Information Networks and Inflation Expectations: Evidence from Professional Forecasters and Conflicts*

Francesco Grigoli
Georgetown University

Bertrand Gruss
International Monetary Fund

This draft: October 6, 2026

Abstract

How are inflation expectations affected by foreign supply shocks? We show that, beyond bilateral exposure, a second network matters: the one linking forecasters to the countries they cover. Using forecaster-level expectations around conflict episodes—precisely dated supply shocks that raise inflation through trade and energy linkages—we find that expectations are revised upward gradually in the days following a conflict. Within the same country and day, forecasters who cover the conflict country revise by 0.2 to 0.3 percentage points less than those who do not, and the differential roughly triples where bilateral exposure is highest. We rationalize this in a model where a public narrative overstates spillovers in proportion to bilateral exposure, and only forecasters covering the conflict country can discount it. The consensus therefore inherits a bias largest where exposure is high and coverage thin, and a central bank reading it at face value tightens more than fundamentals warrant.

JEL Codes: D84, D85, E31, F51
Keywords: inflation expectations, professional forecasters, information networks, conflict shocks, cross-country spillovers, Bayesian learning

*Francesco Grigoli, francesco.grigoli@georgetown.edu; Bertrand Gruss bgruss@imf.org. The views expressed in this paper are those of the authors and do not necessarily represent those of the IMF, its Executive Board, or its management. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

How are inflation expectations in individual countries affected by supply shocks originating abroad? The question is of first-order relevance for the calibration of monetary policy. The textbook prescription is to look through a temporary spike in headline inflation. But if long-term expectations drift up, central banks may need to tighten to pre-empt persistent second-round effects, adding an output cost on top of the one the shock itself imposes. And even when long-term expectations have not yet moved, a drift in near-term expectations mechanically lowers the ex-ante real rate, so policy rates must rise simply to hold the stance unchanged, and failing to do so risks the near-term drift spilling into medium-term expectations. The calibration therefore hinges on how far expectations actually move, which central banks typically infer from the professional-forecaster consensus.

What determines how far a country's consensus forecast moves in response to a foreign supply shock? Two literatures offer complementary answers. The network literature points to bilateral exposure—trade, energy dependence, and other propagation channels—which governs the size of the spillover a country actually receives (Gabaix, 2011; Acemoglu *et al.*, 2012, 2016; Carvalho and Tahbaz-Salehi, 2019). The expectations literature points to how individual forecasters update under noisy or sticky information (Carroll, 2003; Maćkowiak and Wiederholt, 2009; Coibion and Gorodnichenko, 2012, 2015; Bordalo *et al.*, 2020). In this paper, we leverage the occurrences of conflicts to identify a third determinant, and a second network: the one linking forecasters to the countries they cover. We show that this network shapes not only cross-forecaster heterogeneity but also the consensus a central bank reads.

A forecaster's revision for a non-conflict country following a conflict shock could reflect either of two channels. The first channel is the second country's economic exposure to the conflict country—the standard *economic* channel of cross-border propagation. The second channel is the fact that the forecaster happens to cover the conflict country and is therefore already processing the shock—the *information* channel. The latter has a natural Bayesian interpretation in the spirit of Shiller (2017) and Andre *et al.* (2022). That is, a public narrative connecting the conflict to a non-conflict country is salient where bilateral exposure is salient, and only forecasters with granular information about the conflict country can strip the narrative bias from their posteriors.[1] Conflict shocks are well-suited to identify this channel because they are precisely dated, exogenous to the pre-existing forecaster network, and produce non-negligible aggregate macroeconomic costs (Caldara and Iacoviello, 2022; Federle *et al.*, 2024) that vary across countries through bilateral exposure. Thus, the central question of the paper is whether the information channel exists, what its sign and magnitude are, and how it interacts with the standard economic channel.

Using the forecaster-level Continuous Consensus Forecasts (CCF) from Consensus Economics, which cover 48 economies and 611 forecasting institutions, we build the bipartite forecaster-country network and study how it shapes inflation expectations across conflict episodes between 2014 and

[1] The public narrative is a reduced-form representation of the coverage decisions and editorial incentives logic studied in media economics (Mullainathan and Shleifer, 2005; Eisensee and Strömberg, 2007). This literature argues that outlets emphasize the bilateral links most salient to their audience, so empirical coverage of cross-border events scales with the underlying magnitude of the link.

2026.[2] We document three empirical patterns. Our first finding is that post-conflict inflation expectations for non-conflict countries revise gradually upward over the days following the shock. The estimated jump in the inflation forecast at the cutoff date is essentially zero, while the change in its daily trend is positive and statistically significant for horizons out to four or five quarters ahead, reaching 0.03 to 0.04 percentage points per day at the one- to three-quarter horizons. We document this using a regression-discontinuity-in-time design both on the most recent episode (the Middle East escalation of February 28, 2026) and on a pooled stack of all conflict episodes between 2014 and 2026. The estimates confirm that forecasters do not interpret the conflict as a one-time, fully-revealed inflation shock; instead, they walk their expectations upward as information accumulates over time.

The second finding is the central one. Within the same non-conflict country on the same date, forecasters who already cover the conflict country revise their forecasts by 0.2 to 0.3 percentage points less, in absolute value, than forecasters who do not. We identify this within-pair difference from a stacked regression of post-cutoff revisions on an indicator for whether the forecaster also covered the conflict country in the two years prior to the conflict. Identification compares forecasters within the same country, the same submission day, and the same conflict episode. The comparison absorbs every revision common to forecasters submitting on the same day for the same country in the same episode—e.g., trade exposure, energy prices, financial contagion, central-bank communication, common risk sentiment—leaving only the cross-forecaster variation attributable to differential information. The forecaster-coverage variation is predetermined relative to any individual conflict shock, with coverage assignments typically set by employer mandates and not by the events we study. The dampening is consistently negative, statistically significant, and robust across episode-selection rules, panel-construction choices, and a battery of composition diagnostics.

The third finding sharpens the second. The dampening differential roughly triples at country pairs where bilateral exposure to the conflict country is highest, concentrating at the energy and intermediate-goods-trade margins. Low-exposure pairs show a baseline dampening of about 0.2 percentage points; high-exposure pairs show roughly 0.5 percentage points at the energy margin, with non-overlapping confidence intervals. A finer decomposition shows the differential loading on oil, gas, and intermediate-goods trade, but essentially absent at capital- and consumption-goods margins. The pattern fits a mechanism in which the public narrative overstates spillovers most where bilateral linkages are strongest—energy and intermediate-input dependence are widely reported and compared across countries—while connected forecasters' granular private read on the conflict country lets them strip the bias out.

A Bayesian model with two coupled networks rationalizes these findings. The first is the bilateral exposure network—the standard primitive of international macroeconomics—which determines the size of cross-country spillovers. The second is a bipartite information network linking forecasters to the countries they cover. The information network rations who can observe a precise private signal about the conflict shock. Forecasters outside the information network condition only on a public narrative whose magnitude scales with bilateral exposure. The model delivers a closed-form

[2] A "forecaster" in the CCF panel is a forecasting institution—a global investment bank, a multinational corporate-economics team, an international body, or a domestic research house—rather than a named individual analyst. The information channel we identify is therefore an intra-firm signal across country desks of the same institution.

expression for aggregate inflation expectation bias as the product of the two networks. Specifically, bilateral exposure scales the size of the spillover, and the share of forecasters who do not cover the conflict country scales how much of the public-narrative bias survives in the consensus. The two networks compose multiplicatively. Bilateral exposure alone generates no aggregate bias when every forecaster covering the non-conflict country also covers the conflict country, and sparse forecaster coverage alone generates none when bilateral linkages are zero. The patterns identified empirically appear in the model as direct consequences of Bayesian updating with two unequally precise signals.

Closing the model with a New Keynesian Phillips curve generates three macroeconomic implications, which we take to the data. First, the cross-country dispersion of post-conflict consensus revisions decomposes into a fundamentals component, driven by heterogeneity in bilateral exposure, and a network-bias component, driven by the joint distribution of bilateral exposure and forecaster-coverage gaps. The network-bias component contributes more than an order of magnitude more to dispersion than the fundamentals component does in our eight-episode panel. Second, the bilateral inflation wedge—the gap between realized inflation and what full-information rational expectations would imply—scales with the same product of two networks. Calibrating from the within-pair empirical estimate, the wedge runs to 25–60 basis points at the most-exposed economies in our sample, a substantial fraction of central-bank inflation targets. Filling the top fifteen forecaster-country coverage gaps eliminates more than half of the cross-country total wedge. Third, public data quality and forecaster network density are substitutes in determining aggregate expectation bias. Better real-time bilateral data on the conflict country reduces it in the same way broader analyst coverage does, and the model gives the trade-off a precise rate.

Our results speak to the policy tradeoff that a central bank faces after a geopolitical event. The severity of the tradeoff between containing the rise in inflation expectations and cushioning output losses depends on how far the consensus moves in response to the shock. While our panel does not extend to long-horizon expectations, our findings suggest that the consensus overstates the true movement at the near-term horizons, one to four quarters out, which are the expectations that enter the ex-ante real rate. A central bank reading the consensus at face value therefore perceives the tradeoff as sharper than it is and is pushed toward over-tightening to 'buy anchoring' it does not need—roughly 40–90 basis points of excess tightening at the worst-affected economies in our sample. The wedge is largest where bilateral exposure is high and coverage of the conflict country is thin, the configuration typical of small open economies; since the excess output contraction is linear in the wedge and the welfare loss quadratic in it, the cost of misreading the consensus concentrates there. Because the effect we identify operates through the public narrative rather than through fundamentals, a central bank that corrects that narrative—for instance by quantifying the country's actual import share from the conflict country and the implied pass-through to CPI—does publicly what a broader forecaster network would do privately and shrinks the distortion directly.

Literature The paper sits at the intersection of three strands of literature. The first studies how individual forecasters form expectations. Noisy- and sticky-information models (Mankiw and Reis, 2002; Carroll, 2003; Sims, 2003; Maćkowiak and Wiederholt, 2009; Coibion and Gorodnichenko, 2012, 2015;

Maćkowiak and Wiederholt, 2015; Coibion *et al.*, 2021), models of diagnostic over-reaction (Bordalo *et al.*, 2018, 2020), and the literature on forecaster disagreement (Mankiw *et al.*, 2003; Patton and Timmermann, 2010; Andrade and Le Bihan, 2013) and on imperfect macroeconomic expectations (Eusepi and Preston, 2011; Angeletos and La'O, 2013; Angeletos *et al.*, 2021; Andre *et al.*, 2022) all generate heterogeneity across forecasters from information differences that the econometrician does not observe. We identify a source of that heterogeneity in the bipartite information network: which analyst happens to cover which country varies cross-sectionally and is observable in the panel. The premise that coverage confers an informational edge is well established in finance, where analysts who are geographically or organizationally close to what they follow issue sharper forecasts and trade on better information (Coval and Moskowitz, 2001; Malloy, 2005; Bae *et al.*, 2008). Our setting differs from that literature in three respects. First, agents are macroeconomic forecasters covering countries rather than equity analysts covering firms. Second, the advantage we identify operates on a third country's inflation outlook rather than on the covered entity itself, so it is a statement about cross-border transmission rather than about coverage quality. And third, the object of interest is the aggregate bias that the coverage network imprints on the consensus, rather than the individual analyst's forecast performance.

Most closely related on the demand side, Coibion *et al.* (2022) show experimentally that information treatments—about current inflation readings or central-bank communications—shift household inflation expectations significantly. The mechanism in their setting is one in which a public information signal alters posteriors anchored to outdated or imprecise priors. Our finding is the supply-side analogue in a professional-forecaster setting. That is, rather than varying the public information that reaches the agent, we vary the granular private information that some forecasters can use to correct the public-narrative bias, and show that consensus inflation forecasts inherit a bias whenever this private information is rationed by the bipartite network. Together, the two findings establish that the precision of country-level information supply—whether broadcast publicly to households or held privately by a connected forecaster—shapes inflation expectations at quantitatively comparable magnitudes.

The second strand of literature is the one on macroeconomic networks, which takes the bilateral exposure matrix as the central primitive of cross-country transmission (Gabaix, 2011; Acemoglu *et al.*, 2012; Caliendo and Parro, 2015; Acemoglu *et al.*, 2016; Atalay, 2017; Cravino and Levchenko, 2017; Auer *et al.*, 2019; Carvalho and Tahbaz-Salehi, 2019; Forbes, 2019; Pasten *et al.*, 2020; Carvalho *et al.*, 2021), with information-aggregation underpinnings developed by Golub and Jackson (2010); Acemoglu *et al.* (2014). We show that the exposure matrix is one of two networks that compose multiplicatively in aggregate inflation expectations, with the bipartite forecaster-country information network being the second. The two networks together generate the cross-country dispersion and the calibration that we take to the data.

