A Path Is Not a Forecast: Information Structures in Macroeconomic Scenario Analysis

An Application to Oil-Market Disruptions from the Iran Conflict

Aleš Maršál*

September 2026

Abstract

Macroeconomic scenarios are typically defined by paths for conditioning variables. We argue that a path is not a sufficient statistic: the same realised path can generate different decisions depending on whether agents know it in advance or experience it as a sequence of unanticipated innovations, and on how they price uncertainty around its continuation. A simple analytical model shows that, when impact resources are predetermined, information primarily changes their allocation. We quantify these mechanisms in a nonlinear DSGE model of an oil-exporting economy during the 2026 Iran conflict. The conflict combined a large oil-market disruption with conflicting and repeatedly revised signals about its duration and resolution, so the same current outcome could support very different beliefs about its continuation. Holding the realised oil-price and oil-production paths fixed, we cross two information structures: *i*) full-path information and *ii*) sequential surprise with linear and nonlinear decision rules, and add higher-order risk-pricing and stochastic-volatility counterfactuals. This design isolates the value of full information and separates it from nonlinear curvature, approximation error, constant-volatility risk, and time-varying uncertainty. We find that information element is negligible in the impact output but large in consumption, investment, sectoral allocation, and the external account. Full information can substantially discipline investment, while nonlinearities make the information effect strongly state-dependent. Higher perceived volatility induces investment retrenchment and external saving even when the realised oil path is unchanged. A macroeconomic scenario should therefore specify a realised path, an information structure, and a conditional risk law.

JEL: C63, E37, F41, Q43. **Keywords:** conditional forecasting; information structure; perturbation; stochastic volatility; oil windfall.

*King Abdullah Petroleum Studies and Research Center (KAPSARC), Riyadh, Saudi Arabia. E-mail: [ales.marsal@kapsarc.org]. This research was conducted while the author was at KAPSARC. The views expressed in this paper are solely those of the author and do not represent the views or policies of KAPSARC. All remaining errors are the author's own.

1 Introduction

Macroeconomic scenarios are usually described by paths. An oil-market scenario specifies a trajectory for prices and production; a monetary scenario specifies a trajectory for interest rates; a fiscal scenario specifies taxes or public expenditure. Once that trajectory has been imposed, the model is asked to fill in the rest of the economy. In a forward-looking model, however, a path answers only one part of the question. It states what happens, but not when agents learn that it will happen. A trajectory known at the initial date and the same trajectory discovered one realisation at a time are not equivalent economic objects. They generate different expectations, and therefore different decisions, even though the conditioned variables eventually trace exactly the same path.

This paper treats the information structure as part of the economic specification of a macroeconomic scenario. We hold the realised conditioning path fixed and vary when agents learn its continuation, using the standard Waggoner–Zha conditional forecast as the natural benchmark. In that benchmark, future conditioning innovations arrive as surprises and expectations evolve according to the model's law of motion. We contrast this with full-path information, under which agents know the entire trajectory from the initial date and optimise against it. In the nonlinear sequential-surprise economy, agents instead perfectly foresee the deterministic continuation implied by the current state, conditional on no further innovations; the next innovation needed to keep the economy on the analyst's path then arrives unexpectedly, and they reoptimise. Computationally, these two environments correspond to full perfect foresight and to rolling perfect-foresight solutions with one-period MIT shocks. Thus the same realised path can represent an economy that plans around its entire future from the outset or one that repeatedly plans around a continuation that subsequent surprises overturn. By repeating this path-identical experiment under linear and nonlinear decision rules, and then allowing agents to price uncertainty around the continuation, we separate the effects of information, nonlinear adjustment, and risk.

The 2026 conflict between the United States, Israel, and Iran makes this distinction economically salient. The conflict generated a large repricing of oil, a sharp increase in oil-market volatility, and a severe disruption to flows through the Strait of Hormuz. At the same time, public assessments of the likely duration and resolution of the conflict changed repeatedly. Early statements pointed to a campaign lasting several weeks; subsequent announcements alternated between imminent resolution, renewed military operations, and further negotiations, while negotiated arrangements were announced and later reversed. Section 3 documents this sequence using dated and archived sources. The conflict therefore created an unusual information environment in which the current state of the oil market was observable, but its continuation was repeatedly redefined by new political and military signals. The same observed price and production outcomes could consequently support materially different beliefs about how long the underlying disturbance, and hence its macroeconomic effects, would persist. For any forward-looking economy, those beliefs matter for decisions taken today. The episode therefore provides a natural stress test for the information structure embedded in a conditional forecast.

This perspective builds on a literature that has progressively enriched the information con-

tent of conditional forecasts. Waggoner and Zha (1999) provide the natural benchmark, in which conditioning is implemented through a sequence of innovations consistent with a prescribed path. Maih (2010) extends this framework to DSGE models and allows conditioning information to be anticipated by agents; Antolín-Díaz et al. (2021) gives conditional scenarios an explicit structural interpretation; and Montes-Galdón et al. (2022) moves from conditioning on point paths to conditioning on distributions. Our analysis brings this literature into direct contact with nonlinear DSGE solution methods. At first order, decision rules are local and certainty-equivalent; exact nonlinear solutions allow large shocks to move the economy through curved policy functions, while higher-order perturbations introduce state-dependent risk and volatility effects (Schmitt-Grohé and Uribe, 2004; Kim et al., 2008; Andreasen et al., 2018; Fernández-Villaverde et al., 2011). This allows us to ask not only whether information about the continuation matters, but whether a linear conditional forecast assigns the correct macroeconomic value to that information when the scenario is large and uncertain.

Why can the same realised path generate different macroeconomic outcomes? The key margin is expected duration. A temporary shock changes not only current resources but also beliefs about how long the associated change in demand or income will persist. We make this mechanism transparent in a simple two-period model in which current resources are fixed, while households allocate them between consumption, investment, and foreign assets. Investment depends on the expected payoff to capital after the current period and is therefore particularly sensitive to beliefs about the continuation. Consumption instead reflects perceived lifetime wealth, and the external account absorbs the difference between current resources and domestic absorption. When the continuation implied by the current state is more persistent than the path ultimately realised, sequentially surprised agents invest against a tail of demand that full-information agents know will disappear. The same realised shock can therefore generate little difference in current output but large differences in consumption, investment, and external saving. The quantitative model embeds this mechanism in a richer general-equilibrium environment and allows us to measure how strongly it operates for large scenarios.

We take this mechanism to a nonlinear real-business-cycle model of an oil-exporting economy. The model adds the general-equilibrium margins needed to study the economic allocation and is particularly well suited to our experiment because the conditioning variables, oil prices and oil production, enter through an exogenous oil block with invertible laws of motion. A prescribed path for prices and quantities can, therefore, be inverted into the sequence of innovations. Because the mapping from paths to innovations is explicit we can transparently isolate the differences in information and nonlinear adjustment.

The first result is that information about the continuation matters much more for the allocation of resources than for contemporaneous output. In the benchmark scenario, giving agents full-path information rather than exposing them to sequential surprise changes impact GDP by only 0.07 percentage points. Yet consumption rises by about 5.9 percentage points, investment falls by about 5.2 percentage points, and the trade balance declines by about 4.4 percentage points of GDP. The realised oil path therefore delivers almost the same current resource envelope under the two information structures, but agents use those resources very differently. Knowing that the windfall will be short-lived raises current consumption, discourages invest-

ment in capacity that would outlast the shock, and shifts the corresponding adjustment to the external account. The information effect is particularly pronounced in non-tradable investment, whose return depends most directly on the expected persistence of domestic demand.

The second result is that the source of forecast disagreement is strongly state-contingent. Our path-identical decomposition separates the effect of information from the effect of nonlinear adjustment and, crucially, identifies their interaction. Close to the steady state, the linear and nonlinear forecasts agree and the value of information is small. As the disruption grows, the decomposition changes sharply. In the benchmark scenario, the Waggoner–Zha forecast raises impact GDP by about 3.1 per cent, while its nonlinear sequential-surprise counterpart raises it by only about 0.5 per cent, so nonlinear adjustment dominates the aggregate response. For allocation variables, however, information remains quantitatively large even when its effect on GDP is small. Moving further through the scenario space changes the balance again: nonlinearities can attenuate the information effect, reverse its sign, or amplify it sharply. In the most severe scenarios, the nonlinear value of information reaches roughly ten times its linear counterpart for GDP and nineteen times for investment. There is therefore no single correction to the linear conditional forecast: which margin matters depends on where the scenario places the economy in the state space.

The third result is that uncertainty constitutes a distinct, and equally state-contingent, margin of adjustment. Holding both the realised oil path and the sequential-surprise information structure fixed, the higher-order solutions allow us to separate deterministic curvature from the pricing of risk and from changes in the volatility path itself. At constant volatility, risk pricing has little effect on aggregate GDP but materially changes the composition of expenditure, particularly investment. When volatility is allowed to rise along the scenario, the response becomes much larger: relative to the same realised oil path under flat volatility, cumulative investment falls by about 16 per cent and the external surplus rises by about 10 percentage points of GDP over the conditioning horizon. The effect is again state dependent.

When the economy is projecting a sizable investment response, for example, expanding non-tradable capacity in anticipation of persistent domestic demand, a rise in uncertainty makes waiting more valuable and can sharply reduce investment. In the most severe states, by contrast, the realised disruption has already weakened the case for those projects, leaving less investment for additional uncertainty to postpone. The marginal effect of uncertainty can therefore be smaller precisely when the underlying shock is larger.

Taken together, these results show why disagreement across conditional forecasts is economically informative rather than merely numerical. Forecasts can differ because agents know different continuations of the same realised path, because that path is propagated through nonlinear rather than linear decision rules, or because agents price uncertainty around the continuation. Our decomposition assigns these differences to distinct economic margins and shows that their relative importance is itself state dependent. Near the steady state the conventional linear forecast is often sufficient; farther away, information, nonlinear adjustment, and uncertainty can dominate in different combinations.

The paper sits at the intersection of three strands of macroeconomic research. First, it contributes to conditional forecasting and structural scenario analysis, building on Waggoner

and Zha (1999), Maih (2010), Antolín-Díaz et al. (2021), and Montes-Galdón et al. (2022). Our contribution is to treat the information structure as an explicit characteristic of the scenario and to compare alternative information structures while holding the realised path itself fixed. Second, the paper connects to the literature on news, foresight, and expectations (Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009; Barsky and Sims, 2011; Leeper et al., 2013; Chahrour and Jurado, 2018). That literature shows that the timing with which information about future fundamentals reaches agents can move current decisions; our experiment applies this insight to conditional forecasting and identifies its quantitative importance for the same realised macroeconomic path. Third, the paper builds on nonlinear and higher-order solution methods (Schmitt-Grohé and Uribe, 2004; Kim et al., 2008; Andreasen et al., 2018; Fernández-Villaverde et al., 2011) by using them to decompose forecast differences into information, nonlinear propagation, risk pricing, and time-varying uncertainty. The resulting decomposition turns differences across conditional forecasts into economically interpretable objects and shows that their relative importance varies with the state of the economy.

Section 2 develops the mechanism in a simple analytical model. Section 3 documents the information environment surrounding the 2026 conflict, and Section 4 presents the quantitative model. Section 5 defines the alternative information structures, solution methods, and path-identical conditioning experiment. Section 6 presents the benchmark decomposition, while Sections 8 and 7 study uncertainty and variation across the severity of the shock. Section 9 concludes.

2 Information and allocation along a fixed path

The purpose of this section is to isolate the economic effect of information about the continuation of a scenario. The realised path is held fixed. At any decision date, the two economies observe the same current windfall and face the same current resource constraint; they differ only in what they know about what comes next. We embed this distinction in the standard Fisherian consumption–saving problem, augmented with an investment margin. The underlying allocation problem is familiar (Fisher, 1930; Obstfeld and Rogoff, 1996); our interest is in the mapping from beliefs about the continuation to current consumption, investment, and foreign saving. We therefore state the relevant Fisherian decision rules and collect the derivations in Appendix A.

To build the stage, consider two otherwise identical small open economies. At a decision date t, both wake up with the same current resources, $Y_t = \bar{Y} + w_t^*$, where $\bar{Y}$ denotes normal income and w_t^* the windfall realised at that date. These resources can be consumed, invested in productive capacity, or saved abroad. The two economies eventually experience the same sequence $\{w_\tau^*\}$; what differs is the information they receive about that sequence. One may think of them as reading different editions of the same newspaper. The full-information edition reports the entire remaining path. The sequential-surprise edition reports today's realised windfall and the economy's law of motion. Its reader therefore plans against the continuation implied by the current state, conditional on no further innovations. If the analyst's path requires another innovation next period, that innovation arrives unexpectedly and the economy forms

a new plan.

This distinction already contains the mechanism. At a given decision date the two economies have the same current "pie," so information cannot change current resources; it can only change how those resources are allocated. Their views of the future, however, can be very different. An economy that expects a stronger continuation perceives greater future wealth and a higher payoff to capacity installed today. It therefore has an incentive to consume and invest more and to save less abroad. Moreover, the ranking need not remain fixed as the scenario unfolds. Early in a front-loaded episode, full information may reveal an expansion that is not yet visible in the current state. At the peak, the same information reveals the coming decline, while the sequential-surprise economy continues to extrapolate the boom. The full-information effect can therefore change sign over the course of a single realised path.

2.1 One realised path, two information structures

Let the analyst fix the realised windfall path $\{w_t^*\}_{t=0}^T$. Under *full-path information*, or full information, agents know the entire remaining path at date t,

$$E_t^F w_{t+s} = w_{t+s}^*, \qquad s \geq 1. \tag{1}$$

Under *sequential surprise*, agents observe the current state and know its law of motion, but future conditioning innovations are not in their information set. For the AR(1) representation

$$w_{t+1} = \rho w_t + \varepsilon_{t+1}, \qquad 0 < \rho < 1, \tag{2}$$

their continuation at date t is

$$E_t^S w_{t+s} = \rho^s w_t^*, \qquad s \geq 1. \tag{3}$$

The next point on the analyst's path is realised through the innovation $\varepsilon_{t+1}^* = w_{t+1}^* - \rho w_t^*$. When this innovation arrives, agents observe it and reoptimise from the new state. Sequential surprise is therefore a sequence of internally consistent plans, each conditional on the information available when the plan is made. Full information and sequential surprise generate the same path ex post but generally imply different continuations ex ante.

For any endogenous variable x_t, define the *full-information effect* as

$$\mathcal{A}_{x,t} \equiv x_t^F - x_t^S. \tag{4}$$

Elsewhere in the paper we refer to this variable-specific outcome differential as the macroeconomic value of full information[1].

[1]It is not a welfare measure, and its sign need not be positive. A negative $\mathcal{A}_{x,t}$ simply means that knowing the continuation lowers the variable in question relative to sequential surprise.

2.2 A Fisherian allocation problem

The Fisherian problem translates these alternative continuations into current choices. At date t, realised resources are allocated among consumption c_t, investment in productive capacity i_t, and foreign assets b_{t+1}, which earn the gross world return R:

$$c_t + i_t + b_{t+1} = Y_t. \tag{5}$$

The common-pie argument is contemporaneous rather than intertemporal. The two economies begin date t with the same Y_t, but today's investment and foreign saving affect future resources. Hence identical exogenous windfall paths need not imply identical future endogenous consumption or wealth.

Beliefs about the continuation enter the allocation problem through two objects. The first is perceived future windfall wealth,

$$W_t^j = \sum_{s=1}^{T-t} R^{1-s} E_t^j[w_{t+s}], \tag{6}$$

and the second is the expected return shifter on capacity installed today,

$$Q_t^j = \sum_{s=1}^{T-t} R^{1-s} E_t^j[d(w_{t+s})], \qquad d'(w) > 0, \qquad j \in \{F, S\}. \tag{7}$$

The assumption $d'(w) > 0$ captures the complementarity relevant for the mechanism: a stronger future windfall raises future demand and therefore the payoff to capacity installed in advance. Thus W_t^j summarises perceived future resources, whereas Q_t^j summarises the expected profitability of current investment.

Let $\Phi(i)$ denote the future output of installed capacity, with $\Phi' > 0$ and $\Phi'' < 0$. Fisher separation gives the investment rule as the solution to

$$Q_t^j \Phi'(i_t^j) = R, \tag{8}$$

so that $i_t^j = i(Q_t^j)$ with $i'(Q) > 0$. Define the net value of the optimally chosen investment programme in next-period units by

$$H(Q) \equiv \max_i \{Q\Phi(i) - Ri\}. \tag{9}$$

Since $H'(Q) = \Phi(i(Q)) > 0$, both perceived windfall wealth and the profitability of investment raise perceived lifetime resources. With log utility, the standard Fisher consumption rule can be written as

$$c_t^j = \frac{1}{1+\beta}\left[Y_t + \frac{W_t^j + H(Q_t^j)}{R}\right], \qquad b_{t+1}^j = Y_t - c_t^j - i_t^j. \tag{10}$$

These expressions formalise the intuition above. Information does not change the resources already in hand. It changes perceived lifetime wealth and the return to committing part of today's resources to productive capacity. The resulting differences in consumption and

investment are accommodated by the foreign-asset position:

$$\mathcal{A}_{c,t} + \mathcal{A}_{i,t} + \mathcal{A}_{b,t+1} = 0. \tag{11}$$

Lemma 1 (Continuation dominance)**.** *Suppose that, at a given decision date t, the continuation under full information lies weakly above the sequential-surprise continuation at every future horizon, $E_t^F w_{t+s} \geq E_t^S w_{t+s}$ for all $s \geq 1$, with a strict inequality for at least one s. Then*

$$\mathcal{A}_{c,t} > 0, \qquad \mathcal{A}_{i,t} > 0, \qquad \mathcal{A}_{b,t+1} < 0. \tag{12}$$

If the ordering of the two continuations is reversed, all three inequalities are reversed.

The result follows directly from the positive weights in (6)–(7), the monotonicity of d, and the Fisherian decision rules. Its economic content is simple. Conditional on the same current resources, the economy that perceives the stronger continuation is both wealthier and faces a higher expected payoff to installed capacity. It therefore consumes more and invests more today, financing part of the additional absorption through a lower foreign-asset position. The full-information effect has no unconditional sign: it depends on whether the realised script lies above or below the continuation that a sequentially surprised agent extrapolates from the current state.

