Bank of Israel
Research Department

# What Drives the Scenario? Interpreting Conditional Forecasts in Reduced-Form VARs

Occasional Papers Series

Itamar Caspi, Tim Ginker

July 2026

**2026.01**

# What Drives the Scenario? Interpreting Conditional Forecasts in Reduced-Form VARs*

Itamar Caspi[1], Tim Ginker[2]

**Abstract**

This paper develops methods to analyze and quantify the dynamic effects of path conditioning in conditional forecasts from reduced-form vector autoregressive (VAR) models. Building on Kalman filtering and the observation-weight extraction method of Koopman and Harvey (2003), we characterize how imposed future paths of observables influence forecast outcomes. We establish the analytical relationship between observation weights and unscaled generalized impulse response functions (GIRFs): under single-period conditioning, the weights coincide with unscaled GIRFs, while multi-period path conditioning generates a richer, non-local structure of influence. We use these weights to construct measures of overall and marginal variable importance that disentangle model-implied sensitivity from realized sample dynamics. An empirical illustration shows how the proposed framework enhances the interpretability of scenario forecasts by attributing forecast revisions to specific variables and horizons. The approach is model-consistent, structurally agnostic, and well suited for policy-oriented scenario analysis.

**Keywords:** Vector Autoregression; Conditional Forecast; Filter Weights; Impulse Response Analysis.

---

Bank of Israel - http://www.boi.org.il

* We thank Jamel Saadaoui for serving as discussant and for his insightful comments and suggestions. We are grateful to Eyal Argov for his review and advice. We also thank the participants of the Bank of Israel research seminar for helpful comments and discussions.

[1] Bank of Israel Email: Itamar.Caspi@boi.org.il

[3] Bank of Israel Email: ginker.tim@boi.org.il

# מה מניע את התרחיש? מתן פרשנות לתחזיות מותנות במודלי VAR מהצורה המצומצמת

איתמר כספי, טים גינקר


## תקציר

מאמר זה מפתח שיטות לניתוח ולכימות ההשפעות הדינמיות של האילוצים בתחזיות מותנות המיוצרות על ידי מודל vector autoregression (VAR). בהתבסס על מסנן קלמן ועל השיטה לחילוץ משקלי תצפיות של Koopman and Harvey (2003), אנו מראים כיצד האילוצים על ההתפתחות העתידית של המשתנים במערכת משפיעים על התחזית המותנית. בנוסף, אנו מראים את הקשר האנליטי בין משקלי התצפיות המאולצות לבין ה-generalized impulse response functions (GIRFs): כאשר ההתניה היא על תקופה בודדת, המשקלים חופפים לפונקציות GIRF לא מתוקננות, ואילו התניה רב־תקופתית יוצרת מבנה השפעה עשיר יותר. אנו משתמשים במשקלים אלה לבניית מדדים לחשיבות הכוללת והשולית של משתנים בתחזית, המפרידים בין הרגישות הנובעת מן המודל לבין הדינמיקה שהתממשה במדגם. יישום אמפירי ממחיש כיצד המתודולוגיה המוצעת משפרת את יכולת הניתוח של תחזיות מותנות באמצעות ייחוס עדכוני התחזית למשתנים ולאופקי חיזוי ספציפיים. הגישה עקבית עם המודל, אינה נשענת על זיהוי מבני, ומתאימה במיוחד לניתוח תרחישים המכוון לצורכי מדיניות.


**מילות מפתח:** vector autoregression; תחזית מותנית; משקלי מסנן קלמן; ניתוח תגובות לזעזועים.

**הדעות המובעות במאמר זה אינן משקפות בהכרח את עמדתו של בנק ישראל**

# 1 Introduction

In macroeconomic forecasting, conditional projections—often referred to as scenario analyses—are forecasts for a subset of target variables, conditional on assumed future paths for a selected set of other endogenous variables within the system. These scenario-based forecasts play a central role in modern monetary policymaking and communication, and central banks now routinely publish them to inform policy decisions and shape expectations (Binder and Sekkel, 2024).

Conditional forecasts are typically produced using structural models, such as dynamic stochastic general equilibrium (DSGE) models or identified structural vector autoregressions (SVARs), which provide a coherent mapping from economic narratives and expert judgment to forecast outcomes. This mapping is valuable for policy analysis, but it typically depends on identifying restrictions, model structure, and the estimation sample (Antolin-Diaz et al., 2021). Reduced-form vector autoregressions (VARs) are therefore often used as benchmarks or robustness checks (Binder and Sekkel, 2024). Their appeal is that they impose fewer structural restrictions and may be less sensitive to model misspecification, making them useful conservative benchmarks in conditional forecasting exercises. This comes at the cost of interpretability, since the reduced-form model does not make explicit how imposed paths of observables translate into forecast outcomes.

The purpose of this paper is to address this gap by developing a framework for interpreting conditional forecasts in reduced-form VARs. We characterize how imposed future paths of observables propagate through the system using observation weights from the Kalman smoother (Koopman and Harvey, 2003), and show that these weights are formally related to generalized impulse response functions (Pesaran and Shin, 1998). From an applied perspective, the framework provides an economic rationale for why a scenario produces a given forecast revision and helps identify which conditioning assumptions are material for scenario design and monitoring, thereby improving the transparency and communication of reduced-form scenario forecasts.

Various algorithms have been developed to produce conditional forecasts from VAR models. In Bayesian implementations, simulation-based methods introduced by Waggoner and Zha (1999) are widely used, while more computationally efficient state-space methods based on Kalman filtering have been developed to handle larger systems (Bańbura et al., 2015). Notably, Kalman filtering techniques were already applied in early frequentist work by Clarida and Coyle (1984) to compute conditional forecasts. These methods are now standard, and recent empirical applications include Angelini et al. (2019), Bańbura et al. (2015), Giannone et al. (2014), and Bloor and Matheson (2011), among others, who employ conditional VAR forecasts in policy contexts.

A crucial aspect of scenario analysis is the ability to provide an economic interpretation of the effects of different conditioning assumptions. While this is relatively straightforward in SVARs, where identified shocks have clear economic meanings, there has been limited progress in developing analogous interpretative tools for path conditioning in reduced-form VARs. In the absence of structural assumptions, conditional forecasts are typically evaluated based on their out-of-sample accuracy—either using realized values of the conditioning variables or using the scenario paths assumed in real time (Angelini et al., 2019). The former approach evaluates the information content of the conditioning set and the model's ability to translate those assumptions into accurate forecasts, whereas the latter provides a real-time measure of forecast accuracy conditional on the information available at the time.

VAR models are inherently high-dimensional, and researchers typically summarize their dynamics using impulse response functions (IRFs) and forecast error variance decompositions (FEVDs). In the absence of full structural identification, generalized versions—namely the generalized IRF (GIRF) and generalized FEVD (GFEVD) of Pesaran and Shin (1998)—are commonly employed. While useful for tracing the propagation of shocks, standard IRFs and variance decompositions are not designed to evaluate the effects of *path conditioning*. Imposing an assumed path on certain observables influences not only future outcomes but also the inferred historical dynamics of other variables, introducing a bidirectional dependence that GIRFs cannot capture. Moreover, these tools do not quantify the marginal contribution of each element in a conditioning path, nor do they provide a framework for assessing the relative importance of different conditioning variables.

Decompositions and measures of the sensitivity of conditional forecasts to alternative conditioning assumptions are important for several reasons. First, they facilitate the economic interpretation and communication of scenario forecasts by identifying which assumptions are quantitatively driving the results. Second, they help practitioners focus scenario design and monitoring efforts on variables that materially affect the forecast, rather than on assumptions with negligible influence. Finally, sensitivity measures may help assess the adequacy of the forecasting framework itself: economically implausible importance weights, or substantial shifts in their distribution following major shocks or structural changes, may indicate instability, structural breaks, or misspecification in the underlying model.

To address these challenges, this paper develops a framework for interpreting conditional forecasts in reduced-form VAR models. Building on the observation-weight extraction approach of Koopman and Harvey (2003), we characterize how imposed future paths propagate through the system under path conditioning. We further show that the resulting weights are closely related to generalized impulse response functions, thereby linking conditional forecast decompositions to standard tools used in VAR analysis. The framework also provides horizon-specific measures of the contribution and relative importance of conditioning variables, improving the transparency and interpretation of reduced-form scenario forecasts.

Our approach is related to Bańbura and Rünstler (2011), who also exploit observation weights obtained from the Kalman smoother to decompose unconditional forecasts in a dynamic factor model into contributions from distinct data blocks. We extend this weight-based decomposition to the context of conditional forecasting in reduced-form VARs. Unlike the real-time data-flow setting in Bańbura and Rünstler (2011), our analysis focuses on multi-period path restrictions. We provide an explicit analytical characterization of how imposed future trajectories transmit across horizons and quantify their dynamic influence using observation weights that are formally linked to generalized impulse response functions.

