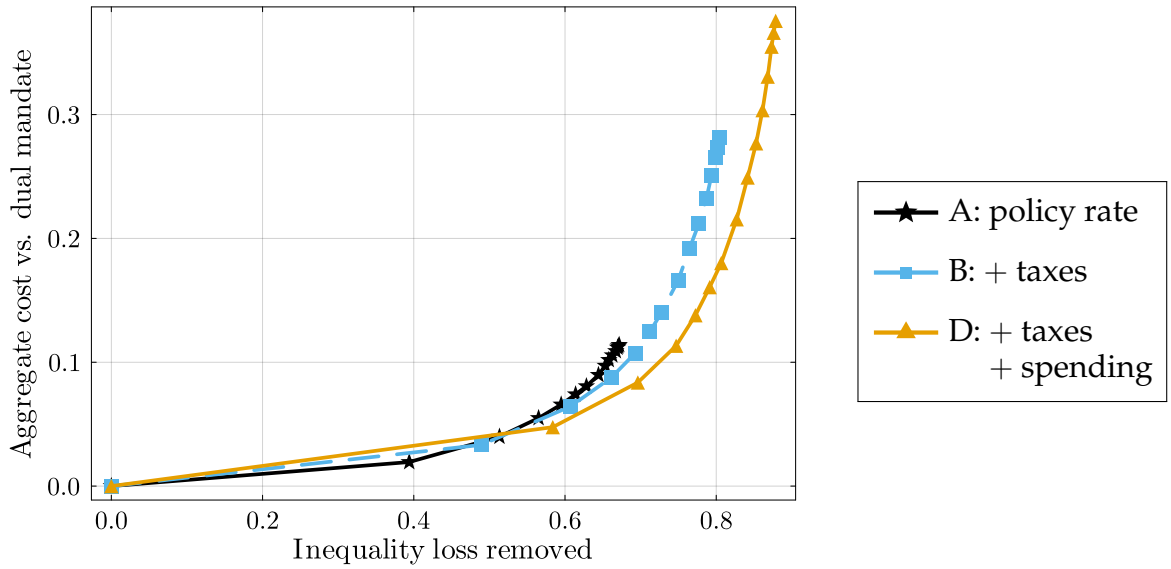
I measure the price of distributional stabilization in two ways. Table 2 fixes the weight on the spread, $\omega_g \in \{0.2, 0.5, 1\}$, and reports two ratios, both computed within each posterior draw d and reported as medians over draws:

$$\text{Removed}_d = 1 - \frac{L_{g,d}^{\omega_g}}{L_{g,d}^{\text{dual}}}, \qquad \text{Cost}_d = \frac{L_{\pi u,d}^{\omega_g}}{L_{\pi u,d}^{\text{dual}}} - 1,$$

where $L_{g,d}$ is the consumption-spread loss and $L_{\pi u,d}$ the aggregate (π, u) loss of draw d under the mandate with weight ω_g or under the dual mandate $\omega_g = 0$. Figure 3 instead abstracts from ω_g and asks: for a given amount of spread loss removed, which span achieves it at the lowest aggregate cost?

The measured aggregate cost is modest relative to the improvement in the consumption spread, at small and at large weights. At $\omega_g = 0.2$, a small weight relative to the benchmark

Figure 3: The price of distributional stabilization by span



Notes: The price paid from moving from the dual to the triple mandate, for each span. Each point is a triple mandate with weight ω_g on the consumption spread. The x-axis is the share of each span's own dual-mandate spread loss removed; the y-axis is the increase in the aggregate (π, u) loss relative to that dual mandate, both medians of per-draw ratios. $T_c = 12$.

Table 2: The price of stabilizing the consumption spread

	$\omega_g = 0.2$		$\omega_g = 0.5$		$\omega_g = 1$	
Span	Removed	Cost	Removed	Cost	Removed	Cost
A	0.59	+0.07	0.65	+0.10	0.67	0.11
B	0.69	+0.11	0.76	+0.19	0.80	0.28
D	0.77	+0.14	0.84	+0.25	0.88	0.38

Notes: Table reports for three weights $\omega_g \in \{0.2, 0.5, 1.0\}$ two statistics. *Removed* is the share of the dual mandate's consumption-spread loss removed and *Cost* the increase in the aggregate (π, u) loss relative to the dual mandate, both computed within each posterior draw and reported as the median over draws.

of one given how the targets are scaled, the mandate removes 59 to 77% of the dual mandate's spread loss for an aggregate cost of 7 to 14% across the three spans; at $\omega_g = 1$ it removes 67 to 88% for 11 to 38%. In the table, wider spans remove more at every weight and also show a higher cost. What the fiscal levers buy in the end is *capacity*: their ability to remove part of the spread loss that the policy rate could not remove under this set of mandates (a removed share of 0.67 in Span A against 0.88 in Span D).

Figure 3 shows the second exercise, where I hold the amount of consumption inequality stabilization fixed (any point on the x-axis). Up to a removed share of about 0.6, all curves are similar, beyond which only the spans with a fiscal lever continue. This speaks to the point emphasized by McKay and Wolf (2026): what matters for a distributional objective is the incidence of the available instruments, not their number; the rate alone has a large incidence on

Table 3: The price of distributional stabilization decomposed

Span	Total	Inflation	Unemployment	Inflation share
A	+0.11	+0.01	+0.10	12%
B	+0.28	+0.07	+0.21	26%
D	+0.38	+0.13	+0.24	35%

Notes: *Total* is the increase in the aggregate (π, u) loss when moving from the dual to the triple mandate ($\omega_g = 1$), relative to the dual mandate's aggregate loss, reported as the median over posterior draws, as in Table 2. *Inflation share* is the share of that increase attributable to inflation, calculated as $\sum_d (L^{\text{tri}}_{\pi,d} - L^{\text{dual}}_{\pi,d}) / \sum_d (L^{\text{tri}}_{d} - L^{\text{dual}}_{d})$. *Inflation* and *Unemployment* decompose the total increase using this share.

the consumption distribution, and the fiscal levers add a second direction rather than one with a lower measured cost.

Decomposing the aggregate price The aggregate price as shown includes both inflation and unemployment costs, masking whether distributional stabilization is paid for through weaker inflation or weaker unemployment stabilization. Table 3 decomposes the cost of moving from the dual to the triple mandate. Results indicate most of it falls on unemployment: inflation accounts for 12% of the cost with the policy rate alone and 26 to 35% when fiscal levers are added to the span.

This is not to say there is no trade-off between price stability and distributional stability. A single mandate on the consumption spread raises the inflation loss to 1.17 in Span A, and in every span the triple mandate's inflation loss is above one. But *around* the dual mandate, the relevant trade-off after an oil supply shock is between the distribution and real activity: the triple mandate can only *give back* stabilization that the dual mandate had *achieved*, and what the dual mandate achieves is unemployment stabilization, not inflation stabilization.

It is for the same reason that the inflation share of the price increases with the span: the fiscal levers let the dual mandate stabilize inflation more (an inflation loss of 1.28 in Span A against 1.02 in Span D), and the triple mandate gives up some of that. The price-stability–distribution trade-off emphasized by Bilbiie (2025) and Bhandari et al. (2021) therefore becomes operative, in the estimated spans and around the dual-mandate benchmark, only once the span gives the dual mandate some capacity to stabilize inflation in the first place.

5.4 Whose consumption does policy stabilize? TBD

Thus far, only the consumption 90/10 spread has been considered as a targeted measure of inequality; however, the measure only describes the tails and does not bring an understanding to *which* households are actually experiencing the stabilization achieved through policy. This

Table 4: Stabilizing the 90/10 spread vs. Household Deciles

	Span A		Span B		Span D	
Mandate	90/10	Risk sharing	90/10	Risk sharing	90/10	Risk sharing
Dual	0.67	0.77	0.86	0.95	0.94	1.04
Triple (spread)	0.22	0.53	0.16	0.49	0.11	0.46
$\text{Triple}_{\text{rs}}$	0.30	0.50	0.26	0.39	0.22	0.31

Notes: *90/10*: the consumption 90/10 spread in asinh points. *Risk sharing*: the ten deciles' log consumption responses net of the aggregate consumption response. The dual mandate targets inflation and unemployment; Triple (spread) adds the consumption spread; $\text{Triple}_{\text{rs}}$ adds risk sharing instead. Span A: policy rate; B: policy rate and taxes; D: policy rate, taxes and government spending. Span C is in Appendix C.

speaks directly to heterogeneous agent literature, which emphasizes that aggregate shocks transmit differentially through the *joint* distribution of individual states. This begs the question: *whose* consumption policy stabilizes? Because the VAR carries the joint distribution of consumption, income, and wealth, I can answer this question directly, estimating consumption dynamics of different household income and wealth groups together.

5.5 The limits of distributional stabilization under policy uncertainty

The counterfactuals so far let the planner find one bundle per posterior draw: the intervention is chosen knowing that draw's impulse responses, including the effects of the policy levers and the responses to the oil shock. This section instead considers one common policy bundle under posterior parameter uncertainty. The planner must commit to only *one* vector of time-zero lever surprises m, to be used in every draw, before knowing which draw is the true economy, and chooses it to minimize the posterior expected loss

$$m^* = \arg\min_m \ \mathbb{E}_{\text{post}}\big[L_d(m)\big],$$

where $L_d(m)$ is the mandate's weighted loss in draw d given some policy choice m. Table 5 compares the results so far with those under one bundle; both are evaluated draw by draw and reported as medians of per-draw ratios.