2.3 A front-loaded scenario and the sign of information

This comparison becomes especially informative when the realised path is front-loaded. Sequentially surprised agents do not carry one fixed forecast through the episode; they form a new continuation after every realised innovation. Before the peak, the full-information agent may see a future expansion that is absent from the extrapolation of the current state. At the peak, the ordering can reverse: the sequentially surprised agent extrapolates the high current state with persistence ρ, whereas the full-information agent already knows the subsequent decline. The same information structure can therefore raise a decision early in the episode and lower it later, even though the analyst never changes the realised path.

Proposition 1 (Sign reversal along a front-loaded path)**.** *Let t_0 denote the first decision date and t^* the peak of the realised windfall. Suppose that the realised continuation initially dominates the extrapolation from the current state,*

$$w^*_{t_0+s} \geq \rho^s w^*_{t_0} \quad \textit{for all } s \geq 1, \tag{13}$$

with strict inequality at some horizon, while after the peak it lies below the continuation implied by the same law of motion,

$$w^*_{t^*+s} \leq \rho^s w^*_{t^*} \quad \textit{for all } s \geq 1, \tag{14}$$

again with strict inequality at some horizon. Then the full-information effect on consumption and investment is positive at t_0 and negative at t^*, while the effect on foreign assets has the opposite sign.

Proposition 1 is an immediate application of Lemma 1 at two different decision dates. Its economic content is the key result. At the beginning of a front-loaded episode, full information reveals the increase that is still to come. The informed economy therefore perceives more future resources and a higher payoff to capacity than an economy extrapolating from the relatively small current disturbance. At the peak, the comparison reverses. The sequential-surprise economy projects the peak into the future, while the informed economy sees the decline already written into the path. Full information therefore raises consumption and investment relative to sequential surprise before a foreseen peak, but lowers them once extrapolation of the current state implies a continuation that is too persistent.

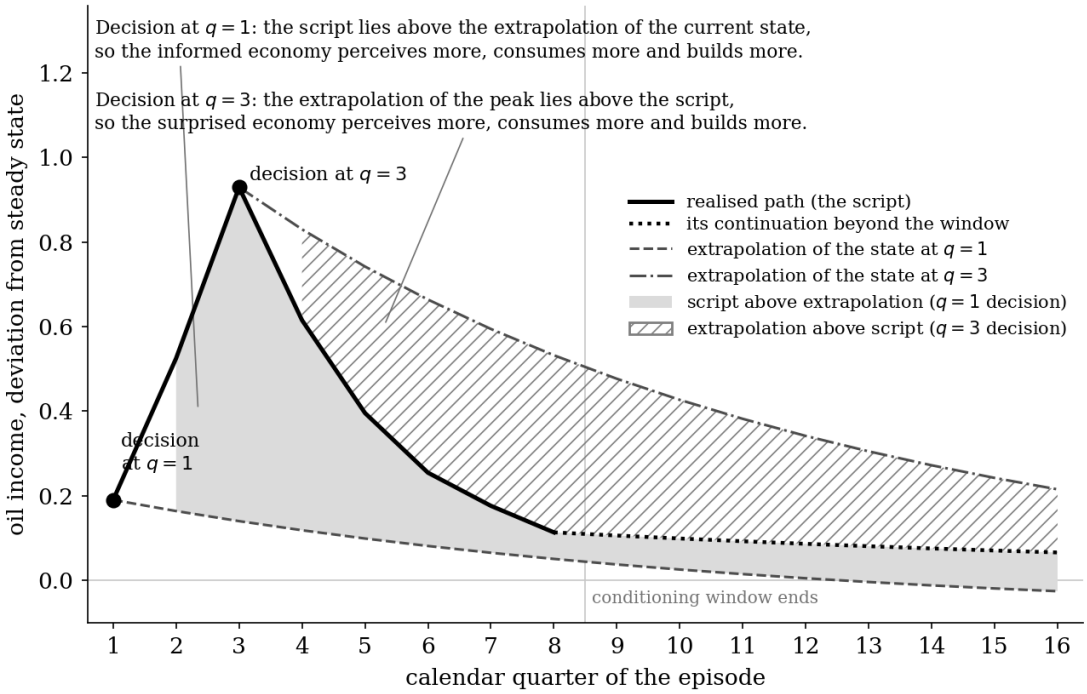


Figure 1: **One realised path, two beliefs about its continuation.** The solid line is the realised windfall path. The dashed line is the sequential-surprise continuation formed from the state at the first decision, and the dash-dotted line is the continuation formed from the state at the peak. At the first decision the script lies above the extrapolated continuation, so full information implies higher perceived wealth and a higher return to capacity. At the peak the ordering reverses: extrapolation places more weight on a persistent tail than the realised path. The shaded regions visualise the two dominance conditions in Proposition 1.

Figure 1 clarifies the role of duration. What matters is not persistence in isolation, but the difference between the persistence agents infer from the current state and the persistence of the path that is actually realised. Holding the realised post-peak path fixed, a larger ρ raises the sequential-surprise continuation from a positive peak and widens the post-peak gap between the two information structures. Conversely, a faster realised reversal widens the same gap for a given ρ. Information is most consequential when the historical law of motion and the realised duration of the episode point to different continuations.

The model therefore delivers a concise result. A common realised path need not imply a common allocation because the relevant continuation is information-set dependent. Moreover,

the full-information effect is itself state dependent: as a front-loaded path unfolds, the same information can first raise and later lower consumption and investment relative to sequential surprise. The economic object is the gap between the script and the continuation agents would otherwise infer from the state they currently observe.

3 The information environment of the 2026 conflict

Section 2 distinguished two ways of experiencing the same realised path. Under full information, agents know its continuation from the outset. Under sequential surprise, they repeatedly form plans from the state and information currently available, only for subsequent innovations to overturn those plans. The 2026 conflict between the United States, Israel, and Iran provides a particularly informative real-world analogue of the second environment. The current state of the oil market was observable, but authoritative statements about the continuation repeatedly supplied concrete horizons that were subsequently revised, extended, or reversed. Agents therefore faced not one uncertain forecast formed at the beginning of the episode, but a sequence of provisional continuations against which decisions could be made and subsequently reconsidered as new information arrived.

This section documents that information environment. Three features are central. First, the conflict was preceded by a visible military build-up but still generated a sharp repricing of oil when hostilities began. Second, public statements repeatedly changed the expected horizon of the conflict and of negotiations intended to end it. Third, the physical disruption to energy flows evolved on a considerably slower clock than many of those announcements. Together, these features make the episode particularly useful for studying the difference between a path known in advance and a path navigated through repeated revisions of its expected continuation.

3.1 A visible build-up, but an unresolved onset

The military build-up was substantial and public. From late December 2025, the United States moved air, naval, and missile-defence assets into the region on a scale that the Congressional Research Service describes as the largest since 2003, including a carrier strike group; the deployment continued for roughly two months (Congressional Research Service, 2026). The build-up made the possibility of conflict increasingly salient but left its timing, scale, and duration unresolved.

When hostilities began on 28 February 2026, nearly nine hundred strikes were reported in the first twelve hours, the Iranian Supreme Leader was among the dead, and the stated objective was to "eliminate imminent threats from the Iranian regime" (Congressional Research Service, 2026). Brent rose by about eight per cent on the first trading day, while oil implied volatility subsequently climbed from a pre-war average near 47 to a peak above 120.[2] The market response is consistent with substantial information about the realised severity of the conflict arriving at the onset despite the prolonged and visible build-up.

[2]Daily Brent (FRED `DCOILBRENTEU`) and OVX. All market figures in this section are descriptive; Appendix E provides the date-by-date verification ledger.

3.2 A moving horizon

The continuation remained difficult to infer after the war had begun. On 2 March, three days into the conflict, the announced horizon was four to five weeks, with the capability to continue for longer.[3] One week later, the war was described in a telephone interview as essentially complete.[4] By 1 April, however, a national address announced military operations for a further two to three weeks.[5] Other dated statements in the intervening period placed the remaining horizon anywhere from a few days to as long as necessary.

For forward-looking decisions, these were not merely different descriptions of the same current state. Each implied a different continuation against which households, firms, financial markets, and policymakers could plan. A conflict expected to end within days, one expected to last several weeks, and one with no clear terminal date imply different expected paths for energy prices, trade flows, income, costs, and demand even when today's observed oil-market conditions are identical. When new information superseded the previous horizon, plans formed against that continuation had to be reconsidered.

Ultimatums followed a similar pattern. A 48-hour deadline for reopening the Strait of Hormuz, announced on 21 March, was successively extended to five days, then ten days, subsequently delayed again, folded into a two-week suspension, and eventually extended indefinitely.[6] The economic content lies in the sequence rather than in any individual announcement: a seemingly precise continuation repeatedly ceased to be the relevant one before the underlying conflict had been resolved.

The pattern was particularly pronounced around negotiations. Between late March and mid-June, public statements repeatedly placed a settlement within days or weeks: five days on 23 March; an agreement described as largely negotiated on 23 May; a final determination within days on 29 May; victory within two weeks on 8 June; a signing within two or three days on 9 June; a settlement that weekend on 11 June; and a deal described as sufficiently close to be signed on 13 June.[7] On 1 June, the divergence was contemporaneous: Iranian state media reported that Tehran had suspended the talks, while Washington described the same negotiations on the same day as proceeding rapidly.[8]

The authority of these signals is part of their economic relevance. Statements about the likely end of a conflict by the governments conducting it are salient inputs into expectations about its continuation. Agents can therefore form internally coherent plans around the currently announced horizon and still be surprised when subsequent political or military information renders that horizon obsolete. This repeated plan revision is the feature of the episode that most closely parallels the sequential-surprise economy studied in the paper.

[3] White House ceremony, 2 March 2026; Al Jazeera, 2 March 2026 (accessed 17 July 2026); corroborated by Military Times, 2 March.

[4] CBS News, 9 March 2026 (accessed 17 July 2026).

[5] Address to the nation, 1 April 2026; PolitiFact/PBS timeline of duration statements, 21 April 2026 (accessed 17 July 2026).

[6] ABC News, 22 April 2026 (accessed 18 July 2026). ABC characterised the sequence as "a pattern of setting unenforced deadlines."

[7] Dated compilation: CBS News, 11 June 2026 (accessed 17 July 2026), with the 13 June entry from NPR, 13–14 June 2026 (accessed 18 July 2026).

[8] CNBC, 1 June 2026, citing Iran's state-affiliated Tasnim agency (accessed 18 July 2026); the same-day US statement is reported in the CBS compilation.

Some revisions reflected new fundamental events, whereas others primarily revised the announced horizon itself. Table 1 records selected commitment–supersession pairs and flags cases in which a discrete intervening event accompanied the reversal. A sixty-day memorandum of understanding signed in mid-June illustrates the distinction. The naval blockade was lifted on 18 June under its terms; twenty-five days into the sixty-day horizon, after Iranian fire on a commercial vessel, the blockade was reinstated.[9]

Table 1: Selected revisions to publicly stated horizons during the 2026 conflict. The first statement in each row records the initial horizon; the second records the subsequent statement or event that superseded it. The full ledger, including URLs, access dates, verification tags, and market context, is reported in Appendix E.

dimension	initial statement	superseding statement or event	days	intervening event?
military duration	2 Mar: projected horizon of four to five weeks	1 Apr: a further two to three weeks of operations announced	30	no discrete event
military status	9 Mar: conflict described as essentially complete	1 Apr: further military operations announced; fighting continued thereafter	23	continued fighting
negotiation	18 Jun: 60-day memorandum; blockade lifted	13 Jul: blockade reinstated on day 25 of the agreement	25	vessel attack on 11 Jul

3.3 Public horizons and physical persistence

The horizon communicated through political statements moved much faster than the physical disruption to energy flows. Shipping data compiled for the House of Commons Library show crude movements through the Strait of Hormuz falling by roughly 95 per cent and liquefied-gas movements by 99 per cent over the conflict. The Emirati state oil company estimated that full flows would not return before 2027 even under a rapid agreement (House of Commons Library, 2026). The political horizon and the physical horizon therefore diverged sharply.

This distinction matters for any economy exposed to the disruption. The macroeconomic persistence of an oil shock depends on the path of prices and quantities, not solely on the announced horizon of the conflict producing them. An announcement of imminent resolution may therefore imply a very different economic continuation depending on whether energy flows, production constraints, and prices are expected to normalise immediately or only with a lag.

The repeated revisions also had precedent. During the June 2025 Iran episode, a decision announced as forthcoming "within the next two weeks" was followed by strikes two days later.[10] This earlier episode does not determine how agents interpreted the 2026 announcements, but it formed part of the public information set against which new statements were received.

3.4 From the information environment to the model experiment

The public record maps naturally into the distinction developed in Section 2. Full-path information describes the limiting case in which agents know from the outset how the episode

[9] Al Jazeera, 18 June 2026; NPR, 13 July 2026 (both accessed 17 July 2026). The verification ledger records the 11 July vessel attack as an intervening event.

[10] NPR, 20 June 2025 (text edition accessed 17 July 2026).

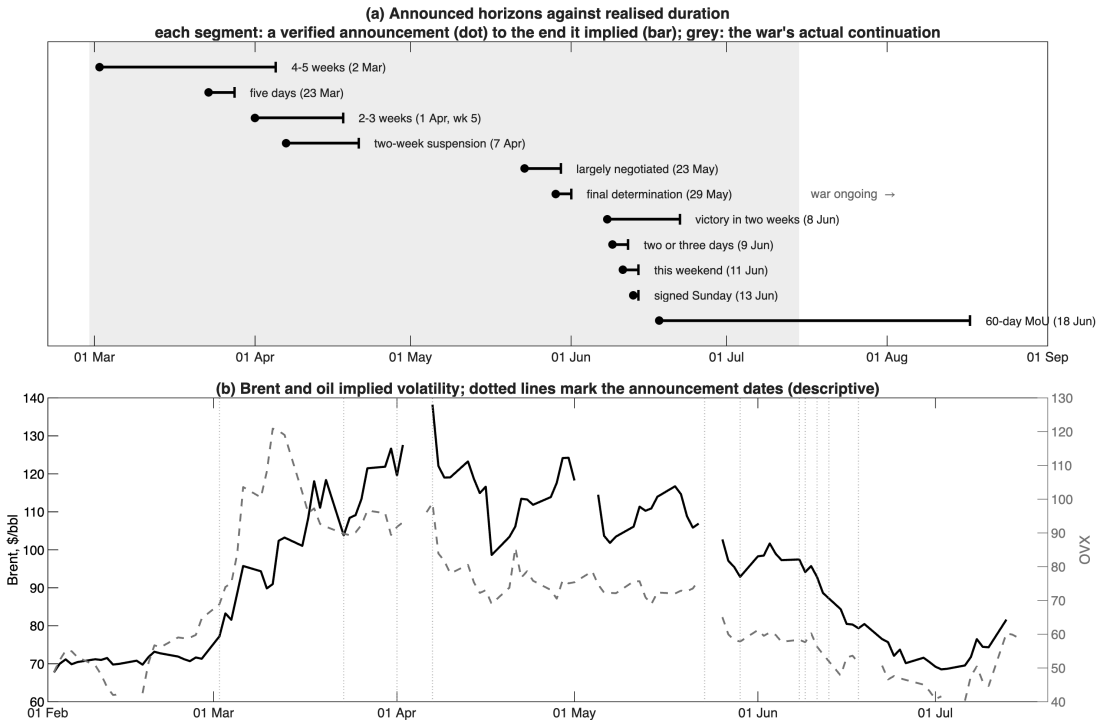


Figure 2: **Public horizons and oil-market conditions during the 2026 conflict.** Panel (a) plots each verified announcement together with the horizon it implied and the subsequent realised continuation of the conflict. Marker types distinguish revisions accompanied by an identifiable intervening event from revisions without one. Panel (b) reports daily Brent prices and oil implied volatility, with announcement dates marked.

will unfold. Sequential surprise describes an economy in which agents instead form a coherent continuation from the state and information currently available, only to revise that plan when new information arrives. The empirical episode illustrates why this distinction is economically meaningful: a plan can be entirely reasonable given the information available when it is made and nevertheless be overturned by subsequent news.

In the quantitative implementation, sequential surprise is represented by rolling perfect-foresight solutions with one-period MIT shocks. At each decision date, agents know the model and perfectly foresee the deterministic continuation implied by the current state conditional on no further innovations. The next innovation required to keep the economy on the analyst's scenario path then arrives unexpectedly, and agents reoptimise. The conflict does not imply that private agents literally followed this forecasting rule; rather, the repeated sequence of apparently informative horizons followed by new political and military surprises provides the economic intuition for the rolling-plan experiment. Full perfect foresight provides the corresponding counterfactual in which the entire continuation is known at the initial date.

The episode also motivates the third dimension of the analysis. Repeated revisions change not only the expected continuation but potentially the confidence attached to it. The sharp rise in oil implied volatility during the conflict provides a direct market counterpart to this second margin. We therefore distinguish throughout the quantitative analysis between what eventually happens, when agents learn about that continuation, and the risk they attach to it. The realised oil-price and oil-production paths define the first object; the information structure defines the second; and the higher-order experiments introduce the third. This separation allows the same economic scenario to be studied under full information, sequential surprise, and priced uncertainty without changing the realised path itself.

4 A model of an oil exporter

The laboratory is a calibrated small open economy built around three facts of an oil exporter under a fixed exchange rate: oil income accrues to the state and reaches the private sector through the fiscal accounts; the private economy divides into a tradable sector whose prices the world sets and a non-tradable sector whose prices domestic demand sets; and the peg removes nominal adjustment, so relative prices move through non-tradable inflation. The model has 85 endogenous variables; this section presents each block's economics with the minimum notation needed to read the results, and Appendix B states every equation, the steady-state construction, and the calibration.

4.1 Households, firms, and the peg

A representative national household consumes a bundle of goods and energy services, supplies labour to both production sectors, accumulates the two sectors' capital stocks subject to investment adjustment costs, and trades a foreign bond at a rate that carries a small debt-elastic premium (which closes the open economy in the usual way). Utility is non-separable between consumption and energy services — energy is a complement to consumption spending, which matters because energy prices are administered and stable, making energy services a stabilising

component of household welfare during price shocks. An expatriate-labour block captures a structural feature of Gulf labour markets: a fraction of employment is internationally mobile, consumes a fixed share of its wage domestically, and remits the rest; this mutes the wage channel of domestic demand booms.