The remainder of the paper is organized as follows. Section 2 presents the econometric methodology. Section 3 reports our empirical application, and Section 4 concludes.

## 2 Dynamic Influence and Variable Importance in Conditional Forecasting

Consider the following vector autoregressive (VAR) model for the $n \times 1$ vector

$$y_t = \left(y_{1t}, y_{2t}, \ldots, y_{nt}\right)', \quad t = 1, \ldots, T,$$

given by

$$y_t = \sum_{i=1}^{p} \Phi_i \, y_{t-i} + u_t, \quad t = 1, \ldots, T, \tag{1}$$

where each $\Phi_i$ is an $n \times n$ coefficient matrix, and $u_t$ is an $n \times 1$ innovation vector.

We impose the following standard assumptions to ensure weak stationarity and identification of the VAR model:

**Assumption 1.** *$\{u_t\}$ is a white noise process with $E[u_t] = \mathbf{0}$, $E[u_t u_t'] = \Sigma$, where $\Sigma$ is positive definite, and*

$$E[u_t u_s'] = \mathbf{0}, \quad \forall \, t \neq s.$$

**Assumption 2.** *All roots of the characteristic equation $\left|I_n - \sum_{i=1}^{p} \Phi_i z^i\right| = 0$ lie outside the unit circle.*

**Assumption 3.** *The set $\{y_t, y_{t-1}, \ldots, y_{t-p}\}$ is linearly independent for all $t = 1, \ldots, T$.*

Under Assumptions 1 and 2, the process $\{y_t\}$ is (weakly) stationary and admits the following infinite-order moving average (MA) representation:

$$y_t = \sum_{i=0}^{\infty} A_i \, u_{t-i}, \quad t = 1, \ldots, T, \tag{2}$$

where the $n \times n$ matrices $\{A_i\}$ are determined recursively as:

$$A_i = \Phi_1 A_{i-1} + \Phi_2 A_{i-2} + \cdots + \Phi_p A_{i-p}, \quad i \geq 1,$$

with $A_0 = I_n$ and $A_i = \mathbf{0}$ for $i < 0$.

To study shock propagation without relying on structural identification, Pesaran and Shin (1998) introduced the generalized impulse response function (GIRF). Let $\Omega_t$ denote the filtration generated by $\{y_i : i \leq t\}$. The GIRF at horizon $s$, following a one-time shock $\delta_j$ to the innovation at time $t$, is defined as:

$$GIRF(s, \delta_j, \Omega_{t-1}) \equiv E\big[y_{t+s} \mid u_{jt} = \delta_j, \Omega_{t-1}\big] - E\big[y_{t+s} \mid \Omega_{t-1}\big]. \tag{3}$$

Assuming $u_t \overset{\text{iid}}{\sim} \mathcal{N}(0, \Sigma)$, standard properties of the multivariate normal distribution imply:

$$E\big[u_t \mid u_{jt} = \delta_j, \Omega_{t-1}\big] = \Sigma^{(j)}\Sigma_{j,j}^{-1}\delta_j = \Sigma e_j \Sigma_{j,j}^{-1}\delta_j, \tag{4}$$

where $\Sigma^{(j)}$ denotes the $j$th column of $\Sigma$, and $e_j$ is the $j$th canonical basis vector of dimension $n$. Substituting equations (2)–(4), the unscaled GIRF for a shock to the $j$th component of $u_t$ is given by:

$$GIRF(s, \delta_j, \Omega_{t-1}) = A_s \Sigma e_j \Sigma_{j,j}^{-1}\delta_j. \tag{5}$$

In conditional forecasting, rather than imposing constraints on the path of reduced-form innovations, assumptions are placed directly on the future trajectory of a subset of observable variables in $y$. To formalize this, consider the mechanism underlying the construction of conditional forecasts. Let the forecast periods be $\mathcal{T} = \{T+1, \ldots, T+h\}$, and suppose that a subset of $n_J \geq 1$ variables is constrained over this horizon. The stacked path restriction can then be written as

$$\mathcal{S}\mathbf{y}_{\mathcal{T}} = \boldsymbol{\delta}_y^c, \tag{6}$$

where $\mathbf{y}_{\mathcal{T}} = (y'_{T+1}, \ldots, y'_{T+h})' \in \mathbb{R}^{nh}$, $\mathcal{S} \in \mathbb{R}^{n_J h \times nh}$ is a known selection matrix, and $\boldsymbol{\delta}_y^c \in \mathbb{R}^{n_J h}$ contains the imposed values for the constrained variables.

Bańbura et al. (2015) show that conditional forecasts can be computed using Kalman filtering techniques by recasting the VAR model into a state-space form and treating the constraints as observed realizations for a subset of the variables. In this framework, the conditional forecast corresponds to the smoothed estimate of the components of the state vector. As demonstrated by Koopman and Harvey (2003), the smoothed state vector can be expressed as a weighted average of observed data, with the weights obtained via a recursive algorithm. Consequently, the conditional expectation of $y_{T+s}$, given the imposed future path constraints, can be expressed as:

$$E\big[y_{T+s} \mid \mathcal{S}\mathbf{y}_{\mathcal{T}} = \boldsymbol{\delta}_y^c, \Omega_T\big] = \sum_{i=1}^{T+h} W_{T+s,i}(T+h)\, y_i, \tag{7}$$

where $W_{T+s,i}(T+h)$ denotes the weight assigned to observation $y_i$ in the estimate of $y_{T+s}$, conditional on the full information set comprising both the historical data $\Omega_T$ and the imposed constraints $\mathcal{S}\mathbf{y}_{\mathcal{T}} = \boldsymbol{\delta}_y^c$.[1]

We define the conditional response function (CRF) for the $j$th variable in the system at forecast horizon $s$ to the constraint at time $T+l$, as the weight assigned to that constrained observation in equation (7):

$$CRF_j(s, l, \delta_y^c, \Omega_T) \equiv W_{T+s,l}^{(j)}(T+h), \tag{8}$$

where $W_{T+s,l}^{(j)}(T+h)$ denotes the $j$th row of the weight matrix $W_{T+s,l}(T+h)$ associated with the forecast at time $T+s$. Notably, $l$ may exceed $s$, giving rise to a bidirectional influence of constraints—whereby anticipated future observations can affect not only forward-looking projections but also inferred past dynamics through the smoothing mechanism.

This distinction underscores a fundamental conceptual difference from standard impulse response analysis. Whereas generalized impulse response functions (GIRFs) trace the dynamic effects of shocks to reduced-form innovations $u_s$, conditional forecasts impose weaker—yet operationally feasible and often more interpretable—restrictions directly on future observable outcomes. Viewed through the moving average representation in (2), these path-based constraints implicitly restrict the joint distribution of future shocks and thus reshape the entire forecast trajectory. Importantly, the influence of each constrained observation on future outcomes can be quantified using the conditional response function in (8).

We begin by establishing the relationship between the conditional response function, $CRF_j(s, l, \delta_y^c, \Omega_T)$, and the generalized impulse response function, $GIRF(s, \delta_j, \Omega_T)$, under single-period conditioning. We then extend the analysis to general multi-period path conditioning and use the resulting framework to evaluate the relative importance of variables included in the conditioning set.

In the following Propositions we establish the relation between CRF and GIRF.

**Proposition 1.** ***(i)*** *For a single-period conditioning on $y_{j,T+1}$, the conditional response function $CRF_i(s, 1, \delta_y^c, \Omega_T)$, defined in equation* (8)*, coincides with the unscaled generalized*

[1] Weights corresponding to missing observations are set to zero.

*impulse response function (GIRF) at horizon $s-1$, that is, with $A_{s-1}\Sigma^{(j)}\Sigma_{j,j}^{-1}$, as given in equation* (3).

***(ii)*** *For a general multi-period conditioning over $\mathcal{T} = \{T+1, \ldots, T+h\}$, the CRF weight assigned to $y_{j,T+1}$ in the smoothed forecast of $y_{i,T+s}$ differs from the corresponding unscaled GIRF at all horizons $s$.*

**Proof.** See Appendix A.

Beyond clarifying the economic interpretation of the conditional response function in (8), Part (i) of Proposition 1 provides a formal justification for the common practice of comparing conditional and unconditional forecasts (see, e.g., Bańbura et al., 2015): under single-period conditioning, this difference coincides exactly with the generalized impulse response function at horizon $s-1$.