The third strand of literature is the one on geopolitical risk and narrative economics. Aggregate measures of geopolitical risk (Caldara and Iacoviello, 2022; Federle *et al.*, 2024) document that conflict shocks have non-negligible macroeconomic costs, but typically at the aggregate level and without separating the channels through which the costs propagate. The narrative economics tradition of

Shiller (2017) provides the conceptual framing for the public narrative bias term in our model. We identify the forecaster network channel separately from the standard economic exposure channels and from common risk co-movement, and we calibrate the model's narrative bias parameter directly from the within-pair empirical estimate.

Outline Section 2 introduces the CCF panel and identifies the conflict episodes. Section 3 establishes the three empirical patterns: post-conflict inflation expectations revise gradually upward following the shock; forecasters connected to the conflict country dampen their revisions relative to unconnected ones within country and date; and the dampening sharpens at country pairs with high bilateral exposure. Section 4 develops the two-network model and traces it through to cross-country dispersion, Phillips-curve pass-through, the implied inflation wedge, and the policy error that follows from reading the consensus at face value. Section 5 concludes.

2 Data

Our central data source is the Continuous Consensus Forecasts (CCF) from Consensus Economics, which spans the period 2012–2026 and covers 48 countries and 611 forecasting institutions with 216,513 forecast submissions in total.[3] The dataset is at the country–forecaster–date level and includes CPI inflation forecasts at three annual horizons (current year, 1 year ahead, 2 years ahead) and eight quarterly horizons (current quarter through 7 quarters ahead). Each "forecaster" in the panel is an institution—a global investment bank (e.g., Goldman Sachs, JP Morgan), a multinational corporate-economics team (e.g., Santander, BBVA), an international body (e.g., the Institute of International Finance), or a domestic research house—rather than a named individual analyst. For the purpose of the empirical analysis, we focus on the quarterly forecasts because they let us observe how each institution's view evolves along the forecast horizon, rather than collapsing the response into a single annual average. The forecasts are reported in real time on a continuous (or replacement) basis, so that each observation reflects an institution's view of country i's outlook at a specific submission date. The country coverage spans G7 and Western Europe, Eastern Europe, and Latin America.

We identify candidate conflict episodes from the Uppsala Conflict Data Program (UCDP) Georeferenced Event Dataset (Sundberg and Melander, 2013), which records geocoded violent events at the day level for state-based, non-state, and one-sided violence. We restrict the universe to events occurring during the years with CCF coverage and require an episode to satisfy at least one of the following two thresholds, both involving a CCF-covered country as a primary armed actor:

(i) **Battle-event threshold.** A discrete onset or sharp escalation, on the conflict country's own territory, defined as a month with at least 50 deaths recorded in UCDP for events occurring in the conflict country, and at least three times the rolling 24-month median of monthly deaths there.

[3] Appendix A-I reports the sources and variable definitions.

(ii) **Expeditionary military threshold.** The CCF country is the lead state-level armed actor in a discrete onset of a coalition or unilateral military intervention abroad, with at least 1,000 cumulative deaths recorded in UCDP across the operation's first thirty days. This admits episodes where the CCF country is unambiguously the source of the news flow but where the violence itself is geographically displaced.[4]

Episodes must additionally satisfy that (iii) the conflict country must be in our CCF dataset and have a sufficient stock of forecasters covering it in the year before the cutoff, and (iv) episodes must be separated by at least three months on each side of the cutoff to avoid contamination across event windows. Table 1 lists the eight episodes that satisfy these criteria. Seven qualify under threshold (i); the 2015 Saudi-led intervention in Yemen qualifies under threshold (ii).

All eight episodes carry supply-side content.[5] The conflict country is either a major energy producer (Russia and Saudi Arabia), a key pipeline-transit country (Ukraine), or located in a region whose energy supply is exposed to the conflict (the Israel-anchored episodes). The public narrative around each episode emphasizes upward inflation pressure on non-conflict economies through bilateral energy and intermediate-goods exposure to the conflict country. Appendix A-II shows that media coverage of each conflict-country pair does scale with that exposure. The realized energy-market response nonetheless varies considerably across episodes.[6] Yet, we do not screen episodes on that response. This is because the response is measured inside the estimation window and moves with the forecast revisions the paper seeks to explain, so conditioning on it would select episodes on an outcome rather than on a characteristic fixed before the shock.

Table 1: Conflict episodes in the sample

Episode	Cutoff date	Conflict country	Description
Crimea 2014	February 27, 2014	Russia, Ukraine	Russian seizure of Crimea
Gaza 2014	July 8, 2014	Israel	Operation Protective Edge
Saudi–Yemen 2015	March 26, 2015	Saudi Arabia	Saudi-led coalition Operation Decisive Storm
RU–UA 2022	February 24, 2022	Russia, Ukraine	Russian full-scale invasion of Ukraine
Hamas 2023	October 7, 2023	Israel	Hamas attack; start of Gaza war
Iran 2024	April 13, 2024	Israel	First direct Iranian missile strike on Israel
Hezbollah 2024	September 17, 2024	Israel	Pager attacks; Operation Northern Arrows
ME 2026	February 28, 2026	Israel	Middle East escalation

Notes: The table reports the episodes selected from the UCDP Georeferenced Event Dataset (events in 2012 or later) using the two-pronged criterion described in the text. The "Conflict country" column lists only the countries that are themselves covered by CCF forecasters.

[4] An example of this is the 2015 Saudi-led intervention in Yemen, where Saudi Arabia is the principal belligerent and Saudi-coverage forecasters retain their informational advantage even though most fatalities occur on Yemeni soil.

[5] The supply-side character matters for the sign of the prediction. A conflict with predominantly demand-side macroeconomic content would imply the opposite sign on bilateral exposure, so pooling the two types would offset rather than average. None of the eight episodes admitted by the criteria is of that type.

[6] Over the twenty trading days after each cutoff, Brent rises by 70 percent for ME 2026, 28 percent for RU–UA 2022 and 16 percent for Saudi–Yemen 2015, while the front-to-twelve-month futures spread widens by $27.4, $13.7 and $2.4 respectively, the signature of a near-term supply disruption. The European gas price (TTF, available from 2017) rises by 70 percent for ME 2026, 32 percent for RU–UA 2022, 26 percent for Hamas 2023 and 19 percent for Hezbollah 2024. Crimea 2014, Gaza 2014 and Iran 2024 show no positive response in any series available.

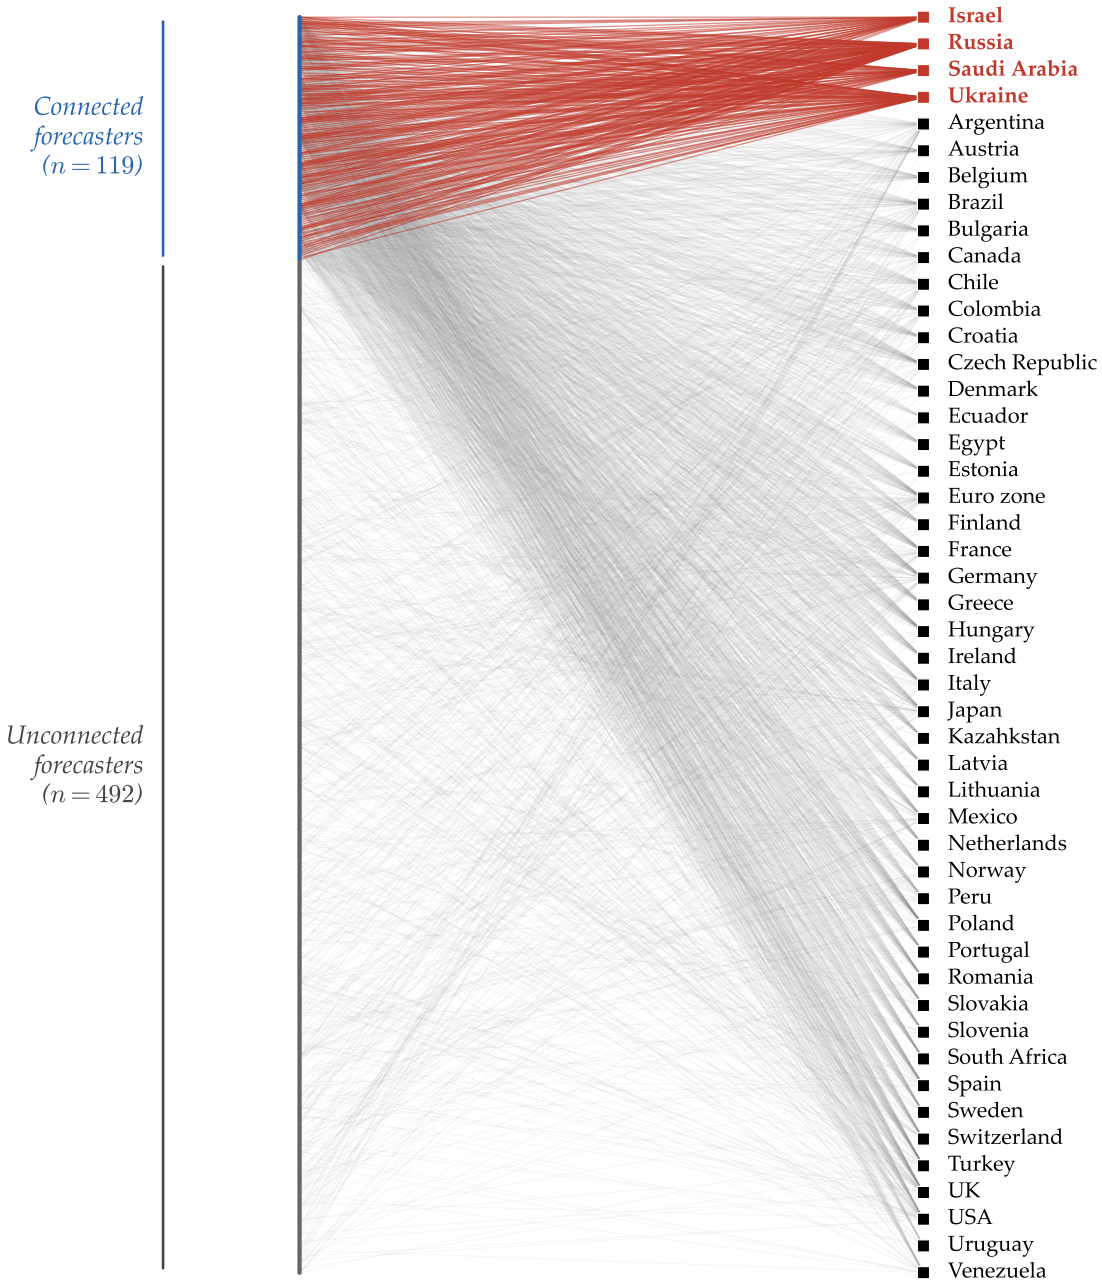
Using the CCF data and the identified conlfict episodes, we construct the bipartite forecaster–country network. The network is bipartite in that it has two types of nodes—forecasting institutions and countries—with edges running only between types and recording which institution covers which country. Figure 1 visualizes the network as institutions on the left and countries on the right, with an edge drawn whenever the institution submitted at least one forecast for that country during the sample. The network is sparse, with a realized density of 0.097, or roughly one in ten possible forecaster–country pairs. It is also stratified, as 119 institutions cover at least one of the conflict countries that anchor our episodes and form the connected group used in the identification, while the remaining 492 institutions cover only non-conflict countries.[7]

Because a forecaster in the panel is an institution rather than an individual analyst, the connectedness we exploit below rests on an intra-firm information-sharing assumption. When a multi-country institution that already covers the conflict country revises its forecast for a non-conflict country, we assume that its research department draws on knowledge held elsewhere in the firm, across country desks.[8]

[7] Among the 492 unconnected forecasters, 80 percent cover only one country—the panel's tail of local and national research houses; the maximum is 16. Connected forecasters have median coverage of 12 countries. A small set of large multi-country institutions sit on the unconnected side because their country mandate happens to skip Russia, Ukraine, Israel, and Saudi Arabia.

[8] This is, in our view, institutionally reasonable. Forecasting institutions in the CCF panel typically operate integrated research departments that share country macroeconomic views across geographic teams, through internal databases, model inputs, and regular cross-desk briefings; published institution-level reports likewise tend to take a unified house view across the country coverage.

Figure 1: Forecaster–country bipartite network in the CCF data



Notes: The figure plots the realized bipartite forecaster–country network. Forecasters (left, dots) are sorted by whether they cover at least one of the four conflict countries that anchor the eight episodes (Russia, Ukraine, Israel, Saudi Arabia): connected forecasters (blue, $n = 119$) cover at least one of these, while unconnected forecasters (gray, $n = 492$) cover none. Countries (right, squares) are sorted with the four conflict countries (red) at the top and the remaining 44 countries listed alphabetically below. An edge is drawn whenever a forecaster submitted at least one forecast for that country at any date in the sample (2,847 edges in total). Edges to a conflict country are drawn in red.

Table 2 reports descriptive statistics for the panel dataset.[9] Country coverage per institution is highly skewed (mean 4.7, standard deviation 9.4), spanning from single-country domestic research houses to global banks tracking dozens of economies. Every country is, however, well-represented in the cross-section. The mean country is followed by 59 institutions, supplying ample within-country–date variation across forecasters for our identification strategy. The panel is also persistent. The mean institution is active in 8.4 sample years, 62 percent of institutions have at least six active years, and the count of active institutions in any given year is stable around 344. For all empirical exercises we exclude countries with a mean current-year CPI inflation forecast at or above 10 percent in the analysis window, whose forecast dynamics are dominated by chronic inflation rather than by the conflict shock.