Production is Cobb–Douglas–CES in both sectors: labour combines with a capital–energy composite, so energy enters as an input alongside capital, not only as a consumption good. The tradable sector sells at world prices — the law of one price holds and the peg fixes the nominal anchor — so a domestic demand boom cannot raise the tradable price; it raises the non-tradable price and the real exchange rate appreciates through that channel alone. This asymmetry — *world-pinned tradable returns, demand-driven non-tradable returns* — is load-bearing for everything that follows: it is why beliefs and risk, which operate through domestic demand, express themselves almost entirely in the non-tradable sector's investment. The mechanism is generic to pegged oil exporters by construction, which is what makes the calibration representative of the GCC class rather than idiosyncratic.

The state receives oil and gas export income and domestic energy-sales income, spends on public consumption, invests in oil-sector capital at a fixed rate, issues domestic bonds held by households, and transfers the residual to households. This fiscal routing is the windfall's front door: an oil price increase arrives as transfers and public demand, not as private factor income so the private sector's response to the windfall is a response to *expected* transfers, which is precisely where the information structure bites.

4.2 The oil block

Oil is the exogenous driver, and its specification is deliberately exact in logarithms because Section 5 exploits that exactness. The log oil price deviation $\hat{p}_t$ follows an autoregression whose innovation is scaled by a stochastic volatility state:

$$\hat{p}_t \;=\; \rho_P\,\hat{p}_{t-1} \;+\; \sigma_{t-1}\,\varepsilon_t^P, \qquad \log(\sigma_t/\bar{\sigma}) \;=\; \rho_\sigma \log(\sigma_{t-1}/\bar{\sigma}) \;+\; \sigma_\sigma\,\varepsilon_t^\sigma, \tag{15}$$

with the price level $P_t^O = \bar{P}^O e^{\hat{p}_t}$. In words: the oil price mean-reverts ($\rho_P = 0.937$ quarterly), and the *size* of its surprises is itself a persistent, shock-driven state — the channel through which announced-path volatility will enter in Section 8. The timing convention matters and recurs later: the volatility state dated $t-1$ scales the innovation at t, so an uncertainty shock today first amplifies tomorrow's price surprise. Oil production follows a log-linear law that mean-reverts to a level implied by oil-sector capital; because oil-sector investment is a fiscal *parameter*, that capital path is deterministic, and the production law is, like the price law, exactly invertible in its innovation. Gas production and prices follow parallel processes. Two conversion constants map model units to data: one model price unit is \$36.38 per barrel (steady-state price \$63.78) and one production unit is 80.33 million barrels per day of flow equivalent (steady-state Saudi production 9.5 mb/d).

4.3 Calibration

The model is calibrated at quarterly frequency to Saudi national accounts and oil-market data; Appendix B tabulates all parameter values with sources. Three calibration choices matter most for the results. First, the oil price process ($\rho_P = 0.937$, unconditional innovation scale $\bar{\sigma} = 0.15$) matches quarterly real Brent dynamics over 1992–2026 (Appendix E); at these values a K3-sized price path is a 0.96 log-point displacement — large enough that linear and exact representations of the same shock differ materially. Second, the volatility-of-volatility parameter σ_σ, to which Section 8 is devoted, is set to 0.27, the value implied by matching the dispersion of oil implied volatility (OVX); the calibration inherited a placeholder here, and Appendix E documents the estimation that replaced it. Third, capital adjustment costs and the debt-elastic premium are standard ($h_T = 0.22$, $h_{NT} = 0.15$, $\xi = 0.05$); they discipline how fast the composition effects at the heart of the paper can play out, and none of the qualitative results depends on their fine values.

5 Information structures and decomposition

The previous sections separated three objects that are often bundled together in a conditional forecast: the realised path, the information agents have about its continuation, and the risk surrounding that continuation. This section turns that distinction into a quantitative design. The principle is simple. We keep the economic scenario fixed and change one feature of the forecast at a time. The first comparison changes what agents know. The second changes whether their decisions are governed by linear or nonlinear policy rules. A separate higher-order branch then introduces priced risk while holding the sequential-surprise information structure fixed.

This organisation is useful because differences across conditional forecasts otherwise have no unique economic interpretation. A forecast may change because agents know a different continuation, because the same continuation is propagated through a nonlinear economy, or because uncertainty affects current decisions. The objects defined below assign these differences to separate margins.

5.1 Four conditional forecasts

The deterministic experiment crosses the two information structures introduced in Section 2 with two solution representations. Along the information dimension, agents either face *sequential surprise* or have *full-path information*. Along the solution dimension, decisions are computed from the first-order approximation or from the exact nonlinear deterministic system. This gives the four cells in Table 2.

The upper-left cell, **WZ**, is the natural benchmark. It is the first-order conditional forecast in the spirit of Waggoner and Zha (1999). The sequence of conditioning innovations required to reproduce the prescribed path is recovered in the linear model, but future conditioning innovations do not enter agents' information sets before they occur. Agents therefore plan from the current state and the model's law of motion.

Table 2: **Information crossed with solution method.** All four conditional forecasts are constructed to reproduce the same realised oil-price and oil-production paths.

	Sequential surprise	Full-path information
Linear	WZ	PF_{lin}
Exact nonlinear deterministic	MIT	PF

The lower-left cell, **MIT**, keeps this information structure and changes the solution method. At each decision date agents observe the current conditioning innovation and perfectly foresee the nonlinear deterministic continuation that would follow if no additional innovation arrived. The model is solved by perfect foresight, the first-period decision is retained, and the next conditioning innovation then arrives unexpectedly. Repeating this procedure generates the rolling sequence of plans described in Sections 2 and 3.

The right-hand column instead gives agents the entire continuation at the initial date. $\mathbf{PF}_{\text{lin}}$ solves this full-information experiment on the first-order approximation, while **PF** solves the same experiment using the exact nonlinear deterministic equations. The labels PF and MIT are therefore computational implementations. The economic objects remain full information and sequential surprise.

At first order, the solution is certainty-equivalent: conditional variances do not enter current decision rules (Schmitt-Grohé and Uribe, 2004). This removes risk pricing from the upper row, but not the information effect. WZ and PF_{lin} can differ because the same future conditioning innovations enter agents' information sets at different dates.

5.2 Holding the scenario fixed

Comparing the four cells requires a precise definition of what is held fixed. Our object is the realised economic path, not the numerical innovations used to generate it. Let $\mathcal{P}^*$ denote the prescribed level paths for the conditioning variables. For a solution representation m, let $\mathcal{R}^m$ denote the recovery operator that maps those level targets into the innovations required by that representation, and let $\mathcal{S}_t^{m,j}$ propagate those innovations under information structure $j \in \{S, F\}$:

$$u^m = \mathcal{R}^m(\mathcal{P}^*), \qquad y_t^{m,j} = \mathcal{S}_t^{m,j}(x_0, u^m). \tag{16}$$

This distinction matters away from the steady state. A first-order representation and a nonlinear representation need not require the same numerical innovations to reproduce the same path in economic levels. Holding the innovations fixed would therefore change the scenario when the solution method changes. We instead recover the innovations appropriate to each representation and require each forecaster to reproduce the same prescribed oil-price and oil-production paths. Section C describes the inversion in the oil block; Appendix D gives the corresponding operator notation and derivatives.

This path-identical convention makes the comparisons below economically meaningful. When two forecasts differ, they disagree about the endogenous response to the same realised oil-market outcome rather than about the oil-market outcome itself. It also means that a linear–nonlinear difference should be read as a *linearisation gap in the conditional-forecast op-*

erator: both propagation through the DSGE model and the innovations needed to recover the fixed level path may contribute to that gap.

5.3 Information, linearisation, and their interaction

The four deterministic cells form a square, and each side of the square changes one assumption. Section 2 defined the full-information effect as the difference between the informed and sequential-surprise economies. Here it can be measured under both solution representations:

$$\mathcal{A}_t \equiv y_t^{PF} - y_t^{MIT}, \qquad \mathcal{A}_t^{\text{lin}} \equiv y_t^{PF_{\text{lin}}} - y_t^{WZ}. \tag{17}$$

The first object is the exact nonlinear full-information effect; the second is the value assigned to the same information change by the linear model. Both are variable-specific outcome differentials rather than welfare measures, and either may be positive or negative.

The vertical comparisons measure the effect of replacing the first-order conditional forecast with the nonlinear one while preserving the information structure:

$$\mathcal{L}_t^S \equiv y_t^{MIT} - y_t^{WZ}, \qquad \mathcal{L}_t^F \equiv y_t^{PF} - y_t^{PF_{\text{lin}}}. \tag{18}$$

We refer to these as the sequential-surprise and full-information *linearisation gaps*. Their difference has a particularly useful interpretation:

$$\mathcal{I}_t \equiv \mathcal{A}_t - \mathcal{A}_t^{\text{lin}} = \mathcal{L}_t^F - \mathcal{L}_t^S. \tag{19}$$

The equality is an exact identity. $\mathcal{I}_t$ measures the interaction between information and nonlinearity: it is the change in the full-information effect generated by moving from linear to nonlinear decision rules, or, equivalently, the difference between the nonlinear corrections under the two information structures. This is distinct from "nonlinearity" in general. Both vertical gaps can be large while $\mathcal{I}_t$ is small if nonlinear adjustment affects the two information structures similarly.

The same definitions apply date by date. During the conditioning window, MIT is re-solved every period, so

$$\mathcal{A}_t = y_t^{PF} - y_t^{MIT,(t)}, \qquad t = 1, \ldots, H, \tag{20}$$

where $y_t^{MIT,(t)}$ denotes the first-period outcome of the rolling solve in which the t-th conditioning innovation has just been observed. This is the quantitative counterpart of the sign reversal in Proposition 1: as the realised path unfolds, the gap between the script and the continuation inferred from the current state changes with the decision date, and so can $\mathcal{A}_t$.

The square also gives a local benchmark for interpreting the results. Scale the scenario away from the steady state as $\mathcal{P}(\lambda) = \bar{\mathcal{P}} + \lambda\Delta\mathcal{P}$, with $\lambda = 0$ at the steady state. Under differentiability of the recovery and solution maps, the linear and nonlinear full-information effects have the same first-order term, so

$$\mathcal{I}_t(\lambda) = O(\lambda^2). \tag{21}$$

Hence information and nonlinearity separate locally. Their interaction is a higher-order object that emerges as the scenario moves the economy away from the expansion point. This result provides the formal benchmark for the state-dependent pattern documented in Section 7. The derivation, including the separate roles of curvature in shock recovery and in DSGE propagation, is in Appendix D.

5.4 Risk as a separate branch

Risk is not a third information structure. The higher-order experiments retain sequential surprise and ask how the forecast changes when decisions respond to the distribution around the expected continuation. The state-space notation of Andreasen et al. (2018) is useful for making this distinction explicit. Write the DSGE solution as

$$y_t = g(x_t; \sigma), \qquad x_{t+1} = h(x_t; \sigma) + \sigma \eta \varepsilon_{t+1}, \tag{22}$$

where σ scales uncertainty. At first order, $g_\sigma = h_\sigma = 0$, so the decision rules are certainty-equivalent (Schmitt-Grohé and Uribe, 2004). At higher order, curvature and derivatives with respect to σ generate corrections for uncertainty; at third order these corrections can also depend on the state of the economy. The pruned state-space representation of Kim et al. (2008) and Andreasen et al. (2018) allows these effects to be simulated without introducing the explosive terms generated by naive iteration of a higher-order polynomial approximation.

We use three higher-order forecasts. $\mathbf{HO}_{\mathrm{CE}}$ retains the third-order curvature of the policy functions but removes the terms that price risk. **HO** restores those risk corrections at constant volatility. $\mathbf{HO}_{\mathrm{SV}}$ additionally allows the conditional volatility of the oil process to follow a prescribed path. All three retain sequential surprise, so their differences have no information-structure content.

These forecasts extend the sequential-surprise comparison through the chain

$$y_t^{HO} - y_t^{WZ} = \underbrace{\left(y_t^{MIT} - y_t^{WZ}\right)}_{\mathcal{L}_t^S} + \underbrace{\left(y_t^{HO_{\mathrm{CE}}} - y_t^{MIT}\right)}_{\mathcal{T}_t} + \underbrace{\left(y_t^{HO} - y_t^{HO_{\mathrm{CE}}}\right)}_{\mathcal{R}_t^{\mathrm{pure}}}. \tag{23}$$

The first term is the linearisation gap under sequential surprise. $\mathcal{T}_t$ is an approximation diagnostic: it compares the third-order certainty-equivalent representation with the exact nonlinear deterministic solution. $\mathcal{R}_t^{\mathrm{pure}}$ isolates risk pricing at a common approximation order. Finally,

$$\mathcal{R}_t^{\mathrm{tv}} \equiv y_t^{HO_{\mathrm{SV}}(\mathrm{spike})} - y_t^{HO_{\mathrm{SV}}(\mathrm{flat})} \tag{24}$$

isolates the effect of the volatility path itself.

Table 3 summarises the resulting objects. The deterministic square identifies information, linearisation, and their interaction. The higher-order branch then separates approximation error from risk pricing and from time variation in volatility.

The decomposition is deliberately constructed so that each difference changes one economically interpretable margin. For example, HO−WZ is not itself a measure of risk because it also contains the nonlinear conditional-forecast gap and the third-order approximation resid-

Table 3: **Objects in the decomposition.**

Object	Definition	Interpretation
$\mathcal{A}$	PF − MIT	full-information effect, nonlinear
$\mathcal{A}^{\text{lin}}$	PF_{lin} − WZ	full-information effect, linear
$\mathcal{L}^S$	MIT − WZ	linearisation gap under sequential surprise
$\mathcal{L}^F$	PF − PF_{lin}	linearisation gap under full information
$\mathcal{I}$	$\mathcal{A} - \mathcal{A}^{\text{lin}} = \mathcal{L}^F - \mathcal{L}^S$	information–nonlinearity interaction
$\mathcal{T}$	HO_{CE} − MIT	third-order approximation diagnostic
$\mathcal{R}^{\text{pure}}$	HO − HO_{CE}	risk pricing at constant volatility
$\mathcal{R}^{\text{tv}}$	HO_{SV}(spike) − HO_{SV}(flat)	time-varying volatility

ual. Likewise, PF−WZ does not isolate information because it changes both information and solution method. The square and the higher-order chain prevent these margins from being conflated.

6 Benchmark results: one path, different forecasts

The analytical model showed that information about a fixed path acts primarily through allocation, and that the sign of the full-information effect can change as a front-loaded episode unfolds. We now measure those forces in the quantitative model. The benchmark is the three-quarter Hormuz-closure scenario introduced in Section 3, labelled K3 below. Oil prices rise sharply over the first three quarters, peaking at 162 per cent above steady state, while oil production is 26 per cent below steady state; production is restored in the fourth quarter and the oil price then declines. The scenario is front-loaded in exactly the sense of Section 2: early in the episode the realised script lies above the continuation inferred from the current state, whereas around the peak the historical law of motion implies a more persistent continuation than the path that is actually realised.

All forecasts below are constructed on the path-identical convention of Section 5.2. Thus each forecaster experiences the same realised oil-price and oil-production paths; only the information structure, solution representation, or treatment of risk changes. We organise the results around three questions. How much does full information matter for the allocation of a given shock? How does that effect evolve as the script unfolds? And which part of the disagreement across forecasts is due to information, linearisation, or risk?

6.1 Similar output, different allocation

Table 4 reports the impact response under the four deterministic forecasts together with the higher-order risk-pricing forecast. The first fact is that the nonlinear forecasts imply very similar current output and very different private decisions. MIT, PF, and HO place the impact response of real GDP between 0.4 and 0.6 per cent, but they imply sharply different paths for consumption and investment.

This is the quantitative counterpart of the common-resource argument in Section 2. The current oil path is imposed by construction and the capital stock is predetermined, leaving relatively little scope for beliefs about future quarters to change impact output. Those beliefs

Table 4: **Impact responses on the benchmark path.** Quarter-one deviations from steady state, in per cent. All columns condition on the same realised oil-price and oil-production paths.

	WZ (linear)	MIT (surprised)	HO (risk pricing)	PF (informed)
Real GDP	+3.10	+0.48	+0.62	+0.41
Consumption	+2.02	−1.54	−1.50	+4.33
Investment	+9.09	+6.34	+7.02	+1.16
Non-oil GDP	+0.49	+0.31	+0.38	+0.14

can, however, change how current resources are used. Moving from sequential surprise to full information changes impact real GDP by only −0.07 percentage points, but raises consumption by 5.87 percentage points, lowers aggregate investment by 5.17 points, and lowers net exports by 4.4 percentage points of steady-state GDP.

The consumption and external-account responses follow directly from the first decision-date comparison in Section 2. In quarter one the full script dominates the continuation that a sequentially surprised agent infers from the current state. The informed economy therefore perceives more future resources, consumes more immediately, and finances part of that front-loading through the external account.

Investment reveals an additional general-equilibrium margin. At impact, full information raises non-tradable investment from roughly 7 to 64 per cent of its own steady state, while tradable investment moves from a 6 per cent increase under sequential surprise to a decline of about 70 per cent under full information. Expressed in units of steady-state aggregate investment, the two sectoral effects are approximately +30 and −36 percentage points, so the aggregate response is negative. The informed economy therefore does not simply invest less; it reallocates capital sharply across sectors.

The rotation reflects the different exposure of sectoral returns to the domestic boom. Tradable output prices are tied to world markets, while stronger domestic demand raises the non-tradable price and the economy-wide wage. The latter compresses the return to tradable capacity at the same time that higher domestic expenditure raises the return to non-tradable capacity. The wealth effect reinforces this movement: relative to the surprised economy, the informed household supplies about 2.9 percentage points less labour on impact, putting additional upward pressure on the wage. These equilibrium responses explain why the simple model is informative about the direction of the wealth and continuation effects, while the sectoral composition of aggregate investment requires the richer quantitative environment.