While the equivalence between CRFs and GIRFs holds under single-period conditioning, Part (ii) of Proposition 1 shows that this correspondence fails under multi-period path restrictions. To build intuition for this result, consider a setting in which the forecast is conditioned on a multi-period path for a single variable $y_j$. At horizon $s = 1$, the difference between the conditional and unconditional forecasts of $y_{T+1}$ reflects the cumulative information embedded in the full assumed path, as captured by equation (7). Unlike the single-period case—where the influence of the constraint is localized—the multi-period restriction induces a bidirectional dependence across time, altering the weighting structure and breaking the equivalence with the unscaled GIRF. The same logic applies for any $s$: the forecast difference reflects the joint effect of the entire conditioning path.

We now turn to the overall impact of conditioning on forecast outcomes. This effect depends jointly on the dynamic properties of the system and the contemporaneous correlations among the variables. These influences are summarized by the observation weights in equation (7).

As noted above, the effectiveness of a given conditioning set is typically assessed through out-of-sample forecast evaluation. However, standard accuracy metrics—such as reductions in forecast errors—are influenced by two distinct factors: the underlying importance of the conditioning information, and the realized dynamics of the variables in the evaluation sample. Disentangling these effects is critical. Doing so allows one to (i) anticipate, ex ante, which variables are likely to have the strongest influence on forecast performance, and (ii) assess whether transitory shocks have materially altered the structure of influence embedded in the model.

To this end, we propose a set of ex ante measures of forecast sensitivity for each horizon $s$. Both the marginal and overall importance of each conditioning variable can be assessed by examining the relative magnitude of the (standardized) observation weights associated with past information and future scenarios. Specifically, we define two measures: the *overall variable importance* (VI) and the *marginal importance* (MI) of variable $j$ at forecast horizon $s$, conditional on the full information set comprising both the historical data $\Omega_T$ and the imposed constraints $\mathcal{S}\mathbf{y}_\mathcal{T} = \boldsymbol{\delta}_y^c$ , given by:

$$VI_j(s) = \frac{\sum_{k=1}^{T+h} |w_{jk}(s)\sigma_j|}{\sum_{i=1}^{n} \sum_{k=1}^{T+h} |w_{ik}(s)\sigma_i|}, \tag{9}$$

$$MI_j(s) = \frac{\sum_{k=T+1}^{T+h} |w_{jk}(s)\sigma_j|}{\sum_{i=1}^{n} \sum_{k=1}^{T+h} |w_{ik}(s)\sigma_i|}, \tag{10}$$

where $w_{jk}(s)$ represent the smoothing weights for variable $j$, at time $k$, for horizon $s$, and $\sigma_j$ denotes the standard deviation of variable $j$, used to normalize for scale. The measure $VI_j(s)$ captures the total (standardized) contribution of variable $j$ to the conditional forecast, combining both historical data and imposed future paths, while $MI_j(s)$ isolates the influence attributable solely to the conditioning assumptions.

# 3 Conditional Inflation Forecasting under Financial and Energy Price Scenarios

This section provides an empirical illustration of the framework developed in Section 2. The aim is not to offer a structural interpretation of the scenario, but to show how conditional response functions (CRFs) and the associated importance measures can be used to (i) assess ex ante which conditioning assumptions are likely to matter for a target forecast, and (ii) attribute an ex post scenario-induced forecast revision to specific elements of a multi-period conditioning path.

We focus on a two-input scenario that is common in policy and stress-testing applications: a tightening in financial conditions (via corporate credit spreads) combined with a temporary oil-price disruption. The key empirical takeaway is that, in this reduced-form VAR, the core inflation revision is driven primarily by the early quarters of the credit-spread path, while the oil-price path contributes much less overall and mainly affects longer horizons. The CRF attribution makes this statement precise by identifying the specific quarters of the imposed paths that account for the revision.

## 3.1 Data and model

We use U.S. quarterly data from FRED over 1986Q2–2015Q4.[2] The VAR includes five variables: (i) real GDP growth (quarter-on-quarter percent change of real GDP), (ii) core PCE inflation (quarter-on-quarter percent change of the core PCE price index excluding food and energy), (iii) the federal funds rate (quarterly average, percent), (iv) a corporate credit spread (Moody's Baa yield relative to the 10-year Treasury, percentage points), and (v) the WTI crude oil price (quarterly average, USD per barrel).

We estimate a reduced-form VAR($p$) with a constant. The lag length $p$ is selected using the Bayesian Information Criterion (BIC), which yields a parsimonious VAR(2). We emphasize that this is a *reduced-form* exercise: the scenario is stated as a path for observables rather than as a sequence of identified shocks. Accordingly, the conditional forecast is interpreted as the most likely joint evolution of the system consistent with the imposed observable paths.

## 3.2 Scenario design and conditional forecasting

We consider a forecast horizon of $h = 20$ quarters (five years) starting at $T + 1 = 2016Q1$. We impose a two-input scenario on the conditioning set

$$\mathcal{J} = \{\text{credit spread},\ \text{WTI oil price}\}.$$

The scenario is designed to be economically coherent and typical of stress-testing narratives:

[2]Series IDs: GDPC1, PCEPILFE, FEDFUNDS, BAA10YM, and DCOILWTICO. The data vintage is January 27, 2026. The dataset is also available in the `cforecast` package (Ginker, 2026), which is used to produce the results reported in the empirical application.

1. **Risk-off financial tightening.** The corporate credit spread widens by roughly 200 basis points relative to baseline, remains elevated for several quarters, and then gradually normalizes.

2. **Temporary energy-price disruption.** The WTI oil price rises sharply above baseline (peaking around 60% above the baseline within the first year) and then mean-reverts.

Figure 1 plots the imposed scenario paths against the model-implied baseline paths (unconditional forecasts).

