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

Itamar Caspi   Tim Ginker

Bank of Israel

June 18, 2026

Conditional Forecasting in Modern Macroeconomic Policy

Conditional forecasts are standard in policy

Targets (e.g., inflation) conditional on assumed paths (e.g., oil, interest rate)

Binder & Sekkel (2024)

Structural models

- ✓ Coherent economic narratives (identified shocks)
- × Sensitive to identification assumptions & sample variation (Antolín-Díaz et al., 2021)

Reduced-form (B)VARs

- ✓ Are often used to cross-check macroeconomic outlooks (Angelini et al., 2025)
- ✓ Robust, easy to implement
- × No causal interpretation

Out-of-Sample Evaluation and Interpretability Limits

In reduced-form VARs, conditional forecasts are typically evaluated out of sample in two main ways (Angelini et al., 2019):

- **Realized-path evaluation.** Conditional forecasts are generated using the *realized* future values of the conditioning variables and evaluated against observed outcomes. This assesses the information content of the conditioning set.
- **Scenario-path evaluation.** Forecasts are conditioned on assumed (real-time) scenario paths and evaluated relative to the same assumed paths, providing a real-time measure of accuracy.

Limitation. Forecast accuracy alone is vulnerable to the evaluation period: inferring that a variable is "important" conflates

1. the *weight* the model assigns to that variable, and
2. the *realized dynamics* of the variable during the evaluation window.

Research Contributions

1. **A new interpretability framework for conditional forecasts.** Using the Kalman filter weights, we develop model-consistent tools that decompose conditional forecasts into contributions from specific elements of the imposed scenario paths, enabling a transparent economic interpretation of reduced-form VAR scenarios.

2. **Linking scenario analysis to impulse response theory.** We formally establish when the weights coincide with generalized impulse response functions (single-period conditioning) and when this equivalence breaks down under multi-period path restrictions.

3. **Weight-based measures of variable importance.** We introduce horizon-specific measures of overall and marginal variable importance that separate intrinsic model sensitivity from realized sample dynamics, allowing ex ante assessment of which variables truly drive conditional forecasts.

4. **Practical implementation.** All methods are implemented in the open-source R package `cforecast`, designed for applied scenario analysis with reduced-form VARs.

Related: Banbura and Runstler (2011) decompose forecasts using observation weights in factor models. We extend this idea to conditional VAR forecasting with multi-period paths and link it to GIRFs.

Scenario: Risk-off shock with energy disruption

Data and model

- U.S. quarterly data (1986–2015)
- Five-variable reduced-form VAR(2)

Variables: GDP growth, core inflation, policy rate, credit spread (Baa–10Y), oil price

Scenario

- **Credit spread:** +200 bps, persistent, gradual normalization
- **Oil price:** sharp spike, peak in year 1, mean reversion

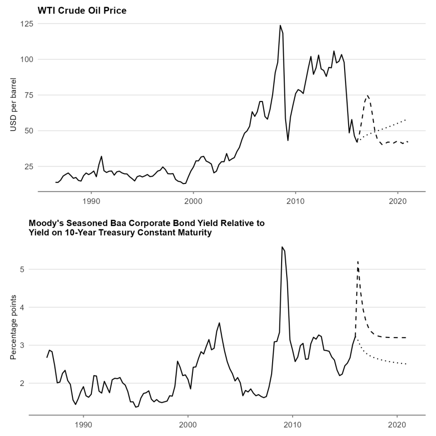

Imposed scenario paths vs. unconditional baseline forecasts.