6.2 The information effect changes as the script unfolds

The full-information effect is a sequence, not a single number. Figure 3 reports $\mathcal{A}_t = PF - MIT$ at each decision date of the conditioning window. It is the quantitative counterpart of Figure 1 and Proposition 1. The correspondence with the analytical mechanism is shown directly in Appendix Figure **??**, which places the continuation beliefs of Section 2 above the resulting full-information effects in the quantitative model.

At the first decision date, the informed economy sees the increase that is still to come.

Consumption and non-tradable investment are therefore higher under full information, while the external account moves toward deficit. By the third quarter, the comparison has reversed. The oil price is at its peak and the sequential-surprise economy extrapolates that state with the estimated persistence of the oil process, whereas the informed economy already knows that the realised path is about to fall. Consumption is then 3.2 percentage points lower under full information and non-tradable investment about 74 points lower; in the fourth quarter the corresponding differences are roughly 5.8 and 58 points. The external account moves in the opposite direction, as the informed economy saves relative to the surprised one.

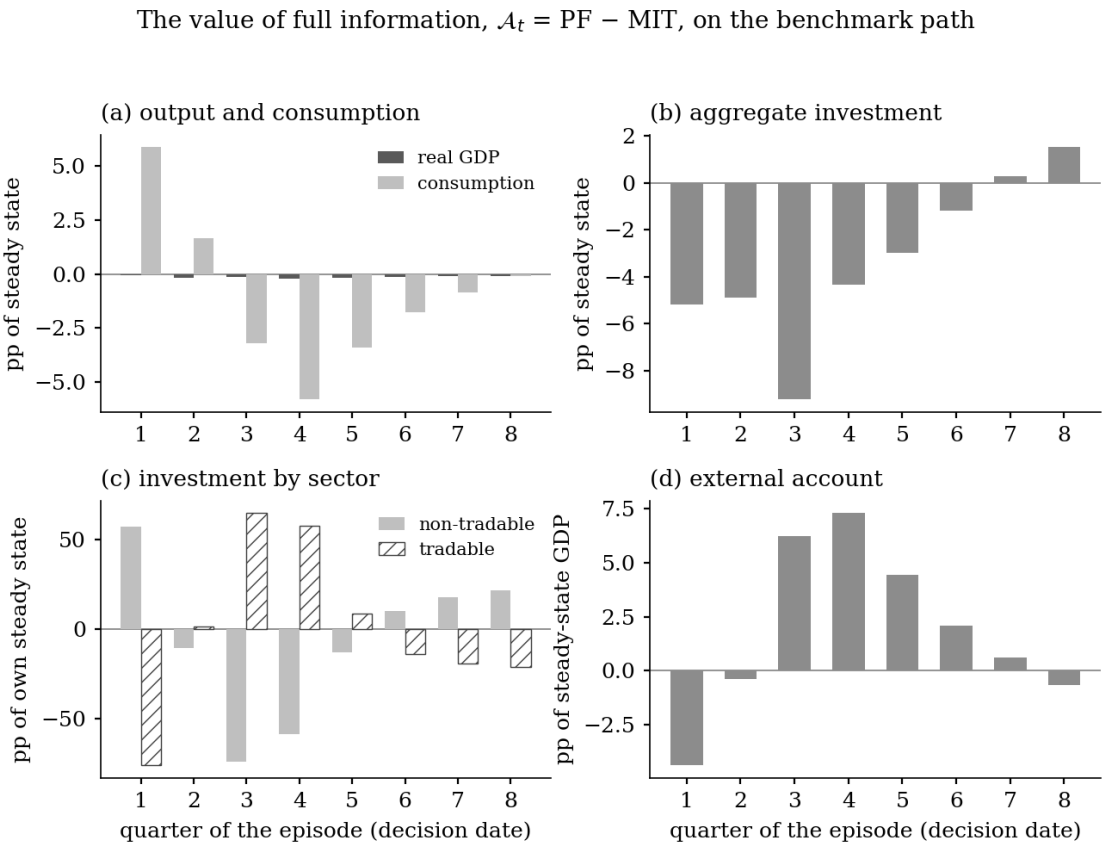


Figure 3: **The full-information effect along the benchmark path.** Bars report $\mathcal{A}_t = PF - MIT$ at each decision date: (a) real GDP and consumption, in percentage points of steady state; (b) aggregate investment; (c) investment by sector, in percentage points of each sector's own steady-state investment; and (d) net exports, in percentage points of steady-state GDP. The sign reversal in consumption and non-tradable investment mirrors the change in the ranking of the two continuations in Figure 1.

The figure also shows where the two-period mechanism reaches its limit. In the second quarter, the two continuations cross, so the dominance condition of Lemma 1 does not apply and the information effects are correspondingly small. Later in the episode, inherited capital matters. From the sixth quarter onward, non-tradable investment is again higher under full information because the sequential-surprise economy is running down capacity accumulated when it extrapolated the peak. A model without a capital stock cannot capture this post-boom adjustment. Appendix A compares the sign prediction of the analytical mechanism with the quantitative model decision date by decision date.

The cumulative effects are dominated by the peak and post-peak regime. Over the eight conditioning quarters, full information lowers cumulative real GDP by 1.1 percentage points, consumption by 7.6 points, and investment by 26.0 points relative to sequential surprise, while the external account is 15.2 percentage points of GDP higher. The sectoral decomposition is again important: cumulative non-tradable investment is 49 percentage points of its steady state lower under full information, whereas tradable investment is 2.4 points higher. The front-loaded consumption response is therefore reversed over the episode, while the surprised economy's extrapolation of the peak leaves a persistent investment footprint.

6.3 What drives the disagreement across forecasts?

The decomposition of Section 5 identifies which assumption is responsible for these differences. Table 5 reports two complementary comparisons. The left block decomposes the difference between the higher-order risk-pricing forecast and the linear WZ benchmark into the sequential-surprise linearisation gap $\mathcal{L}^S$, the third-order approximation diagnostic $\mathcal{T}$, and constant-volatility risk pricing $\mathcal{R}^{\text{pure}}$. The right block compares the full-information effect under exact and linear rules and reports their interaction $\mathcal{I}$.

Table 5: **Decomposition on the benchmark path.** The left block decomposes HO−WZ. The right block reports the nonlinear and linear full-information effects and their interaction. Entries are percentage points of steady state, except for variables expressed relative to steady-state GDP.

	HO−WZ	$\mathcal{L}^S$	$\mathcal{T}$	$\mathcal{R}^{\text{pure}}$	$\mathcal{A}$	$\mathcal{A}^{\text{lin}}$	$\mathcal{I}$
Impact (quarter 1)							
Real GDP	−2.49	−2.62	+0.01	+0.13	−0.07	+0.12	−0.18
Consumption	−3.51	−3.56	−0.01	+0.05	+5.87	+5.30	+0.57
Investment	−2.07	−2.76	+0.05	+0.63	−5.17	+2.83	−8.00
Non-oil GDP	−0.11	−0.18	+0.00	+0.06	−0.18	+0.20	−0.38
Cumulative, quarters 1–8							
Real GDP	−15.66	−16.03	+0.14	+0.23	−1.12	+0.81	−1.93
Consumption	−28.19	−26.87	+0.19	−1.51	−7.56	−11.95	+4.39
Investment	−30.42	−39.06	+0.66	+7.99	−25.99	+20.82	−46.81
Non-oil GDP	−0.48	−0.13	+0.03	−0.39	−2.26	+2.24	−4.50

The first result is that the benchmark disagreement with WZ is primarily a linearisation problem rather than a risk-pricing problem. WZ implies an impact GDP response of 3.10 per cent, compared with 0.48 per cent under the exact sequential-surprise solution. In the decomposition, $\mathcal{L}^S$ accounts for −2.62 percentage points of the −2.49 point HO–WZ impact wedge and −16.0 points of the −15.7 point cumulative wedge. This should be read as a linearisation gap in the conditional-forecast operator: away from the steady state, both the recovery of the innovations required to reproduce the fixed level path and the propagation of those innovations through the model are nonlinear. At the severity of the benchmark disruption, that combined gap is the dominant source of the aggregate forecast difference.

The second result is that nonlinearity changes the economic value assigned to information, not merely the level of the forecast. For consumption, the linear and nonlinear full-information effects are similar on impact: +5.30 and +5.87 percentage points. For investment, however, the

sign reverses. Under linear rules, full information raises impact investment by 2.83 points and cumulative investment by 20.8 points. Under the exact nonlinear solution, it lowers investment by 5.17 points on impact and by 26.0 points cumulatively. The information–nonlinearity interaction is therefore −8.0 points on impact and −46.8 points cumulatively. For real and non-oil GDP, the information effects are much smaller, but the linear and nonlinear measures again differ in sign over the eight-quarter window. The interaction is therefore not a small correction to a common information effect; for some margins it changes the economic conclusion.

The source of the investment reversal is the state dependence suppressed by the first-order approximation. The linear model evaluates marginal responses around the steady state. In the exact model, a large anticipated boom moves the economy into a region where installation costs, wages, relative prices, and sectoral returns change with the state. The informed economy consequently rotates strongly toward non-tradable capacity and away from tradable capacity, and the latter response dominates aggregate investment. Section 7 shows that this interaction itself changes systematically with the severity of the scenario.

The third result is that constant-volatility risk pricing is quantitatively different from these two margins. For aggregate GDP, the cumulative risk-pricing term is only 0.23 percentage points, of the same order as the third-order approximation diagnostic; we therefore treat its sign as below the resolution of the method. For investment, in contrast, the cumulative risk-pricing effect is about 8.0 percentage points, far larger than the corresponding diagnostic. Risk is therefore visible in the allocation of the windfall even when its effect on aggregate output cannot be resolved reliably. Section 8 develops this margin and then allows volatility itself to vary along the scenario.

Benchmark takeaway. The benchmark establishes three facts that the next sections generalise. First, information about the continuation has little effect on current output but large effects on consumption, investment, sectoral allocation, and the external account. Second, the sign of that information effect changes as a front-loaded path moves from build-up to peak and reversal. Third, the disagreement across conditional forecasts is economically interpretable: linearisation dominates the aggregate WZ discrepancy at the benchmark, nonlinearity can reverse the value assigned to information, and risk pricing mainly affects allocation. The next section asks how these margins change as the scenario moves the economy through different regions of the state space.

7 State dependence across scenarios

The benchmark established that information about the continuation acts on allocation rather than on current output, that its sign changes as a front-loaded path unfolds, and that the linear and exact solutions assign different values to the same information. Section 5 also supplied a local benchmark for that last fact: near the steady state the information–nonlinearity interaction is a second-order object, $\mathcal{I}_t(\lambda) = O(\lambda^2)$ in equation (21), so the linear and exact measures of the value of information coincide for small scenarios and can diverge only as the scenario moves the economy away from the expansion point. This section varies the realised

path across seven scenarios of increasing severity and traces that divergence. The unifying result across scenarios is as follows. The exact value of full information moves monotonically with the severity of the disruption; the linear value does not. Between the two lie four regimes: local agreement in mild scenarios, attenuation and a "wrong" sign in short partial closures, the sign reversal of the benchmark, and amplification by an order of magnitude in the most severe scenarios.

7.1 Seven realised paths

Table 6 and Figure 4 define the scenarios. Three are partial closures of the Strait of Hormuz lasting one, two, and three quarters (K1–K3). Their oil-price paths are those of Kilian et al. (2026), who derive them from a global oil-market model: a first-quarter price of $103 per barrel, rising to $132 and $167 if the closure lasts into the second and third quarters, then falling back towards the steady state of $63.78. Oil exporter production is 26 per cent below steady state while the closure lasts, the Strait-dependent share of production (the remainder is exported through Yanbu on the Red Sea), and is restored the quarter after it ends. K3 is the benchmark of Section 6. Two scenarios are mild. A1 is a rapid de-escalation: a 15 per cent price increase that vanishes within three quarters and no change in production, mirroring the June 2025 episode described in Section 3. A2 is a strike on the Kharg Island loading terminal: the same initial price increase decaying over eight quarters, with oil exporter production 2 to 5 per cent above steady state as the kingdom substitutes for lost Iranian exports. Two scenarios are severe. A4 is a Gulf-wide escalation in which energy infrastructure is targeted directly: the price doubles at impact and remains elevated for two years, and production falls by 40 per cent before recovering at the pace of infrastructure repair. A3 is the full closure: the K3 price path combined with a 63 per cent production cut for three quarters, the case in which the Yanbu route is also lost. A3 and K3 therefore share the realised price path exactly and differ only in the depth of the production cut.

Table 6: **The seven scenarios.** Conditioning paths for the oil price and Saudi oil production, quarters 1–8, as deviations from the model steady state ($63.78 per barrel; 9.5 million barrels per day). The last column sums the production shortfall over the window in percentage points of steady-state production. Scenarios are listed in the order used throughout the section.

	scenario	peak price (% above ss)	production (% of ss, quarters)	cumulative production loss
A1	rapid de-escalation	+15	unchanged	0
A2	Kharg Island strike	+15	+2 to +5, q1–q7	−20
K1	one-quarter closure	+61	−26, q1	26
K2	two-quarter closure	+107	−26, q1–q2	53
K3	three-quarter closure (benchmark)	+162	−26, q1–q3	79
A4	Gulf escalation	+100	−40 recovering to −2, q1–q8	145
A3	full closure, Yanbu cut	+162	−63, q1–q3	189

The scenarios are ordered by the size of the production shortfall, which is also the order of the price displacement within each family. A4 and K3 are comparable in overall distance from the steady state; A4 is listed after K3 because its volume loss is nearly twice as large and persists through the window. Every scenario is conditioned on the path-identical convention

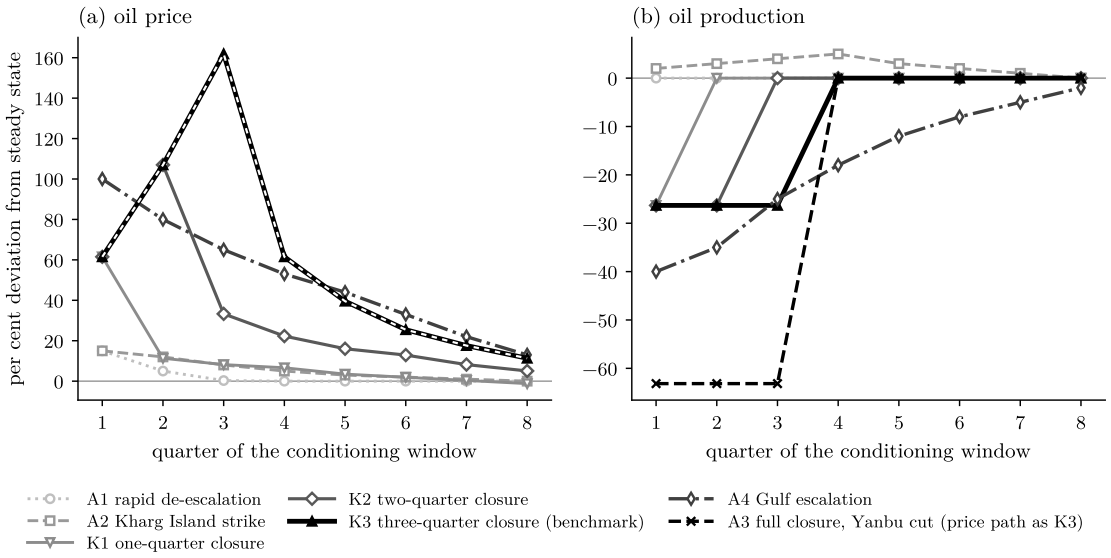


Figure 4: **The conditioning paths.** Oil price (a) and oil production (b), per cent deviation from steady state over the eight-quarter conditioning window, for the seven scenarios of Table 6. A3 shares the K3 price path (white dashes over the K3 line in panel (a)) and deepens the production cut. Source: `define_scenarios.m`; K-family price paths from Kilian et al. (2026), Table 1.

of Section 5.2, and the four deterministic forecasts and the risk-pricing forecast are computed for each.

7.2 Agreement on output near the steady state, not away from it

Table 7 reports real GDP under the linear sequential-surprise forecast WZ, the exact sequential-surprise forecast MIT, the full-information forecast PF, and the risk-pricing forecast HO, at impact and cumulatively. Three facts organise the table.

Table 7: **Real GDP across scenarios and forecasts.** Impact (quarter-one per cent deviation from steady state) and cumulative (sum over quarters 1–8, percentage points) responses under the four forecasts that retain the realised oil path. WZ and MIT differ by the solution method under sequential surprise; PF gives agents full-path information; HO prices risk at constant volatility under sequential surprise.

	impact				cumulative, q1–8			
	WZ	MIT	PF	HO	WZ	MIT	PF	HO
A1	+1.88	+1.87	+1.88	+1.87	+2.8	+2.8	+3.4	+2.8
A2	+2.23	+2.26	+2.30	+2.25	+9.8	+10.0	+11.1	+9.9
K1	+3.10	+0.48	+0.41	+0.62	+7.9	+5.0	+5.0	+5.1
K2	+3.10	+0.48	+0.43	+0.62	+26.5	+18.4	+17.9	+18.7
K3	+3.10	+0.48	+0.41	+0.62	+50.9	+34.8	+33.7	+35.2
A4	+5.53	−0.88	−1.14	−0.64	+28.9	+10.8	+7.8	+11.5
A3	−3.35	−9.43	−10.18	−9.07	+30.8	−5.0	−13.4	−3.7

First, in the mild scenarios the four forecasts agree. Under A1 the impact responses lie within 0.01 percentage points of one another and the cumulative responses within 0.6 points,

the latter being the value of full information itself. The linear forecast is adequate here, and so is the assumption that information about the continuation can be ignored for output.

Second, the three partial closures are indistinguishable at impact to every sequentially surprised forecaster. At the first decision date WZ, MIT, and HO observe the same state, a price of \$103 and a 26 per cent production cut, and cannot know how long the closure will last; their impact responses are therefore identical across K1, K2, and K3, at 3.10, 0.48, and 0.62 per cent. Only the informed forecast distinguishes the three, and it does so barely, at 0.41 to 0.43 per cent, which is the output-pinning fact of Section 6.1 once more. What the duration of the closure changes is the cumulative response, and there the linear overstatement grows with the closure: 7.9 against 5.0 points for K1, 26.5 against 18.4 for K2, and 50.9 against 34.8 for K3.