The first column covers results reported from the previous tables. The second column under Inflation shows that under one bundle, the inflation gains substantially diminish: the inflation-only mandate's loss moves between 0.94 to 0.99 in every span, against the 0.66 to 0.94 with full information. This is again because the sign of the levers' inflation effects differs across draws and a bundle that lowers inflation in one draw raises it in another. Unemployment stabilization survives but shrinks: the dual mandate's unemployment loss rises from 0.16–0.32 to 0.40–0.43, still more than half of the realized loss removed.

Table 5: Full information vs. one bundle under parameter uncertainty

		Inflation		Unemployment		Cons. spread		Risk sharing	
Span	Mandate	full	one bundle	full	one bundle	full	one bundle	full	one bundle
A	Single_π	0.94	0.99	1.08	1.13	1.33	1.26	1.16	1.16
	Dual	1.28	1.32	0.32	0.43	0.67	0.47	0.77	0.64
	Triple	1.17	1.18	0.56	0.54	0.22	0.45	0.53	0.64
B	Single_π	0.78	0.95	1.01	1.04	1.53	1.49	1.37	1.34
	Dual	1.14	1.26	0.21	0.40	0.86	0.56	0.95	0.74
	Triple	1.19	1.21	0.45	0.57	0.16	0.43	0.49	0.62
D	Single_π	0.66	0.94	0.99	1.04	1.62	1.58	1.45	1.41
	Dual	1.02	1.26	0.16	0.40	0.94	0.56	1.04	0.75
	Triple	1.20	1.22	0.36	0.56	0.11	0.43	0.46	0.62

Notes: Medians of per-draw loss ratios. *Full*: one bundle m_{aux} is chosen per draw (the planner knows the economy). *One bundle*: the same single bundle is chosen instead across all draws. Minimizes the posterior expected loss.

Under the Cons. spread header, second column, you can see distributional gains also shrunk under the triple mandate: with one bundle its spread loss is 0.43–0.45 in every span, against 0.11–0.22 with the full information case. Interestingly, under the single-bundle setting, the dual mandate's consumption-spread loss falls to 0.47–0.56, despite consumption inequality remaining untargeted. Explicit distributional targeting provides further gains where fiscal instruments are available, reducing spread loss from 0.56 to 0.43 in Spans B and D. However, these gains are smaller than when policy is optimized separately for each posterior draw. Thus, uncertainty about policy transmission substantially limits the additional stabilization attainable through an explicit distributional mandate, but improves that of the dual.

6 Conclusion

This paper asks what monetary and fiscal policy can achieve after an oil supply shock when both aggregate and distributional outcomes matter. I combine a VAR that tracks the joint distribution of household consumption, income and wealth with the counterfactual method of Caravello, McKay, and Wolf (2026), and I compare nested mandates over inflation, unemployment and consumption inequality, implemented with the policy rate alone or together with tax and spending levers.

The oil supply shock raises inflation and depresses real activity, and it compresses the consumption distribution: consumption falls more at the top than at the bottom. Policy can stabilize the real side of this response but not its inflation. The reason is the shape of the shock: inflation rises and then turns into disinflation within the policy window; one surprise to the levers cannot offset both phases.

Consumption inequality is different. The dual mandate stabilizes it only partially, but an explicit target strongly so. Moving from the dual to the triple mandate removes 67% of the dual mandate's consumption-spread loss with the policy rate alone, 80% with taxes and 88% with taxes and spending. The aggregate loss rises by 11%, 28% and 38%, and most of that increase is in unemployment: inflation accounts for 12% of it with the rate alone and 26 to 35% with fiscal levers. *Around the dual mandate*, conditional on the levers and the horizon for which policy is considered, the trade-off after an oil supply shock is therefore between the distribution and real activity, not between the distribution and inflation.

Fiscal levers matter mainly for how far distributional stabilization can go. Wider spans remove more of the spread loss at every weight on the distributional target, but they do not make it cheaper: the aggregate cost rises with the span, from 11% with the rate alone to 28% with taxes and 38% with all levers. Fiscal policy extends the reach of distributional stabilization more than it lowers its price. Much of the gain is available at small weights: with a weight of 0.2 on the spread, the triple mandate already removes 59 to 77% of the dual mandate's spread loss at a cost of 7 to 14%.

The joint distribution also shows whose consumption is stabilized. Stabilizing the consumption spread does not stabilize income and wealth inequality to the same degree, so policy's effect on consumption inequality cannot be read off the marginal distributions of income and wealth. Grouping households by income and wealth together shows which households carry the stabilization.

These results are conditional on the planner knowing the economy. When one common bundle must be chosen under posterior parameter uncertainty, the small inflation gains largely vanish, unemployment stabilization shrinks but survives, and the incremental gain from targeting consumption inequality shrinks substantially. It does not disappear where fiscal levers are available: under one bundle the triple mandate's spread loss is 0.43 against the dual mandate's 0.56 in Spans B and D, but only 0.45 against 0.47 with the rate alone.

References

Acharya, Sushant, Edouard Challe, and Keshav Dogra (2023). "Optimal Monetary Policy According to HANK". In: *American Economic Review* 113.7, pp. 1741–1782.

Andersen, Asger Lau et al. (2023). "Monetary Policy and Inequality". In: *Journal of Finance* 78.5, pp. 2945–2989.

Antolín-Díaz, Juan and Paolo Surico (2025). "The Long-Run Effects of Government Spending". In: *American Economic Review* 115.7, pp. 2376–2413.

Aruoba, S Borağan and Thomas Drechsel (2024). *Identifying monetary policy shocks: A natural language approach*. Tech. rep. National Bureau of Economic Research.

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

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2024). "The intertemporal keynesian cross". In: *Journal of Political Economy* 132.12, pp. 4068–4121.

— (2025). "Fiscal and Monetary Policy with Heterogeneous Agents". In: *Annual Review of Economics* 17.

Auclert, Adrien et al. (2024). "Managing an Energy Shock: Fiscal and Monetary Policy". In: *Heterogeneity in Macroeconomics: Implications for Monetary Policy*. Ed. by Sofía Bauducco, Andrés Fernández, and Giovanni L. Violante. Santiago: Central Bank of Chile, pp. 39–108.

Auerbach, Alan J and Yuriy Gorodnichenko (2012). "Measuring the Output Responses to Fiscal Policy". In: *American Economic Journal: Economic Policy* 4.2, pp. 1–27.

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

Barnichon, Régis et al. (2026). "Humble VARs". Paper prepared for the NBER Macroeconomics Annual Conference, July 2026.

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

Baumeister, Christiane (2025). "Discussion of "Local Projections or VARs? A Primer for Macroeconomists" by Montiel Olea, Plagborg-Møller, Qian, and Wolf". Prepared for the NBER Macroeconomics Annual 2025.

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

Bayer, Christian, Benjamin Born, and Ralph Luetticke (2024). "Shocks, Frictions, and Inequality in US Business Cycles". In: *American Economic Review* 114.5, pp. 1211–1247.

Bayer, Christian, Luis Calderon, and Moritz Kuhn (2025). *Distributional Dynamics*. Working paper. University of Bonn.

Bayer, Christian et al. (2025). *Being and consciousness: Fiscal attitudes according to HANK*. Tech. rep. ECONtribute Discussion Paper.

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

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

Bilbiie, Florin O. (2025). "Monetary Policy and Heterogeneity: An Analytical Framework". In: *Review of Economic Studies* 92.4, pp. 2398–2436.

Bjørnland, Hilde C., Jamie L. Cross, and Jonas Hölz (2025). *Re-visiting the Relationship Between Oil Prices and Monetary Policy*. Norges Bank Working Paper 4/2025. Norges Bank.

Bobasu, Alina, Michael Dobrew, and Amalia Repele (2025). "Energy price shocks, monetary policy and inequality". In: *European Economic Review* 175, p. 104986.

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

Calderon, Luis (2025). *The distributional effects of oil supply shocks*. Working paper. Companion paper to "Does Monetary Policy Matter for Oil Supply Shocks?" University of Bonn.

Caravello, Tomás E., Alisdair McKay, and Christian K. Wolf (2026). *Evaluating Monetary Policy Counterfactuals: (When) Do We Need Structural Models?* Working Paper 32988. Revised March 2026. Earlier version circulated as "Evaluating Policy Counterfactuals: A VAR-Plus Approach". National Bureau of Economic Research.

Castelnuovo, Efrem, Lorenzo Mori, and Gert Peersman (2024). *Commodity Price Shocks and Global Cycles: Monetary Policy Matters*. CAMA Working Paper 36/2024. Centre for Applied Macroeconomic Analysis, ANU.

Chan, Jenny, Sebastian Diz, and Derrick Kanngiesser (2024). "Energy prices and household heterogeneity: Monetary policy in a Gas-TANK". In: *Journal of Monetary Economics* 147, p. 103620.

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

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

Chang, Minsu and Frank Schorfheide (2026). "On the Effects of Monetary Policy Shocks on Income and Consumption Heterogeneity". In: *American Economic Journal: Macroeconomics*. Forthcoming.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico (2020). "Monetary policy when households have debt: new evidence on the transmission mechanism". In: *The Review of Economic Studies* 87.1, pp. 102–129.

Dávila, Eduardo and Andreas Schaab (2023). *Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy*. Working Paper 30961. National Bureau of Economic Research.

De Graeve, Ferre and Andreas Westermark (2025). "Long-lag VARs". In: *Journal of Monetary Economics* 156, 103831.

Del Canto, Felipe N. et al. (2025). "Are Inflationary Shocks Regressive? A Feasible Set Approach". In: *Quarterly Journal of Economics* 140.4, pp. 2685–2747.

Del Negro, Marco et al. (2025). "Tradeoffs for the poor, divine coincidence for the rich". In: *Journal of Economic Dynamics and Control*, p. 105221.

Gagliardone, Luca and Mark Gertler (2026). "Oil Prices, Monetary Policy and Inflation Surges". In: *American Economic Journal: Macroeconomics*. Forthcoming.