Table 2: Descriptive statistics for the CCF data

	Mean	SD
Network structure		
Countries per forecaster	4.7	9.4
Forecasters per country	59.3	19.5
Submissions per forecaster–country–year	10.5	7.7
Forecaster persistence		
Years active per forecaster	8.4	5.3
Share with 6+ years active	0.62	
Share first observed by 2015	0.72	
Share still active in 2026	0.49	
Yearly dynamics		
New forecasters entering per year	40.7	87.6
Active forecasters per year	343.7	21.7

Notes: The table reports the descriptive statistics computed on the 611-forecaster CCF panel covering 48 countries from March 2012 to March 2026 (216,513 forecast submissions in total). *Network structure* reports the distribution of countries per forecaster, forecasters per country, and submission frequency at the forecaster–country–year level; a forecaster–country pair is counted whenever the forecaster submitted at least one forecast for that country in the sample, generating 2,847 such pairs. *Forecaster persistence* measures activity by the number of distinct calendar years in which a forecaster submitted at least one forecast. *Yearly dynamics* reports the per-year flow of new entrants and the count of forecasters active in any given year.

We finally complement our data by merging in other datasets. First, we obtain conflict event data from UCDP for the timing of armed-conflict episodes. Second, we draw the bilateral economic linkages from CEPII's BACI (HS6 bilateral trade), the IEA (energy flows), and the OECD TiVA (production-network exposure). And third, we use GDELT for media coverage of conflict events.

[9] The dataset starts two years before the earliest conflict episode we study because the connectedness indicator is built from a two-year coverage lookback ending the day before each episode.

3 Empirical evidence on the information channel

This section presents our empirical findings. Subsection 3.1 provides evidence on forecasters' gradual learning based on their response in the aftermath of conflict onsets. Subsection 3.2 identifies the within-country–date information channel by comparing connected and unconnected forecasters submitting on the same day for the same country in the same episode. Subsection 3.3 tests whether the information differential varies with bilateral exposure linkages. Subsection 3.4 presents the results of the robustness checks.

3.1 Gradual updating of inflation expectations

To estimate the effect of conflicts on inflation expectations in the days following the onset of these episodes, we rely on a regression-discontinuity-in-time design (Davis and Kahn, 2008), which isolates the conflicts' effect on forecasts within a narrow window around the cutoff date. Identification rests on the assumption that, within a narrow window around the cutoff, unobserved factors affecting forecasts are smooth in time. We first present the evidence based on the Middle East escalation of February 2026. It is both the best-covered episode around its cutoff and the largest supply shock in the sample, so it is where a change in the daily trajectory is easiest to see. We then estimate the same design on all eight episodes pooled.

For each CPI inflation forecast horizon h we estimate:

$$y^h_{i,f,t} = \alpha + \beta(t - t^*) + \tau \mathbf{1}(t \geq t^*) + \delta(t - t^*)\mathbf{1}(t \geq t^*) + \gamma_i + \phi_f + \varepsilon_{i,f,t}, \tag{1}$$

where $y^h_{i,f,t}$ is the CPI inflation forecast at quarterly horizon h submitted by forecaster f for country i on (working) day t; t^* denotes the cutoff date February 28, 2026 (the onset of the war); γ_i are country fixed effects; and ϕ_f are forecaster fixed effects. The country fixed effects absorb persistent cross-country differences in the level of inflation forecasts, and the forecaster fixed effects absorb systematic level differences across institutions, so that τ and δ are identified from the within-(country, forecaster) time variation around t^*. The coefficient β captures the baseline time trend, τ captures the discrete level shift at the cutoff, and δ captures the change in the slope after the cutoff. Standard errors are clustered at the day level.

Table 3 reports the estimates of τ and δ for all horizons. Panel A shows that the level shift τ is essentially zero at most horizons. The general absence of immediate level shifts indicates that forecasters did not interpret the onset of the conflict as a one-time, fully-revealed inflation shock. The slope change δ is significant at virtually every horizon, indicating that the war's impact on inflation expectations accumulated progressively over the days following the onset. The slope is positive and significant from the current quarter through 4 quarters ahead, with the largest revisions at the 1- to 3-quarter horizons. At those horizons, forecasters revise their inflation forecasts upward by 0.03 to 0.04 percentage points per day, accumulating to about half a percentage point over the three weeks following the war onset. This gradual revision is consistent with forecasters updating their inflation outlook as the conflict persisted and its pass-through to prices became clearer.

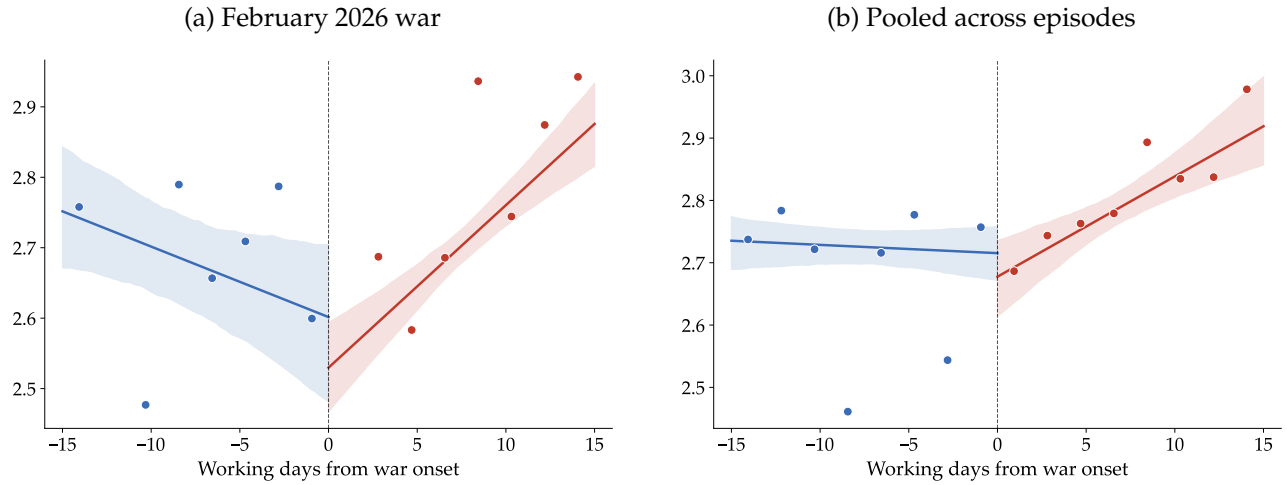
To assess whether this gradual upward revision is specific to the February 2026 episode or reflects a general feature of how professional forecasters react to conflict shocks, Panel B of Table 3 stacks the ± 15 working-day windows around the onsets of all eight conflict episodes and re-estimates Equation 1 with episode×country and episode×forecaster fixed effects. The episode-by-country fixed effects absorb country-level differences within each episode window, and the episode-by-forecaster fixed effects absorb forecaster-level differences within each episode window, so that τ and δ are identified from the within-(episode, country, forecaster) time variation around each cutoff. The pooled estimates confirm that the level shift τ remains close to zero at most horizons, while the slope change δ is positive and statistically significant from the current quarter through 5 quarters ahead, with the largest effects at the 1- to 4-quarter horizons. Figure 2 illustrates the patterns for the recent conflict and across episodes at the 4-quarter-ahead horizon, where the slope change peaks.

Table 3: Effect of conflicts on inflation expectations

	Forecast horizon							
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
Panel A: Middle East war (February 28, 2026)								
Level shift (τ)	0.142	0.208	−0.034	−0.016	−0.072	−0.124*	−0.276	0.045
	(0.164)	(0.217)	(0.122)	(0.119)	(0.078)	(0.068)	(0.260)	(0.110)
Slope change (δ)	0.018*	0.048***	0.057***	0.049***	0.033***	0.005	−0.020	0.001
	(0.010)	(0.017)	(0.015)	(0.013)	(0.009)	(0.006)	(0.019)	(0.008)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Forecaster FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	524	563	561	565	516	490	425	411
Countries	31	31	31	31	31	31	31	31
Forecasters	157	156	157	157	150	149	146	145
R^2	0.031	0.111	0.095	0.094	0.055	0.012	0.003	0.001
Panel B: Pooled across 8 conflict episodes								
Level shift (τ)	−0.088	0.079	0.006	0.002	0.031	−0.047	0.007	0.128
	(0.096)	(0.107)	(0.071)	(0.089)	(0.081)	(0.065)	(0.083)	(0.088)
Slope change (δ)	0.012*	0.056***	0.056***	0.048***	0.041***	0.017*	−0.003	−0.001
	(0.006)	(0.013)	(0.016)	(0.016)	(0.011)	(0.010)	(0.008)	(0.006)
Episode × country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Episode × forecaster FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,784	1,958	2,695	2,726	2,455	2,261	1,917	1,799
Countries	32	32	32	32	32	32	32	32
Forecasters	244	244	263	264	259	255	252	252
R^2	0.093	0.113	0.050	0.042	0.026	0.005	0.000	0.002

Notes: The table reports the results from a regression discontinuity around the onset of the conflict. The dependent variable is the forecast at the indicated horizon. Qh denotes h quarters ahead (Q0 = current quarter). The running variable is the working-day distance from each episode's cutoff. Each specification includes a baseline day trend. Panel A includes country and forecaster fixed effects; Panel B stacks the ± 15 working-day windows around the cutoffs of all eight episodes and replaces the country and forecaster fixed effects with episode×country and episode×forecaster fixed effects. In each window, countries with a mean current-year CPI inflation forecast at or above 10 percent are excluded. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Figure 2: Inflation expectations around conflict onset, 4 quarters ahead



Notes: The figure shows the binned means of the residualized inflation forecast at the 4-quarter-ahead horizon around the beginning of the conflict. The horizontal axis is in working days from the cutoff. The sample is restricted to a ±15 working-day window. Panel (a) shows the February 28, 2026 episode, with country and forecaster fixed effects absorbed. Panel (b) pools the same window across all conflict episodes, with episode×country and episode×forecaster fixed effects absorbed. Linear fits and bootstrapped 90 percent confidence bands are displayed on each side.

3.2 Network estimates of the information channel

We have established in Subsection 3.1 inflation forecasts for non-conflict countries revise gradually upward in the days following a conflict. The revision is consistent with two channels operating in parallel: aa standard economic channel, through which the conflict spills over to country j via bilateral linkages, and an information channel through which the revision depends on whether the forecaster making it also covers the conflict country.

The two channels yield observationally similar aggregate revisions, since both raise the average forecast for country j. Yet, they differ in how revisions vary across forecasters looking at j on the same date. Under the economic channel, the shock reaches j through its bilateral linkages and all forecasters covering j face the same news, so all should revise by a similar amount whether or not they also cover the conflict country. Under the information channel, forecasters who cover both the conflict country and j revise differently from those covering only j, because they are already processing the shock at its source. The direction, however, is not determined a priori. A forecaster with a direct read on the conflict country could revise by more, if the disruption is larger than the public account suggests, or by less, if it is smaller.

To empirically identify the role of the information channel, for each conflict episode e with cutoff date t_e^* and conflict country $i(e)$, and for each non-conflict country $j \neq i(e)$, we estimate the stacked regression:

$$\Delta y^h_{j,f,t,e} = \alpha_{j,t,e} + \beta \, \text{Connected}_{f,i(e)} + \varepsilon_{j,f,t,e}, \tag{2}$$

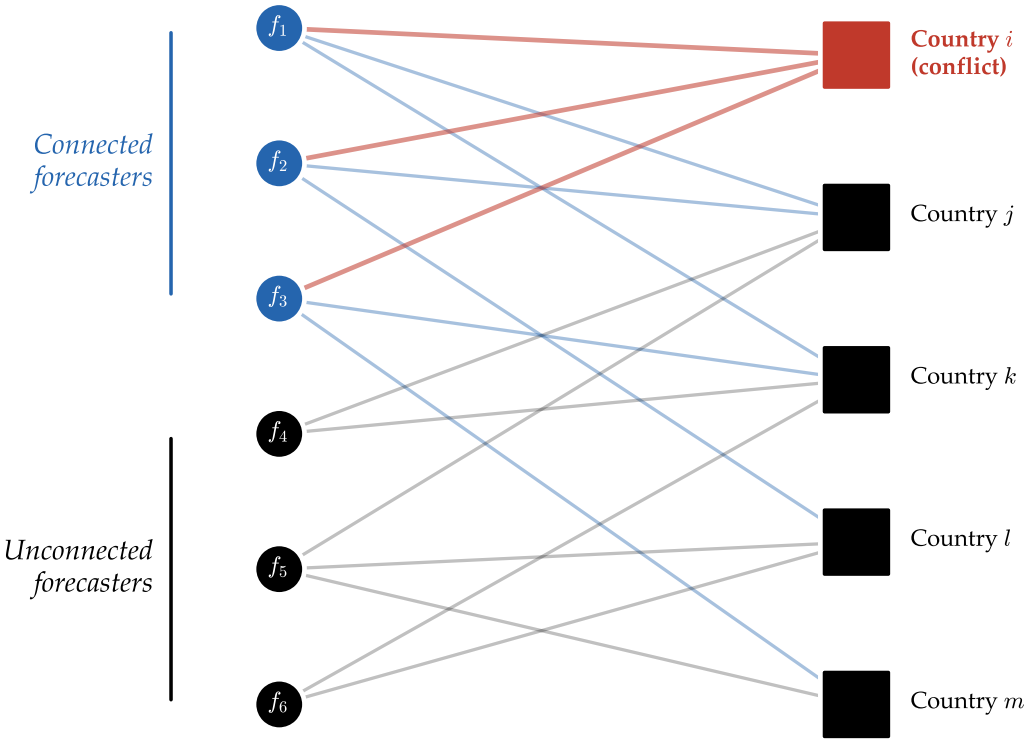
where the date subscript t refers to the calendar day on which forecaster f submitted the post-cutoff forecast. The unit of observation is therefore a (forecaster, country, submission date, episode) tuple, and the same forecaster–country pair contributes multiple rows to the regression—one for every distinct submission date in $[t_e^* + 1,\ t_e^* + 30]$.