Third, in the severe scenarios the linear forecast is wrong in sign or in magnitude. Under A4 it reports an impact expansion of 5.5 per cent where every exact forecast reports a contraction, of 0.6 to 1.1 per cent. Under A3 it reports a contraction of 3.3 per cent where the exact forecasts report 9.1 to 10.2 per cent; cumulatively it reports an expansion of 31 points where the exact forecasts report contractions of 4 to 13 points. A reader of linear conditional forecasts is told that the partial closures are more expansionary than they are and that the full closure is far milder than it is. The error is not conservative in either direction. All of this is the sequential-surprise linearisation gap $\mathcal{L}^S$ of Section 6.3 traced across the ladder; the remainder of the section concerns the value of information and how the linear model measures it.

7.3 The value of full information along the ladder

Table 8 reports, for real GDP, consumption, and aggregate investment, the full-information effect under exact rules, $\mathcal{A} = \text{PF} - \text{MIT}$, and under linear rules, $\mathcal{A}^{\text{lin}} = \text{PF}_{\text{lin}} - \text{WZ}$, together with their difference, the interaction $\mathcal{I} = \mathcal{A} - \mathcal{A}^{\text{lin}}$ of equation (19), and the ratio $|\mathcal{A}|/|\mathcal{A}^{\text{lin}}|$ of the cumulative effects. Figure 5 draws the cumulative effects for GDP and investment. The ratio is marked where the two measures have opposite signs, because a ratio of magnitudes is then a statement about a sign error rather than about scale.

Local agreement. In the mild scenarios the linear and exact measures coincide to within three per cent for every variable: cumulative GDP +0.57 against +0.56 points under A1 and +1.10 against +1.09 under A2; investment +12.9 against +12.5 and +24.7 against +24.5. The interaction never exceeds 0.34 points. This is equation (21) realised: at these displacements the second-order term is negligible, and the agreement doubles as an internal check that the four cells nest one another where they should. Two remarks follow. The value of information is far from negligible even here, since cumulative investment is 12 to 25 points higher when the path is known, but the linear model measures it correctly. And in both mild scenarios full information raises investment and, at impact, lowers consumption, the opposite of the benchmark pattern; the sign of the information effect is a property of the path, as Lemma 1 implies, not of the information structure.

Table 8: **The value of full information under exact and linear rules.** $\mathcal{A}$ = PF − MIT and $\mathcal{A}^{\text{lin}}$ = PF_{lin}− WZ at impact (quarter-one per cent deviation) and cumulatively (quarters 1–8, percentage points); the interaction $\mathcal{I} = \mathcal{A} - \mathcal{A}^{\text{lin}}$ and the ratio $|\mathcal{A}|/|\mathcal{A}^{\text{lin}}|$ refer to the cumulative effects. †: the two cumulative measures have opposite signs.

	impact		cumulative, q1–8			
	$\mathcal{A}^{\text{lin}}$	$\mathcal{A}$	$\mathcal{A}^{\text{lin}}$	$\mathcal{A}$	$\mathcal{I}$	ratio
Real GDP						
A1	+0.01	+0.01	+0.57	+0.56	−0.01	0.99
A2	+0.04	+0.04	+1.10	+1.09	−0.01	0.99
K1	−0.03	−0.07	+0.74	−0.03	−0.77	0.04†
K2	+0.08	−0.05	+0.99	−0.50	−1.49	0.51†
K3	+0.12	−0.07	+0.81	−1.12	−1.93	1.38†
A4	−0.06	−0.26	−0.31	−3.03	−2.71	9.66
A3	−0.10	−0.75	−3.43	−8.38	−4.95	2.44
Consumption						
A1	−1.58	−1.53	+0.95	+0.93	−0.02	0.98
A2	−1.23	−1.20	−2.82	−2.86	−0.04	1.02
K1	−0.49	+2.15	+0.16	−1.28	−1.44	8.29†
K2	+2.37	+4.01	−3.00	−3.37	−0.37	1.12
K3	+5.30	+5.87	−11.95	−7.56	+4.39	0.63
A4	+0.64	+5.33	+1.47	+2.12	+0.65	1.45
A3	+11.55	+13.47	−8.44	−3.27	+5.17	0.39
Aggregate investment						
A1	+1.71	+1.59	+12.87	+12.54	−0.34	0.97
A2	+3.49	+3.43	+24.73	+24.49	−0.25	0.99
K1	−1.66	−4.08	+18.93	+2.14	−16.79	0.11
K2	+2.71	−3.52	+25.24	−8.05	−33.29	0.32†
K3	+2.83	−5.17	+20.82	−25.99	−46.81	1.25†
A4	−3.98	−13.54	−3.24	−62.36	−59.12	19.26
A3	−15.31	−36.76	−69.23	−176.60	−107.38	2.55

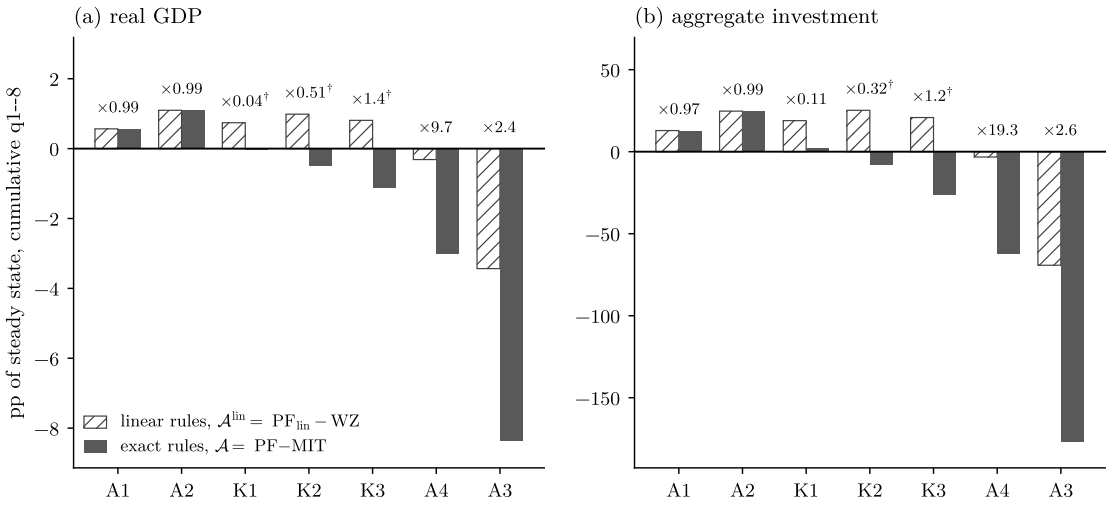


Figure 5: **The value of full information under linear and exact rules across the ladder.** Cumulative effects over quarters 1–8 for real GDP (a) and aggregate investment (b); hatched bars are the linear measure $\mathcal{A}^{\text{lin}}$, filled bars the exact measure $\mathcal{A}$. The annotation is the ratio $|\mathcal{A}|/|\mathcal{A}^{\text{lin}}|$, marked † where the two measures have opposite signs. The exact measure moves monotonically down the ladder; the linear measure does not. Source: `HO_stage3d_results.mat`.

Attenuation and a wrong sign. In the short partial closures the linear model overstates the value of information for investment and mis-signs it for output. Under linear rules full information raises cumulative investment by 18.9 points in K1 and 25.2 points in K2; under exact rules it raises it by 2.1 points and lowers it by 8.1 points. For GDP the linear model reports that knowing the path adds 0.7 to 1.0 cumulative points; the exact model reports that it subtracts 0.03 to 0.5. At impact the discrepancy is already visible in consumption: the linear informed economy in K1 consumes 0.5 per cent less than the surprised one, the exact informed economy 2.2 per cent more. The economics is that of Section 6.3. Linear anticipation front-loads a windfall at zero curvature cost; the exact informed economy runs into installation costs, wage increases, and the rotation between sectoral returns, and discounts a short boom accordingly.

The sign reversal at the benchmark. K3 is the case documented in Section 6.3: cumulative investment +20.8 points under linear rules against −26.0 under exact rules, GDP +0.8 against −1.1. The ratios of 1.25 and 1.38 compare quantities of opposite sign and should not be read as amplification. The three closures together display a regularity that the benchmark alone cannot show. As the closure lengthens from one to three quarters, the exact value of information moves monotonically, from +2.1 to −8.1 to −26.0 points for investment and from −0.03 to −0.50 to −1.12 for GDP, and the impact consumption effect rises from 2.2 to 4.0 to 5.9 per cent as the informed economy sees a longer windfall ahead, exactly as the ordering of continuations in Lemma 1 predicts. The linear value of information barely responds to the duration of the closure: 18.9, 25.2, and 20.8 points for investment, 0.74, 0.99, and 0.81 for GDP. The linear model assigns a nearly duration-invariant value to knowing the path; the

exact model assigns a value that scales with what there is to know.

Amplification. In the severe scenarios the two measures agree in sign and the exact one is much larger. For GDP the exact value of information is 2.4 times the linear value under A3 (-8.4 against -3.4 cumulative points) and 9.7 times under A4 (-3.0 against -0.3); for investment the factors are 2.6 under A3 (-177 against -69) and 19.3 under A4 (-62 against -3). The very large A4 factors have a specific origin that the signed values make plain: the linear model assigns almost no value to information for output and investment in A4, because the linear measure crosses zero between K3 and A4, while the exact measure continues to grow. Amplification is a statement about output and investment. For consumption the exact measure is larger than the linear one at impact in both severe scenarios (5.3 against 0.6 per cent in A4; 13.5 against 11.6 in A3) but smaller cumulatively in A3 (-3.3 against -8.4 points), so the consumption effect is attenuated over the window even where the investment effect is amplified.

What full information does in the severe scenarios is discipline investment. Under sequential surprise the economy keeps investing through the disruption: aggregate investment is 7.8 per cent above steady state at impact in A4 and 3.6 per cent in A3, and 60.6 and 27.9 points cumulatively. Under full information it falls by 5.7 and 33.2 per cent at impact and by 1.8 and 148.7 points cumulatively. The surprised economy extrapolates each observed state, a high price and a deep production cut, with the long half-lives of the two oil processes; the informed economy knows that the price will fall and, in A3, that production will be restored in the fourth quarter. It therefore perceives a very different continuation, consumes against it (consumption at impact is 13.5 per cent higher under full information in A3, -4.0 against -17.5 per cent), and withdraws the capacity that the surprised economy is building against a state that the path will not sustain. This is the intuition of the introduction: sequentially surprised agents invest against a continuation that full-information agents know will disappear.

7.4 Why the interaction grows: two linearisation gaps

The identity $\mathcal{I} = \mathcal{L}^F - \mathcal{L}^S$ of equation (19) locates the interaction in the difference between the linearisation gap under full information and the gap under sequential surprise. Table 9 reports both. In every disruption scenario the full-information gap is larger, and for investment it is between two and three times larger: -24 against -7 points in K1, -86 against -39 in K3, -107 against -48 in A4, and -203 against -95 in A3. For GDP the excess is 10 to 30 per cent of the sequential-surprise gap. The interaction is therefore of the same order as the linearisation gap itself for investment, which is why it changes the economic conclusion rather than correcting it at the margin.

The asymmetry has a simple source, made precise in Appendix D. The two economies are evaluated at different points of the same curved policy functions. The sequentially surprised economy responds at each date to the current innovation and to the deterministic continuation of the observed state; its decisions are, in every quarter, a response to one surprise. The informed economy responds at the first date to the entire anticipated sequence and undertakes the large intertemporal reallocation that the sequence warrants. It is the economy that makes

Table 9: **Linearisation gaps under the two information structures.** Cumulative effects over quarters 1–8, percentage points of steady state. $\mathcal{L}^S$ = MIT − WZ and $\mathcal{L}^F$ = PF − PF_{lin}; their difference is the interaction $\mathcal{I}$ of Table 8.

	real GDP		aggregate investment	
	$\mathcal{L}^S$	$\mathcal{L}^F$	$\mathcal{L}^S$	$\mathcal{L}^F$
A1	−0.01	−0.02	+0.1	−0.2
A2	+0.22	+0.21	+1.0	+0.8
K1	−2.88	−3.65	−7.3	−24.1
K2	−8.03	−9.51	−19.8	−53.1
K3	−16.03	−17.96	−39.1	−85.9
A4	−18.06	−20.77	−47.6	−106.7
A3	−35.78	−40.72	−95.5	−202.8

the large move, and the large move is where linearisation fails: installation costs, the wage response, and the rotation of sectoral returns all change with the size of the reallocation. Information and curvature are complements in this precise sense. The value the linear model assigns to information is computed for an economy that never moves far from the steady state, and it is reliable only where the informed economy does not either.

Takeaway. The ladder gives the second result of the paper its shape. The exact value of full information grows monotonically with the severity of the disruption, from a fraction of a point of cumulative GDP in the short closures to 8 points in the full closure, and from 2 points of cumulative investment to 177; the linear value does not track it. Close to the steady state the two agree and the linear conditional forecast can be used with confidence, for the level of the forecast and for the value of information alike. In the partial closures the linear model overstates the value of information for investment and mis-signs it for output; at the benchmark it reverses the sign for both; in the severe scenarios it understates it by a factor of two and a half (A3) to an order of magnitude (A4). There is no single correction to apply, because the correction is itself a function of where the scenario places the economy. The uncertainty margin studied next is state dependent in the same sense, though not in the same direction: Section 8 shows that the volatility channel is larger at the benchmark than in the full closure, because uncertainty postpones investment only where investment is still being considered.

8 Risk and uncertainty

The deterministic experiments hold one thing constant that the 2026 episode did not: the confidence agents attach to the continuation. Section 3 documented that oil implied volatility rose from a pre-war level near 47 to a peak above 120 while the price path itself was being revised. Section 5.4 separated the corresponding margin into two objects, both computed under sequential surprise so that neither has information-structure content. The first, $\mathcal{R}^{\text{pure}}$ = HO − HO_{CE}, asks what changes when agents price the risk around the continuation at constant volatility. The second, $\mathcal{R}^{\text{tv}} = \text{HO}_{\text{SV}}(\text{spike}) - \text{HO}_{\text{SV}}(\text{flat})$, asks what changes when the volatility of the oil process is itself allowed to follow the path observed during the conflict, holding the

realised oil-price and oil-production paths fixed to machine precision. This section reports both, assembles them with the deterministic margins into a single accounting for the benchmark and the full closure, and shows that the uncertainty margin is state dependent in a direction opposite to the information margin of Section 7.

8.1 Pricing risk at constant volatility

Table 10 reports the third-order approximation diagnostic $\mathcal{T}$ and the risk-pricing term $\mathcal{R}^{\text{pure}}$, cumulatively over the window, for real GDP and aggregate investment across the seven scenarios of Section 7.1. Two features stand out. For GDP the risk-pricing term is small and, at the benchmark, unresolved: it amounts to 0.23 cumulative points, against an approximation diagnostic of 0.14, so that the ratio $|\mathcal{T}|/|\mathcal{R}^{\text{pure}}|$ reaches 0.61. We treat this ratio as the resolution limit of the method and read the GDP entries of the K scenarios as a pattern rather than as point estimates. The pattern follows the severity ordering, and in the severe scenarios the term is resolved: under A3 it is 1.35 points against a diagnostic of 0.03. For investment the term is an order of magnitude larger than the diagnostic in every disruption scenario, rising from 1.8 points in K1 to 8.0 at the benchmark and 8.3 in the full closure. Risk pricing is therefore visible in the allocation of the windfall even where its effect on output cannot be resolved reliably, which is the third fact of Section 6.3 traced across the ladder.

Table 10: **Approximation diagnostic and risk pricing across the ladder.** $\mathcal{T} = \text{HO}_{\text{CE}} - \text{MIT}$ and $\mathcal{R}^{\text{pure}} = \text{HO} - \text{HO}_{\text{CE}}$, cumulative over quarters 1–8, percentage points of steady state. The last column reports the resolution ratio $|\mathcal{T}|/|\mathcal{R}^{\text{pure}}|$ for real GDP; 0.61 at the benchmark is the binding verdict of Appendix **??**.

	real GDP		aggregate investment		
	$\mathcal{T}$	$\mathcal{R}^{\text{pure}}$	$\mathcal{T}$	$\mathcal{R}^{\text{pure}}$	$\lvert\mathcal{T}\rvert/\lvert\mathcal{R}^{\text{pure}}\rvert$, GDP
A1	−0.00	−0.01	−0.01	+0.37	–
A2	−0.00	−0.13	−0.04	+1.03	0.02
K1	+0.01	+0.13	+0.06	+1.77	0.07
K2	+0.05	+0.19	+0.23	+4.79	0.25
K3	+0.14	+0.23	+0.66	+7.99	0.61
A4	+0.07	+0.65	+0.45	+6.96	0.11
A3	−0.03	+1.35	+0.24	+8.30	0.02

What risk pricing does to the allocation is to build a buffer in capacity. At the benchmark it raises cumulative investment by 8.0 points and lowers cumulative consumption by 1.5 points; the sectoral harvest of the same experiment attributes the investment response almost entirely to the non-tradable sector, roughly 13 points of its own steady state against under one point for tradable investment. The composition of the term explains why it is state dependent. Appendix D splits $\mathcal{R}^{\text{pure}}$ into the constant correction that shifts the stochastic steady state, the state-dependent correction that changes with the position of the economy, and the innovation-dependent correction. For investment at the benchmark the three contribute +2.9, +5.4 and −0.3 points; in the full closure +0.3, +7.8 and +0.2. Pricing risk at constant volatility is, at conflict severity, mostly a state-dependent object, and its growth down the ladder is the growth of that term.

8.2 The conditioned volatility path

The second margin conditions on the volatility path itself. The volatility-of-volatility parameter is $\sigma_\sigma = 0.27$, set to match the dispersion of oil implied volatility at the persistence $\rho_\sigma = 0.64$, which survives an audit against realised-volatility and OVX-based estimates (Appendix E).[11] The conditioned path is the 2026 episode: during the three closure quarters the OVX stood at 1.88, 1.98 and 1.20 times its pre-war level, and we impose exactly those ratios on the conditional standard deviation of the oil price, letting it decay freely thereafter under its own law (1.12, 1.08, 1.05 times steady state in quarters four to six). Volatility roughly doubles for two quarters and then subsides.