Giannone, Domenico, Michele Lenza, and Giorgio E Primiceri (2015). "Prior selection for vector autoregressions". In: *Review of Economics and Statistics* 97.2, pp. 436–451.

Gornemann, Nils, Keith Kuester, and Makoto Nakajima (2022). *Doves for the Rich, Hawks for the Poor? Distributional Consequences of Systematic Monetary Policy*. Discussion Paper 089. Revise and resubmit, *Review of Economic Studies*. ECONtribute.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson (2005). "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements". In: *International Journal of Central Banking* 1.1, pp. 55–93.

Hack, Lukas, Klodiana Istrefi, and Matthias Meier (2025). *The Systematic Origins of Monetary Policy Shocks*. Working Paper 1021. Also CEPR Discussion Paper 19063. Banque de France.

Hamilton, James D (1983). "Oil and the macroeconomy since World War II". In: *Journal of political economy* 91.2, pp. 228–248.

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

Hofmann, Boris, Cristina Manea, and Benoît Mojon (2025). *Targeted Taylor rules: some evidence and theory*. BIS Working Papers 1234. Bank for International Settlements.

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek (2021). "The transmission of monetary policy under the microscope". In: *Journal of Political Economy* 129.10, pp. 2861–2904.

Ider, Gökhan et al. (2026). *Friend, Not Foe? Monetary Policy and Energy Prices*. DIW Discussion Paper. Revised May 2026. DIW Berlin.

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

Känzig, Diego R. (2026). "The Unequal Economic Consequences of Carbon Pricing". In: *American Economic Review*. Forthcoming.

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

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

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

Kilian, Lutz and Xiaoqing Zhou (2023). "The Econometrics of Oil Market VAR Models". In: *Essays in Honor of Joon Y. Park: Econometric Methodology in Empirical Applications*. Vol. 45B. Advances in Econometrics. Emerald, pp. 65–95.

Kuttner, Kenneth N. (2001). "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market". In: *Journal of Monetary Economics* 47.3, pp. 523–544.

Le Grand, François, Alaïs Martin-Baillon, and Xavier Ragot (2026). "Should Monetary Policy Care about Redistribution? Optimal Monetary and Fiscal Policy with Heterogeneous Agents". In: *Review of Economic Studies*. Advance access.

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

Luetticke, Ralph (2021). "Transmission of monetary policy with heterogeneity in household portfolios". In: *American Economic Journal: Macroeconomics* 13.2, pp. 1–30.

McCracken, Michael W. and Serena Ng (2021). "FRED-QD: A Quarterly Database for Macroeconomic Research". In: *Review* 103.1, pp. 1–44.

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

McKay, Alisdair and Christian K. Wolf (2026). "Optimal Policy Rules in HANK". In: *Review of Economic Studies*. Forthcoming.

Mertens, Karel and Morten O Ravn (2012). "Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks". In: *American Economic Journal: Economic Policy* 4.2, pp. 145–181.

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

Pallotti, Filippo et al. (2024). "Who bears the costs of inflation? Euro area households and the 2021–2023 shock". In: *Journal of Monetary Economics* 148, p. 103671.

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

Ramey, Valerie A (2011). "Identifying government spending shocks: It's all in the timing". In: *The Quarterly Journal of Economics* 126.1, pp. 1–50.

Romer, Christina D and David H Romer (2010). "The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks". In: *American Economic Review* 100.3, pp. 763–801.

Sims, Christopher A (1992). "Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy". In: *European Economic Review* 36.5, pp. 975–1000.

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

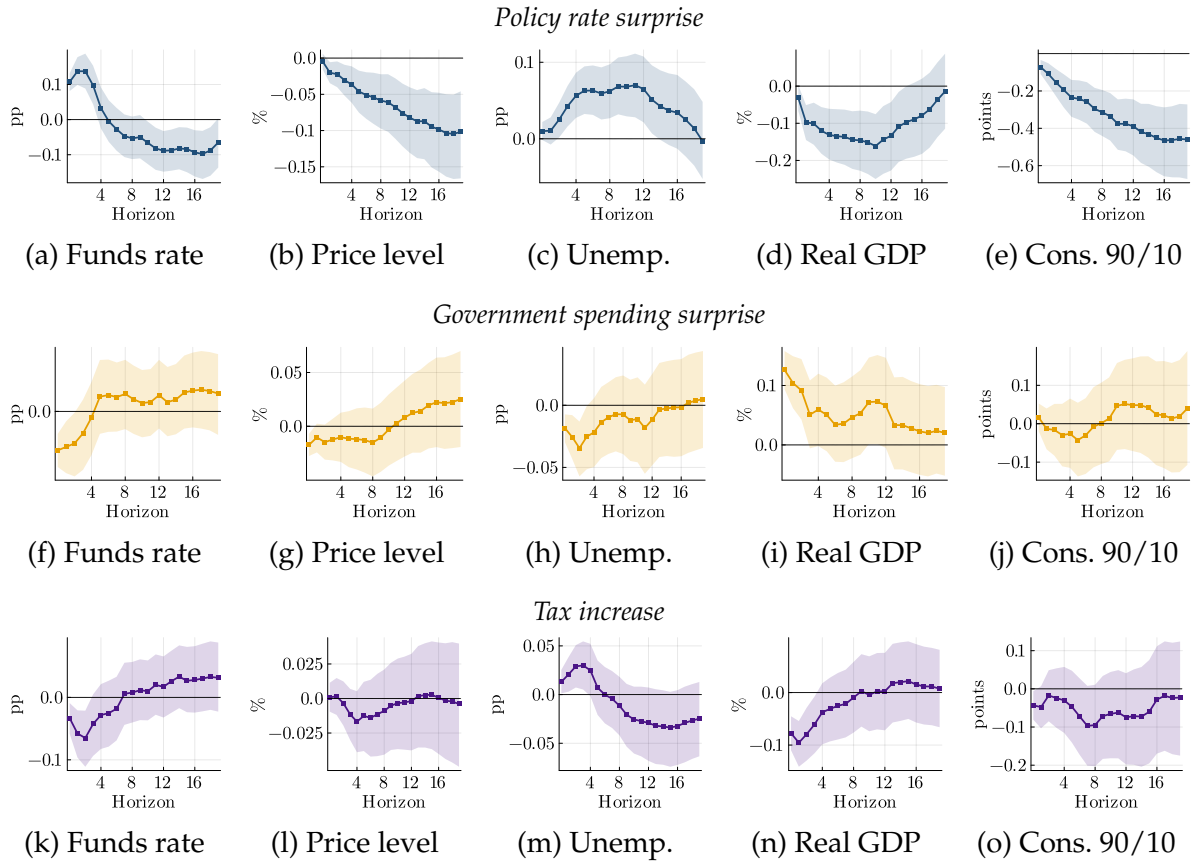
Sklar, Abe (1959). "Vol. 8 of Fonctions de Répartition à n Dimensions et Leurs Marges, 229–231". In: *Paris: Publications de l'Institut de statistique de l'Université de Paris*.

— (1973). "Random variables, joint distribution functions, and copulas". In: *Kybernetika* 9.6, pp. 449–460.

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

Wolf, Christian K (2025). "Interest rate cuts versus stimulus payments: An equivalence result". In: *Journal of Political Economy* 133.4, pp. 1235–1275.

Figure A.1: Responses to a one-standard-deviation surprise in each policy lever



Notes: Rows: the policy rate (Hack, Istrefi, and Meier's Aruoba and Drechsel purged narrative series to 2007Q4 spliced with Bauer–Swanson surprises from 2008Q1), the government spending surprise (Auerbach–Gorodnichenko forecast error, purged of SPF macro news) and legislated tax changes (Romer–Romer; a positive surprise is a tax increase). Columns: the federal funds rate (pp), the price level (GDP deflator, %), unemployment (pp), real GDP (%) and the consumption 90/10 spread, $100 \times [\operatorname{asinh}(Q_{90}/\mu) - \operatorname{asinh}(Q_{10}/\mu)]$ in points, in quarters after the surprise. Posterior median and 68% band. Units are one historical standard deviation of each surprise, not the 10% oil-price normalization of the counterfactuals.

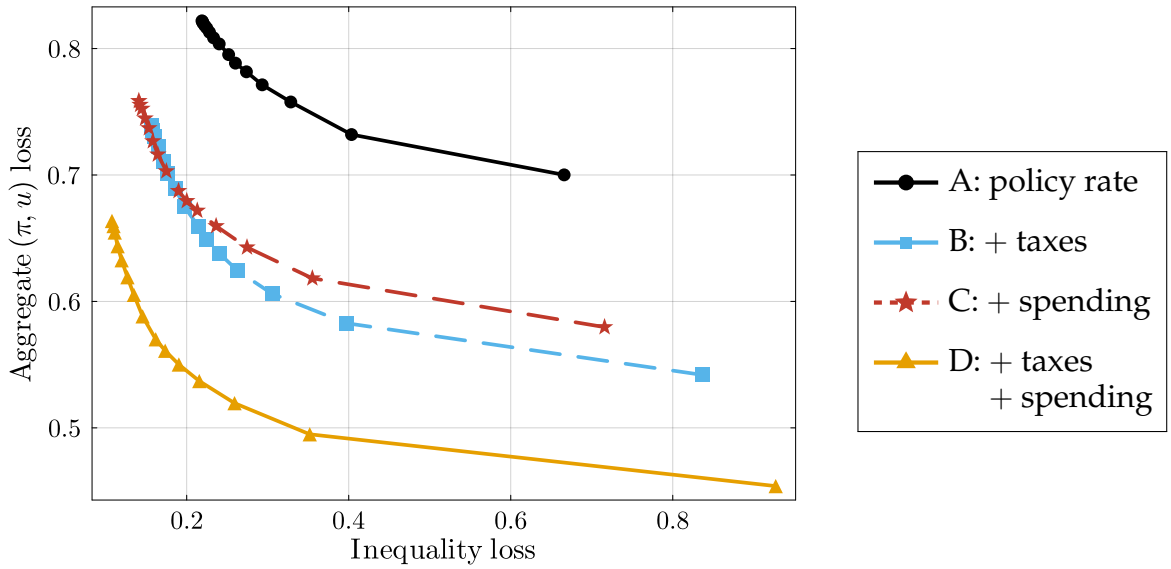
A Policy levers

Table A.1: Correlations of the policy levers, 1983Q1–2019Q3

	Policy rate	Gov. spending	Taxes
Policy rate	1.00	−0.04	0.09
Gov. spending		1.00	−0.05
Taxes			1.00