The dependent variable $\Delta y^h_{j,f,t,e}$ is the difference between the post-cutoff forecast on date t and the same forecaster's most recent pre-cutoff forecast for the same country, restricted to be within seven days before t_e^* (a single pre-baseline per forecaster–country pair, common to all of that pair's post-cutoff rows). A short pre window (7 days) ensures the baseline reflects the forecaster's current view at the time of the shock. A forecaster–country pair contributes to the regression only if a pre-cutoff baseline exists. Each forecast that the same forecaster submitted for the same country in $[t^* + 1,\ t^* + 30]$ is a different observation. Using every post-cutoff submission within 30 days exploits the within-forecaster–country variation in revision timing while acknowledging that monthly Consensus update cycles often span more than a week.

Connected$f, i(e)$ is an indicator equal to one if forecaster f also covers the conflict country $i(e)$, defined as having submitted at least one forecast for it at any point during the two years before the cutoff. Critically, the fixed effects $\alpha j, t, e$ absorb any revision common to all forecasters updating country j on date t in episode e—whether driven by trade exposure, energy prices, or financial contagion—and so removes the economic channel. The coefficient β is therefore identified from cross-forecaster variation within each (j, t, e) group, and it compares revisions made for the same country on the same day by forecasters who do and do not cover the conflict country.[10] Under this identification strategy, because both groups are updating the same country on the same day, any difference between them comes from something other than the shock's economic transmission to j. Figure 3 illustrates the design. The identifying assumption is that, within a country–date–episode, connectedness is uncorrelated with any determinant of the size of the revision other than information about the conflict country; and that connected forecasters would not have revised differently for reasons unrelated to their coverage, as they would if institutions covering many countries were systematically different revisers.

[10] As an example, if Goldman Sachs, JPMorgan, and Morgan Stanley all submit Germany forecasts on March 8, 2022, the group $\alpha_{\text{DEU,,Mar 8 2022,,RU–UA 2022}}$ averages their three revisions, and β measures whether the forecasters in this group who also covered Russia revised differently from those who did not.

Figure 3: Identification of the information channel



Notes: The figure illustrates the identification of the information channel. Stylized bipartite network linking forecasters (left) to countries (right). Connected forecasters (blue) cover the conflict country and at least one non-conflict country. Unconnected forecasters (gray) cover only non-conflict countries. Red edges link to the conflict country. The identification compares, within the same non-conflict country on the same date, revisions by connected vs. unconnected forecasters.

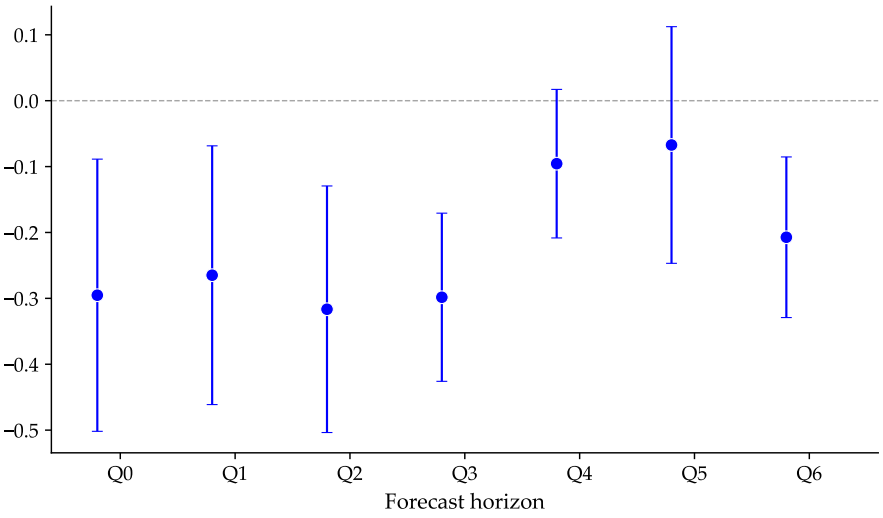
We now turn to the results. A priori, the sign of β is not pinned down. If connected forecasters react more aggressively to the conflict shock than unconnected forecasters, β should be positive for CPI inflation (assuming wars push inflation up). If connected forecasters react less aggressively—because they better understand the geographic limits of the shock—β should be negative. Figure 4 plots the stacked network coefficient $\hat{\beta}$ for CPI inflation across quarterly horizons, along with 90 percent confidence intervals.[11] The estimates point to the existence of the network information channel. Specifically, connected forecasters revise CPI inflation forecasts downward at the near-horizon quarters (Q0–Q3 in particular), relative to unconnected forecasters.

Since the underlying CPI inflation revisions are positive on average (as a result of the negative supply shock), the negative $\hat{\beta}$ implies that connected forecasters revised by a smaller amount in the same direction as unconnected forecasters. That is, connected forecasters dampen the inflation spillover predictions for non-conflict countries. Forecasters with direct knowledge of the conflict country appear to act as informational anchors, while unconnected forecasters over-extrapolate the

[11] Table A2 in Appendix A-III reports the underlying numerical estimates.

conflict's inflationary implications to non-conflict countries.[12]

Figure 4: Network coefficient across horizons



Notes: The figure plots the estimated network coefficient isolating the information channel. The dependent variable uses an asymmetric pre/post window (pre 7 days, post 30 days). Qh denotes h quarters ahead. The dots denote the coefficient estimates, the vertical lines denote the 90 percent confidence intervals from standard errors clustered at the forecaster level.

Figure 4 displays the regression results over quarterly horizons. However, not every forecaster constructs a full quarter-by-quarter path of inflation projections. Many report only annual averages, and those annual numbers move little within a month of a shock—updating an annual average is mechanically less responsive to a recent event than updating a near-term quarterly path. Table 4 reports the results from two diagnostic tests, suggesting that the information channel is identified from forecasters who produce a quarterly inflation path, not from those who only update an annual average.

The first diagnostic splits the sample on the granularity of the projection. We define "quarterly reporters" as forecasters who submit at least one near-horizon quarterly observation for the relevant country and episode—forecasters who construct an internally consistent quarter-by-quarter path of inflation, rather than only an annual average. The complementary group reports only annual horizons. The first three columns of Table 4 re-estimate the annual CPI inflation coefficient on each subsample. Among quarterly reporters, $\hat{\beta}$ on the annual horizons turns sharply negative, matching in magnitude the quarterly-horizon estimates in Figure 4. Among annual-only reporters, $\hat{\beta}$ is indistinguishable from zero. That is, the two groups generate the same observable forecast variable (an annual CPI inflation revision) but reveal the information channel only in the first.

[12] Estimated on the February 2026 escalation alone, the coefficient is -0.36 at the two-quarter horizon and -0.50 at three quarters, both statistically significant and larger in absolute value than the pooled estimates. The shorter horizons' coefficients cannot be estimated, since a single episode does not supply enough country–date cells containing both a connected and an unconnected forecaster.

The second diagnostic asks whether the per-horizon quarterly estimates are internally coherent enough to survive pooling. For each (forecaster, country, date) observation we compute the average revision across the four short quarterly horizons (Q0 through Q3) and re-estimate Equation 2 on this single aggregated outcome. The last two columns of Table 4 report the pooled coefficient. The CPI inflation estimate is sharply negative and statistically significant, and of the same magnitude as each individual quarterly point estimate—evidence that the dampening is consistent across quarterly horizons rather than the result of one or two horizons happening to align in sign. Extending the pool to all quarterly horizons preserves the result.

Read together, the two diagnostics deliver a sharper reading of the headline result. The information channel operates through forecasters who construct a quarter-by-quarter inflation path; their annual numbers, internally consistent with that path, inherit the same network coefficient. Forecasters who report only annual averages do not produce a near-term path that responds to a shock in our 30-day window, and their annual revisions are essentially flat. The relevant heterogeneity in the CCF panel is therefore not in the information-processing frequency of different forecasters but in the granularity of the projection they actually publish.

Table 4: Composition and pooled-horizon diagnostics for the network coefficient

	Annual vs quarterly forecasters			Pooling	
	Baseline forecasters	Quarterly reporters	Annual-only forecasters	Q0–Q3 average	Q0–Q7 average
CPI, current year	0.021	-0.151^{**}	0.026		
	(0.017)	(0.061)	(0.018)		
	[1,309]	[162]	[1,147]		
CPI, 1 year ahead	-0.002	-0.193^{***}	0.015		
	(0.025)	(0.061)	(0.024)		
	[1,126]	[120]	[1,006]		
CPI, pooled near-horizon				-0.292^{***}	-0.212^{***}
				(0.065)	(0.059)
				[248]	[258]
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the results of composition and pooled-horizon diagnostics for the network coefficient. The first three columns re-estimate each annual horizon on three samples: the baseline (all forecasters); "quarterly reporters", defined as forecasters who report at least one Q0–Q3 observation for that country in that episode; and the complementary subsample of annual-only forecasters. The last two columns replace the dependent variable with the (forecaster, country, date) average of revisions across Q0–Q3 and Q0–Q7 respectively. Coefficients are reported with cluster-robust standard errors in parentheses and sample sizes in square brackets. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

3.3 Heterogeneity by exposure

Figure 4 reports the estimates of the average information channel across all forecaster–country pairs. The mechanism we have in mind—a public narrative whose intensity scales with bilateral exposure—makes a sharper prediction, which is that the gap between connected and unconnected forecasters should be larger for non-conflict countries with greater bilateral exposure to the conflict country.

The idea is that headlines portraying war in country i as an economic catastrophe scale with bilateral linkages—a country importing 50 percent of its gas from i generates a more dramatic public narrative than one importing 5 percent. Unconnected forecasters, who rely on this public narrative, should over-extrapolate exactly where the headline is most pronounced. Connected forecasters, who use granular information about the actual disruption in i, should instead temper the extrapolation by more wherever the public-narrative-driven baseline is most aggressive. The information channel should therefore be concentrated at high-exposure pairs.

To test this, we extend Equation 2 as:

$$\Delta y^{h}_{j,f,t,e} = \alpha_{j,t,e} + \beta_0 \,\text{Connected}_{f,i(e)} + \beta_1 \,\text{Connected}_{f,i(e)} \times \text{HighExp}_{j,e} + \varepsilon_{j,f,t,e}. \tag{3}$$

where $\text{HighExp}_{j,e}$ is an indicator equal to one if country j's bilateral exposure to the conflict country in episode e is above the within-episode median (the conflict country $i(e)$ is implicit since it is pinned by the episode). We construct three complementary exposure measures, all computed at the (country j, episode e) level using bilateral flows in the calendar year before the cutoff.[13] The first is the energy share, defined as j's primary energy imports sourced from $i(e)$, scaled by j's total energy imports. The second is the trade share, which corresponds to j's total goods imports sourced from $i(e)$, scaled by j's total goods imports. And the last is the value-added share, which is the foreign value added consumed in j's final demand that originates in $i(e)$, scaled by total foreign value added in j's final demand. Captures bilateral exposure through global value chains rather than only direct gross-trade flows.

We use the within-episode median split for the energy and value-added shares, and a fixed absolute threshold of 2 percent for the trade share.[14] The $\text{HighExp}_{j,e}$ indicator is constant within each (j, e) pair and is therefore absorbed by the episode×country×week fixed effect, where the week is the week of the submission date. Hence, only the interaction is identified. The coefficient β_0 measures the information channel at low-exposure pairs, while $\beta_0 + \beta_1$ measures it at high-exposure pairs. The information channel predicts that both coefficients are negative. That is, connected forecasters dampen their non-conflict-country revisions at low-exposure pairs, and the dampening is even larger in absolute value at high-exposure pairs.