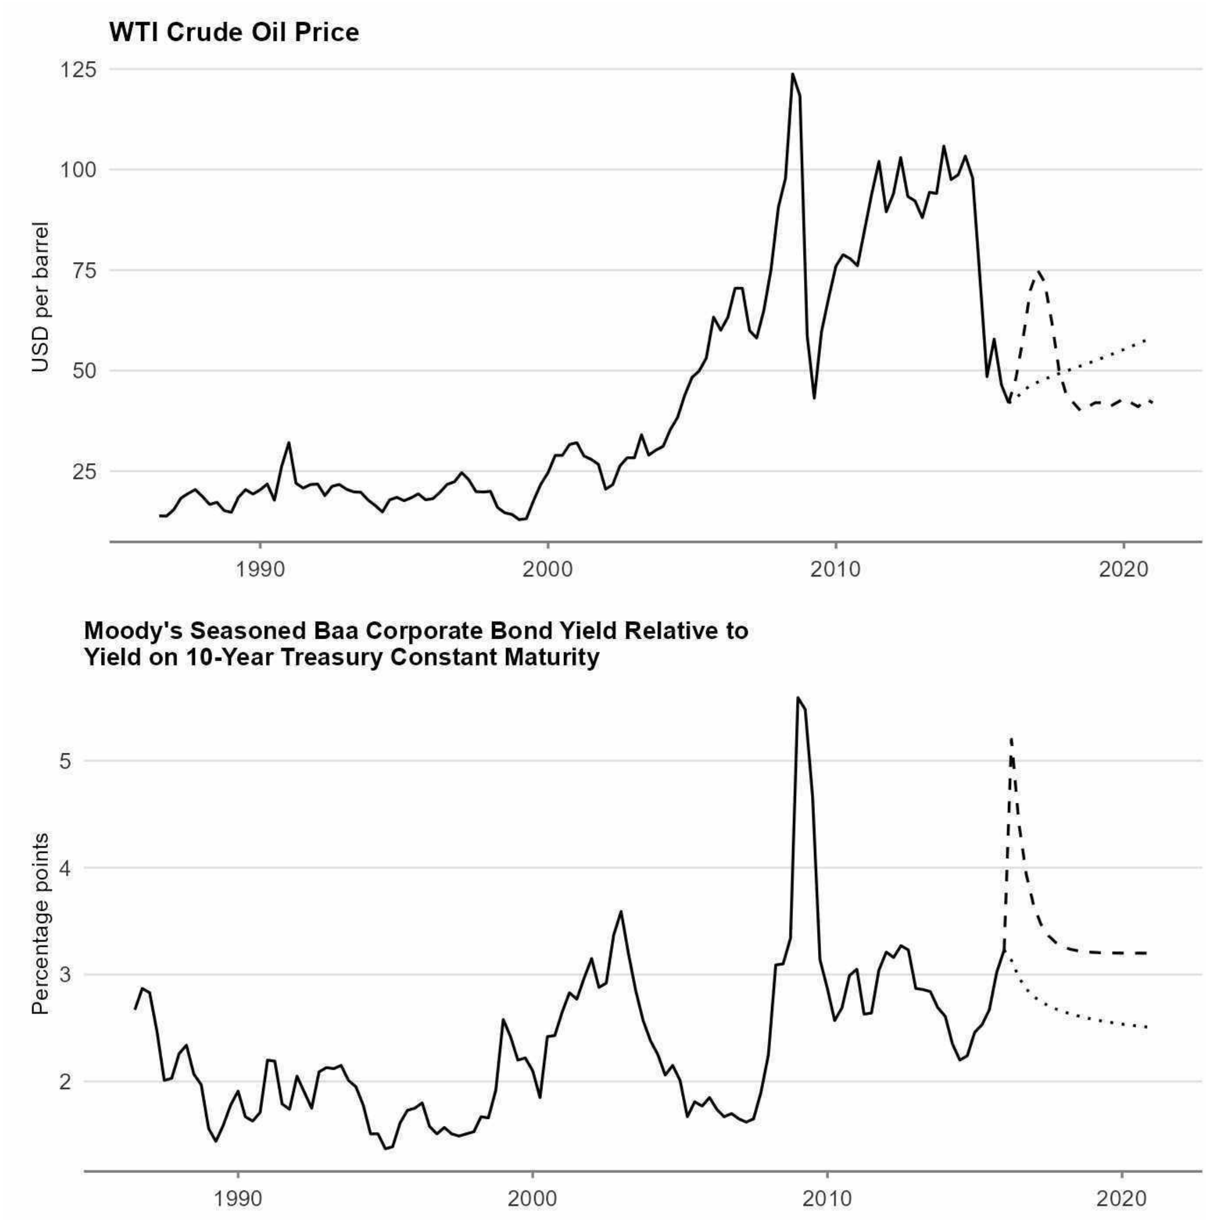


Figure 1: **Scenario inputs (conditioning paths).** The figure shows the conditioning variables: a corporate credit spread and the WTI crude oil price.

*Notes:* Solid lines denote historical data. Dotted lines show model-implied baseline paths (unconditional forecasts). Dashed lines indicate the imposed scenario paths, featuring a persistent risk-off widening in credit spreads and a temporary, mean-reverting increase in oil prices.

### 3.3 Conditional Forecast Transmission and Variable Importance

We begin by examining overall variable importance in the conditional forecast of core inflation, as reported in Figure 2. These measures are computed ex ante, given the conditioning design: we specify that the future paths of the Moody's Baa–10Y Treasury yield spread and WTI crude oil prices are constrained, while the remaining variables are left unconstrained. The decomposition therefore quantifies the relative contribution of (i) the imposed conditioning paths and (ii) past information contained in the historical data, without requiring realized future values. The results show that the importance of the credit spread increases steadily with the forecast horizon, indicating that financial conditions—particularly corporate credit risk—play an increasingly prominent role in shaping medium-term inflation expectations. This finding is consistent with evidence that credit spreads have predictive content for demand-side pressures (López-Salido et al., 2017; Caldara and Herbst, 2019).

In contrast, the oil price contributes modestly to core inflation forecasts, with a slight increase in importance at longer horizons—an intuitive outcome given that core inflation excludes energy components and oil affects inflation only indirectly, through production costs and aggregate demand.

The federal funds rate, which is unconstrained in the forecast scenario, contributes through its past values. Its importance peaks with a lag of around five quarters, consistent with standard estimates of monetary policy transmission (Christiano et al., 1996).

Finally, GDP growth receives minimal weight, reflecting its weak marginal signal once financial and commodity variables are accounted for. This finding is consistent with Stock and Watson (2007), who show that real activity indicators offer limited predictive content for inflation.

Figure 3 focuses on the marginal importance of the imposed future constraints and corroborates the earlier results. The decomposition is dominated by the corporate credit spread, with oil prices gaining importance at longer horizons because of their indirect and lagged effects on core inflation. The results further indicate that the imposed constraints account for approximately 24% to 32% of the total forecast across horizons, suggesting that the conditioning assumptions play a quantitatively important role in shaping the forecast path. Since the measure incorporates both contemporaneous and dynamic interactions within the VAR system, it provides an additional indication of the sensitivity of the forecast to the imposed constraints. Taken together, Figures 2 and 3 show that while overall variable importance reflects both historical information and imposed paths, marginal importance isolates the contribution of future constraints, highlighting the dominant role of the credit spread and the delayed influence of oil prices.

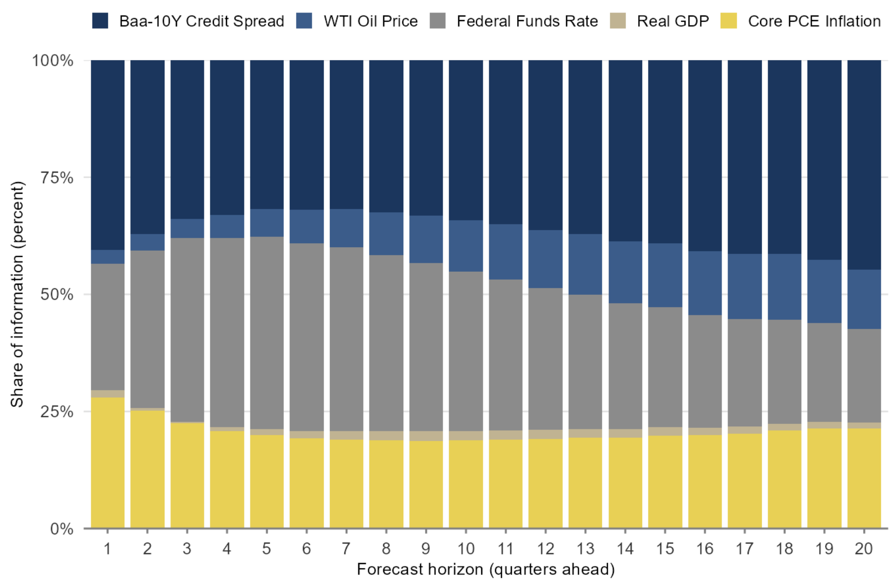


Figure 2: **Overall variable importance in the conditional CPI forecast.**

*Notes:* The figure decomposes the conditional CPI forecast weights into variable- and forecast-horizon-specific information shares.

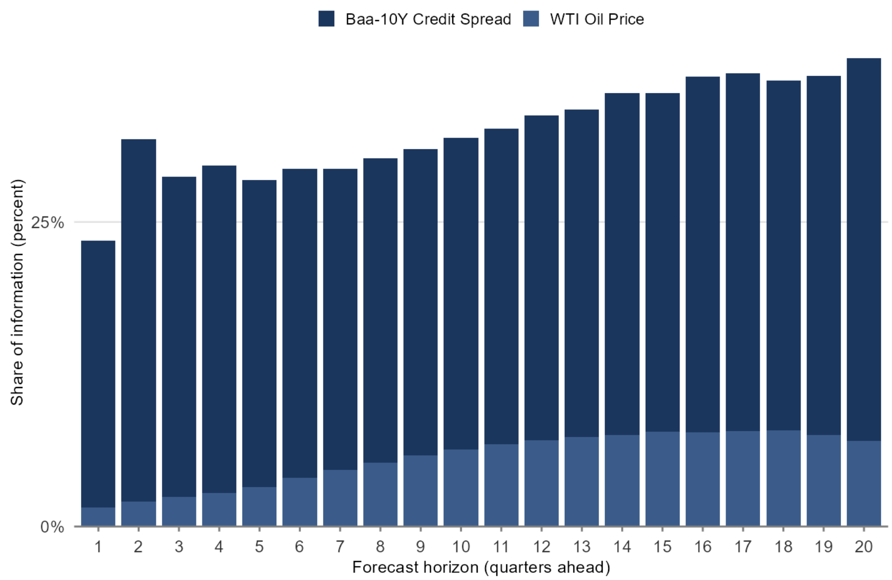


Figure 3: **Marginal variable importance in the conditional CPI forecast.**
*Notes:* The figure decomposes the conditional forecast weights assigned to the scenario constraints into variable- and horizon-specific information shares for core inflation.

Figure 4 illustrates the conditional forecast of core inflation under the assumed scenario, while Figure 5 clarifies the sources of these dynamics. The imposed widening in the corporate credit spread generates a pronounced disinflationary impulse at short horizons, and the reduced-form system implies a corresponding accommodative policy-rate response. The credit spread continues to exert downward pressure on inflation as it stabilizes at a persistently elevated level relative to the baseline scenario. Oil prices affect the forecast with a delay through the gradual transmission of energy cost shocks to production costs and aggregate demand. However, because oil prices normalize below the baseline path following the initial spike, their contribution eventually becomes negative. Despite these dynamics, the scenario remains dominated by the credit-spread assumption at all horizons, consistent with its high variable-importance measures. More generally, these measures help discipline scenario design and interpretation by distinguishing between effects arising from strong model linkages and those driven primarily by large imposed deviations. In turn, this helps identify the economically central components of the scenario and highlights which imposed assumptions are largely irrelevant for the forecast variables of interest.