Notes: Pairwise correlations of the quarterly lever series as they enter the VAR, each set to zero outside its coverage. *Policy rate*: the narrative surprises of Aruoba and Drechsel (2024) up to 2007Q4, spliced with the high-frequency surprises of Bauer and Swanson (2023) from 2008Q1, each standardized over its own window. *Gov. spending*: the Auerbach and Gorodnichenko (2012) forecast error of federal government spending, purged of the macroeconomic news received by forecasters within the quarter (Section 3.2.2). *Taxes*: the exogenous legislated tax changes of Romer and Romer (2010), which end in 2007.

Figure B.1: Attainable outcomes by policy span



Notes: Both axes are losses relative to the *realized path*, medians of per-draw ratios as in the loss tables. The x-axis is the consumption-spread loss; the y-axis is the aggregate (π, u) loss. Each curve runs from the dual mandate ($\omega_g = 0$, right end) to the triple mandate ($\omega_g = 1$, left end) over the grid $\omega_g \in \{0, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.4, \ldots, 0.9, 0.95, 1\}$. Every span includes the policy rate; the legend lists what is added.

B Results by policy span

C Adding government spending alone (span C)

Table C.1: Losses by mandate, span C: policy rate and government spending

	Aggregates			90/10 spreads		
Mandate	Inflation ⌖	Unemployment ⌖	GDP	Consumption ⌖	Income	Wealth
Single$_\pi$	0.81	1.06	1.13	1.42	1.01	0.91
Dual	1.18	0.25	0.41	0.74	0.65	1.07
Triple	1.24	0.45	0.52	0.14	0.75	1.05

Notes: Span C is the policy rate with the government spending surprise. See Table 1 for other details.

Table C.2: Span C (policy rate and government spending): the price of distributional stabilization

	$\omega_g = 0.2$		$\omega_g = 0.5$		$\omega_g = 1$	
	Removed	Cost	Removed	Cost	Removed	Cost
C	0.70	+0.09	0.78	+0.17	0.81	+0.24

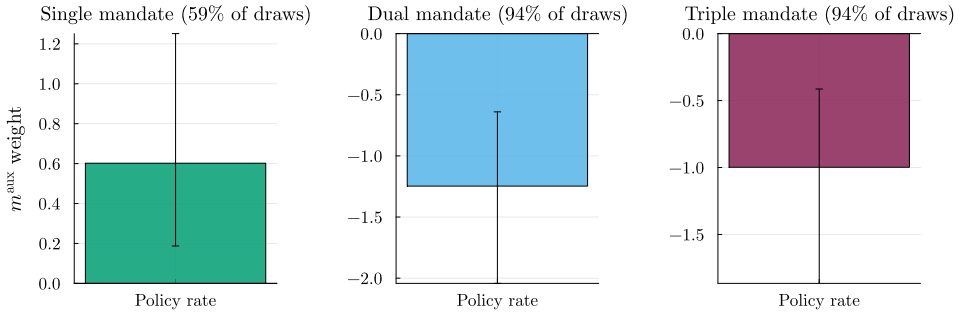
Notes: As Table 2. Decomposition of the cost at $\omega_g = 1$ (posterior-mean losses, as Table 3): total +0.35, inflation +0.10, unemployment +0.25; median of the per-draw total +0.24.

Table C.3: Span C: full information vs. one bundle under parameter uncertainty

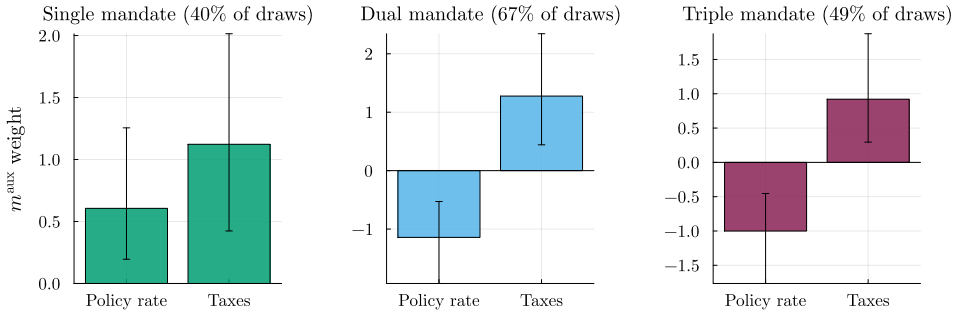
	Inflation		Unemployment		Cons. spread		Risk sharing	
Mandate	full	one bundle	full	one bundle	full	one bundle	full	one bundle
Single$_\pi$	0.81	0.98	1.06	1.12	1.42	1.33	1.24	1.22
Dual	1.18	1.33	0.25	0.44	0.74	0.46	0.84	0.64
Triple	1.24	1.20	0.45	0.52	0.14	0.44	0.47	0.64

Notes: As Table 5.

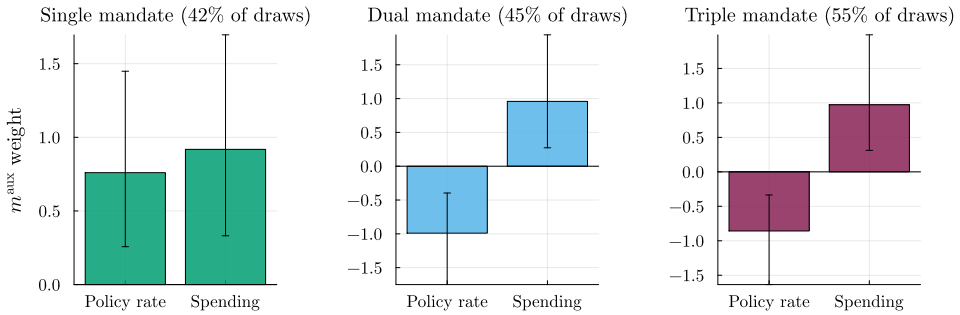
Figure D.1: The modal policy bundle under each mandate, by span



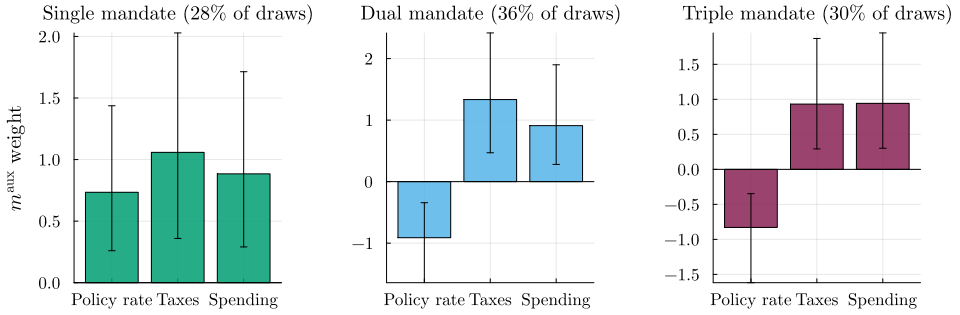
(a) Span A: policy rate



(b) Span B: policy rate and taxes



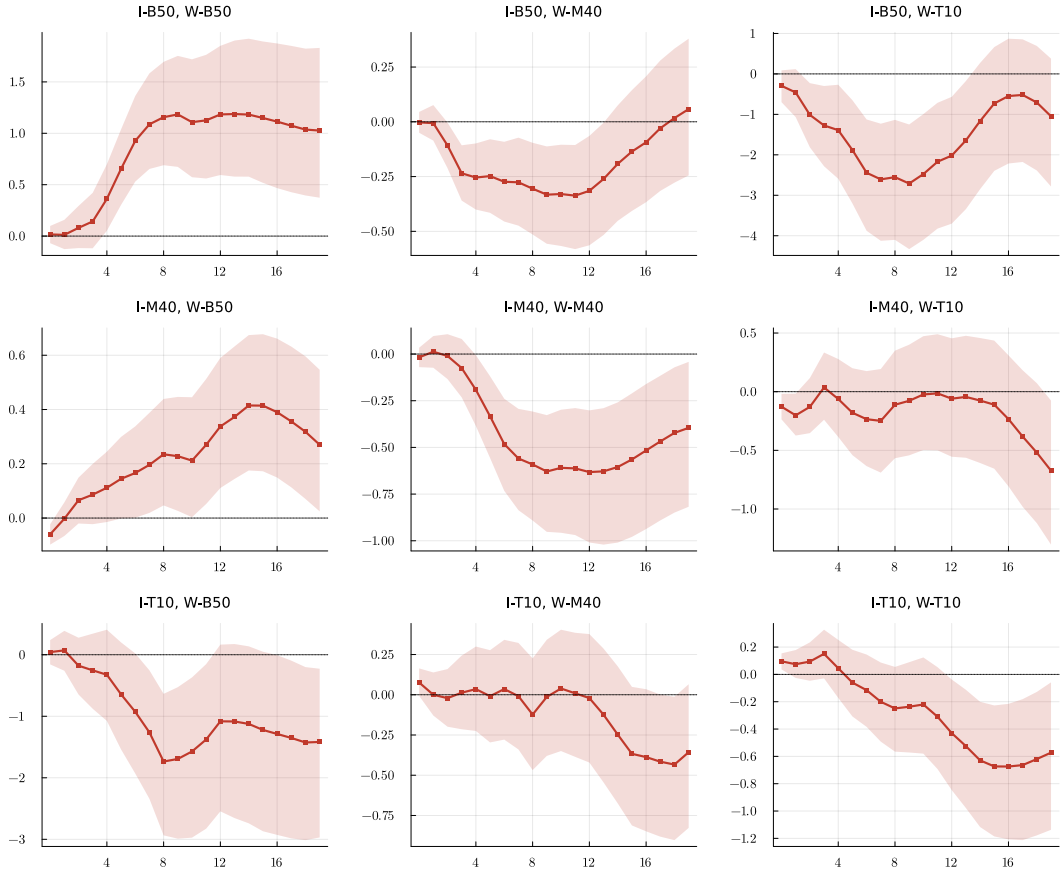
(c) Span C: policy rate and government spending