Table 5 reports the regression results for the three exposure measures on the pooled CPI inflation Q0–Q7 revisions. Low-exposure pairs show the same baseline dampening already documented in Figure 4. High-exposure pairs show roughly three times as much dampening at the energy margin, with non-overlapping confidence intervals. The differential is also statistically significant at the value-added margin—production-network exposure, which transmits supply-side conflict shocks through intermediate-input prices. Composite bilateral goods trade, which averages supply-side

[13] When the calendar year before the cutoff falls outside the data coverage we fall back to the most recent available year. For multi-conflict episodes ($|I(e)| > 1$, i.e., Crimea 2014 and Russia–Ukraine 2022) the bilateral numerators are summed across $i \in I(e)$.

[14] The bilateral total-trade-share distribution is highly skewed, with most country pairs sitting well below 1 percent, so the within-episode median falls in that near-zero tail and a median split puts both "low" and "high" on essentially the same region. We use a fixed 2 percent threshold instead, which sits in the middle of a stable range (1.5 to 4 percent) over which the estimates do not move materially, and corresponds to a non-trivial bilateral trade dependence.

and demand-side sub-categories as we show below, shows no significant differential.

Table 5: Information channel heterogeneity by bilateral exposure

	Energy share (1)	Trade share (2)	Value-added share (3)
Connected ($\hat{\beta}_0$)	-0.161^{***}	-0.163^{***}	-0.116^{**}
	(0.050)	(0.040)	(0.054)
Connected × HighExp ($\hat{\beta}_1$)	-0.305^{***}	-0.150	-0.130^{*}
	(0.065)	(0.111)	(0.076)
Episode × country × week FE	Yes	Yes	Yes
Observations (N)	164	235	235
R^2 (within)	0.072	0.042	0.043

Notes: The table reports the information channel estimates for the heterogeneity by bilateral exposure. The dependent variable is the (forecaster, country, date) average CPI inflation revision across horizons. HighExp is the within-episode median split for each share. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Table 6 re-estimates the same regression after splitting the energy and trade margins into the sub-shares: oil and gas; and intermediate, capital, and consumption goods (classified at the BEC Rev. 5 broad-category level). The decomposition resolves the composite-trade null and sharpens the supply-vs-demand reading. On the energy side, both oil and gas exposures generate a sizable additional dampening at high-exposure pairs, comparable in magnitude to the composite energy estimate; both are supply-side channels through which a disruption in the conflict country raises imported input costs in the non-conflict country.

On the trade side, the heterogeneity loads on the intermediate-goods margin only, consistent with intermediate inputs being the segment of trade that most directly transmits supply-side conflict shocks to non-conflict countries' production costs. Capital and consumption goods, which are closer to demand-side channels with ambiguous or negative inflation transmission, show no statistically significant differential. With no clear inflationary direction for those channels, the public narrative has nothing salient and unambiguously inflationary to overstate.

These results clarify the average estimates in Figure 4 and pin the information channel to the supply-side interpretation of the conflict shocks in our sample. The dampening differential loads where the bilateral narrative points unambiguously toward higher inflation—oil and gas dependence and intermediate-goods supply-chain exposure—and is absent where the underlying economic transmission is ambiguous or runs in the opposite direction (capital and consumption goods). We conclude that the dampening loads precisely where supply-side narrative bias predicts it should, and is null where it should not.

Table 6: Bilateral exposure decomposition

	Energy		Trade		
	Oil share (1)	Gas share (2)	Interm. share (3)	Capital share (4)	Cons. share (5)
Connected ($\hat{\beta}_0$)	-0.161^{***}	-0.217^{***}	-0.094^{*}	-0.192^{***}	-0.082
	(0.050)	(0.047)	(0.056)	(0.044)	(0.096)
Connected × HighExp ($\hat{\beta}_1$)	-0.305^{***}	-0.281^{***}	-0.140^{*}	0.006	-0.165
	(0.065)	(0.066)	(0.074)	(0.067)	(0.117)
Episode × country × week FE	Yes	Yes	Yes	Yes	Yes
Observations (N)	164	164	235	235	235

Notes: The table reports the information channel estimates for the heterogeneity by bilateral exposure, using sub-components of the exposure metrics. Energy is split into oil and gas (both bilateral imports as a share of the importer's total energy imports). Trade is split at the BEC Rev. 5 broad-category level into intermediate, capital, and consumption goods, classified from BACI HS6 codes and divided by total goods imports. HighExp is the within-episode median split for each share. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

3.4 Robustness

The dampening effect documented in Subsection 3.2 is subject to some measurement and identification challenges. The first challenge is episode selection. We therefore examine the robustness of the results to alternative episode samples. Table A3 in Appendix A-III re-estimates Equation 2 on a more restrictive seven-episode sample that applies the on-territory battle-event threshold only (dropping the 2015 Saudi-led intervention in Yemen) and on a more permissive fourteen-episode sample that relaxes both thresholds and admits wartime sovereign-authority shocks (the September 2022 Russian partial-mobilization decree, the July 2016 Turkish coup attempt, several smaller Israel–Gaza flare-ups, and Turkey's 2019 cross-border operation in Syria). The dampening effect at the near-horizon quarters is preserved across the three samples, while the magnitude attenuates somewhat in the aggressive sample because the additional events generated smaller forecaster revisions on average. Table A4 repeats the heterogeneity-by-exposure regression of Equation 3 on the same three samples. The energy-margin differential remains negative and significant across all three samples, and the value-added differential survives in the restrictive and headline samples.

Second, we turn to the construction of the forecast revision. Table A5 in Appendix A-III re-estimates Equation 2 with a single observation per forecaster–country pair (the closest post-cutoff submission within 30 days) rather than stacking every post-cutoff submission. Sample sizes are roughly two-thirds smaller and standard errors widen accordingly, but the sign and approximate magnitude of the coefficient of interest track the baseline estimates closely. The baseline defines a forecaster as "connected" if they submitted at least one forecast for the conflict country in the two-year lookback. As shown in columns (2) and (3) of Table A6, tightening this to at least five submissions and then to twenty leaves the coefficient at the peak horizon unchanged or even larger in

absolute value.[15]

Third, connectedness may be proxying for the breadth of an institution's coverage rather than for information about the conflict country. Covering a conflict country implies a mandate beyond a single domestic market, so connected forecasters are mechanically broader (80 percent of unconnected institutions cover one country, while the median connected institution covers twelve). Columns (4) and (5) of Table A6 restrict the sample to forecasters covering at least five and then at least ten countries—which holds coverage breadth roughly constant across the connected and unconnected groups—and leaves the coefficient at the peak horizon unchanged or larger in absolute value.[16]

Fourth, the gap may be a standing difference in revision style rather than something the conflict opens. Two placebo tests speak to it. If the dampening genuinely reflects an information channel that opens at the moment of the conflict, it should disappear at dates when no shock occurs and at events that are equally salient to all forecasters regardless of coverage. Thus, the first test shifts each of the eight conflict cutoffs backward by 90 calendar days—snapping to the closest active date in the CCF panel if the naive shift lands on a weekend or holiday—and holds the connectedness assignment fixed. The placebo cutoff is a non-event date at which no conflict shock occurs. The second test uses pre-scheduled real events in the conflict countries—presidential elections in Russia (2018, 2024) and Israeli legislative elections (2015, 2019$\times$2, 2021, 2022)—events that are anticipated months in advance and equally salient to all forecasters, so that covering the country confers no informational advantage.

Table 7 reports the results. The baseline coefficient on Connected is negative and significant at every near-horizon quarter (Q0–Q3), ranging from -0.27 to -0.32 percentage points. At shifted cutoffs the coefficient is small and statistically indistinguishable from zero at every horizon, with mixed signs. At anticipated events the coefficient is also indistinguishable from zero at every horizon, with mostly positive point estimates. Neither placebo reproduces the negative dampening pattern. This is consistent with the information-channel interpretation in which connected forecasters dampen their non-conflict-country revisions only when the conflict country is the source of an active, unanticipated shock—not when nothing is happening, and not when the event is one everyone could see coming. The shifted-cutoff test also speaks to the identifying assumption of Subsection 3.2. Because it holds the connectedness assignment fixed and moves only the date, a standing difference in revision style between connected and unconnected institutions would have produced the same gap at the placebo cutoffs. It does not, which is difficult to reconcile with an explanation resting on firm scale or forecasting style rather than on information.

[15] The results are largely unchanged also for other horizons.

[16] At the shortest horizons the estimate weakens, and at the current quarter it cannot be estimated at all, because too few country–date cells retain both a connected and an unconnected forecaster once the restriction is imposed.

Table 7: Placebo tests of the information channel

	Forecast horizon						
	Q0	Q1	Q2	Q3	Q4	Q5	Q6
Headline (8 conflict episodes)	-0.295^{**}	-0.265^{**}	-0.317^{***}	-0.298^{***}	-0.096	-0.067	-0.207^{***}
	(0.126)	(0.119)	(0.114)	(0.078)	(0.069)	(0.109)	(0.074)
N	101	164	181	195	137	104	56
Shifted cutoffs ($t^* - 90$ days)	0.027	0.058	-0.056	0.032	0.098	0.080	-0.066
	(0.130)	(0.078)	(0.078)	(0.064)	(0.118)	(0.118)	(0.052)
N	50	115	149	206	211	203	162
Anticipated events	-0.145	0.012	0.073	0.055	0.090	-0.009	0.068
	(0.133)	(0.077)	(0.044)	(0.070)	(0.058)	(0.046)	(0.050)
N	105	170	166	182	164	137	141
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the placebo estimates for the information channel. "Shifted cutoffs ($t^* - 90$ days)" takes each of the baseline cutoffs and shifts them backward by 90 calendar days—snapping to the closest active date in the CCF panel if the naive shift lands on a weekend or holiday—and holds the connectedness assignment fixed (connectedness is computed from the two-year lookback before the original conflict cutoff). "Anticipated events" uses seven pre-scheduled events in the conflict countries: Russian presidential elections of March 2018 and March 2024, and Israeli legislative elections of March 2015, April and September 2019, March 2021, and November 2022. Qh denotes h quarters ahead. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

4 A two-network model of inflation expectation formation

This section presents a model that rationalizes the three empirical facts of Section 3. First, post-conflict inflation expectations revise gradually upward after the shock. Second, within country and date, forecasters who already cover the conflict country dampen their non-conflict-country revisions relative to forecasters who do not. And third, this dampening differential is larger at country pairs where bilateral exposure to the conflict country is highest.

We build a model with two networks.[17] A bilateral exposure network—energy, trade, production linkages—scales the size of cross-country spillovers. A bipartite information network governs who can observe a precise private signal about the conflict shock, with forecasters outside the network relying exclusively on a public narrative. The two networks are independent primitives of the economy. A running example may help. Consider the day Russia invaded Ukraine in February 2022 and a forecaster issuing a Hungary CPI inflation forecast a few days later. Two forecasters in our panel, say Goldman Sachs and a small local Hungarian research firm, both produce Hungary forecasts. Goldman has covered Russia for years (i.e., it is connected to Russia in the bipartite information network); the local firm has not and therefore is unconnected. Both read the same financial news, which emphasizes Hungary's deep gas dependence on Russia and projects severe inflation for Hungary. Goldman Sachs, however, also has its own real-time Russia desk that tells it more precisely how disrupted Russia's gas exports actually are. The model tracks how each forecaster combines these two information

[17] The formal derivations behind the predictions, propositions, and corollaries below are collected in Appendix A-IV.

streams, why their post-cutoff Hungary forecasts will differ, and how those individual differences add up to a wedge in aggregate Hungarian inflation expectations.

4.1 Setup

The economy has two networks. The exposure network E assigns a bilateral weight $x_{ij} \in [0,1]$ to every ordered pair (i,j) of countries, which are the share of j's primary energy imports sourced from i, the share of j's intermediate-goods imports sourced from i, and so on. We treat E as an observable primitive of the international economy. The information network $G_e \subset F \times N$ is bipartite, linking forecasters F to countries N, and indexed by episode e. In this network, an edge $(f,i) \in G_e$ means forecaster f has covered country i in the lookback window prior to the cutoff. The episode-specific connectedness indicator is

$$z_{f,e} \;=\; \mathbf{1}\{(f, i(e)) \in G_e\},$$

where $i(e)$ is the conflict country in episode e. We will refer to forecasters with $z_{f,e} = 1$ as connected and to those with $z_{f,e} = 0$ as unconnected.

A conflict shock $\xi_e \sim \mathcal{N}(0, \sigma_\xi^2)$ realizes at the cutoff date t_e^*. The true cross-country spillover from $i(e)$ to a non-conflict country j is

$$\Delta y_{j,e} \;=\; \gamma_{ij}\,\xi_e \;+\; u_{j,e}, \qquad \gamma_{ij} \;=\; \bar{\gamma} + \lambda\, x_{ij}, \tag{4}$$

with $u_{j,e} \sim \mathcal{N}(0, \sigma_u^2)$ idiosyncratic and $\lambda > 0$, so that countries with greater bilateral exposure inherit larger spillovers.[18]

Two signal streams arrive on each day t after the cutoff. Both convey information about the shock ξ_e, but they differ both in who can see them and in what they say. The first is a *public narrative* that every forecaster observes:

$$s_t^{\text{pub}} \;=\; \xi_e \;+\; b\, x_{ij} \;+\; \eta_t^{\text{pub}}, \qquad \eta_t^{\text{pub}} \sim \mathcal{N}(0, \sigma^2_{\text{pub}}), \tag{5}$$

where $b > 0$ is the narrative-bias slope, which governs how strongly the public signal overstates the spillover in proportion to that exposure. The second is a *private granular* signal that only connected forecasters observe:

$$s_t^{\text{priv}} \;=\; \xi_e \;+\; \eta_t^{\text{priv}}, \qquad \eta_t^{\text{priv}} \sim \mathcal{N}(0, \sigma^2_{\text{priv}}), \tag{6}$$

with $\sigma^2_{\text{priv}} < \sigma^2_{\text{pub}}$.