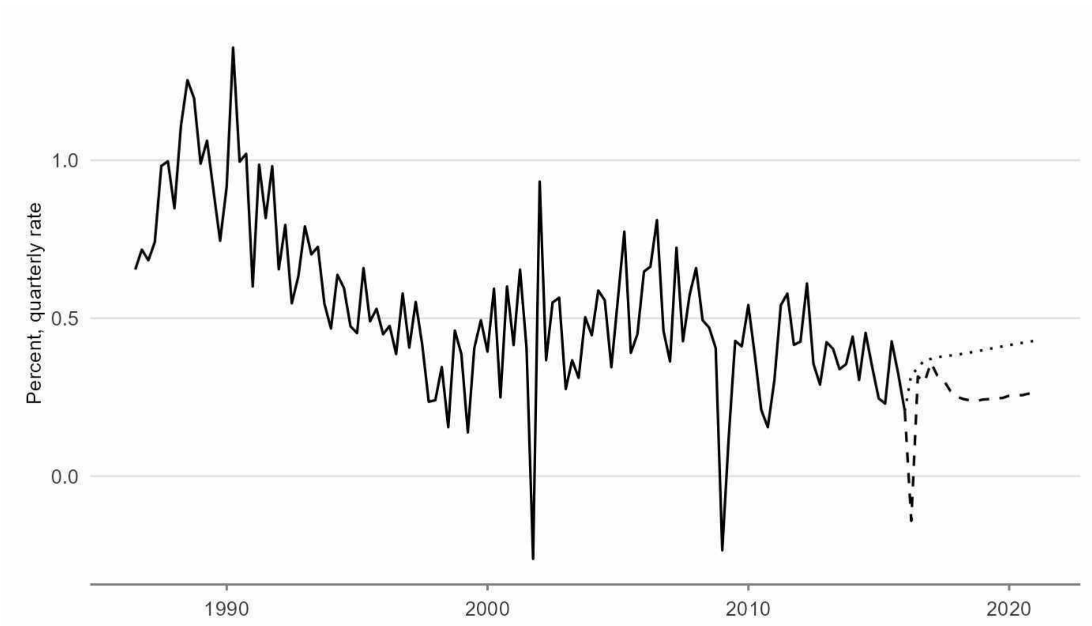


Figure 4: **Core inflation forecasts.** The figure presents historical core inflation (solid line) together with the unconditional baseline forecast (dotted line) and the conditional scenario forecast (dashed line). *Notes:* Core inflation is measured as the Personal Consumption Expenditures Price Index excluding food and energy.

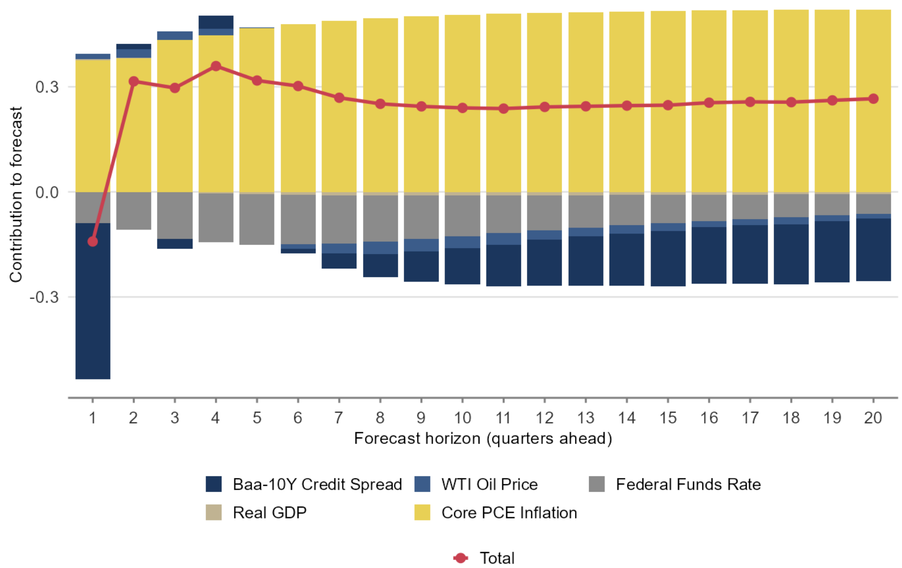


Figure 5: **Composition of the conditional forecast.** The figure decomposes the conditional forecast into the contributions of its underlying variables across forecast horizons. *Notes:* Bars show the contribution of each variable at a given forecast horizon; the stacked height of each bar equals the forecast value.

## 3.4 What drives the scenario? Constraint-by-constraint CRF attribution

The central interpretability question in scenario analysis is: which elements of the imposed conditioning paths drive the forecast revision? Because the conditional expectation in the linear-Gaussian state-space representation can be written as a weighted sum of observed and imposed data (Section 2), the revision in core inflation at horizon $s$ admits an additive decomposition across constraint points.

Let $\delta_{j,l}^{\text{scen}}$ and $\delta_{j,l}^{\text{base}}$ denote the imposed scenario and baseline (unconditional forecast) values of constrained variable $j \in \mathcal{J}$ at constraint horizon $l = 1, \ldots, h$. Let $CRF_{j\to\pi}(s,l)$ denote the CRF weight mapping the constraint on variable $j$ at $T+l$ into the forecast of core inflation $\pi$ at $T+s$. Then

$$\Delta\hat{\pi}_{T+s} = \sum_{j\in\mathcal{J}}\sum_{l=1}^{h} CRF_{j\to\pi}(s,l)\left(\delta_{j,l}^{\text{scen}} - \delta_{j,l}^{\text{base}}\right), \qquad s = 1, \ldots, h. \tag{11}$$

We refer to each term in the double sum as the contribution of constraint point $(j,l)$ to the revision at forecast horizon $s$.

Figure 6 visualizes these contributions as a horizon-by-horizon heatmap for each constrained variable. The credit-spread heatmap displays large-magnitude contributions concentrated near the diagonal, indicating that the imposed spread path is the dominant driver of the core inflation revision and that each quarter of the spread assumption primarily affects inflation at comparable horizons. The oil-price heatmap shows contributions that are an order of magnitude smaller; oil contributes modestly to the revision, with a pattern consistent with a temporary spike and subsequent mean reversion that matters mostly at longer horizons.

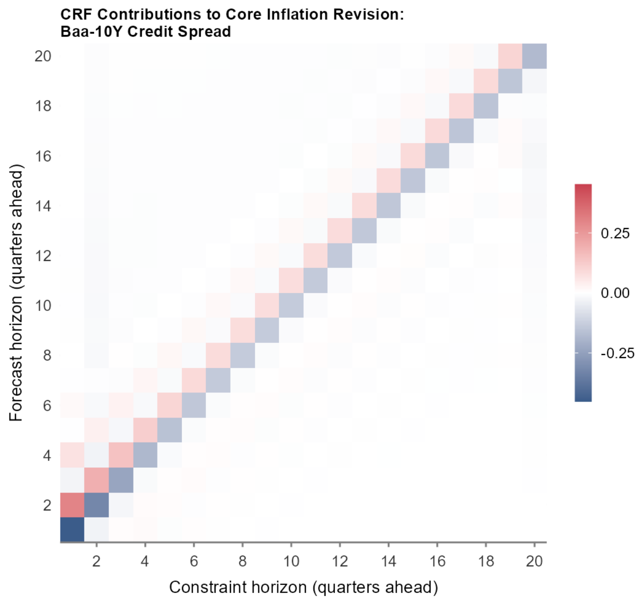


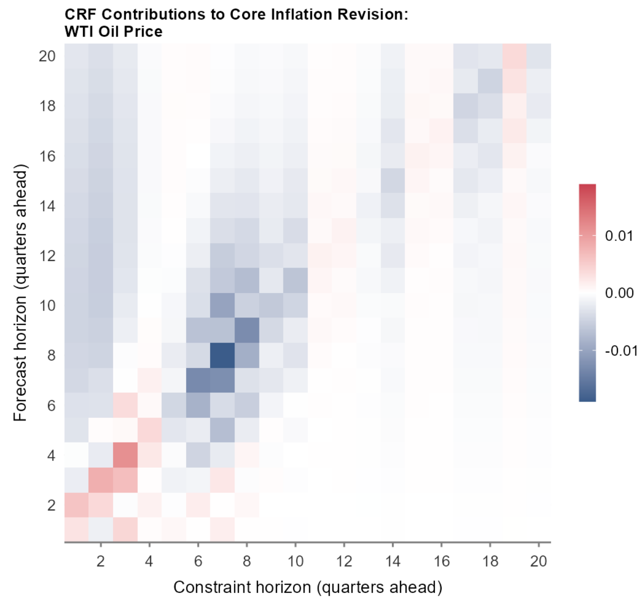


Figure 6: **CRF contributions to the core inflation revision (constraint-by-constraint).** The figure shows the contribution of each constraint point (horizontal axis: constraint horizon $l$) to the core inflation revision at each forecast horizon $s$ (vertical axis), as in equation (11). The left panel presents the contributions associated with credit-spread constraints, while the right panel reports the contributions of oil-price constraints.