Ex-ante Variable Importance

Key results

- **Credit spread:** dominant, rising with horizon (Lopez-Salido et al., 2017; Caldara and Herbst, 2019)
- **Oil price:** modest, delayed effect
- **Policy rate:** peaks at 4–6 quarters (Christiano et al., 1996)
- **GDP:** negligible (Stock and Watson, 2007)

Importance computed ex ante (no realized data)

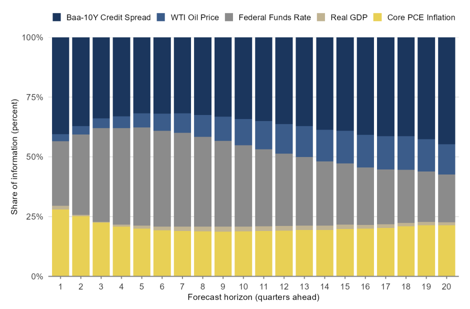

Overall variable importance across horizons.

Ex-ante Marginal Variable Importance

Key results

- **Scenario constraints:** account for about 25–30% of forecast information
- **Credit spread:** dominates the contribution of imposed future paths
- **Oil price:** smaller contribution, increasing at longer horizons
- **Scenario inputs:** effects are concentrated in the constrained variables

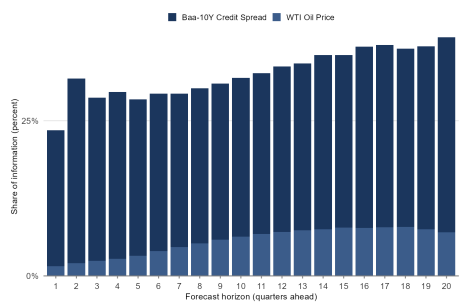

Marginal variable importance from imposed scenario paths.

What Drives the Scenario Forecast?

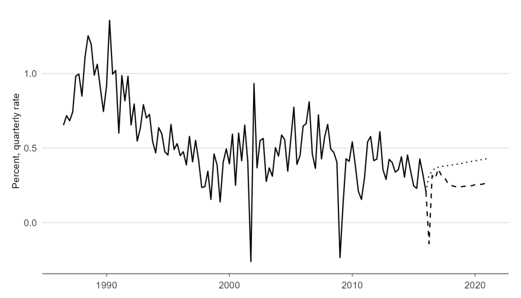

Conditional vs. baseline forecast

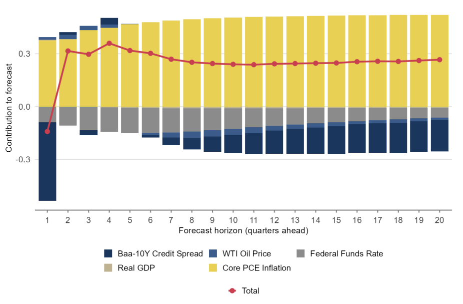

Forecast decomposition

Key takeaways

- Credit spreads are the dominant driver of the forecast revision.
- Oil matters mainly through the imposed path magnitude, not model sensitivity.
- Decomposition identifies which scenario assumptions are pivotal.

Pruned-Scenario Validation: Core Inflation Revision

Idea

Rank all imposed constraint points by their absolute contribution, then keep only the most important ones.

Validation exercise

- Full scenario: 40 imposed constraints
- Pruned scenario: top-10 constraints only. Remaining constraints set back to NA

Interpretation: If the pruned scenario closely tracks the full revision, many imposed path details are redundant.

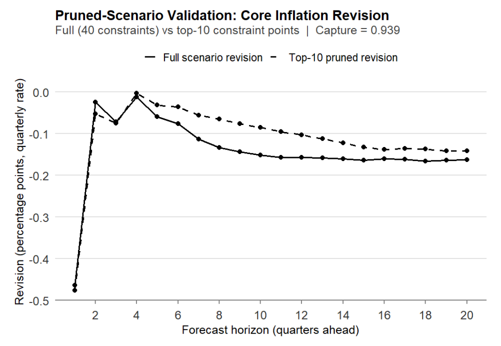

Full scenario revision vs. top-10 pruned revision.

Additional Tools for Scenario Interpretation

Further empirical exercises in the paper

Constraint-by-constraint attribution	Identifies which imposed path elements drive forecast revisions.
Constraint-horizon rankings	Ranks which quarters of a scenario are pivotal.
Absolute impact profiles	Summarizes how influence is distributed across horizons.
Pruned-scenario validation	Shows a small subset of constraints can reproduce much of the full scenario.
Capture curves	Quantifies redundancy and supports parsimonious scenario design.