The difference between Equation 5 and Equation 6 encodes two ideas. The first idea is that the public narrative is biased such that it overstates how much the conflict matters for country j in proportion to its bilateral exposure to the conflict country.[19] This is consistent with the notion that the

[18] Equation 4 is the standard linear-in-exposure transmission rule. It could be derived from any of several structures (input–output spillover, financial-flow propagation, monetary-union pass-through). The model is agnostic about which.

[19] The narrative interpretation we adopt is the one developed by Shiller (2017), formalized in noise-and-bias terms by

narrative is salient where bilateral linkages are salient, and the salience of the narrative outruns the actual size of the spillover.[20] Outlets emphasize the bilateral links most salient to their audience, and empirical coverage of cross-border events scales with the underlying magnitude of the link.[21] The over-extrapolation tendency of public-narrative-driven baselines is documented by Coibion and Gorodnichenko (2015); Bordalo *et al.* (2020), and rational learning aggregated along network primitives is studied by Golub and Jackson (2010); Acemoglu *et al.* (2014). The second idea is that the private signal is unbiased but rationed. That is, only forecasters who already cover the conflict country—for reasons unrelated to the current shock, namely employer mandate—can read the granular country-level data that allows them to back out ξ_e directly. The bipartite network G_e encodes who has access.

A forecaster's information set after t days is

$$\mathcal{I}_t^f \;=\; \{s_1^{\text{pub}}, \ldots, s_t^{\text{pub}}\} \;\cup\; z_{f,e} \cdot \{s_1^{\text{priv}}, \ldots, s_t^{\text{priv}}\}.$$

where connected forecasters condition on both streams, and unconnected forecasters condition on the public stream only. Both update Bayesian-rationally and report their posterior mean.

4.2 Forecaster expectations

The forecaster's problem is a standard Gaussian filtering exercise. Define cumulative signal averages $\bar{s}_t^{\text{pub}} = t^{-1}\sum_{\tau\le t} s_\tau^{\text{pub}}$ and $\bar{s}_t^{\text{priv}} = t^{-1}\sum_{\tau\le t} s_\tau^{\text{priv}}$. The posterior mean of ξ_e given $\mathcal{I}_t^f$ is the precision-weighted average of the prior and the realized signals:

$$\mathbb{E}[\xi_e \mid \mathcal{I}_t^U] \;=\; w_t^U\, \bar{s}_t^{\text{pub}}, \tag{7}$$

$$\mathbb{E}[\xi_e \mid \mathcal{I}_t^C] \;=\; w_t^{C,\text{pub}}\, \bar{s}_t^{\text{pub}} \;+\; w_t^{C,\text{priv}}\, \bar{s}_t^{\text{priv}}, \tag{8}$$

where the Kalman weights satisfy

$$w_t^U \;=\; \frac{t/\sigma^2_{\text{pub}}}{1/\sigma^2_\xi + t/\sigma^2_{\text{pub}}}, \qquad w_t^{C,\text{pub}} + w_t^{C,\text{priv}} \;=\; \frac{t/\sigma^2_{\text{pub}} + t/\sigma^2_{\text{priv}}}{1/\sigma^2_\xi + t/\sigma^2_{\text{pub}} + t/\sigma^2_{\text{priv}}},$$

with the connected forecaster placing relatively more weight on the more precise private signal.

Each posterior is the precision-weighted average of the prior (mean zero, precision $1/\sigma^2_\xi$) and the accumulated signals (precision t/σ^2_{pub} for the public stream and, for connected forecasters, t/σ^2_{priv}

Eusepi and Preston (2011); Andre *et al.* (2022).

[20] The modeling choice that the bias scales with bilateral exposure is the natural reduced form of the salience-and-audience logic of the media-economics literature (Mullainathan and Shleifer, 2005; Eisensee and Strömberg, 2007).

[21] The bias is an empirical regularity in media coverage. Headlines about a conflict in country i emphasize the link to country j exactly where the link is most visible: Germany's gas dependence on Russia, Central Europe's reliance on Russian intermediate inputs, the energy-importing periphery of an energy-exporting conflict country. As a direct test in our setting, the log of GDELT mentions of each (conflict country, non-conflict country) pair across the eight episode windows is positively and significantly related to the log of the bilateral trade share, with the slope positive in every episode individually. The relationship is large in size: a country with twice the bilateral trade share to the conflict country receives roughly fifty percent more media mentions of the link. Appendix A-II reports these regressions.

for the private stream). This is in line with the standard normal-normal updating used in the noisy-information literature (Sims, 2003; Maćkowiak and Wiederholt, 2009; Coibion and Gorodnichenko, 2015). The Kalman weight w_t^U is bounded strictly below one for every finite t, because a posterior anchored to a finite-variance prior never collapses entirely onto the observed signal.

Two intuitions for the comparative statics of w_t^U are worth spelling out. First, a tighter prior over the conflict shock (smaller σ_ξ^2) shrinks w_t^U further. This means that a forecaster who a priori believes that conflict shocks rarely cause large cross-country spillovers discounts a strong public signal claiming otherwise, because the prior dominates the noisy data. Second, the noisier the public stream relative to the prior (larger σ_{pub}^2), the smaller w_t^U at any given t. Thus, a noisy public narrative is also discounted. As $t \to \infty$, $w_t^U \to 1$, while the connected forecaster's public-stream weight $w_t^{C,\text{pub}}$ converges to $\sigma_{\text{pub}}^{-2}/(\sigma_{\text{pub}}^{-2} + \sigma_{\text{priv}}^{-2}) < 1$ because the connected forecaster always allocates positive weight to the more precise private stream. The private signal is therefore a substitute for the public one in the connected forecaster's filter, and a closer substitute the larger the precision ratio $\sigma_{\text{pub}}^2/\sigma_{\text{priv}}^2$.[22]

From these expressions, three predictions follow that map directly to the three empirical patterns of Section 3:

Prediction 1 (gradual revision). For each forecaster type $\tau \in \{C, U\}$, the expected spillover $\gamma_{ij}\,\mathbb{E}[\xi_e \mid \mathcal{I}_t^\tau]$ has level zero at the cutoff ($t = 0$) and an interior positive slope thereafter as t grows. This is because at $t = 0$ no signals have arrived, the posterior equals the prior (mean zero), and the expected spillover therefore equals zero. As signals accumulate, the posterior weight on data grows linearly in t, and the expected spillover walks upward at a rate proportional to $\gamma_{ij} \cdot dw_t/dt$.

This matches the RDD pattern in Table 3. The level shift τ is essentially zero at the cutoff, while the slope δ is positive and statistically significant from the current quarter through 4–5 quarters ahead. Forecasters do not interpret the conflict as a fully revealed shock to be priced in instantly, rather they walk their expectations upward as information accumulates. **Prediction 2 (information channel).** The within-pair difference between connected and unconnected forecasters' expected spillovers is non-positive:

$$\beta_t \;\equiv\; \mathbb{E}^C[\Delta y_{j,e} \mid \mathcal{I}_t] - \mathbb{E}^U[\Delta y_{j,e} \mid \mathcal{I}_t] \;=\; -\,\gamma_{ij}\, w_t^U\, b\, x_{ij} \;\le\; 0, \tag{9}$$

with equality only in the degenerate cases $b = 0$ or $x_{ij} = 0$.

The intuition is the following. The public narrative overstates the spillover by bx_{ij}, and the connected forecaster nets out the bias because she has a second, unbiased signal of ξ_e. The unconnected forecaster has no such reference and absorbs the bias into her posterior with weight w_t^U. The cross-section difference is therefore the spillover loading γ_{ij}, multiplied by the bias term bx_{ij}, multiplied

[22] Both forecaster types' posterior expectations of ξ_e are shrunk toward the prior mean of zero by the standard Bayesian-shrinkage logic: connected forecasters have $w_t^{C,\text{pub}} + w_t^{C,\text{priv}} < 1$ and so the absolute level of $\mathbb{E}^C[\xi_e|\mathcal{I}_t]$ is an attenuated version of the true ξ_e. This is a property of any normal-normal Bayesian update with a finite-variance prior and is the same setting used in Lucas (1972); Morris and Shin (2002), where agents combine prior beliefs with one or more imperfectly informative signals. The within-(forecaster, country, date) regression in Equation 2 absorbs the common shrinkage that connected and unconnected forecasters share, through the episode×country×date fixed effect; the coefficient β is identified from the differential weight that the two groups place on the biased public signal, not from the absolute level of either group's expectation.

by the Kalman weight the unconnected forecaster places on the biased public signal. This is the dampening estimate documented in Figure 4 for which connected forecasters revise inflation in non-conflict countries upward by less than unconnected forecasters do. The model recovers the sign and the within-pair logic without further structure. **Prediction 3 (heterogeneity by exposure).** The information-channel differential is more negative at higher-exposure pairs:

$$\beta_{1,t} \equiv \frac{\partial \beta_t}{\partial x_{ij}} = -b\, w_t^U\, (\bar{\gamma} + 2\lambda x_{ij}) < 0. \tag{10}$$

The bias term bx_{ij} in Equation 5 is itself increasing in exposure. So is the spillover loading γ_{ij}. Both forces pull the connected-vs-unconnected gap in the same direction at high-exposure pairs. Equation 10 reproduces the heterogeneity result of Table 5, for which the dampening differential $\hat{\beta}_1$ is large and statistically significant at high-exposure margins. The components decomposition of Table 6 indicates that the differential loads on the energy and intermediate-goods margins. The model interprets this as b being largest precisely at margins where the public narrative is most easily organized around bilateral linkages—oil and gas dependence is a media staple, intermediate-input dependence is the natural object of supply-chain reporting—while capital- and consumption-goods margins generate a weaker b and therefore a weaker differential. Predictions 1–3 absorb the empirical content of Section 3. A single Bayesian filter operating on the bipartite information network G delivers, jointly, the gradual upward revision of post-conflict expectations (Prediction 1, matching Figure 4 and Table 3), the within-pair dampening of connected forecasters (Prediction 2, matching Table A2), and the concentration of the differential at high-exposure pairs with a precise pattern across components (Prediction 3, matching Table 5 and Table 6). The two-network primitive does the work jointly: exposure E supplies the bias term bx_{ij}, information G supplies the connectedness indicator $z_{f,e}$. Neither network alone would generate the patterns.

4.3 Aggregation and macroeconomic implications

We move from individual forecasters to the consensus, and from there to aggregate inflation. Let $\rho_{j,e} \in [0,1]$ be the share of forecasters covering country j in episode e who are also connected to the conflict country $i(e)$ in G_e. This share is a primitive of the bipartite network and is directly observable in our data. The consensus expectation of the spillover for country j in episode e, after t days of post-cutoff signals, is the weighted average of the two types' posterior means:

$$\mathbb{E}^{\text{cons}}_{e,t}[\Delta y_{j,e}] = \rho_{j,e}\, \mathbb{E}^{C}_{t}[\Delta y_{j,e}] + (1-\rho_{j,e})\, \mathbb{E}^{U}_{t}[\Delta y_{j,e}]. \tag{11}$$

Taking expectations over signal realizations in the long-t limit (where $w_t^U \to 1$ and $w_t^{C,\text{pub}} \to w_\infty^{C,\text{pub}} < 1$) yields a clean closed form for the aggregate bias relative to the rational-expectations benchmark $\gamma_{ij}\xi_e$:

$$\mathbb{E}\big[\mathbb{E}^{\text{cons}}_{e,\infty}[\Delta y_{j,e}] - \gamma_{ij}\xi_e\big] = (1-\rho_{j,e})\ \gamma_{ij}\ b\, x_{ij}\ w_\infty^U. \tag{12}$$

Proposition 1 (network mediation). Aggregate inflation expectation bias factors as the product of two distinct network primitives: bilateral exposure x_{ij} in the exposure network E, and the complement of the connectedness share $1 - \rho_{j,e}$ in the information network G.

Equation 12 is the central macroeconomic statement of the model. A conflict shock translates into biased aggregate inflation expectations through the joint operation of two networks. Bilateral exposure alone is insufficient: when $\rho_{j,e} = 1$ (every forecaster covering j also covers i), even high-exposure countries have unbiased aggregate forecasts, because the granular private signal absorbs the narrative bias before it reaches the consensus. The bipartite information network alone is also insufficient: when $x_{ij} = 0$ (no bilateral linkage), the public narrative carries no bias term to magnify and the consensus is unbiased regardless of ρ. The two networks compose multiplicatively, mirroring—for the macro-expectations primitive—the multiplicative structure of input–output spillover networks formalized by Acemoglu *et al.* (2012) and Carvalho and Tahbaz-Salehi (2019).