Mechanically, the same price path now requires smaller price innovations where volatility is high, together with an explicit sequence of volatility innovations; both are recovered analytically by the inversion of Appendix C, respecting the timing of equation (15) under which the volatility state dated $t-1$ scales the innovation at t. The price innovation of quarter one is unchanged, those of quarters two and three roughly halve, and the volatility innovations are 2.3, 1.0 and -0.9 standard deviations in the first three quarters and zero thereafter. A consistency gate protects the exercise: with a flat volatility path the stochastic-volatility solution reproduces the constant-volatility results of Section 6 to 6×10^{-14}, so everything below is attributable to the volatility path alone.

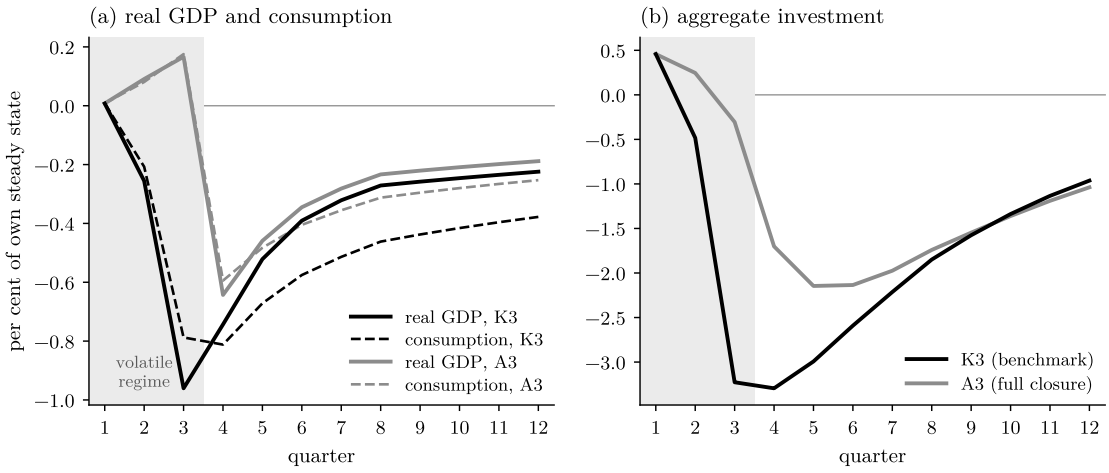


Figure 6: **The effect of the conditioned volatility path.** $\mathcal{R}^{\text{tv}} = \text{HO}_{\text{SV}}(\text{spike}) - \text{HO}_{\text{SV}}(\text{flat})$ for real GDP and consumption (a) and aggregate investment (b), in per cent of each variable's own steady state, for the benchmark (black) and the full closure (grey). The shaded band marks the quarters in which conditional volatility is elevated. The realised oil-price and oil-production paths are identical across the two forecasts being compared. Source: `HO_stage3e_results.mat`.

Figure 6 shows the response on the benchmark path. At impact almost nothing happens: real GDP moves by 0.01 percentage points and investment by 0.5, because the volatility state dated $t-1$ scales the next quarter's surprise and because output is pinned by the conditioned path regardless. The retrenchment then builds while the volatile regime lasts and peaks in the

[11] A sensitivity run at $\sigma_\sigma = 0.40$, the value implied by realised rather than implied dispersion, is not yet part of the results; the risk-pricing component of $\mathcal{R}^{\text{tv}}$ scales approximately with σ_σ^2.

third and fourth quarters, at −1.0 point of GDP and −3.3 points of investment. Cumulatively over the window the volatility path subtracts 3.5 points of GDP, 4.0 of consumption and 16.2 of investment relative to the same realised oil path under calm, and the resources that are not committed leave the country: the external account is 10.2 percentage points of steady-state GDP higher over the eight quarters, with the surplus building from the third quarter on. The behavioural reading is precautionary retrenchment. With not a dollar of realised oil income changed, agents who attribute the same windfall to a more volatile process postpone capacity, hold back consumption, and park the difference abroad (Bernanke, 1983; Bloom, 2009). Uncertainty here taxes commitments, not income.

8.3 One accounting for the benchmark and the full closure

Adding $\mathcal{R}^{\text{tv}}$ to the chain of equation (23) completes the accounting for the forecast that prices both risk and the episode's volatility. Table 11 reports the four terms for the benchmark and for the full closure. The sum is exact: the bridge between the stochastic-volatility solution under a flat path and the constant-volatility solution is below 10^{-13}, as the consistency gate implies.

Table 11: **The four-way decomposition, cumulative over quarters 1–8.** The wedge between the stochastic-volatility risk-pricing forecast on the conditioned volatility path and the linear benchmark, $\text{HO}_{\text{SV}}(\text{spike}) - \text{WZ}$, decomposed into the sequential-surprise linearisation gap $\mathcal{L}^S$, the approximation diagnostic $\mathcal{T}$, constant-volatility risk pricing $\mathcal{R}^{\text{pure}}$, and the volatility path $\mathcal{R}^{\text{tv}}$. Percentage points of steady state; the terms add to the total to machine precision.

		$\text{HO}_{\text{SV}}-\text{WZ}$	$\mathcal{L}^S$	$\mathcal{T}$	$\mathcal{R}^{\text{pure}}$	$\mathcal{R}^{\text{tv}}$
K3	real GDP	−19.1	−16.0	+0.1	+0.2	−3.5
	consumption	−32.2	−26.9	+0.2	−1.5	−4.0
	investment	−46.6	−39.1	+0.7	+8.0	−16.2
A3	real GDP	−36.2	−35.8	−0.0	+1.3	−1.7
	consumption	−57.7	−57.5	−0.2	+1.8	−1.9
	investment	−96.2	−95.5	+0.2	+8.3	−9.3

The ordering of channels is itself a finding. Nonlinearity comes first: the linearisation gap accounts for 16 of the 19 cumulative GDP points by which the fully specified forecast departs from the linear benchmark, and for 39 of 47 investment points. The volatility path comes second, at 3.5 GDP points and 16 investment points. Constant-volatility risk pricing is a distant third for GDP and enters investment with the opposite sign, so that the two risk margins partly offset: pricing risk at constant volatility builds capacity, pricing a volatility spike postpones it. The linear certainty-equivalent forecast misses all three, but what it misses after curvature is mostly the volatility path, not the level of risk.

8.4 Uncertainty is state dependent, in the opposite direction

Section 7 showed that the value of information grows monotonically with the severity of the disruption. The volatility margin does not. On the same conditioned volatility path, the

cumulative effect on GDP is −3.5 points at the benchmark and −1.7 in the full closure; for investment, −16.2 against −9.3; for consumption, −4.0 against −1.9. Figure 6 shows where the difference arises. In the full closure the volatility spike has no contractionary effect while production is cut: GDP and consumption are marginally higher in the second and third quarters and investment is essentially unchanged, and the retrenchment sets in only in the fourth quarter, when production is restored and the economy would otherwise begin to rebuild.

The economics is that of an option that is only valuable when there is something to postpone. Uncertainty makes waiting valuable (Bernanke, 1983; Bloom, 2009), so its effect is largest where the economy is in the middle of a sizeable investment programme. At the benchmark the sequentially surprised economy is building capacity against an extrapolated windfall, with cumulative investment 122 points above steady state, and a rise in uncertainty postpones a substantial part of that programme. In the full closure the realised disruption has already cancelled most of it, cumulative investment under the same forecaster being 28 points, and there is correspondingly less for additional uncertainty to postpone. The marginal effect of uncertainty is therefore smaller precisely when the underlying shock is larger. The two risk margins thus move in opposite directions down the ladder. Constant-volatility risk pricing grows with severity, because its state-dependent component grows with the displacement of the economy; the volatility path bites through the option value of the investment that the scenario leaves standing, and that stock shrinks as the disruption deepens.

8.5 What unconditional risk is made of

A third-order economy drifts from its deterministic toward its stochastic steady state even with no innovations at all, here by 3.9 per cent of GDP after twenty quarters, with a deep transition dip in consumption, and it is tempting to read that drift as the model pricing oil risk. It is not. Switching off each shock variance in turn and re-solving attributes essentially the entire drift to sectoral productivity, 67 per cent to non-tradable and 34 per cent to tradable productivity, and none of it to the oil-price variance, whose contribution is slightly negative. The unconditional risk buffer of this economy is a business-cycle object, not an oil object, which is why the conditioned volatility path, and not the unconditional drift, is the correct measurement of the war's uncertainty channel, and why every risk-pricing object in this paper is measured net of its own zero-innovation counterfactual (Appendix D). Whether announcement reversals in particular move implied volatility without moving the price level, the empirical fingerprint of this section's experiment, is a pre-registered test in the companion empirical design described in Appendix E; it is not run here.

Takeaway. The uncertainty margin completes the accounting. A macroeconomic scenario for this economy is characterised by a realised path, by the information agents have about its continuation, and by the risk they attach to it. The third element is not a residual: on the benchmark path the volatility of the episode alone removes 3.5 points of cumulative GDP and 16 points of investment, second only to nonlinear adjustment, and it does so in a state-dependent way that no unconditional risk premium would reproduce.

9 Concluding remarks

Macroeconomic scenarios are usually described by paths, but a realised path is not a sufficient statistic for a forward-looking economy. The same trajectory can generate different decisions depending on what agents know about its continuation and on the uncertainty they attach to it. This paper therefore treats the information structure as part of the economic specification of a scenario. The analytical model isolates the mechanism in its simplest form. When current resources are predetermined, information about the continuation does not change the resources available today; it changes how those resources are allocated between consumption, investment, and foreign assets. Along a front-loaded path, the direction of this effect changes with the decision date. Before the peak, full information may reveal an expansion that is not yet visible in the current state. At the peak, the same information reveals the coming reversal, while a sequentially surprised economy continues to extrapolate the boom.

We quantify this mechanism in a nonlinear DSGE model of an oil-exporting economy during the 2026 conflict. Throughout the deterministic experiments, the realised oil-price and oil-production paths are held fixed while the information structure and solution representation are varied. This produces a path-identical comparison between sequential surprise and full-path information under both linear and exact nonlinear decision rules. A separate higher-order branch then introduces priced risk and time-varying volatility. Because the conditioning variables belong to an invertible exogenous oil block, the innovations required to reproduce any prescribed path can be recovered directly under each solution representation. The resulting decomposition therefore attributes differences across forecasts to distinct economic margins: information, linearisation, their interaction, and risk.

The results show that these margins matter in different ways and in different regions of the state space. First, information about the continuation is largely an allocation margin rather than an impact-output margin. In the benchmark three-quarter closure, moving from sequential surprise to full-path information changes impact GDP by only -0.07 percentage points, but raises consumption by 5.9 points, lowers aggregate investment by 5.2 points, and reduces the external balance by 4.4 percentage points of GDP. The signs then change as the episode unfolds: by the peak, the surprised economy extrapolates a continuation that the informed economy knows will not materialise. Over the eight-quarter conditioning window, the informed economy consumes 7.6 percentage points less, invests 26 points less, and runs an external balance about 15 percentage points of GDP higher than its sequentially surprised counterpart. The aggregate investment response also conceals a large sectoral rotation between tradable and non-tradable capacity.

Second, the macroeconomic value assigned to information is itself state-dependent. Near the steady state, the linear and nonlinear measures are close, as the local approximation implies. As the disruption becomes larger, however, the two separate. The linear conditional forecast can attenuate the information effect, assign it the wrong sign, or substantially understate it. Across the scenario ladder, the nonlinear value of full information reaches roughly ten times its linear counterpart for GDP and nineteen times for investment in the most severe case. The interaction between information and nonlinearity is therefore not a small correction applied

after the fact. It can determine the economic interpretation of the scenario itself.

Third, uncertainty operates through a separate margin. At constant volatility, risk pricing has no reliably resolved effect on aggregate GDP in the benchmark scenario, but it materially changes the allocation of expenditure. When the volatility path observed during the conflict is imposed while the realised oil path is kept unchanged, the effects become much larger: cumulative GDP falls by about 3.5 percentage points, investment by about 16 points, and the external balance rises by about 10 percentage points of GDP relative to the same realised path under flat volatility. This effect is also state-dependent. Uncertainty has the greatest influence when agents are still considering costly commitments that can be postponed; once a sufficiently severe realised shock has already eliminated those projects, additional uncertainty has less to act on.

Taken together, the results show why disagreement across conditional forecasts is economically informative rather than merely numerical. Forecasts can differ because agents know different continuations of the same realised path, because that path is propagated through nonlinear rather than linear decision rules, or because agents price uncertainty around the continuation. Which margin matters most depends on where the scenario places the economy. Close to the steady state, the conventional linear conditional forecast may provide a good approximation. In large temporary disruptions, however, information, nonlinear adjustment, and uncertainty can dominate in different combinations, so there is no single correction that converts a linear scenario into the relevant nonlinear one.

The broader implication is that a macroeconomic scenario should specify more than the path that is eventually realised. It should also make explicit what agents know about the continuation when they make decisions and, where risk matters, the distribution they attach to that continuation. The information environment of the 2026 conflict illustrates why this distinction is substantive rather than semantic: current conditions were observable while the expected continuation was repeatedly revised. In such environments, a path is not yet a forecast.

References

Martin M. Andreasen, Jesús Fernández-Villaverde, and Juan F. Rubio-Ramírez. The pruned state-space system for non-linear DSGE models: Theory and empirical applications. *Review of Economic Studies*, 85(1):1–49, 2018. VERIFY exact pages before submission.

Juan Antolín-Díaz, Ivan Petrella, and Juan F. Rubio-Ramírez. Structural scenario analysis with SVARs. *Journal of Monetary Economics*, 117:798–815, 2021. VERIFY.

Robert B. Barsky and Eric R. Sims. News shocks and business cycles. *Journal of Monetary Economics*, 58(3):273–289, 2011. VERIFY.

Paul Beaudry and Franck Portier. Stock prices, news, and economic fluctuations. *American Economic Review*, 96(4):1293–1307, 2006. VERIFY.

Ben S. Bernanke. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics*, 98(1):85–106, 1983. VERIFY.

Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009. VERIFY.

Ryan Chahrour and Kyle Jurado. News or noise? the missing link. *American Economic Review*, 108(7):1702–1736, 2018. VERIFY.

Congressional Research Service. U.s. conflict with iran. Technical report, CRS Report R48887, March 2026. PDF archived 2026-07-18, data/stage4/statements/.

Jesús Fernández-Villaverde, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martín Uribe. Risk matters: The real effects of volatility shocks. *American Economic Review*, 101 (6):2530–2561, 2011.

Irving Fisher. *The Theory of Interest*. Macmillan, New York, 1930.

House of Commons Library. Israel/us-iran conflict 2026: Reopening the strait of hormuz. Technical report, Research Briefing CBP-10636, 2026. content retrieved 2026-07-18; PDF bot-gated, to be archived manually.

Nir Jaimovich and Sergio Rebelo. Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118, 2009. VERIFY.

Lutz Kilian, Michael D. Plante, Alexander W. Richter, and Xiaoqing Zhou. The impact of the 2026 Iran war on U.S. inflation: A scenario analysis. Working Paper 2609, Federal Reserve Bank of Dallas, April 2026.

Jinill Kim, Sunghyun Kim, Ernst Schaumburg, and Christopher A. Sims. Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models. *Journal of Economic Dynamics and Control*, 32(11):3397–3414, 2008.

Eric M. Leeper, Todd B. Walker, and Shu-Chun Susan Yang. Fiscal foresight and information flows. *Econometrica*, 81(3):1115–1145, 2013. VERIFY.

Junior Maih. Conditional forecasts in DSGE models. *Norges Bank Working Paper*, (2010/07), 2010.

Carlos Montes-Galdón, Joan Paredes, and Elias Wolf. Conditional density forecasting: A tempered importance sampling approach. *ECB Working Paper*, (2754), 2022. VERIFY.

Maurice Obstfeld and Kenneth Rogoff. *Foundations of International Macroeconomics*. MIT Press, Cambridge, MA, 1996.

Glenn D. Rudebusch and Eric T. Swanson. The bond premium in a DSGE model with long-run real and nominal risks. *American Economic Journal: Macroeconomics*, 4(1):105–143, 2012.

Stephanie Schmitt-Grohé and Martín Uribe. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4):755–775, 2004.

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

A The analytical mechanism: derivations and the benchmark check

This appendix supplies what Section 2 states without proof: the two-period programme that generates the Fisherian rules (8)–(10), the proofs of Lemma 1 and Proposition 1, and the comparison of the mechanism's sign predictions with the quantitative model on the benchmark path, decision date by decision date.

A.1 The programme behind the decision rules

At decision date t the economy holds resources $Y_t = \bar{Y} + w_t^*$ and, from date $t+1$ on, receives the windfall sequence $\{w_{t+s}\}_{s=1}^{T-t}$, consumes $\{c_{t+s}\}$, and borrows or lends at the gross world return R, with foreign assets b_τ carried into quarter τ and a terminal condition $b_{T+1} = 0$. Capacity i_t installed at date t produces $\Phi(i_t)$ units in every later quarter, valued at the relative price $d(w_{t+s})$, so that the quarter-by-quarter budget from $t+1$ on is

$$b_{t+s+1} = R\, b_{t+s} + w_{t+s} + d(w_{t+s})\,\Phi(i_t) - c_{t+s}, \qquad s = 1, \ldots, T-t. \tag{25}$$

Iterating (25) forward, imposing $b_{T+1} = 0$ and discounting to date $t+1$ gives a single constraint on the future consumption stream,

$$\tilde{c}_{t+1} \equiv \sum_{s=1}^{T-t} R^{1-s} c_{t+s} = R\, b_{t+1} + \underbrace{\sum_{s=1}^{T-t} R^{1-s} w_{t+s}}_{W_t} + \underbrace{\Big[\sum_{s=1}^{T-t} R^{1-s} d(w_{t+s})\Big]}_{Q_t} \Phi(i_t). \tag{26}$$

The two sums are the windfall-wealth and return statistics of equations (6)–(7); under information structure j they are evaluated at the point beliefs $E_t^j[\cdot]$, which is why each structure solves a deterministic programme. The collapse is exact because the economy borrows and lends freely at R in every future quarter.