(d) Span D: policy rate, taxes and government spending

Notes: For each mandate, the modal combination is the most frequent joint sign pattern of the lever weights across posterior draws; the panel title gives its share of draws. Bars are the median weight of each lever over the draws with that pattern, in one-standard-deviation historical surprises (negative: an easing, a tax cut, or less spending), with 16–84% bars over the same draws. The modal pattern's share ranges from 94% of draws (dual and triple mandates, Span A) to 28–36% (Span D). Where the share is low, the remaining draws use other sign patterns, so the bars describe the most frequent combination, not a typical draw.

Figure E.1: Realized Responses of Mean Consumption by Income and Wealth



Notes: Response of the mean consumption of households in each income–wealth cell to the oil supply shock, in percent of the cell's steady-state mean consumption: posterior median and 68% band. Rows: income rank bottom 50%, 50–90%, top 10%; columns: the same for wealth. Horizontal axis is time horizon after the shock in quarters.

D Policy bundles

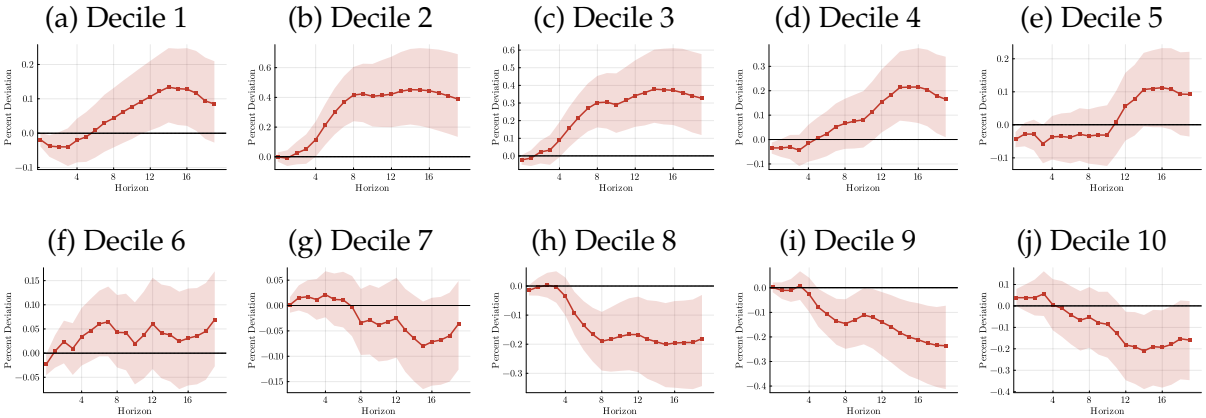
E Inside the distributions

Table E.1: Whose consumption is stabilized? Mean consumption by income and wealth (bottom 40%, 40–80% and top 20%)

	Span A: policy rate			Span D: all levers		
Household cell	Dual	Triple	$\text{Triple}_{\text{rs}}$	Dual	Triple	$\text{Triple}_{\text{rs}}$
Income bottom 40%, wealth bottom 40%	0.45	0.41	0.38	0.58	0.31	0.25
Income bottom 40%, wealth 40–80%	4.52	3.40	3.58	4.59	3.56	3.06
Income bottom 40%, wealth top 20%	1.13	1.08	1.07	1.15	1.10	0.90
Income 40–80%, wealth bottom 40%	0.90	0.87	0.83	1.23	0.99	0.75
Income 40–80%, wealth 40–80%	0.78	0.65	0.60	1.08	0.55	0.40
Income 40–80%, wealth top 20%	0.77	0.72	0.75	0.82	0.78	0.72
Income top 20%, wealth bottom 40%	1.82	1.56	1.61	2.86	2.21	1.86
Income top 20%, wealth 40–80%	2.81	1.41	1.62	3.02	1.65	1.71
Income top 20%, wealth top 20%	1.61	0.69	0.81	1.77	0.65	0.71

Notes: Each entry is the median over posterior draws of $\sum_{h=1}^{12} y_h^2$ for the response of the cell's mean consumption, relative to the realized response within the same draw; 1 means no change and below 1 means more stable. Cells are defined jointly by the household's income and wealth rank: bottom 40%, 40–80% and top 20% of each distribution. Responses are exact reconstructions from the joint distribution of consumption, income and wealth, in percent of the cell's steady-state mean consumption. None of these moments is targeted.

Figure E.2: Realized Responses of Consumption by Decile to an Oil Supply Shock



Notes: Response of the mean consumption of each decile to the oil supply shock, $100 \times \text{asinh}(Q/\mu)$ units (points), with Q the quantile function and μ mean consumption. Posterior median and 68% band. Horizontal axis is time horizon after the shock in quarters.

Table E.2: Inside the distributions: losses of the group means of consumption, income and wealth

		Consumption			Income			Wealth		
Span	Mandate	Bottom 50	Middle 40	Top 10	Bottom 50	Middle 40	Top 10	Bottom 50	Middle 40	Top 10
A	Dual	0.68	1.94	1.33	1.23	1.54	0.98	0.97	0.94	1.03
	Triple	0.35	0.91	1.08	1.12	1.30	0.95	0.88	0.93	1.01
	$\text{Triple}_{\text{rs}}$	0.32	1.03	1.09	1.12	1.35	0.96	0.86	0.90	1.01
B	Dual	0.93	2.38	1.79	1.75	2.85	1.31	1.29	0.91	0.98
	Triple	0.31	1.04	1.16	1.35	1.82	1.21	0.96	0.87	1.04
	$\text{Triple}_{\text{rs}}$	0.23	1.16	0.86	1.14	1.47	0.99	0.87	0.79	0.93
C	Dual	0.80	1.91	1.56	1.52	1.92	1.05	1.21	0.94	1.03
	Triple	0.28	0.93	1.04	1.40	1.67	1.01	1.02	0.90	1.01
	$\text{Triple}_{\text{rs}}$	0.22	1.00	0.89	1.24	1.48	0.96	0.91	0.81	0.95
D	Dual	1.11	2.28	2.09	2.11	3.14	1.38	1.62	0.93	0.96
	Triple	0.27	1.08	1.19	1.62	2.05	1.26	1.12	0.84	0.99
	$\text{Triple}_{\text{rs}}$	0.17	1.08	0.71	1.16	1.41	0.98	0.92	0.71	0.84

Notes: Loss of the response of each group's mean (asinh of the group mean relative to the overall mean, in points) relative to the realized response, median of per-draw ratios over $h = 1, \ldots, 12$. None of these moments is targeted. Realized losses (median over draws, points^2): consumption b50 0.218, m40 0.046, t10 0.191; income b50 0.094, m40 0.174, t10 0.800; wealth b50 0.088, m40 9.274, t10 3.118.

F Robustness

Table F.1: Robustness: rules enforced from the impact quarter

Span	Enforcement window	Single$_\pi$: π	Dual: π / u	Dual: spread	Triple: π / u	Triple: spread	Triple$_{rs}$: risk s
A	rules from $h = 1$ (baseline)	0.94	1.28 / 0.32	0.67	1.17 / 0.56	0.22	0.50
A	rules from impact	0.94	1.22 / 0.37	0.69	1.18 / 0.57	0.24	0.55
B	rules from $h = 1$ (baseline)	0.78	1.14 / 0.21	0.86	1.19 / 0.45	0.16	0.39
B	rules from impact	0.80	1.12 / 0.25	0.89	1.19 / 0.47	0.18	0.42
C	rules from $h = 1$ (baseline)	0.81	1.18 / 0.25	0.74	1.24 / 0.45	0.14	0.38
D	rules from $h = 1$ (baseline)	0.66	1.02 / 0.16	0.94	1.20 / 0.36	0.11	0.31
D	rules from impact	0.65	0.96 / 0.20	1.00	1.14 / 0.38	0.12	0.34

Notes: Medians of per-draw loss ratios, each relative to the realized loss over the rule's own window: $h = 1, \ldots, 12$ (baseline) or $h = 0, \ldots, 11$ (i impact). Span C was run only from $h = 1$.