The wedge in Equation 12 is the long-run steady-state value. Along the time window after the cutoff, the unconnected forecaster's Kalman weight w_t^U grows monotonically from zero toward w_∞^U, so the wedge builds gradually toward this long-run value. These are the same dynamics generate Prediction 1. Returning to our running example: Hungary's post-invasion consensus inflation forecast is biased upward by an amount proportional to (1) Hungary's gas dependence on Russia and (2) the share of Hungary-covering forecasters who do not cover Russia. If most of Hungary's forecasters happen to be Russia-covering global houses like Goldman (high ρ), the consensus is close to truth. If they are mostly local Hungarian firms with no Russia coverage (low ρ), the consensus inherits the public-narrative bias.

The remaining results follow from Equation 12 as comparative statics. The first concerns cross-country dispersion. Conditional on a global shock realization, the cross-country variance of consensus inflation expectations decomposes as

$$\mathrm{Var}_j\big(\mathbb{E}^{\mathrm{cons}}_{e,\infty}[\Delta y_{j,e}] \,\big|\, \xi_e\big) \;=\; \xi_e^2\,\mathrm{Var}_j(\gamma_{ij}) \;+\; (b\,w_\infty^U)^2\,\xi_e^2\,\mathrm{Var}_j\big[(1-\rho_{j,e})\gamma_{ij}x_{ij}\big]. \tag{13}$$

Corollary 1 (cross-country dispersion). The cross-country dispersion of consensus inflation expectations following a global conflict shock is the sum of a fundamentals component (heterogeneity in γ_{ij}) and a network-bias component (heterogeneity in $(1 - \rho_{j,e})x_{ij}$).

The corollary says that two countries with identical bilateral exposures can display very different consensus inflation expectations purely because of the structure of forecaster coverage. Empirically, this rationalizes why two non-conflict countries with comparable exposure to a conflict country—e.g., two European economies with similar pre-2022 gas dependence on Russia—can nonetheless exhibit visibly different post-conflict consensus revisions. The model attributes the gap to differences in $\rho_{j,e}$. Table 8 implements this decomposition empirically. For each non-conflict country j in each episode, we compute the average post-cutoff revision in the current-year CPI inflation forecast (using the multi-post w30 panel construction of Figure 4) and regress it on the bilateral exposure x_{ij} (the fundamentals component) and on the network-bias term $(1 - \rho_{j,e})\,x_{ij}$. We report the R^2 of the fundamentals-only regression and the marginal R^2 from adding the network-bias regressor. We

measure x_{ij} by the IEA bilateral energy import share—the supply-side channel where the heterogeneity analysis in Table 5 identifies the strongest information differential.[23] Pooling across episodes (with episode fixed effects absorbed), the network component contributes substantially more to cross-country dispersion than the bilateral-exposure component does on its own ($R^2_{\text{network}} = 11.8$ percent vs. $R^2_{\text{fundamentals}} = 0.8$ percent). Per-episode, the dominance is most visible in ME 2026, Russia–Ukraine 2022, and Saudi–Yemen 2015, where the cross-country variation in forecaster coverage of the conflict country is wide.

Table 8: Cross-country dispersion decomposition, Corollary 1

Episode	N countries	$\text{Var}_j(\Delta y_{j,e})$	R^2 fundamentals	R^2 + network	Marginal R^2 network
Crimea 2014	23	0.092	0.299	0.312	0.013
Saudi-Yemen 2015	14	0.042	0.199	0.232	0.033
RU-UA 2022	17	1.337	0.011	0.140	0.128
Iran 2024	16	0.014	0.045	0.052	0.006
ME 2026	10	0.027	0.008	0.269	0.261
Pooled (episode FE absorbed)	100	0.252	0.008	0.126	0.118

Notes: The table provides a cross-country dispersion decomposition, following Corollary 1. For each non-conflict country in each conflict episode, we compute the average post-cutoff CPI inflation revision in the current year and average across forecasters and post-cutoff dates. We then a regression on the bilateral exposure x_{ij} alone (the fundamentals component), and another with the network-bias term $(1 - \rho_{j,e})\, x_{ij}$ added. "R^2 fundamentals" is the R^2 of the first regression; "R^2 + network" is the R^2 of the second; "Marginal R^2 network" is the difference. The bilateral exposure x_{ij} is the bilateral energy import share with the exporter set extended to include Iran for the Israel-anchored episodes (Gaza 2014, Hamas 2023, Iran 2024, Hezbollah 2024, ME 2026). The connectedness share $\rho_{j,e}$ is the share of forecasters covering j in the post-cutoff window of episode e who are also connected to the conflict country in the two-year lookback before the conflict. The pooled row absorbs episode fixed effects via demeaning. The Gaza 2014 episode (only nine countries with non-missing exposure) is omitted from the per-episode rows but included in the pooled regression.

The second result closes the macro loop. Closing the model with a country-level New Keynesian Phillips curve, $\pi_{j,t} = \kappa\, mc_{j,t} + \theta\, \mathbb{E}_e^{\text{cons}}[\pi_{j,t+1}]$—where $\pi_{j,t}$ and $mc_{j,t}$ are inflation and real marginal cost in deviations from steady state, $\kappa > 0$ is the New Keynesian slope, $\theta \in (0, 1)$ is the subjective discount factor, and the forward expectation is the (biased) consensus—propagates the bias from Equation 12 to realized inflation:

Proposition 2 (Phillips-curve pass-through). Realized inflation in country j in the post-conflict window inherits a wedge proportional to the network-mediated expectation bias:

$$\pi_{j,t} - \pi_{j,t}^{\text{full info}} \;\approx\; \theta\,(1 - \rho_{j,e})\,\gamma_{ij}\,b\,x_{ij}\,w_\infty^U \;+\; O(\theta^2), \tag{14}$$

where $\pi_{j,t}^{\text{full info}}$ is the counterfactual inflation that would obtain if the consensus expectation in the

[23] Israel exports negligible quantities of energy, so bilateral energy shares computed against Israel alone are uninformative. For the Israel-anchored episodes the energy exporter implicated in the conflict is Iran, and we therefore extend the exporter set to include Iran for Gaza 2014, Hamas 2023, Iran 2024, Hezbollah 2024, and ME 2026. The Russia–Ukraine and Saudi–Yemen episodes are unchanged.

Phillips curve were replaced by the full-information rational expectation of the spillover, and w_∞^U is the long-run weight that unconnected forecasters place on the public signal (the $t \to \infty$ limit of w_t^U). The leading term is the one-period-ahead bias; the $O(\theta^2)$ remainder collects the contributions of horizons two quarters and beyond, which are damped both by additional powers of the discount factor θ and by the gradual public revelation of the true shock ξ_e.

The slope of the cross-country Phillips curve in response to conflict shocks therefore depends not only on standard pricing primitives κ and θ but also on the structure of the forecaster network. A sparser bipartite information network makes the global Phillips curve appear steeper than fundamentals would imply. This means that conflict episodes look more inflationary to a policymaker viewing the cross section than they actually are, and the gap is larger in countries with both more exposure to the conflict country and less coverage by analysts who follow it. The wedge is in the reduced-form Phillips-curve sense: it is the part of inflation that comes from forecasters' upward-biased expectations becoming self-fulfilling through standard NK pricing.[24]

The third result delivers an explicit substitutability between public-data quality and forecaster-network density. Differentiating Equation 12 along the two relevant margins:

Corollary 2 (public data vs. network density). For any country j, public-data quality (proxied by σ_{pub}^{-2}) and forecaster-network density $\rho_{j,e}$ are substitutes in determining aggregate-expectation bias, with marginal rate of substitution

$$\left.\frac{d\rho_{j,e}}{d\sigma_{\text{pub}}^{-2}}\right|_{\text{bias fixed}} = -\frac{1-\rho_{j,e}}{w_\infty^U} \cdot \frac{\partial w_\infty^U}{\partial \sigma_{\text{pub}}^{-2}} > 0.$$

That is, more granular real-time bilateral data on the conflict country reduces the wedge by improving the public signal, in exactly the way that broader analyst coverage does. The model gives this substitution a precise rate. In our running example, doubling the precision of public real-time Russian gas-flow data reduces the wedge in Hungary's consensus inflation forecast by the same amount as adding several Russia-covering forecasters to the Hungary panel. The two margins of improvement act on the same primitive—the precision of the information used to interpret the bilateral spillover—and the corollary tells us they substitute at a known rate. This is the consensus-side counterpart of the experimental information-treatment effects on household expectations documented by Coibion *et al.* (2022). Their treatments improve the precision of the public information that reaches a household, while our corollary characterizes how the same primitive shows up at the consensus level when the public information is granular bilateral data on the conflict country.

Before translating the model's bias into basis points of inflation, we ask how concentrated the aggregate welfare loss from coverage gaps is across (importer, episode) pairs. We define the squared-

[24] A natural follow-up question is whether ex-post forecast errors are smaller for connected or unconnected forecasters. The model's prediction is ambiguous and possibly counterintuitive. If realized inflation tracks the biased consensus through the pass-through in Equation 14, unconnected forecasters—whose revisions sit closer to the consensus—would have smaller ex-post errors than connected forecasters, even though connected forecasters' subjective revisions are the ones tilted toward the true shock ξ_e. The within-pair design in Equation 2 sidesteps this by identifying the information channel from the cross-forecaster gap in revisions on the same country and date, where the realized inflation that the forecasts will be compared against is held fixed.

loss aggregate $\mathcal{L} = \sum_{(j,e,c)}(1-\rho_{j,e})^2 x_{ij,c}^2$ summed over the full set of (importer, episode, channel) tuples in the sample. Setting $\rho_{j,e} = 1$ at the top-N tuples—closing the connectedness gap at the highest-loss pairs by adding forecaster–country edges—removes their contribution to $\mathcal{L}$. Table 9 reports the cumulative percentage reduction in $\mathcal{L}$ at $N = 5, 10, 15, 20, 30, 50, 100$, separately by channel and pooled across all channels. The welfare loss is highly concentrated. Filling the top 15 pooled gaps captures 67 percent of the total welfare loss, and the top 50 capture over 90 percent. The concentration is even sharper on individual channels—79 percent of the energy-channel welfare loss is captured by the top 15 energy gaps, and 81 percent of the trade-channel loss by the top 15 trade gaps.

Table 9: Welfare reduction from filling top-N coverage gaps

	Top N gaps filled, percent reduction in total welfare loss						
Channel	5	10	15	20	30	50	100
Energy	47.1	63.1	79.4	85.7	94.4	99.1	100.0
Trade	50.9	71.2	80.7	86.0	92.7	96.6	99.4
Value-added	37.7	56.6	66.0	72.7	79.7	88.2	96.2
All channels (pooled)	40.2	53.9	66.8	73.5	82.1	91.7	97.5

Notes: Each cell reports the percentage reduction in the total welfare loss from coverage gaps achieved by giving full conflict-country coverage to the N highest-loss (importer, conflict country, episode) tuples. A tuple's contribution to the welfare loss rises with both its coverage gap and its bilateral exposure, so the ranking places the most exposed and least covered pairs first. The per-channel rows rank tuples within a single channel; the "All channels (pooled)" row ranks across the three channels jointly. Both reach 100 percent once N reaches the total number of tuples in scope.

The squared-loss reduction is dimensionless. To translate the model's implications into basis points of inflation, we close Equation 14 with a calibration of the composite parameter $\gamma \cdot B$, where $B \equiv b \cdot w_\infty^U$ bundles the narrative-bias slope b of Equation 5 with the weight the unconnected forecaster ultimately places on the public signal. Only this product is identified. The within-cell empirical relationship $\hat{\beta} = -\gamma \cdot B \cdot \bar{x}$ documented in Figure 4 and Table 5 pins down $\gamma \cdot B$, but not b, w_∞^U, and γ separately, and we normalize γ to a constant throughout. We use the bilateral energy share as x_{ij}, with the Iran-extended exporter set for Israel-anchored episodes (as in Table 8): $\bar{x} = 0.117$ and $\hat{\beta} = -0.161$, so $\gamma \cdot B \approx 1.37$. Setting $\theta = 0.99$, the model's implied inflation wedge for each (importer, episode) pair becomes $\text{Wedge}_{j,e} = \theta \cdot (1-\rho_{j,e}) \cdot \gamma B \cdot x_{j,e}$.