Preferences are logarithmic, $\sum_{s=0}^{T-t} \beta^s \log c_{t+s}$. Given the value $\tilde{c}_{t+1}$ of the future stream, its optimal allocation across quarters yields indirect utility $\tilde{\beta} \log \tilde{c}_{t+1}$ plus a constant, with $\tilde{\beta} = \sum_{s=1}^{T-t} \beta^s$.[12] The date-$t$ decision is therefore the two-period problem

$$\max_{c_t,\, i_t,\, b_{t+1}} \ \log c_t + \tilde{\beta} \log \tilde{c}_{t+1} \quad \text{s.t.} \quad c_t + i_t + b_{t+1} = Y_t, \qquad \tilde{c}_{t+1} = W_t^j + R\, b_{t+1} + Q_t^j\, \Phi(i_t), \tag{27}$$

which is the problem of Section 2.2 with β standing for $\tilde{\beta}$. Maintained assumptions: $\Phi' > 0$, $\Phi'' < 0$, $\Phi'(0) = +\infty$ (interior investment), $d' > 0$, and $Y_t + [W_t^j + H(Q_t^j)]/R > 0$ (interior consumption); b_{t+1} is unrestricted in sign.

[12] Maximising $\sum_{s\geq 1} \beta^s \log c_{t+s}$ subject to $\sum_{s\geq 1} R^{1-s} c_{t+s} = \tilde{c}_{t+1}$ gives $c_{t+s} \propto \beta^s R^{s-1}$, and substituting back gives $\big(\sum_{s\geq 1} \beta^s\big) \log \tilde{c}_{t+1}$ up to terms that do not depend on $\tilde{c}_{t+1}$.

A.2 First-order conditions, Fisher separation, and the rules

With multipliers λ_t and λ_{t+1} on the two constraints of (27), the first-order conditions are

$$\frac{1}{c_t} = \lambda_t, \qquad \frac{\tilde{\beta}}{\tilde{c}_{t+1}} = \lambda_{t+1}, \qquad \lambda_t = R\,\lambda_{t+1}, \qquad \lambda_t = \lambda_{t+1}\,Q_t^j\,\Phi'(i_t). \tag{28}$$

The third condition with the first two is the Euler equation $\tilde{c}_{t+1} = \tilde{\beta} R\, c_t$. The fourth divided by the third is the investment rule (8), $Q_t^j \Phi'(i_t) = R$, from which both consumptions have cancelled. This is Fisher separation (Fisher, 1930): with an unrestricted bond, the capacity decision is a present-value decision that depends on Q_t^j alone and not on W_t^j, Y_t or preferences. Since $\Phi'' < 0$ the rule has a unique solution $i(Q)$, and differentiating it gives

$$i'(Q) = -\frac{\Phi'(i)}{Q\,\Phi''(i)} > 0. \tag{29}$$

Define, as in (9), $H(Q) \equiv \max_i\{Q\Phi(i) - Ri\} = Q\Phi(i(Q)) - R\,i(Q)$. By the envelope theorem, $H'(Q) = \Phi(i(Q)) > 0$. Eliminating b_{t+1} between the two constraints of (27) gives the lifetime constraint $c_t + \tilde{c}_{t+1}/R = Y_t + [W_t^j + Q_t^j\Phi(i_t) - R\,i_t]/R$, whose right-hand side at the optimal $i(Q_t^j)$ equals $Y_t + [W_t^j + H(Q_t^j)]/R$. Combining with the Euler equation,

$$c_t^j = \frac{1}{1+\tilde{\beta}}\left[Y_t + \frac{W_t^j + H(Q_t^j)}{R}\right], \qquad b_{t+1}^j = Y_t - c_t^j - i(Q_t^j), \tag{30}$$

which is (10). Finally, the date-t constraint holds for both information structures with the same Y_t; subtracting the two gives the common-pie identity (11), $\mathcal{A}_{c,t} + \mathcal{A}_{i,t} + \mathcal{A}_{b,t+1} = 0$.

A.3 Proof of Lemma 1

Let $E_t^F w_{t+s} \geq E_t^S w_{t+s}$ for all $s \geq 1$, strictly for some s. The weights R^{1-s} in (6) are positive, so $W_t^F - W_t^S = \sum_s R^{1-s}(E_t^F w_{t+s} - E_t^S w_{t+s}) > 0$. Since d is strictly increasing, $d(E_t^F w_{t+s}) \geq d(E_t^S w_{t+s})$ term by term with strict inequality wherever the beliefs differ, so $Q_t^F > Q_t^S$. By (29), $i(Q_t^F) > i(Q_t^S)$, that is $\mathcal{A}_{i,t} > 0$. By (30), c_t is increasing in W_t and, through $H' = \Phi > 0$, in Q_t, so $\mathcal{A}_{c,t} > 0$. Then $\mathcal{A}_{b,t+1} = -\mathcal{A}_{c,t} - \mathcal{A}_{i,t} < 0$ by the common-pie identity. Reversing the premise reverses every inequality. □

A.4 Proof of Proposition 1

Under sequential surprise the continuation formed at date t is (3), $E_t^S w_{t+s} = \rho^s w_t^*$, and under full information it is the realised path, $E_t^F w_{t+s} = w_{t+s}^*$. Condition (13) therefore states that at t_0 the full-information continuation dominates the sequential-surprise continuation, and the first case of Lemma 1 gives $\mathcal{A}_{c,t_0} > 0$, $\mathcal{A}_{i,t_0} > 0$, $\mathcal{A}_{b,t_0+1} < 0$. Condition (14) states the reverse ordering at t^*, and the second case of the lemma reverses the three signs. □

Comparative statics of the post-peak gap. The discussion following Figure 1 concerns the size of the gap at the peak,

$$W^S_{t^*} - W^F_{t^*} = \sum_{s\geq 1} R^{1-s}(\rho^s w^*_{t^*} - w^*_{t^*+s}), \tag{31}$$

and the analogous expression with $d(\cdot)$ applied to each term for $Q^S_{t^*} - Q^F_{t^*}$. For a positive peak $w^*_{t^*} > 0$, every term of (31) is increasing in ρ and decreasing in the realised post-peak value $w^*_{t^*+s}$; the same holds term by term for the return statistic because d is increasing. Hence, holding the realised post-peak path fixed, a more persistent law of motion widens the gap between the two information structures, and, holding ρ fixed, a faster realised reversal widens it as well. Through (29) and (30) the magnitudes of $\mathcal{A}_{c,t^*}$ and $\mathcal{A}_{i,t^*}$ increase with the gap, which is the sense in which information is most consequential when the historical law of motion and the realised duration point to different continuations.

A.5 The benchmark check, decision date by decision date

The premises of the mechanism are statements about the shape of the realised path relative to the continuation inferred from the current state; its conclusions are sign predictions. Both can be confronted with the quantitative model, whose full endogenous paths under full information (PF) and sequential surprise (MIT) are stored with the results. Section 6.2 refers to this comparison.

The path and the two continuations. The benchmark conditions on log oil-price deviations of $(0.479, 0.727, 0.963, 0.479, 0.333, 0.227, 0.162, 0.107)$ over the eight-quarter window, together with a log production deviation of -0.305 (26 per cent) in the first three quarters and zero thereafter; beyond the window both processes continue by their own laws with zero innovations, which is what the perfect-foresight solver imposes. The windfall that reaches domestic spending is oil revenue, price times exported quantity, so the check uses the single income basis $w_q = \exp(\hat{p}_q + \hat{o}_q) - 1$, the level deviation of oil income from steady state. The quantitative model has two conditioning processes rather than the one of equation (2); the sequential-surprise continuation at a decision in quarter k is the deterministic continuation of each law from the observed state, $\hat{p}_{k+s} = \rho_P^s \hat{p}_k$ and $\hat{o}_{k+s} = \rho_O^s \hat{o}_k$, with $\rho_P = 0.937$ and $\rho_O = 0.974$ (half-lives of 10.6 and 26.6 quarters), combined through the same income basis.

Dominance and the quantitative signs. Figure 7 draws the comparison and Table 12 records it. For each decision date k the table reports which continuation dominates over the twenty quarters that follow the decision ($T - k = 20$ at every k), the two wealth statistics $W^F_k = \sum_{s=1}^{20} R^{1-s} w^*_{k+s}$ and $W^S_k = \sum_{s=1}^{20} R^{1-s} E^S_k w_{k+s}$ at $R = 1.01$, in quarters of steady-state oil income, and the full-information effect on consumption and on investment by sector in that quarter.

At the first decision the realised path dominates and the informed economy consumes more and builds more non-tradable capacity, by 5.9 and 57 percentage points; from the third decision the extrapolation dominates and both signs reverse, as Proposition 1 predicts. The sign of the

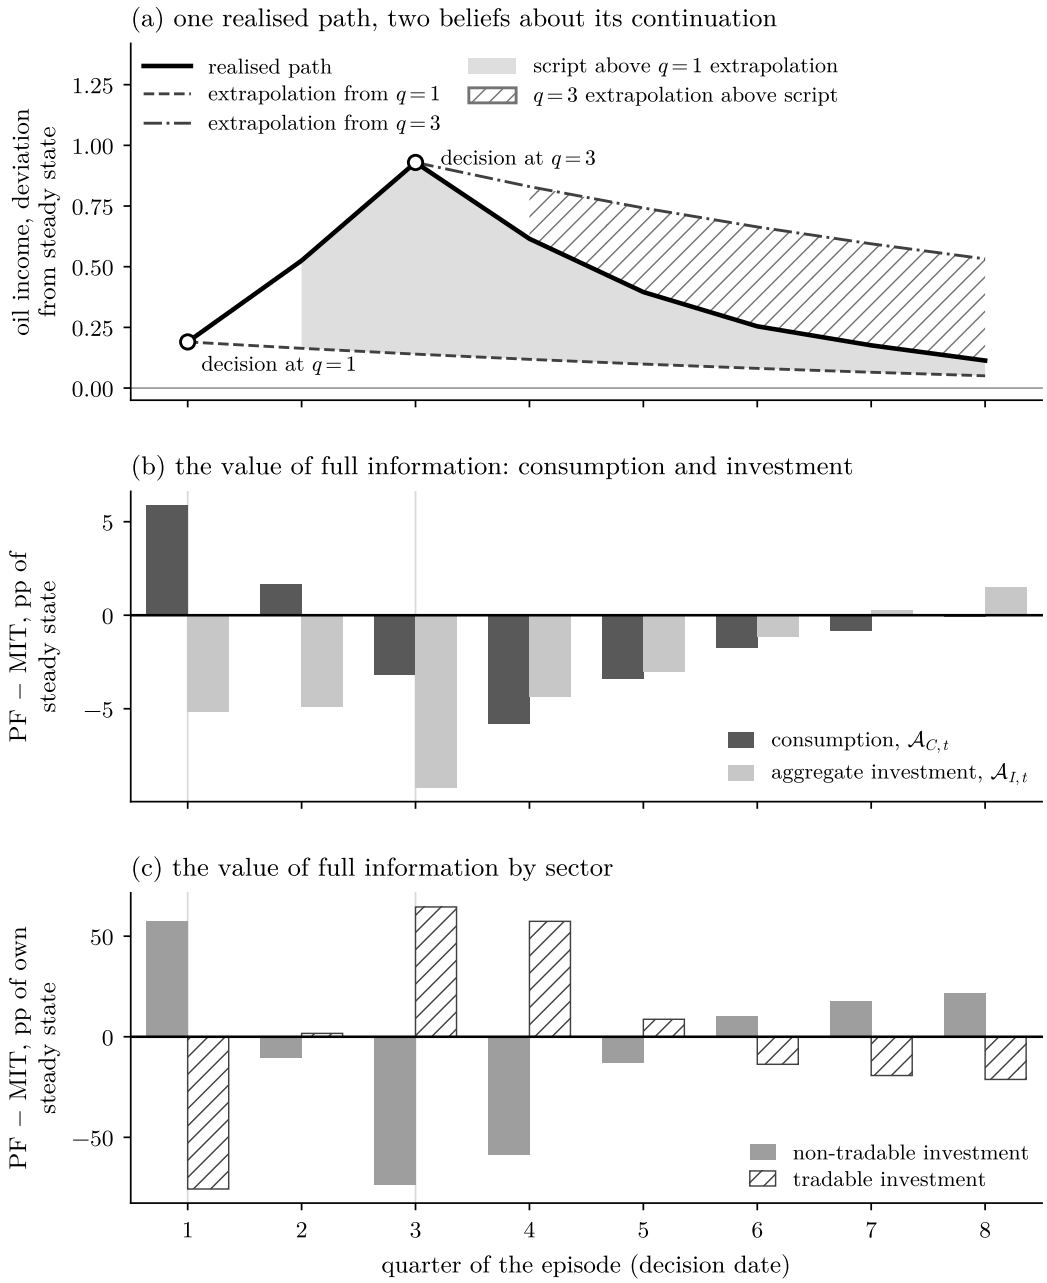


Figure 7: **The mechanism and the quantitative model on the benchmark path.** Panel (a): realised oil income (solid), the sequential-surprise continuation formed at the first decision (dashed) and at the peak (dash-dotted); the shaded area marks the quarters in which the realised path lies above the first continuation, the hatched area those in which the peak continuation lies above the realised path. Panels (b) and (c): the full-information effect $\mathcal{A}_t$ = PF − MIT at each decision date for consumption and aggregate investment, and for non-tradable and tradable investment, in percentage points of each variable's own steady state.

Table 12: **The mechanism at each decision date of the benchmark.** Which continuation dominates over the remaining horizon, the two wealth statistics on the income basis (twenty quarters after each decision, $R = 1.01$, quarters of steady-state oil income), and the full-information effect $\mathcal{A}_t = \text{PF} - \text{MIT}$ in that quarter for consumption, non-tradable investment, tradable investment, and aggregate investment (percentage points of each variable's own steady state).

decision	dominant continuation	W^F	W^S	$\mathcal{A}_C$	$\mathcal{A}_{I_{NT}}$	$\mathcal{A}_{I_T}$	$\mathcal{A}_I$
q1	realised path (informed perceives more)	3.79	0.55	+5.87	+57.3	−75.7	−5.17
q2	neither (the two cross)	3.33	3.40	+1.66	−10.5	+1.7	−4.87
q3	extrapolation (surprised perceives more)	2.46	6.56	−3.20	−73.7	+64.5	−9.21
q4	extrapolation	1.90	5.63	−5.79	−58.5	+57.4	−4.35
q5	extrapolation	1.55	3.73	−3.39	−13.1	+8.7	−3.00
q6	extrapolation	1.33	2.45	−1.76	+10.1	−13.7	−1.17
q7	extrapolation	1.19	1.72	−0.86	+17.7	−19.3	+0.28
q8	the two coincide	1.12	1.12	−0.10	+21.7	−21.2	+1.52
cumulative, quarters 1–8				−7.56	−49.0	+2.4	−26.0

consumption differential is predicted at every date except two near-ties: the second decision, where the two continuations cross and the wealth statistics differ by two per cent, and the eighth, where the two beliefs coincide because the observed state has returned to the law of motion. The sign of the non-tradable investment differential is predicted at the five dates that carry the cumulative account and departs from the prediction in quarters six to eight, where the surprised economy is running down the stock it accumulated while extrapolating the peak; this is the stock effect that Section 6.2 attributes to inherited capital and that a model without a capital stock cannot represent. Tradable investment moves against non-tradable investment at every large date. In units of steady-state aggregate investment the impact effects are −35.8 (tradable) and +30.2 (non-tradable) percentage points, which is why the aggregate effect of −5.2 points carries the tradable sign; Section 6.1 gives the general-equilibrium reason.

Two remarks. First, the labour margin. The informed household supplies 2.9 percentage points less labour than the surprised one at impact and more from the third quarter on, a wealth effect on labour supply that moves the return on both capital stocks and that the two-period economy excludes by Fisher separation. Second, the horizon. On the income basis the extrapolated production cut outlasts the extrapolated price boom, so the surprised economy's perceived income eventually turns negative. Lengthening the horizon to forty or sixty quarters lowers W^S at the first decision (to −0.4 and −1.2) and leaves the ordering of the two wealth statistics at every date in Table 12 unchanged. The one qualification concerns pointwise dominance at the third decision: the extrapolation formed at the peak falls below the realised path after horizon 26, so the premise of Lemma 1 holds there for the first twenty-five quarters rather than for every horizon, which is far beyond the horizons that carry weight in the statistics.

B Model appendix

B.1 Equations

The full system comprises 85 endogenous variables in twelve blocks; we state the economically load-bearing equations and describe the remainder, with the complete listing available in the replication package. **[TODO: expand to the complete equation listing before submission .]**

Households. The national household chooses consumption C_t, energy services ES_t, labour ℓ_t, sectoral investments and next-period capital, and foreign bonds B_t^*. Preferences are non-separable in (C, ES) with curvature σ and weight d, and additively separable in labour with Frisch-type parameter ν. The foreign-bond Euler equation carries a debt-elastic premium with elasticity ξ around steady-state net foreign assets, which pins the long-run external position. Tobin-q investment with quadratic adjustment costs h_T, h_{NT} governs sectoral capital accumulation at common depreciation $\delta = 0.025$. An expatriate household supplies the fraction $x = 0.6$ of employment, consumes fixed wage shares domestically and remits the rest.

Production. Each sector $s \in \{T, NT\}$ produces with labour and a CES capital–energy composite, $y_{s,t} = A_s z_{s,t} L_{s,t}^{\alpha_s} \left[\theta_s K_{s,t-1}^{-\nu_s} + (1-\theta_s) ES_{s,t}^{-\nu_s}\right]^{-(1-\alpha_s)/\nu_s}$, with sectoral TFP following AR(1) processes. The law of one price holds for tradables and the peg fixes the nominal anchor, so the relative price of non-tradables clears domestic demand.

Oil and gas block. As in Section 4.2: the log price deviation follows equation (15); production $\log \bar{O}_t = (1-\rho_O)\log(\alpha_O K_{t-1}^{oil} u) + \rho_O \log \bar{O}_{t-1} + \sigma_O \varepsilon_t^O$ with $u = 1$ and oil-sector investment a fiscal parameter, so K_t^{oil} is deterministic; gas production proportional to oil, gas prices AR(1).

Fiscal block and market clearing. The government budget collects oil/gas export and domestic energy revenue, consumption and factor taxes; spends on public consumption (AR(1) around its steady state), oil-sector investment, and debt service on domestically held bonds (fixed supply); transfers balance the budget. Non-tradable goods and energy-services markets clear domestically; the current account accumulates foreign bonds.