### 3.5 Which constraint quarters matter most? Rankings and horizon profiles

To summarize which parts of the imposed paths matter most, we aggregate the absolute contributions across forecast horizons. For each constrained variable $j$ and constraint horizon $l$, define the absolute impact measure

$$Impact_j(l) \equiv \sum_{s=1}^{h} \left| CRF_{j\to\pi}(s,l) \left(\delta_{j,l}^{\text{scen}} - \delta_{j,l}^{\text{base}}\right) \right|.$$

Figure 7 plots $Impact_j(l)$ for both constrained variables. The profile is sharply front-loaded for the credit spread: early quarters of the spread path account for the largest share of the revision. By contrast, the oil-price impact is smaller throughout, and becomes relatively more important only at longer constraint horizons.

Figure 8 ranks the top constraint points by total absolute impact. The top ten are all credit-spread constraints concentrated in the early part of the forecast window (2016–

2018), reinforcing the conclusion that most of the core inflation revision is driven by the early spread-widening assumptions rather than by details of the longer-run tail of the imposed paths.

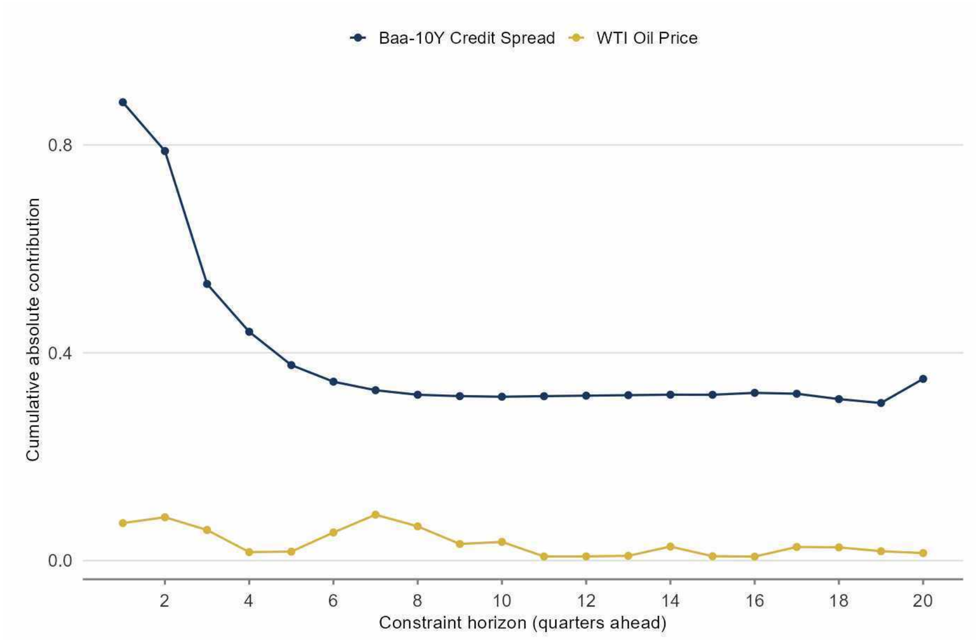


Figure 7: **Absolute impact by constraint horizon.** For each constrained variable and constraint horizon $l$, the figure reports the cumulative absolute contribution to the core inflation revision across forecast horizons.

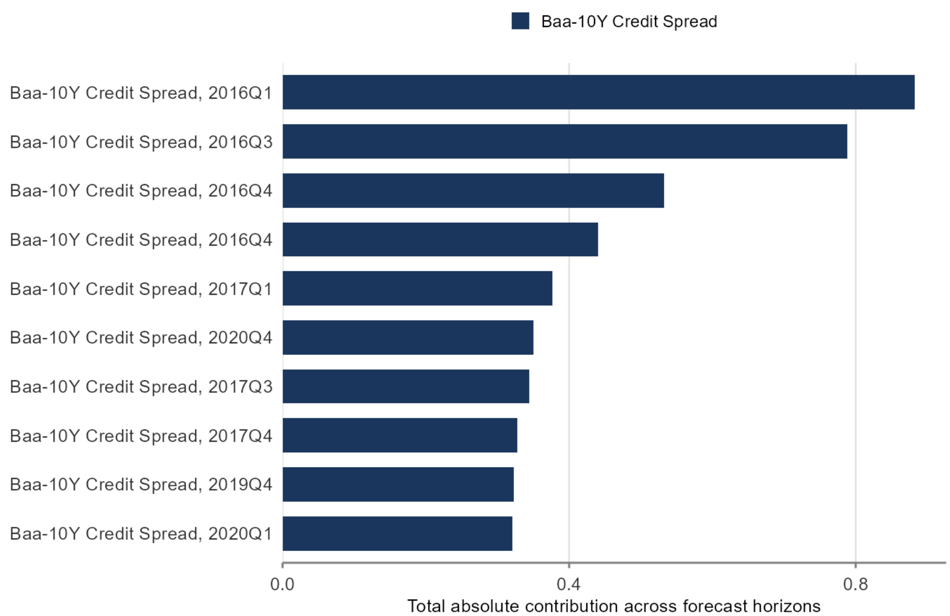


Figure 8: **Top constraint points by absolute impact on core inflation.** The figure ranks constraint points $(j, l)$ by their cumulative absolute contribution to the core inflation revision across forecast horizons.

### 3.6 Pruned-scenario validation and redundancy of imposed details

A practical implication of the attribution is that many imposed path elements may be redundant for a given target: if the revision is driven by a small subset of constraint points, then conditioning on only those points should reproduce much of the scenario forecast.

We implement a pruned-scenario exercise as follows. Rank all $n_J h = 40$ constraint points by absolute impact on core inflation. For a given $K$, construct a *pruned* scenario that imposes only the top-$K$ constraint points and sets all remaining points back to their baseline values. Figure 9 compares the full scenario revision to the revision obtained under the top-10 constraint points. The top-10 points explain about 94% of the variation in the full revision, indicating substantial redundancy in the full 40-point conditioning path.

Figure 10 reports the capture curve: the fraction of revision variance explained as $K$ increases. The curve rises rapidly for small $K$ (driven almost entirely by credit-spread constraints), exceeds 90% after including roughly a quarter of the constraint points, and approaches 1 only once oil-price constraints begin to enter the selected set. This pattern provides an operational answer to “what drives the scenario”: the early credit-spread assumptions are pivotal for the core inflation revision, while oil-price assumptions matter mainly for capturing the remaining longer-horizon features of the revision.

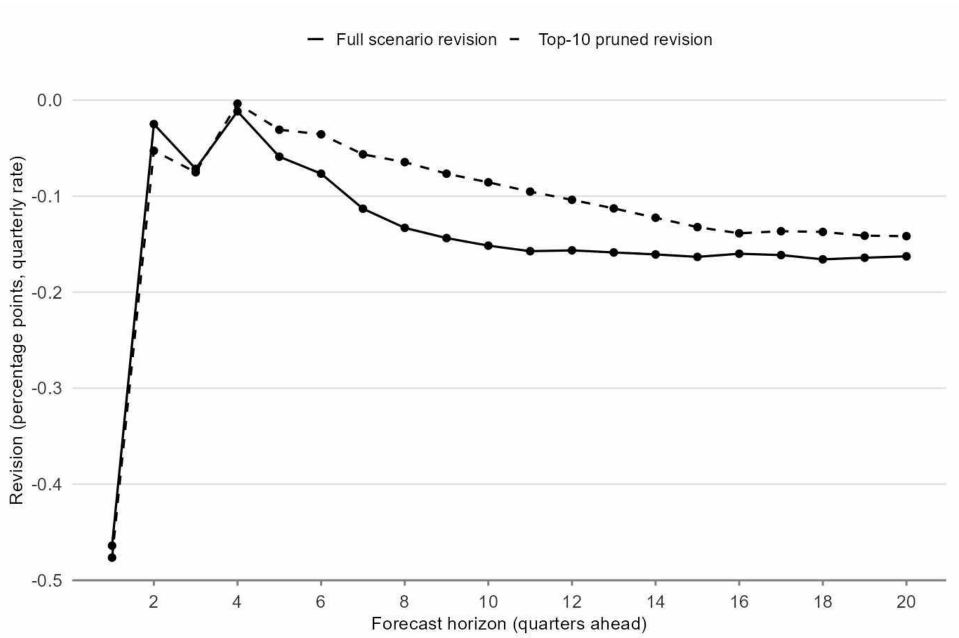


Figure 9: **Pruned-scenario validation: core inflation revision.** The figure compares the full scenario-induced revision in core inflation to the revision obtained when conditioning only on the top-10 constraint points ranked by absolute impact.