G Description of Distributional Data

Description Consider the joint distribution $\Xi_t(\boldsymbol{w})$ over $\boldsymbol{w} \in \mathbb{R}^3$, for consumption, income, wealth. A joint distribution is fully described by two components: its time-varying support in *levels* and the *mass* associated with the time-varying support. Sklar (1959, 1973) formalized this description, and with an adjustment, we can decompose the joint distribution into its marginal quantile functions $\Xi^{-1}_{m,t}$ for $m \in \{c, y, w\}$ and a copula density $c_t : [0,1]^3 \to [0, \infty)$. The copula encodes the dependence structure independently of the marginals and maps to the *mass*, while the quantile functions map to the *levels*, with both functions defined on the now same time-invariant support: the uniform distribution over ranks. In words, the quantile functions tell us how much a household consumes, earns, and has in wealth at each rank, while the copula tells us the mass of households occupying each combination of ranks.

Achieving a finite dimensional representation These distributional objects are infinite dimensional, making their inclusion in the model infeasible. Additionally, we are not interested in partial representations of the distribution e.g., including shares, ginis, level variables, etc in the VAR. The first step in achieving then a *finite* representation of these objects, while maintaining a *functional* approach, is to project each marginal quantile function $\Xi^{-1}_{m,t}$ and the copula density c_t onto an orthonormal basis of Legendre polynomials $\{Q_o\}_{o \geq 0}$, shifted to reside on the same support of the distributional objects, $[0, 1]$:

$$\Xi^{-1}_{m,t}(u_m) = \sum_{o=0}^{\infty} \xi_{m,o,t}\, Q_o(u_m), \qquad u_m \in [0,1], \tag{G.1}$$

$$c_t(u_1, \ldots, u_d) = \sum_{(o_1,\ldots,o_d) \in \mathbb{N}_0^d} \kappa_{(o_1,\ldots,o_d),t} \prod_{m=1}^{d} Q_{o_m}(u_m). \tag{G.2}$$

This means any distributional object is an infinite expansion of polynomials $\{Q_o\}_{o \geq 0}$, weighted

by polynomial coefficients $\xi_{m,o,t}$ and $\kappa_{(o_1,\ldots,o_d),t}$. With the polynomial basis fixed and time-invariant, the distributional data observed at survey dates can be reduced to estimates of the aforementioned coefficients. In that sense, it will be these coefficients that carry the business-cycle fluctuations of the joint distribution of consumption, income, and wealth, so the remainder of the section is dedicated to their estimation and treatment.

Each coefficient is the inner product of some Legendre polynomial Q_o with the corresponding distributional object— be it $c_t(u_1,\ldots,u_d)$ or $\Xi^{-1}_{m,t}(u_m)$— for u_m some rank, *divided* by squared norm of the polynomial $\|Q_o\|^2$. Since this is a projection onto an orthonormal basis, the normality aspect informs us that the squared norm is 1. Concretely, $Q_o(u) = \sqrt{2o+1}\,P_o(2u-1)$, where P_o is the standard Legendre polynomial on $[-1,1]$, so that $\int_0^1 Q_o(u)\,Q_{o'}(u)\,du$ is one if $o = o'$ and zero otherwise: the basis is orthonormal on $[0,1]$ under the uniform measure, not on $[-1,1]$. Assuming the distributional objects are square-integrable, and hence belong to $\mathcal{L}^2$, orthonormality implies that each coefficient is simply the expectation over the product of the object with the corresponding basis polynomial. The corresponding sample analogue is just a sample average over ranked microdata, shown in Equations G.3 and G.4:

$$\hat{\xi}_{m,o,t} = \sum_i \omega_{i,t}\, w_{m,i,t}\, Q_o(u_{m,i,t}) \qquad \text{for the quantile function} \tag{G.3}$$

$$\hat{\kappa}_{(o_1,\ldots,o_d),t} = \sum_i \omega_{i,t} \prod_m Q_{o_m}(u_{m,i,t}) \qquad \text{for the copula density} \tag{G.4}$$

where $u_{m,i,t}$ is the data rank, derived from the cumulative sample weight at observation i after sorting the dimension m in the data and $w_{m,i,t}$ the data level/quantile value corresponding to the rank for dimension m. The survey weights $\omega_{i,t}$ are normalized to sum to one, so the sample averages are population-weighted. Because the ranks are built from the same cumulative weights, the weighted ranks are uniform on $[0,1]$ by construction, up to the discreteness of the sample, so the weighted sums approximate integrals against the uniform measure under which the Q_o are orthonormal. The full derivation is in Bayer, Calderon, and Kuhn (2025). Levels enter in inverse hyperbolic sine units relative to the cross-sectional mean, $w_{m,i,t} = \operatorname{asinh}(x_{m,i,t}/\mu_{m,t})$, so $\Xi^{-1}_{m,t}$ is the quantile function of the transformed variable. Any statistic that is linear in it, such as the consumption spread $100\,[\Xi^{-1}_{c,t}(0.9) - \Xi^{-1}_{c,t}(0.1)]$ targeted in the results, is therefore exactly linear in the coefficients $\xi_{c,o,t}$ and hence in the VAR; statistics such as the Gini coefficient are not.

Still, Equations G.1 and G.2 are infinite sums, implying to track the distribution, we need to track infinitely many coefficients—infeasible. To achieve the finite representation, the sums of Equations G.1 and G.2 are *truncated* to some order O, providing an approximation of the

distributional objects.

$$\hat{\boldsymbol{\xi}}_{m,t} = \left(\hat{\xi}_{m,0,t}, \ldots, \hat{\xi}_{m,O,t}\right)', \tag{G.5}$$

$$\hat{\boldsymbol{\kappa}}_t = \left[\hat{\kappa}_{(o_1,\ldots,o_d),t}\right]_{(o_1,\ldots,o_d)\in\{0,\ldots,O\}^d}, \tag{G.6}$$

This truncation leaves us with two column vectors of finitely-many coefficients G.5 and G.6 to be estimated, achieving the finite representation required to estimate the model. This truncation is not just for feasibility; it is also appropriate. Here, the coefficients $\xi_{m,o,t}$ and $\kappa_{(o_1,\ldots,o_d),t}$ have the property that they non-monotonically become small in size at sufficiently higher orders. Moreover, the squared coefficient gives that polynomial's contribution to the squared L^2 norm of the object. Putting these two properties together implies the contribution of these omitted higher-order components is small in their contribution to the function; similar to an R^2 interpretation. Truncation then avoids estimating coefficients that contribute little to the underlying distributional object. Lastly, by orthogonality of the polynomials, the kept coefficients are not affected by the omission of higher order polynomials.

$$\boldsymbol{\theta}_t \;=\; \left[\; \hat{\boldsymbol{\xi}}'_{1,t},\; \ldots,\; \hat{\boldsymbol{\xi}}'_{d,t},\;\; \mathrm{vec}(\hat{\boldsymbol{\kappa}}_t)' \;\right]' \;\in\; \mathbb{R}^N \tag{G.7}$$

Once vertically concatenated, they form a finite column vector $\boldsymbol{\theta}_t \in \mathbb{R}^N$, with $N = (O+1)^d + d \cdot (O+1)$ entries: quantile coefficients across d dimensions plus copula coefficients.

Representation in BVAR What will be important to know is how these data are represented in the Bayesian VAR. At the moment, to estimate the model with the coefficients $\boldsymbol{\theta}_t$ directly, although finite, would result in far too many equations. To reduce the dimensionality of the system, I exploit the correlational structure of $\boldsymbol{\theta} = [\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_T]'$ and extract the principal components of the coefficient matrix via PCA. From Bayer, Calderon, and Kuhn (2025), I know the data generation process of the joint distribution(s) estimated here and thus also know the first $r = 8$ principal components, $\boldsymbol{F}_t = (F_{1t}, \ldots, F_{rt})'$, capture all the relevant business cycle variation in the consumption, income, and wealth distribution. It is precisely these eight factors that enter the baseline Bayesian VAR. They enter as point estimates; Appendix I states which sources of uncertainty the reported bands carry. The treatment of the distributional data thereafter ensures impulse responses still retain their properties as a distribution.[20]

[20]Specifically, the reconstruction from factor estimates to distributional data preserves the properties of the underlying distributional objects across the impulse-response horizon $h \in 0, \ldots, H$: the reconstructed marginal quantile functions $\hat{\Xi}^{-1}_{m,h}$ are weakly monotone in $u \in [0,1]$ for each m; the reconstructed copula density $\hat{c}_h$ is

H Operationalization of concepts

The distributional factor estimates $\boldsymbol{F_t}$ rely on U.S. micro-data, U.S. macro data, a state-equation dictating the law of motion for the distribution with aggregates, and a measurement equation which optimally weights the contribution of each data source to each estimate at each point in time, and which also reconciles differences in how each underlying micro-source operationalizes consumption, income, and wealth. In that sense, the operationalization of each concept — consumption, income, wealth — is a convex combination of these data sources and thus a consensus estimate for the unobserved distributional object.

Transforming these factors to observables, however, requires reversing some data treatment. For example, there is a de-standardization step, where the means and standard deviations estimated from a *specific* micro-dataset is re-applied back to the consensus estimates. This choice of de-standardization dataset enters only through the steady-state mean and unconditional volatility of each variable; the dynamics of the responses I report below are unaffected, since the measurement equation has already reconciled the operationalization differences across the underlying micro-sources. Since I am interested in the joint consumption, income, and wealth distribution, I use means and standard deviations from the household panel micro-data in Bayer, Calderon, and Kuhn (2025) that measures all three. To nonetheless provide some idea of the underlying variation constituting these estimates, I describe how these data operationalize each concept.

In these data, I measure **consumption** as containing food, rent (for renters) or housing rental equivalence (6% of home market value, for homeowners), utilities, health care, public transportation, education, childcare, and gas/vehicle repairs. **Income** contains labor earnings, public transfers, professional/self-employment income, rental income, dividends, and business/farm income, but is not net of taxes. **Wealth** is total assets minus total debt. Assets comprise liquid assets (checking, savings, CDs, money market, T-bills) plus illiquid assets (housing net of property debt, real estate, automobiles, retirement plans, life insurance, stocks, mutual funds, business equity). Debt comprises housing debt (mortgages, home equity lines of credit) plus personal debt (car loans, student loans, credit cards, medical, legal).