B.2 Steady state

The steady state is constructed by a two-layer strategy: an analytic layer expresses most variables in terms of eleven core unknowns, and a square nonlinear system solves those (residuals verified below 10^{-8}). The steady-state oil price is normalised so that one model price unit is \$36.38/bbl (steady state \$63.78/bbl); production is normalised to 9.5 mb/d.

Table 13: Calibration (quarterly). Values verified against the archived calibration file 2026-07-20.

parameter	value	meaning / target
β	0.9901	discount factor (4% annual rate)
σ_C	1	consumption curvature
σ	$-1/3$	C–ES substitution (complements)
d	0.0598	energy-services weight
ν	1.429	labour supply curvature
θ_C	0.479	tradable share in consumption
h_T, h_{NT}	0.223, 0.149	investment adjustment costs
δ_T, δ_{NT}	0.025	depreciation
$\alpha_T = \alpha_{NT}$	0.582	labour exponents
ν_T, ν_{NT}	1.381	capital–energy CES
x	0.6	expatriate employment share
ξ	0.05	debt-elastic premium
ρ_P	0.9368	oil price AR(1)
$\bar{\sigma}$	0.15	oil price innovation scale
ρ_σ	0.64	volatility persistence (audited)
σ_σ	0.27	vol-of-vol (estimated; App. E)
$\rho_O,\ \sigma_O$	0.9742, 0.032	oil production AR(1), innovation
$\alpha_O,\ \mu_O$	0.0320, 0.0125	production–capital link, depreciation
ρ_{zT}, ρ_{zNT}	0.982, 0.987	sectoral TFP persistence
$\sigma_{zT}, \sigma_{zNT}$	0.110, 0.166	TFP innovations
$\rho_G,\ \sigma_G$	0.7, 0.157	public spending process
conversion	36.38, 80.33	\$ /bbl per price unit; mb/d per flow unit

B.3 Calibration

C Recovery algebra

Setup. A scenario is a pair of level target paths $\{P_t^*, O_t^*\}_{t=1}^H$. The oil block (Section 4.2) is exogenous and exact in logs; recovery of the innovations that make the block trace the targets is triangular.

Block 1 (volatility). If a volatility level path $\{\sigma_t^*\}$ is conditioned on, set $s_t^* = \log(\sigma_t^*/\bar{\sigma})$ and invert the volatility law: $\varepsilon_t^\sigma = (s_t^* - \rho_\sigma s_{t-1}^*)/\sigma_\sigma$, s_0^* the anchor state. With no volatility path (constant volatility), $\varepsilon^\sigma \equiv 0$ and $s_t^* = \rho_\sigma^t s_0$.

Block 2 (price). Under the log (path-identical) operator, $\hat{p}_t^* = \log(P_t^*/\bar{P})$ and, respecting the $t-1$ timing,

$$\varepsilon_t^P \;=\; \frac{\hat{p}_t^* - \rho_P\, \hat{p}_{t-1}^*}{\bar{\sigma}\, e^{s_{t-1}^*}},$$

which is exact because the price law is linear in $\hat{p}$ given the volatility path. Under the hat (linear-representation) operator used for the WZ and PF$_{lin}$ cells, replace logs by fractional deviations and the denominator by $\bar{\sigma}$ (the linear solution carries no volatility–innovation cross term; this constant loading is what reproduces the textbook WZ instruments, Appendix **??**).

Block 3 (production). Oil-sector capital is deterministic (investment is a parameter): $K_t^{oil} = I^{oil} + (1-\mu_O)K_{t-1}^{oil}$, computable from the anchor without simulation. The produc-

tion law then inverts in one pass:

$$\varepsilon_t^O \;=\; \frac{\log O_t^* - (1-\rho_O)\log(\alpha_O K_{t-1}^{oil} u) - \rho_O \log O_{t-1}^*}{\sigma_O}.$$

Proposition (exactness and order-freeness). The recovered sequence $(\varepsilon^\sigma, \varepsilon^P, \varepsilon^O)$ is unique given the targets and the anchor; replaying the recovered innovations through the exact laws reproduces the targets identically (machine precision in every recorded run); and the recovery involves no solution object of any order — the same canonical sequence feeds the deterministic solvers and the pruned third-order simulation. *Remark (the impurity).* Under the third-order solution the realised *level* P_t^O is the Taylor image of $\bar{P}e^{\hat{p}_t}$; with $\hat{p}$ bound exactly, the level gap is truncation error of the representation, bounded on the benchmark path at 1.7 per cent of the price level and 0.05 per cent of production, reported per horizon in the replication package, and not iterated upon (an early fixed-point attempt on this gap diverged, for the stated reason: the gap is not a property of the innovations).

Simulation. Deterministic cells solve the stacked nonlinear system by Newton (tolerance tightened to 10^{-12} residuals after a near-unit-root amplification analysis: per-equation residuals amplify into path gaps by $\approx 1/(1-\rho_O) \approx 39$). MIT is the rolling-window sequence of such solves. Perturbation cells simulate the pruned state space (Andreasen et al., 2018); the risk terms and their role are as in Section 5.4 and Appendix D. All responses are scenario effects: each cell's path minus the same cell's zero-innovation path from the same anchor — for third-order cells this removes the deterministic-to-stochastic steady-state transition (Appendix D).

D Decomposition details

D.1 Operators and the four deterministic cells

Section 5.2 writes each conditional forecast as the composition of a recovery operator and a solution operator, $y_t^{m,j} = \mathcal{S}_t^{m,j}(x_0, \mathcal{R}^m(\mathcal{P}^*))$, where $m \in \{\text{lin}, \text{nl}\}$ indexes the solution representation and $j \in \{S, F\}$ the information structure. The four cells of Table 2 are

$$\begin{aligned} \text{WZ} &: y^{\text{lin},S}, & \text{PF}_{\text{lin}} &: y^{\text{lin},F}, \\ \text{MIT} &: y^{\text{nl},S}, & \text{PF} &: y^{\text{nl},F}. \end{aligned} \tag{32}$$

Two features of the implementation matter for what follows. First, the recovery operators differ by representation. In the exact oil block the level targets are bound in logarithms, $\hat{p}_t^* = \log(P_t^*/\bar{P})$, and the innovations follow from inverting the laws of Section 4.2 (Appendix C); in the first-order representation the same level targets are bound as fractional deviations, $(P_t^* - \bar{P})/\bar{P}$, because that is the object the linear solution propagates. The two innovation sequences u^{lin} and u^{nl} therefore differ away from the steady state even though each reproduces the prescribed level path in its own representation. Second, the solution operators for $j = S$ are rolling objects: at each decision date t within the window the economy is solved forward with the innovations observed up to t and none thereafter, and the date-t outcome is retained.

For $m = \text{lin}$ this rolling solve is the Waggoner–Zha construction; for $m = \text{nl}$ it is the sequence of perfect-foresight solves with one-period MIT shocks.

Remark (certainty equivalence and the WZ cell). At first order the policy function is $g(x, \sigma) = \bar{y} + g_x \hat{x}$ with $g_\sigma = g_{x\sigma} = 0$ (Schmitt-Grohé and Uribe, 2004), so conditional expectations of future states are linear in the current state and do not depend on future shock variances. Under sequential surprise each date's re-optimisation therefore reproduces the linear filtering recursion that the Waggoner–Zha construction implements directly, and the rolling linear solve and the WZ forecast coincide path for path; the pipeline verifies this to 10^{-13} (Appendix **??**). The linear sequential-surprise cell is thus the standard conditional forecast exactly, and $\mathcal{L}_t^S = y_t^{\text{MIT}} - y_t^{\text{WZ}}$ changes the solution representation and nothing else. The equivalence is bought by the first-order truncation: at second order and above, $g_{\sigma\sigma} \neq 0$ and the timing of information acquisition becomes payoff-relevant, which is why the higher-order branch of Section 5.4 holds the information structure fixed at sequential surprise rather than comparing structures across orders.

D.2 The local result: $\mathcal{I}_t(\lambda) = O(\lambda^2)$

Scale the scenario as $\mathcal{P}(\lambda) = \bar{\mathcal{P}} + \lambda \Delta\mathcal{P}$, with $\bar{\mathcal{P}}$ the steady-state path, and write $u^m(\lambda) = \mathcal{R}^m(\mathcal{P}(\lambda))$ and $y^{m,j}(\lambda) = \mathcal{S}^{m,j}(x_0, u^m(\lambda))$, suppressing the date subscript. Assume that the recovery and solution maps are twice continuously differentiable in a neighbourhood of the steady state; the oil laws are smooth and the perfect-foresight system is a smooth square system with a regular Jacobian at the steady state, so this holds for the deterministic cells. At $\lambda = 0$ every cell equals the steady state, $y^{m,j}(0) = \bar{y}$ and $u^m(0) = 0$. Expanding to second order,

$$\begin{aligned} y^{m,j}(\lambda) = \bar{y} &+ \lambda\, D\mathcal{S}^{m,j}\, D\mathcal{R}^m\, \Delta\mathcal{P} \\ &+ \tfrac{1}{2}\lambda^2 \Big\{ D\mathcal{S}^{m,j}\, D^2\mathcal{R}^m[\Delta\mathcal{P}, \Delta\mathcal{P}] + D^2\mathcal{S}^{m,j}\big[D\mathcal{R}^m\Delta\mathcal{P}, D\mathcal{R}^m\Delta\mathcal{P}\big] \Big\} + O(\lambda^3), \end{aligned} \tag{33}$$

where all derivatives are evaluated at the steady state. Two facts identify the first-order terms across representations. First, the first-order perturbation solution is, by construction, the derivative of the exact solution map at the expansion point: $D\mathcal{S}^{\text{nl},j} = D\mathcal{S}^{\text{lin},j}$ for each j, applied solve by solve for the rolling structure. Second, the log and fractional-deviation bindings agree to first order, $\log(1 + x) = x + O(x^2)$, so $D\mathcal{R}^{\text{nl}} = D\mathcal{R}^{\text{lin}} \equiv D\mathcal{R}$. Moreover, the linear cells are exactly linear in λ: the hat recovery is linear in the fractional deviations and the first-order solution is linear in the innovations, so $D^2\mathcal{R}^{\text{lin}} = 0$ and $D^2\mathcal{S}^{\text{lin},j} = 0$. Hence

$$\mathcal{A}^{\text{lin}}(\lambda) = \lambda\big[D\mathcal{S}^{\text{lin},F} - D\mathcal{S}^{\text{lin},S}\big] D\mathcal{R}\, \Delta\mathcal{P} \qquad \text{exactly,} \tag{34}$$

while the nonlinear full-information effect has the same first-order term and a second-order remainder:

$$\begin{aligned}\mathcal{I}(\lambda) &= \mathcal{A}(\lambda) - \mathcal{A}^{\text{lin}}(\lambda) \\ &= \tfrac{1}{2}\lambda^2\Big\{\underbrace{[D\mathcal{S}^F - D\mathcal{S}^S]D^2\mathcal{R}^{\text{nl}}[\Delta\mathcal{P},\Delta\mathcal{P}]}_{\text{recovery curvature}} + \underbrace{D^2\mathcal{S}^{\text{nl},F}[Du,Du] - D^2\mathcal{S}^{\text{nl},S}[Du,Du]}_{\text{propagation curvature}}\Big\} + O(\lambda^3),\end{aligned} \tag{35}$$

with $Du = D\mathcal{R}\,\Delta\mathcal{P}$. This is equation (21). The expansion also separates the two sources of the interaction named in Section 5.2. The first term is the curvature of shock recovery: binding the same level path in logarithms rather than in fractional deviations changes the innovations at second order, and this enters $\mathcal{I}$ only through the difference between the two first-order propagation maps, that is, only to the extent that the informed and surprised economies respond differently to the same innovation sequence. The second term is the curvature of propagation under each information structure. It need not be small relative to the first, and it need not be symmetric across j: the informed economy responds at date t to the entire anticipated sequence of innovations, while the surprised economy responds to the current innovation and the deterministic continuation of the state, so the two economies are evaluated at different points of the same curved policy functions. Section 7 reports that the full-information linearisation gap $\mathcal{L}^F$ exceeds the sequential-surprise gap $\mathcal{L}^S$ in every disruption scenario, which is the empirical face of this asymmetry.

The local result is a statement about the limit $\lambda \to 0$. It says that information and nonlinearity separate near the steady state and that their interaction is a second-order object; it does not bound the interaction at the severities studied in the paper, where Section 7 shows it to be of the same order as the linearisation gaps themselves.

D.3 Curvature and risk in the pruned third-order solution

The higher-order forecasts use the pruned state-space representation of Andreasen et al. (2018). Writing the third-order policy function in the notation of (22), the terms fall into two groups. The deterministic curvature terms g_x, g_{xx}, g_{xxx} (and their counterparts in h) describe how the response to the state bends away from the linear rule; they are present in the exact deterministic solution as well, which is why HO_{CE} can be compared with MIT. The risk terms $g_{\sigma\sigma}$, $g_{\sigma\sigma x}$, and $g_{\sigma\sigma u}$ (the constant, state-dependent, and innovation-dependent corrections; the odd terms g_σ, $g_{\sigma x}$, $g_{\sigma xx}$ and $g_{\sigma\sigma\sigma}$ vanish for symmetric innovations) describe how the level of uncertainty shifts decisions. HO_{CE} is the pruned third-order simulation with the risk terms set to zero; HO restores them at constant volatility; HO_{SV} additionally feeds the recovered volatility innovations through the stochastic-volatility law. Because the three forecasts share the same innovation sequence, information structure, and approximation order, $\mathcal{R}^{\text{pure}} = \text{HO} - \text{HO}_{\text{CE}}$ isolates risk pricing and $\mathcal{T} = \text{HO}_{\text{CE}} -$ MIT isolates the third-order truncation of deterministic curvature, as Table 3 states.

The risk correction can itself be split by term. On the benchmark path the cumulative investment effect of constant-volatility risk pricing, $\mathcal{R}^{\text{pure}} = +8.0$ percentage points, is com-

posed of +2.9 from the constant term, +5.4 from the state-dependent term, and −0.3 from the innovation-dependent term; the three add to the total to machine precision. The state-dependent term is the part of risk pricing that changes with the scenario, which is what makes the risk margin state dependent in Section 8.

D.4 Baseline discipline: the drifting third-order counterfactual

A pruned third-order economy initialised at the deterministic steady state moves with no innovations at all: the constant risk term pushes it toward the stochastic steady state, here by +3.9 per cent of GDP after twenty quarters, with a transition dip in consumption of −13 per cent and an investment surge of about +48 per cent in the first quarter. Every third-order object in the paper is therefore measured net of its own zero-innovation path from the same anchor. The cautionary example: the raw third-order investment response to the benchmark at impact reads +54.7 per cent of steady state; the scenario effect is +7.0. The raw number is the stochastic-steady-state transition wearing the scenario as a disguise, and it is exactly the number an unwary reader would quote. Two facts complete the picture. The drift is almost entirely the constant risk term, so zero-innovation paths at orders two and three are indistinguishable, and its variance budget is a business-cycle object: switching off shock variances one at a time attributes it to sectoral productivity (non-tradable 67 and tradable 34 per cent of the twenty-quarter drift), with the oil-price variance contributing slightly negatively (Section 8). Naive, unpruned third-order simulation did not explode at our shock sizes; it stayed bounded while producing a materially different path, the more dangerous failure mode, and our reason for treating pruning and baseline netting as non-negotiable.

D.5 Rejected decompositions

Risk as HO − WZ. This difference bundles the linearisation gap, the third-order truncation, and risk pricing; calling it a premium relabels the dominant component, curvature, as compensation for risk. The literature's premium is risky minus risk-neutral *within* a solution order (Rudebusch and Swanson, 2012); $\mathcal{R}^{\text{pure}}$ implements exactly that, with HO_{CE} the risk-neutral twin.

Splitting $\mathcal{L}^S$ through PF. Writing MIT − WZ as (MIT − PF) + (PF − WZ) makes the second bracket change information and solution method at once, a confounded object that no single-switch interpretation fits. The square of Section 5.3 and the chain of equation (23) exist precisely to keep one switch per difference.

Anchor dependence (remark). Under path-identical conditioning the experiment itself depends on the anchor: the same level path implies larger innovations from a displaced state with a lower oil price. This is a feature, since conditioning on outcomes means that the severity of the implied shock is anchor-relative, but it is why all results anchor at the steady state, where the experiment is anchor-neutral; anchoring at a filtered current state is listed as an extension in Section 9.

E Empirical appendix

E.1 Estimating the volatility block

The SV block of the oil price (equation 15) is estimated in two steps on quarterly real Brent (monthly Brent deflated by US CPI, quarterly averages, 1992Q1–2026Q1, $n = 136$; FRED series, vintage archived with a re-download script). Step one: OLS of the demeaned log real price on its lag gives $\rho_P = 0.958$ and residual scale 0.140 — close to the calibrated 0.9368 and 0.15, which are retained. Step two: $\log|e_t| = \log\sigma_{t-1} + \log|z_t|$ is AR(1) signal plus i.i.d. noise with known moments; method-of-moments on autocovariances gives $\rho_\sigma = 0.375$ (biased down by the noise; naive OLS gives 0.09) and unconditional dispersion $\mathrm{sd}(\log\sigma) = 0.51$. The cross-check on the CBOE oil implied-volatility index (quarterly log-OVX AR(1), 2007–2026) gives persistence 0.709 — smoothed upward, as an implied measure should be — and dispersion 0.33. The calibration keeps $\rho_\sigma = 0.64$ (inside the bracket) and sets $\sigma_\sigma = 0.27$, matching the OVX dispersion at that persistence; the realised-dispersion alternative is 0.40. **[TODO: sensitivity rerun at 0.40.]** The value the model file historically carried (0.01) implies an essentially constant volatility band and traces to an uncited line predating this project; it is documented, and replaced, in the replication package. Caveats: two-step noise attenuates persistence; time aggregation compresses dispersion; no leverage effect; the sample spans structural breaks. The conditioned volatility path of Section 8 (ratios $1.88, 1.98, 1.20$) is the 2026 OVX episode relative to its 2025H2 pre-war mean, from the same archived vintage.