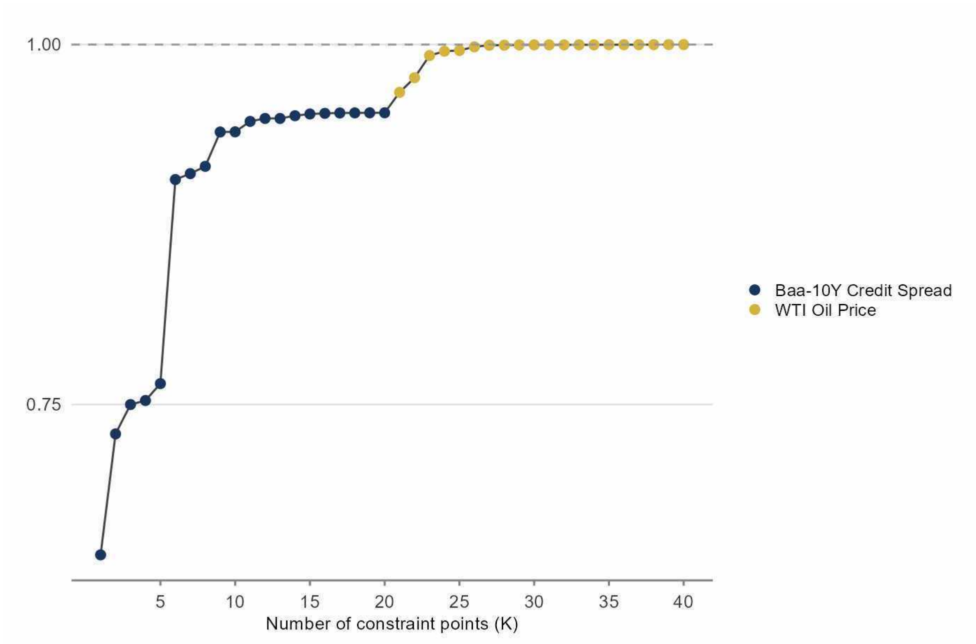


Figure 10: **Capture curve for the core inflation revision.** The figure reports the fraction of revision variance explained by the top-$K$ constraint points (ranked by absolute impact), and indicates whether the $K$-th added constraint is a credit-spread or oil-price point.

# 4 Conclusion

This paper develops a framework for interpreting conditional forecasts in reduced-form VAR models. Building on the observation-weight extraction approach of Koopman and Harvey (2003), we characterize how imposed future paths propagate through the system under path conditioning and show that the resulting conditional response functions are closely related to generalized impulse response functions. While the equivalence with unscaled GIRFs holds under single-period conditioning, multi-period path restrictions generate a richer dynamic structure that standard impulse response analysis cannot capture.

The proposed framework provides horizon-specific decompositions and variable-importance measures that quantify how alternative conditioning assumptions shape forecast outcomes. The empirical illustration shows how these tools can identify the economically central components of a scenario and attribute forecast revisions to specific variables and horizons.

More broadly, the approach improves the transparency, interpretation, and communication of reduced-form scenario forecasts while remaining structurally agnostic and readily applicable in policy-oriented forecasting environments.

# Appendix A

### Technical Preliminaries and Auxiliary Results

Following Koopman and Harvey (2003), we define the linear Gaussian state space model as:

$$y_t = Z_t \alpha_t + G_t \varepsilon_t, \quad \varepsilon_t \sim \mathrm{WN}(0, I), \quad t = 1, \ldots, T, \tag{12}$$

$$\alpha_{t+1} = T_t \alpha_t + H_t \varepsilon_t, \quad \alpha_1 \sim \mathrm{WN}(a, P), \tag{13}$$

where $y_t$ is an $n \times 1$ vector of observed variables, $\alpha_t$ is a $p \times 1$ state vector, and $\varepsilon_t$ is a $q \times 1$ vector of white noise disturbances.

The system in (1) can be rewritten as a first-order VAR using the companion form:

$$\tilde{y}_t = \tilde{\Phi} \tilde{y}_{t-1} + \tilde{\varepsilon}_t, \tag{14}$$

where

$$\tilde{y}_t = \begin{pmatrix} y_t \\ y_{t-1} \\ \vdots \\ y_{t-p+1} \end{pmatrix}, \quad \tilde{\Phi} = \begin{pmatrix} \Phi_1 & \Phi_2 & \cdots & \Phi_{p-1} & \Phi_p \\ I_n & 0 & \cdots & 0 & 0 \\ 0 & I_n & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & I_n & 0 \end{pmatrix}, \quad \tilde{\varepsilon}_t = \begin{pmatrix} u_t \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \tag{15}$$

The system in (12)–(13) can be used to represent the companion-form VAR model defined in (14) by setting:

$$\alpha_t = \tilde{y}_t, \quad Z_t = \begin{pmatrix} I_n & 0_n & \cdots & 0_n \end{pmatrix}_{n \times np}, \quad T_t = \tilde{\Phi} \equiv T, \quad G_t = 0_{n \times 1},$$

and

$$H_t = \begin{pmatrix} \Sigma^{0.5} & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} \equiv H,$$

where $\Sigma$ is the innovation covariance matrix from the VAR.

To accommodate missing observations, we follow Harvey (1990, p. 144) and introduce a sequence of selection matrices $\{S_t\}$, each composed of rows of $I_n$ corresponding to the observed components of $y_t$. The transformed observation equation is then given by:

$$y_t^* = \underbrace{S_t Z_t}_{Z_t^*} \alpha_t + S_t G_t \varepsilon_t, \tag{16}$$

where $y_t^* \equiv S_t y_t$ denotes the observed subset of $y_t$. The dimensions of $y_t^*$, $F_t$, and $V_t$ vary over time depending on the number of observed entries in $y_t$.

The Kalman filter provides recursive estimates of the state vector $\alpha_t$ based on the available information up to time $t$. Let $Z_t^*$ denote the modified observation matrix that accounts for missing observations, as described above. Then, the one-step-ahead forecast error and its variance are given by:

$$v_t = y_t - Z_t^* \alpha_{t|t-1}, \tag{17}$$

$$F_t = Z_t^* P_{t|t-1} Z_t^{*\prime} + H_t G_t', \tag{18}$$

where $v_t$ is the forecast residual, and $F_t$ is its conditional variance.

The filtered state and its covariance matrix are then updated as:

$$\alpha_{t+1|t} = T_t \alpha_{t|t-1} + K_t v_t, \tag{19}$$

$$P_{t+1|t} = T_t P_{t|t-1} T_t' + H_t H_t' - K_t M_t, \tag{20}$$

where the Kalman gain matrix is defined by

$$K_t = M_t F_t^{-1}, \quad \text{with} \quad M_t = T_t P_{t|t-1} Z_t^{*\prime}.$$

If $Z_t^*$ is empty (i.e., no observations are available at time $t$), then $K_t$ is set to zero. The recursion proceeds for $t = 1, \dots, T+h$, initializing with $\alpha_{1|0} = a$ and $P_{1|0} = P$, where $a$ and $P$ are the prior mean and covariance of the initial state.

Consider the effect of conditioning on of the form of (6). Within the state space framework, the effect of each constraint in (6) can be quantified by evaluating the weight of $y_{j,T+i}$ in the smoothed estimate of the state $\alpha_{T+s}$, and, consequently, in the conditional forecast of $y_{i,T+s}$.

More formally, the filtered vector at horizon $T+s$, conditional on conditional on the full information set comprising both the historical data $\Omega_T$ and the imposed constraints $\mathcal{S}\mathbf{y}_{\mathcal{T}} = \boldsymbol{\delta}_y^c$ (i.e. the entire set of observed values $\{y_1^*, \ldots, y_{T+h}^*\}$), admits the following linear representation:

$$\alpha_{T+s|T+h} = \sum_{j=1}^{T+h} w_j\big(\alpha_{T+s|T+h}\big) y_j^*,$$

where $w_j\big(\alpha_{T+s|T+h}\big)$ denotes the vector of observation weights associated with the $j$th observed value. These weights capture the marginal influence of each conditioning observation on the filtered or smoothed state vector. In particular, they provide a natural metric for assessing the impact of specific elements of the conditioning set on forecasted quantities of interest.