I Reconstruction of Distributional Objects

This appendix details how impulse responses estimated in factor space are mapped back to distributional objects—quantile functions, copula densities, and household-level moments—

non-negative on the trimmed cube $[\varepsilon, 1-\varepsilon]^d$; and the copula integrates to unity, $\int [0,1]^d \hat{c}_h = 1$. Uniform marginals follow from restrictions on the coefficients. Monotonicity and non-negativity are verified ex post on a G^d tensor grid; deviations from the strict constraints are below 10^{-6} and treated as numerical noise.

following the methodology of Bayer, Calderon, and Kuhn (2025).

I.1 From Factor IRFs to Coefficient IRFs

The VAR is estimated on a small number of distributional factors $\boldsymbol{F}_{D,t} \in \mathbb{R}^r$, with $r = 8$ in the baseline specification. Let $\hat{\boldsymbol{\theta}}_t \in \mathbb{R}^N$ denote the raw Legendre polynomial coefficient vector, with $N = (O+1)^d + d \cdot (O+1)$ collecting copula and marginal quantile-function coefficients (cf. Section 3). The factors are obtained via principal components analysis on the *standardized* coefficients $\boldsymbol{\Sigma}_\theta^{-1/2}\big(\hat{\boldsymbol{\theta}}_t - \bar{\boldsymbol{\theta}}\big)$, so the implied factor model for the raw coefficients is

$$\hat{\boldsymbol{\theta}}_t \;=\; \bar{\boldsymbol{\theta}} \;+\; \boldsymbol{\Sigma}_\theta^{1/2}\big(\boldsymbol{\Lambda}\,\boldsymbol{F}_{D,t} + \boldsymbol{\gamma}\,\boldsymbol{f}_t\big), \tag{I.1}$$

where $\bar{\boldsymbol{\theta}}$ is the sample mean, $\boldsymbol{\Sigma}_\theta^{1/2}$ is the diagonal matrix of per-coefficient standard deviations, $\boldsymbol{\Lambda} \in \mathbb{R}^{N\times r}$ is the factor loading matrix, and $\boldsymbol{f}_t$ captures idiosyncratic variation orthogonal to the common factors.

Mixed-frequency measurement and the four-quarter projection A subtle point governs the loading used for reconstruction. Bayer, Calderon, and Kuhn (2025) estimate the factor model in a mixed-frequency state-space framework: the latent factors $\boldsymbol{F}_t$ live at quarterly frequency, but the coefficient observations from the household panel survey are annual averages. The measurement equation for that dataset takes the form

$$\tilde{\boldsymbol{\theta}}_t^{\text{micro}} \;=\; \frac{1}{4}\,\boldsymbol{\Lambda}\left(\boldsymbol{F}_t + \boldsymbol{F}_{t-1} + \boldsymbol{F}_{t-2} + \boldsymbol{F}_{t-3}\right) \;+\; \boldsymbol{\nu}_t^{\text{micro}}, \tag{I.2}$$

where $\boldsymbol{\nu}_t^{\text{micro}}$ captures sampling and operationalisation noise. Equation (I.2) is the source of a reconstruction issue when the smoothed factors enter the VAR as observables. Applying $\boldsymbol{\Lambda}$ alone to a quarterly factor IRF produces an over-attenuated coefficient response, because each quarter of the IRF activates a single factor rather than the four-quarter sum the loadings $\boldsymbol{\Lambda}$ were estimated to fit. To recover the quarterly-frequency loading that maps a single quarter of factor variation into its associated coefficient impact, I fit by ordinary least squares

$$\hat{\boldsymbol{\Lambda}}_{4q} \;=\; \arg\min_{\boldsymbol{\Lambda}} \sum_t \left\| \hat{\boldsymbol{\theta}}_{t|T}^{\text{micro}} \;-\; \boldsymbol{\Lambda}\,\tfrac{1}{4}\big(\hat{\boldsymbol{F}}_{t|T} + \hat{\boldsymbol{F}}_{t-1|T} + \hat{\boldsymbol{F}}_{t-2|T} + \hat{\boldsymbol{F}}_{t-3|T}\big) \right\|^2, \tag{I.3}$$

where $\hat{\boldsymbol{\theta}}_{t|T}^{\text{micro}}$ and $\hat{\boldsymbol{F}}_{t|T}$ are the BCK smoother's posterior-mean reconstructions of the survey coefficient and the latent quarterly factor at quarter t, respectively. The four-quarter rolling sum on the right-hand side enforces the BCK measurement equation (I.2); OLS thus recovers the loading $\hat{\boldsymbol{\Lambda}}_{4q}$ that closes the measurement-equation residual. Two points matter for reading

the quarterly responses. First, the scale of the factors does not matter: $\hat{\mathbf{\Lambda}}_{4q}$ is estimated by least squares on the same smoothed factors that enter the VAR, so any rescaling of the factors is undone by the loading, and (I.4) returns coefficient responses in raw units. Second, as written, (I.2) treats every coefficient of the survey, marginal and copula alike, as a four-quarter average, and (I.4) applies $\hat{\mathbf{\Lambda}}_{4q}$ to the whole coefficient vector; the correction is therefore the same for all coefficient series.

Mapping factor IRFs to coefficient IRFs Quarterly factor responses $\mathbf{\Psi}_h^F$ are mapped to raw-coefficient responses by

$$\Delta\hat{\boldsymbol{\theta}}_h \;=\; \mathbf{\Sigma}_\theta^{1/2}\,\hat{\mathbf{\Lambda}}_{4q}\,\mathbf{\Psi}_h^F, \tag{I.4}$$

which undoes the per-coefficient standardization $\mathbf{\Sigma}_\theta^{1/2}$ from the PCA. The result $\Delta\hat{\boldsymbol{\theta}}_h$ is the IRF of the polynomial coefficients in raw units; the steady-state mean is not reintroduced.

Reconstructing nonlinear functionals To evaluate distributional functionals ϕ (quantile values, copula densities, household-group means, inequality indices) at horizon h, the reconstruction is non-linear in the coefficient vector and therefore requires the level coefficient

$$\hat{\boldsymbol{\theta}}_h \;=\; \bar{\boldsymbol{\theta}} + \Delta\hat{\boldsymbol{\theta}}_h, \tag{I.5}$$

where $\bar{\boldsymbol{\theta}}$ is the time-invariant steady-state coefficient (including any trend component removed prior to estimation). The reconstructed objects are evaluated via (I.7)–(I.8) below at the level $\hat{\boldsymbol{\theta}}_h$, and the IRF of any functional ϕ is reported as

$$\Psi_h^\phi \;=\; \phi(\hat{\boldsymbol{\theta}}_h) \;-\; \phi(\bar{\boldsymbol{\theta}}). \tag{I.6}$$

For functionals that are linear in $\boldsymbol{\theta}$ (the marginal quantile values evaluated pointwise via (I.7)), this reduces to direct evaluation at $\Delta\hat{\boldsymbol{\theta}}_h$. For copula-based functionals (household-group means, conditional moments, Ginis), the level coefficient must be used because monotonicity and non-negativity constraints couple the reconstruction nonlinearly across coefficients. The same applies to the counterfactuals. The consumption spread targeted by the distributional mandates is linear in the quantile coefficients (Appendix G), so choosing the policy bundle involves no approximation for it. The outcomes that are not targeted and are nonlinear, such as the mean consumption of an income–wealth cell, which is a ratio of copula integrals, are not approximated to first order: for each posterior draw, the counterfactual coefficient path is passed through (I.5)–(I.6) and the functional is recomputed exactly.

I.2 Reconstructing Quantile Functions and Copula Densities

The coefficient vector $\hat{\theta}_h$ partitions into copula coefficients $\kappa_h \in \mathbb{R}^{(O+1)^d}$ and marginal quantile-function coefficients $\xi_h^m \in \mathbb{R}^{O+1}$ for each dimension $m \in \{\text{consumption, income, wealth}\}$. These reconstruct the distributional objects via the Legendre polynomial basis $\{Q_o\}_{o=0}^{O}$:

Marginal quantile functions For dimension m evaluated on a grid $u \in [0, 1]$:

$$\hat{\Xi}_{m,h}^{-1}(u) = \sum_{o=0}^{O} \xi_{o,h}^m \, Q_o(u). \tag{I.7}$$

In practice, I evaluate this on a grid of $G = 20$ equally spaced quantile points (vigintiles), giving the income, consumption, or wealth level at each vigintile of the distribution.

Copula density The joint dependence structure is reconstructed as:

$$d\hat{C}_h(u_1, u_2, u_3) = \sum_{o_1,o_2,o_3=0}^{O} \kappa_{(o_1,o_2,o_3),h} \prod_{m=1}^{3} Q_{o_m}(u_m). \tag{I.8}$$

This is evaluated on the $G^3 = 8{,}000$ grid points of the three-dimensional unit cube, giving the joint density at each combination of consumption, income, and wealth vigintiles.

I.3 Household Groups and Conditional Moments

A household group is defined by a range of quantiles in each dimension. For a group occupying quantile range $U_i^c \times U_i^y \times U_i^w$ (e.g., the bottom two income vigintiles crossed with all wealth vigintiles), I compute:

Mass (population share) The fraction of households in the cell:

$$\omega_{i,h} = \iiint_{(u_c,u_y,u_w)\in U_i^c \times U_i^y \times U_i^w} d\hat{C}_h(u_c, u_y, u_w). \tag{I.9}$$

Conditional mean consumption The average consumption of households in the cell:

$$\bar{c}_{i,h} = \frac{1}{\omega_{i,h}} \iiint_{U_i^c \times U_i^y \times U_i^w} \hat{\Xi}_{c,h}^{-1}(u_c) \, d\hat{C}_h(u_c, u_y, u_w). \tag{I.10}$$

Income and wealth conditional means are defined analogously by replacing $\hat{\Xi}_{c,h}^{-1}$ with the corresponding marginal quantile function.