Practical implications

- Interprets what drives the scenario, not merely what the scenario implies.
- Supports simpler, more transparent scenario design and communication.

Methodology

Our approach builds on Kalman filtering and the observation-weight extraction method of Koopman and Harvey (2003).

Consider a VAR(p) for the $n \times 1$ vector y_t:

$$y_t = \sum_{i=1}^{p} \Phi_i \, y_{t-i} + u_t, \qquad u_t \sim (0, \Sigma), \quad t = 1, \ldots, T.$$

A conditional forecast s periods ahead, given an imposed future path and information Ω_T, can be written as[1]:

$$E[y_{t+s} | \overbrace{\mathcal{S}\mathbf{y}_\tau = \delta_y^c}^{\text{Future constraints}}, \overbrace{\Omega_T}^{\text{Observed data}}] = \sum_{i=1}^{T+h} W_{T+s,i}(T+h) y_i,$$

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 $n_J \geq 1$ constrained variables.

[1] Missing observations have zero weights.

GIRF recap (unscaled)

Assuming that the process $\{y_t\}$ is (weakly) stationary, it admits the following infinite-order moving average (MA) representation:

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

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 $i < 0$. For a shock to innovation component u_{jt}:

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

Under Gaussianity:

$$\mathbb{E}[u_t \mid u_{jt} = \delta_j, \Omega_{t-1}] = \Sigma e_j \Sigma_{jj}^{-1} \, \delta_j,$$

so

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

Conditional Response Function (CRF)

Definition. The conditional response function (CRF) measures the influence of an imposed constraint on a specific forecasted outcome.

For variable j, forecast horizon s, and a constraint imposed at time $T+l$, the CRF is defined as:

$$CRF_j(s, l, \mathcal{S}\mathbf{y}_T = \boldsymbol{\delta}_y^c, \Omega_T) \equiv W_{T+s,\,T+l}^{(j)}(T+h),$$

where $W_{T+s,\,T+l}^{(j)}(T+h)$ is the observation weight assigned to the constrained value $y_{j,T+l}$ in the conditional forecast of y_{T+s}.

- Unlike GIRFs, which trace shocks to reduced-form innovations, CRFs capture the effects of **path restrictions on future observables**.
- Path conditioning imposes weaker—but operationally feasible and often more interpretable—restrictions than structural shock identification.
- Through the MA representation, imposed paths implicitly restrict the joint distribution of future shocks and reshape the entire forecast trajectory.
- CRFs provide a direct, quantitative measure of how each constrained observation influences future outcomes.

CRF and GIRF: Main Results

Proposition 1

(i) For a single-period conditioning on $y_{j,T+1}$, the conditional response function $CRF_i(s, 1, \delta_y^c, \Omega_T)$ coincides with the unscaled generalized impulse response function (GIRF) at horizon $s-1$, that is, with $A_{s-1}\Sigma^{(j)}\Sigma_{j,j}^{-1}$.

(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.

Implication: CRFs generalize impulse response analysis to path-based scenarios and form the basis for new scenario analysis tools.

Variable importance for conditioning sets

Let $w_{jk}(s)$ be smoothing weights for variable j, time k, horizon s. Scale by σ_j.
Overall importance (VI):

$$VI_j(s) = \frac{\sum_{k=1}^{T+h} |w_{jk}(s)\sigma_j|}{\sum_{j=1}^{n} \sum_{k=1}^{T+h} |w_{jk}(s)\sigma_j|}.$$

Marginal importance (MI) from imposed future paths only:

$$MI_j(s) = \frac{\sum_{k=T+1}^{T+h} |w_{jk}(s)\sigma_j|}{\sum_{j=1}^{n} \sum_{k=1}^{T+h} |w_{jk}(s)\sigma_j|}.$$

Thank you!

Questions and comments are welcome.