Koopman and Harvey (2003) show that the observation weights associated with the smoothed state vector $\alpha_{T+s|T+s}$ can be computed recursively using the following backward recursion:

$$w_j\big(\alpha_{T+s|T+s}\big) = B_{T+s,j} K_j, \tag{21}$$

$$B_{T+s,j-1} = B_{T+s,j} T_j - w_j\big(\alpha_{T+s|T+s}\big) Z_j, \quad j = T+s-1, T+s-2, \ldots, 1, \tag{22}$$

where $K_j$ is the Kalman gain at time $j$, $T_j$ and $Z_j$ are the system matrices as defined in the state space model, and $B_{T+s,j}$ is a matrix tracking the recursive influence of earlier observations.

The recursion is initialized at $j = T+s$ using:

$$w_{T+s}\big(\alpha_{T+s|T+s}\big) = P_{T+s|T+s-1} Z'_{T+s} F_{T+s}^{-1}, \tag{23}$$

$$B_{T+s,T+s-1} = I - P_{T+s|T+s-1} Z'_{T+s} F_{T+s}^{-1} Z_{T+s}. \tag{24}$$

This formulation expresses the filtered state as a linear combination of past and future observations, with the weights $w_j(\cdot)$ capturing the marginal contribution of each $y_j^*$ to the state at horizon $T+s$.

In the context of VAR models, several of the Kalman filter recursions admit simplifications. First, note that given $\Omega_t$, the state vector $\alpha_t$ is known without error, since it is composed of lagged endogenous variables. As a result, the conditional covariance matrix

$P_{t|t-1}$, which can equivalently be expressed as

$$P_{t|t-1} = T_t P_{t-1|t-1} T_t' + H_t H_t',$$

simplifies to

$$P_{t|t-1} = HH',$$

because $P_{t-1|t-1} = 0$ for all $t$.

As a result, the forecast error variance and the auxiliary matrix used in the Kalman gain expression also simplify. Specifically, we have

$$F_t = Z_t^* HH' Z_t^{*\prime}, \quad \text{and} \quad M_t = THH' Z_t^{*\prime}.$$

**Proof of Proposition 1**

(i)

We now derive the observation weights under single-period conditioning on $y_{j,T+1}$. This corresponds to the case of $h = 1$. First, consider the case $l = 1, s = 1$, the observation matrix becomes

$$Z_{T+1}^* = \begin{pmatrix} e_j & 0 & \cdots & 0 \end{pmatrix}_{1 \times np},$$

where $e_j$ is the $j$th canonical basis vector of dimension $n$. Consequently, $F_{T+1}$ simplifies to:

$$F_{T+1} = Z_{T+1}^* HH' Z_{T+1}^{*\prime} = \Sigma_{j,j},$$

and the observation weight becomes:

$$w_{T+1}(\alpha_{T+1|T+1}) = HH' Z' e_j' \Sigma_{j,j}^{-1}.$$

It follows that the first $n$ elements of $w_{T+1}(\alpha_{T+1|T+1})$ are given by:

$$\Sigma^{(j)} \Sigma_{j,j}^{-1},$$

where $\Sigma^{(j)}$ denotes the $j$th column of $\Sigma$.

This expression is precisely the unscaled GIRF to a shock in the $j$th variable, evaluated at horizon 0 (the impact response), since $\Sigma^{(j)} \Sigma_{j,j}^{-1} = A_0 \, \Sigma^{(j)} \Sigma_{j,j}^{-1}$. Thus, the observation weight under single-period path conditioning recovers the GIRF as a special case.

Similarly, for the case $l = 1, s = 2$, the observation weight at horizon $T + 2$ is zero:

$$w_{T+2} = 0,$$

since no information beyond $T + 1$ is available to influence $\alpha_{T+2|T+2}$ through smoothing. The recursion then starts with

$$B_{T+1} = I,$$

and the Kalman gain simplifies to:

$$K_{T+1} = T \begin{pmatrix} \Sigma^{(j)}\Sigma_{j,j}^{-1} \\ \vdots \\ 0 \end{pmatrix},$$

where the first $n$ rows of the vector correspond to the influence of the shock in the $j$th variable on the current state, and the remaining rows are zero due to the companion form structure.

As a result, the observation weight at time $T + 1$ becomes:

$$w_{T+1} = \begin{pmatrix} \Phi_1\Sigma^{(j)}\Sigma_{j,j}^{-1} \\ \vdots \\ 0 \end{pmatrix},$$

where again the first $n$ elements coincide with the unscaled GIRF evaluated at horizon $s = 1$.

More generally, since no information beyond time $T+1$ is available to influence $\alpha_{T+s|T+s}$ through smoothing, it follows that Since $w_{T+i} = 0$ for all $i > 1$, the recursion reduces to a power of the transition matrix and

$$w_{T+1} = T^{s-1} \begin{pmatrix} \Sigma^{(j)}\Sigma_{j,j}^{-1} \\ \vdots \\ 0 \end{pmatrix}.$$

Writing $J = (I_n \ 0 \ \cdots \ 0)'$ for the $np \times n$ selection matrix, the companion structure gives the identity $J'T^kJ = A_k$, the horizon-$k$ MA coefficient of (2); hence the top $n$-block of $w_{T+1}$ equals $A_{s-1}\Sigma^{(j)}\Sigma_{j,j}^{-1}$. The first $n$ rows of this expression coincide with the unscaled generalized impulse response function (GIRF) at horizon $s-1$, confirming that single-period path conditioning at $T+1$ reproduces the GIRF structure across horizons through the propagation dynamics encoded in the VAR companion form.

(ii) The result follows directly from Part (i). Since the observation matrices remain nonzero under the imposed multi-period conditioning, the recursive structure of the weights does not simplify to the generalized impulse response functions (GIRFs).

# References


Angelini, E., M. Lalik, M. Lenza, and J. Paredes (2019). Mind the gap: A multi-country BVAR benchmark for the Eurosystem projections. *International Journal of Forecasting 35*(4), 1658–1668.

Antolin-Diaz, J., I. Petrella, and J. F. Rubio-Ramírez (2021). Structural scenario analysis with SVARs. *Journal of Monetary Economics 117*, 798–815.

Bańbura, M. and G. Rünstler (2011). A look into the factor model black box: publication lags and the role of hard and soft data in forecasting GDP. *International Journal of Forecasting 27*(2), 333–346.

Bańbura, M., D. Giannone, and M. Lenza (2015). Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections. *International Journal of Forecasting 31*(3), 739–756.

Binder, C. C. and R. Sekkel (2024). Central bank forecasting: A survey. *Journal of Economic Surveys 38*(2), 342–364.

Bloor, C. and T. Matheson (2011). Real-time conditional forecasts with Bayesian VARs: An application to New Zealand. *The North American Journal of Economics and Finance 22*(1), 26–42. Symposium on Nowcasting and Model Combination.

Caldara, D. and E. Herbst (2019). Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs. *American Economic Journal: Macroeconomics 11*(1), 157–192.

Christiano, L., M. Eichenbaum, and C. Evans (1996). The effects of monetary policy shocks: Evidence from the flow of funds. *The Review of Economics and Statistics 78*(1), 16–34.

Clarida, R. H. and D. Coyle (1984). Conditional projection by means of Kalman filtering.

Giannone, D., M. Lenza, D. Momferatou, and L. Onorante (2014). Short-term inflation projections: A Bayesian vector autoregressive approach. *International journal of forecasting 30*(3), 635–644.

Ginker, T. (2026). *cforecast: Conditional Forecasting and Scenario Analysis Using VAR Models*. R package version 0.1.1.

Harvey, A. C. (1990). *Forecasting, structural time series models and the Kalman filter*. Cambridge university press.

Koopman, S. J. and A. Harvey (2003). Computing observation weights for signal extraction and filtering. *Journal of Economic Dynamics and Control 27*(7), 1317–1333.

López-Salido, D., J. C. Stein, and E. Zakrajšek (2017). Credit-market sentiment and the business cycle. *Quarterly Journal of Economics 132*(3), 1373–1426.

Pesaran, H. H. and Y. Shin (1998). Generalized impulse response analysis in linear multivariate models. *Economics letters 58*(1), 17–29.

Stock, J. H. and M. W. Watson (2007). Why has US inflation become harder to forecast? *Journal of Money, Credit and banking 39*, 3–33.

Waggoner, D. F. and T. Zha (1999). Conditional forecasts in dynamic multivariate models. *Review of Economics and Statistics 81*(4), 639–651.