Numerical implementation The continuous integrals (I.9)–(I.10) are discretized on the $G \times G \times G$ vigintile grid. Let $\mathcal{G}_i \subseteq \{1, \ldots, G\}^3$ denote the grid indices belonging to household group i, and write $u_g = (g - 1/2)/G$ for the midpoint of vigintile g. The discrete analogues of (I.9) and (I.10) are

$$\hat{\omega}_{i,h} = \frac{1}{G^3} \sum_{\mathbf{g} \in \mathcal{G}_i} d\hat{C}_h(u_{g_c}, u_{g_y}, u_{g_w}), \qquad \hat{\bar{c}}_{i,h} = \frac{1}{G^3\, \hat{\omega}_{i,h}} \sum_{\mathbf{g} \in \mathcal{G}_i} \hat{\Xi}^{-1}_{c,h}(u_{g_c})\, d\hat{C}_h(u_{g_c}, u_{g_y}, u_{g_w}), \tag{I.11}$$

with $\mathbf{g} = (g_c, g_y, g_w)$. The copula density is rescaled across the full grid so that $\sum_{\mathbf{g}} d\hat{C}_h(u_{g_c}, u_{g_y}, u_{g_w}) = G^3$, ensuring $\sum_i \hat{\omega}_{i,h} = 1$ when groups partition the unit cube. This rescaling is applied at every horizon and for every draw; non-negativity and monotonicity are checked on the same grid rather than imposed by a projection (see the footnote in Appendix G). The household cells in the paper, the bottom 50%, 50–90% and top 10% of income and of wealth, and the bottom 40%, 40–80% and top 20% in the appendix, have boundaries at multiples of 0.05. Each cell is therefore a union of whole vigintiles and no interpolation within a vigintile is needed; the remaining error is that of the midpoint rule within each vigintile.

I.4 Inequality Measures

Inequality statistics are computed from the reconstructed marginal quantile functions. For a given dimension m at horizon h:

Gini coefficient Evaluated from the quantile function on a fine grid of $n = 200$ equally spaced points $\{u_j\}_{j=1}^n$:

$$G_{m,h} = 1 - 2 \int_0^1 L_{m,h}(p)\, dp, \qquad L_{m,h}(p) = \frac{\int_0^p \hat{\Xi}^{-1}_{m,h}(u)\, du}{\int_0^1 \hat{\Xi}^{-1}_{m,h}(u)\, du}, \tag{I.12}$$

where $L_{m,h}(p)$ is the Lorenz curve. In practice, the integrals are approximated by cumulative sums over the grid: $L_{m,h}(p_j) = \sum_{k=1}^{j} Q_k / \sum_{k=1}^{n} Q_k$, where $Q_k = \hat{\Xi}^{-1}_{m,h}(u_k)$.

Percentile shares The share of dimension m held by a quantile range $[a, b] \subseteq [0, 1]$ is the corresponding slice of the Lorenz integral,

$$S^{[a,b]}_{m,h} = \frac{\int_a^b \hat{\Xi}^{-1}_{m,h}(u)\, du}{\int_0^1 \hat{\Xi}^{-1}_{m,h}(u)\, du}, \tag{I.13}$$

so the top-10% share is $S^{[0.9,1]}$ and the bottom-50% share is $S^{[0,0.5]}$, both evaluated on the same $n = 200$ grid. The bottom-50 consumption share is the target row of the bottom-50 pin.

Consumption spread (the policy target). The statistic that enters the policy loss is the 90/10 spread in inverse hyperbolic sine units, $100\,[\text{asinh}(Q_{90}/\mu) - \text{asinh}(Q_{10}/\mu)] = 100\,[\hat{\Xi}^{-1}_{c,h}(0.9) - \hat{\Xi}^{-1}_{c,h}(0.1)]$, with Q_p the p-th percentile of consumption and μ its mean. It is linear in the quantile coefficients (Appendix G). The log ratios below are a different statistic, reported for description only; none of them enters a loss.

Log interquantile ratios The 90/10 and 90/50 ratios are:

$$\log 90/10_{m,h} = \log \hat{\Xi}^{-1}_{m,h}(0.9) - \log \hat{\Xi}^{-1}_{m,h}(0.1), \qquad \log 90/50_{m,h} = \log \hat{\Xi}^{-1}_{m,h}(0.9) - \log \hat{\Xi}^{-1}_{m,h}(0.5). \tag{I.14}$$

Standard deviation of logs

$$\text{sdlog}_{m,h} = \text{sd}\left(\left\{\log \hat{\Xi}^{-1}_{m,h}(u_j)\right\}_{j=1}^{n}\right). \tag{I.15}$$

All inequality measures are reported as percent deviations from their steady-state values: $100 \times (G_{m,h} - G_{m,\text{ss}})/G_{m,\text{ss}}$.

I.5 Computation Summary

Table I.1: Reconstruction parameters

Parameter	Value
Polynomial order (O)	11
Distribution dimensions (d)	3 (consumption, income, wealth)
Total coefficients ($N = (O+1)^d + d \cdot (O+1)$)	1,764
Copula coefficients ($(O+1)^d$)	1,728
Marginal coefficients ($d \cdot (O+1)$)	36
Distributional factors (r)	8
Variance explained by r factors	$> 95\%$
Evaluation grid (G)	20 (vigintiles)
Grid points for copula (G^3)	8,000
Gini evaluation grid (n)	200
Household groups	3 income $\times$ 3 wealth $= 9$ cells, in two partitions
Posterior draws	5,000

For each of the 5,000 posterior draws, the full reconstruction pipeline (factor IRF $\rightarrow$ coefficient IRF $\rightarrow$ quantile functions + copula density $\rightarrow$ household moments + inequality measures)

is executed at each horizon $h = 1, \ldots, 20$. The reported impulse responses are posterior medians with 68% credible sets computed pointwise across draws. These bands carry the posterior uncertainty of the BVAR only. The distributional factors enter the VAR as the smoother's posterior means, and the loading $\hat{\mathbf{\Lambda}}_{4q}$, the mean $\bar{\boldsymbol{\theta}}$ and the scale $\mathbf{\Sigma}_\theta^{1/2}$ are held at their point estimates in every draw, so the estimation uncertainty of the first-stage factors and of the copula and quantile coefficients is not propagated. Bands for the household cells are therefore conditional on the first stage. The tensor contraction over the G^d grid is the computational bottleneck; the implementation exploits the separable structure of the Legendre basis—each factor $Q_{o_m}(u_m)$ depends on only one dimension—to reduce the $O\big(G^d(O+1)^d\big)$ naïve evaluation to $O\big(G^d(O+1)\big)$ operations per draw via successive 1D contractions along each dimension.

Formally, the reconstruction proceeds as follows. Given the estimated factors $\hat{\boldsymbol{F}}_t$, I first recover the standardized coefficients, then undo normalization and trend adjustments. In the notation of (I.1)–(I.4), $\hat{\Lambda}$ below is the four-quarter loading $\hat{\mathbf{\Lambda}}_{4q}$, $\mu_n + g(t)_n$ is the n-th entry of $\bar{\boldsymbol{\theta}}$, and $\sigma_{\zeta(n)}$ is the corresponding diagonal entry of $\mathbf{\Sigma}_\theta^{1/2}$:

$$\hat{\tilde{\theta}}_t = \hat{\Lambda}\,\hat{\boldsymbol{F}} \qquad \text{(Project factors)} \tag{I.16}$$

$$\hat{\theta}_{nt} = \sigma_{\zeta(n)}\hat{\tilde{\theta}}_{\zeta(n)t} + \mu_n + g(t)_n \qquad \text{(Unstandardize and add trend)} \tag{I.17}$$

$$\hat{\theta}_t = (\hat{\xi}_{m,o_1,t}, \ldots, \hat{\kappa}_{(o_1,\ldots,o_d),t}) \qquad \text{(Decomposition)} \tag{I.18}$$

$$\tag{I.19}$$

See Table I.2 for details on notation. To generate economically interpretable objects—quantile functions, copula densities, Gini coefficients, percentile shares, and other moments—these reconstructed objects are evaluated on $u \in [0,1]^d$, with boundary points excluded in practice (the grid uses $[10^{-6}, 0.9995]$ for numerical stability) and then integrated to obtain the distributional statistics reported in the body.

Table I.2: Reconstruction of Distributional Objects: Notation

Symbol	Description
$\hat{\boldsymbol{F}}_t$	Estimated high-frequency factors
$\hat{\Lambda}$	Factor loadings mapping factors to coefficients; the four-quarter loading $\hat{\boldsymbol{\Lambda}}_{4q}$ of (I.3)
$\hat{\tilde{\theta}}_t$	Standardized coefficient representation
$\hat{\theta}_{nt}$	Final coefficient vector for coefficient n
$\zeta(n)$	denotes a distributional object—either a quantile function or a copula, which consists of n coefficients
$\sigma_{\zeta(n)}$	Standard deviation *per* dist. object
μ_n	Mean adjustment *per* coefficient
$g(t)_n$	Trend component per coefficient
$\hat{\xi}_{m,o,t}$	Marginal polynomial coefficients
$\hat{\kappa}_{(o_1,\ldots,o_d),t}$	Copula coefficients
$Q_o(u_m)$	Legendre polynomial basis function

Notes: The table summarizes the mapping from factor estimates to reconstructed distributional objects.