Table 10 reports the top fifteen pairs by implied wedge. The 59 bps figure for Finland during Crimea 2014 says that, had every Finland-covering forecaster also been a Russia-covering forecaster at that time, Finland's consensus inflation forecast in early 2014 would have been about 0.6 percentage points lower; through the standard New Keynesian pass-through to wages and prices, realized Finnish inflation would itself have been about 0.6 percentage points lower in the immediate post-conflict window. More generally, the model puts inflation wedges of roughly 25–60 basis points at the most exposed (importer, episode) pairs. These are calibrated magnitudes rather than directly estimated ones, and they inherit the sampling uncertainty of the dampening coefficient they are built from. Because the wedge is linear in $\gamma \cdot B$, the 95 percent confidence interval around $\hat{\beta}$ scales every wedge by a factor between 0.4 and 1.6, so the Finland figure ranges from 23 to 95 basis points and the

25–60 range from about 10 to 95. The point estimates should be read as a central case within a band of that width. To benchmark these magnitudes, 25–60 bps is comparable to the four-quarter cumulative pass-through of a moderate (10–20 percent) energy-price shock to advanced-economy CPI inflation documented by Choi *et al.* (2018), or it is 12–30 percent of a typical central-bank inflation target. The list is dominated by the Eastern European gas-importing periphery during Crimea 2014 and Russia–Ukraine 2022 where bilateral energy dependence on Russia is largest. The average wedge across (importer, episode) pairs in our sample is 7 basis points but is highly concentrated. The top fifty pairs account for over 90 percent of the cross-country total, and filling the top fifteen forecaster–country coverage gaps eliminates 53 percent of it.[25]

Table 10: Model-implied inflation wedge: top 15 (importer, conflict country, episode) tuples

	Importer (j)	Conflict country	Episode	x_{ij}	$\rho_{j,e}$	Wedge (bps)
1	Finland	Russia, Ukraine	Crimea 2014	0.753	0.42	59.3
2	Finland	Russia, Ukraine	RU-UA 2022	0.658	0.46	48.2
3	Lithuania	Russia, Ukraine	Crimea 2014	0.883	0.60	48.0
4	Bulgaria	Russia, Ukraine	Crimea 2014	0.840	0.65	40.3
5	Germany	Russia, Ukraine	RU-UA 2022	0.440	0.44	33.8
6	Greece	Russia, Ukraine	Crimea 2014	0.410	0.42	32.3
7	Poland	Russia, Ukraine	Crimea 2014	0.824	0.72	31.5
8	Latvia	Russia, Ukraine	Crimea 2014	0.486	0.55	30.1
9	Hungary	Russia, Ukraine	RU-UA 2022	0.632	0.68	27.6
10	Poland	Russia, Ukraine	RU-UA 2022	0.596	0.69	25.3
11	Latvia	Russia, Ukraine	RU-UA 2022	0.510	0.64	25.2
12	Hungary	Russia, Ukraine	Crimea 2014	0.860	0.79	25.1
13	Estonia	Russia, Ukraine	Crimea 2014	0.388	0.55	24.0
14	Germany	Russia, Ukraine	Crimea 2014	0.332	0.47	23.8
15	Czech Republic	Russia, Ukraine	Crimea 2014	0.678	0.76	22.1

Notes: The table reports the model-implied inflation wedge for the top 15 (importer, conflict country, episode) tuples. For each tuple, the implied wedge is the product of the forecaster-coverage gap (i.e., share of forecasters covering the importer in the window following the episode's cutoff who are not also connected to the conflict country), the importer's bilateral energy-import share from the conflict country, and a single calibration constant. The calibration constant is recovered from the baseline within-cell coefficient of Table 5: the absolute coefficient on the energy margin (0.161) divided by the mean bilateral energy share across pairs (0.117) gives 1.37, and the Phillips-curve pass-through is set to 0.99. The bilateral energy share extends the exporter set to include Iran for the Israel-anchored episodes (Gaza 2014, Hamas 2023, Iran 2024, Hezbollah 2024, ME 2026), reflecting the supply-side adversary in those conflicts. The wedge is reported in basis points of inflation.

[25]Summed across the (importer, episode) cells in our sample, the cumulative network-mediated inflation excess is 9.5 percentage points.

4.4 Implications for monetary policy

The inflation wedge in Equation 14 is not directly observable. What a central bank observes is the professional-forecaster consensus, which embeds it, so a policy rule conditioned on that consensus inherits the bias. This has a direct consequence for the supply-shock tradeoff. Consider a central bank with the standard objective $\frac{1}{2}\,\mathbb{E}\left[\pi_{j,t}^2 + \vartheta\, y_{j,t}^2\right]$, where $y_{j,t}$ is the output gap and $\vartheta > 0$ is the weight on output stabilization, facing the Phillips curve $\pi_{j,t} = \kappa\, y_{j,t} + \theta\, \mathbb{E}_t[\pi_{j,t+1}] + u_{j,t}$ with $u_{j,t}$ the conflict-driven cost-push shock. The observed consensus exceeds the full-information expectation by the network-mediated expectation bias of Equation 12, $\widehat{\mathbb{E}}_t[\pi_{j,t+1}] = \mathbb{E}_t^*[\pi_{j,t+1}] + W_{j,e}/\theta$, where $W_{j,e}$ is the inflation wedge of Equation 14 reported in Table 10 and the two differ by the discount factor carried by the Phillips curve. The only non-standard ingredient is that the forward expectation entering the bank's problem is the observed consensus rather than the full-information expectation. This is not an assumption that the bank is naive. Removing the bias would require knowing $\rho_{j,e}$ and b, neither of which is observable in real time, so even a bank that suspects the consensus is too high cannot net the bias out of it.

Corollary 3 (policy wedge). A central bank that conditions on the observed consensus opens an output gap that is too negative, by an amount proportional to the inflation wedge,

$$\hat{y}_{j,t} - y_{j,t}^* \;=\; -\,\frac{\kappa}{\vartheta + \kappa^2}\, W_{j,e} \;<\; 0, \tag{15}$$

and incurs an excess loss that is quadratic in it, $\Delta L_{j,e} = \frac{1}{2}\,\kappa^2\, W_{j,e}^2/(\vartheta + \kappa^2)$.

Two considerations follow. First, because the excess tightening is increasing in $(1 - \rho_{j,e})\, x_{ij}$, it is largest exactly where bilateral exposure is high and forecaster coverage of the conflict country is thin—the configuration typical of small open economies. Second, because the loss is quadratic in the wedge, which is itself highly concentrated, the welfare cost is more concentrated still, and the returns to closing the coverage gaps of Table 9 are convex. By Corollary 2 the same reduction is available from more granular public data on the conflict country, or from communication that quantifies the bilateral exposure directly. These act on the measurement rather than on the economy, which is why they are inexpensive relative to the tightening they would replace.

To gauge the magnitude, a bank following a simple rule $i_{j,t} = r^n + \phi_\pi \widehat{\mathbb{E}}_t[\pi_{j,t+1}]$ tightens by $\phi_\pi\, W_{j,e}/\theta$ more than the fundamentals warrant. At $\phi_\pi = 1.5$, $\theta = 0.99$, and a wedge of 25–60 basis points, that is roughly 40–90 basis points of excess tightening at the most exposed pairs. This figure inherits the calibration uncertainty of the wedge itself, and should be read as an order of magnitude rather than a point estimate.[26]

To summarize, Predictions 1–3 deliver an internally consistent account of the empirical findings in Section 3, and Propositions 1–2 with Corollaries 1–3 turn the same primitives into a set of macroeconomic statements about cross-country inflation dispersion, Phillips-curve pass-through, the substi-

[26] Two caveats bound the exercise. Central banks form their own staff forecasts and do not mechanically adopt the consensus; if the consensus receives weight $\omega \in (0, 1]$ in the bank's assessment, every expression above scales by ω, so the magnitudes are upper bounds. And the policy algebra itself is standard: the contribution here is the measurement of $W_{j,e}$, not the mapping from a biased expectation to a policy error.

tutability between public-data quality and forecaster connectedness, and the policy error that follows when the wedge is read at face value. The empirical implementations of these statements in Table 8, Table 9, and Table 10 translate the corollaries into observable numbers. The bottom line is that the network channel accounts for more than an order of magnitude more of the cross-country dispersion in post-conflict consensus revisions than bilateral exposure does on its own; the welfare loss from coverage gaps is highly concentrated at a small set of high-exposure pairs; and at the most exposed (importer, episode) pairs, the model's implied inflation wedge is in the range of 0.25–0.6 percentage points.

5 Conclusion

Inflation expectations are usually studied as a property of the individual forecaster. This paper has examined a complementary dimension: the network linking forecasters to the countries they cover, which is in place before any given shock. Our findings indicate that, across eight conflict episodes between 2014 and 2026, post-conflict inflation expectations for non-conflict countries revise gradually upward over the days following the shock rather than jumping at the cutoff. Within that adjustment, forecasters who already cover the conflict country revise by 0.2 to 0.3 percentage points less, in absolute value, than forecasters who do not. The dampening differential roughly triples at country pairs where bilateral exposure to the conflict country is highest, concentrating at the energy and intermediate-goods-trade margins.

A Bayesian model with two coupled networks rationalizes these findings. Bilateral exposure determines the size of cross-country spillovers and the magnitude of the public-narrative bias, while the bipartite information network rations who observes a precise private signal about the shock; aggregate inflation expectation bias is the product of the two. Three macroeconomic implications follow. First, the network-bias component accounts for an order of magnitude more of the cross-country dispersion in post-conflict revisions than the fundamentals component. Second, closed with a New Keynesian Phillips curve, the model implies an inflation wedge of 25–60 basis points at the most exposed economies; filling the top fifteen forecaster–country coverage gaps eliminates more than half of the cross-country total. Third, public data quality and forecaster network density are substitutes, in that better real-time bilateral data on the conflict country, or communication that quantifies the exposures at issue, reduces the wedge in the same way than as broader forecaster coverage would do.

The results imply that the cross-sectional and dynamic properties of inflation expectations depend on the network structure of information producers, not only on how individual forecasters process news. From the policy perspective, a central bank's tradeoff after a conflict depends on how far the inflation expectations it observes have moved. If the central bank reads the consensus forecasts at face value and these overstate the true movement, monetary policy may contract by more than it needs to. The results emphasize the importance of clear communication of the central bank's own reading of the shock and its transmission channels where exposure to the conflict country is high and forecasters' coverage of it is thin. Still, since the central bank cannot tell in real time how much

of the consensus movement is narrative rather than fundamentals, it will need to balance the risk of discounting it and facing higher inflation against the risk of taking it at face value and contracting more than needed.

References

ACEMOGLU, D., AKCIGIT, U. and KERR, W. (2016). Networks and the macroeconomy: An empirical exploration. *NBER Macroeconomics Annual*, **30**, 273–335.

—, CARVALHO, V. M., OZDAGLAR, A. and TAHBAZ-SALEHI, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, **80** (5), 1977–2016.

—, OZDAGLAR, A., PARANDEHGHEIBI, A. and TAHBAZ-SALEHI, A. (2014). Dynamics of information exchange in endogenous social networks. *Theoretical Economics*, **9** (1), 41–97.

ANDRADE, P. and LE BIHAN, H. (2013). Inattentive professional forecasters. *Journal of Monetary Economics*, **60** (8), 967–982.

ANDRE, P., PIZZINELLI, C., ROTH, C. and WOHLFART, J. (2022). Subjective models of the macroeconomy: Evidence from experts and representative samples. *Review of Economic Studies*, **89** (6), 2958–2991.

ANGELETOS, G.-M., HUO, Z. and SASTRY, K. A. (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, **35**, 1–86.

— and LA'O, J. (2013). Sentiments. *Econometrica*, **81** (2), 739–779.

ATALAY, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, **9** (4), 254–280.

AUER, R. A., LEVCHENKO, A. A. and SAURÉ, P. (2019). International inflation spillovers through input linkages. *Review of Economics and Statistics*, **101** (3), 507–521.

BAE, K.-H., STULZ, R. M. and TAN, H. (2008). Do local analysts know more? a cross-country study of the performance of local analysts and foreign analysts. *Journal of Financial Economics*, **88** (3), 581–606.

BORDALO, P., GENNAIOLI, N., MA, Y. and SHLEIFER, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, **110** (9), 2748–2782.

—, — and SHLEIFER, A. (2018). Diagnostic expectations and credit cycles. *Journal of Finance*, **73** (1), 199–227.

CALDARA, D. and IACOVIELLO, M. (2022). Measuring geopolitical risk. *American Economic Review*, **112** (4), 1194–1225.

CALIENDO, L. and PARRO, F. (2015). Estimates of the trade and welfare effects of NAFTA. *Review of Economic Studies*, **82** (1), 1–44.

CARROLL, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics*, **118** (1), 269–298.

CARVALHO, V. M., NIREI, M., SAITO, Y. U. and TAHBAZ-SALEHI, A. (2021). Supply chain disruptions: Evidence from the Great East Japan earthquake. *Quarterly Journal of Economics*, **136** (2), 1255–1321.

— and TAHBAZ-SALEHI, A. (2019). Production networks: A primer. *Annual Review of Economics*, **11**, 635–663.

CHOI, S., FURCERI, D., LOUNGANI, P., MISHRA, S. and POPLAWSKI-RIBEIRO, M. (2018). Oil prices and inflation dynamics: Evidence from advanced and developing economies. *Journal of Interna-*

tional Money and Finance, **82**, 71–96.

COIBION, O. and GORODNICHENKO, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, **120** (1), 116–159.

— and — (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, **105** (8), 2644–2678.

—, —, KUMAR, S. and RYNGAERT, J. (2021). Do you know that I know that you know? higher-order beliefs in survey data. *Quarterly Journal of Economics*, **136** (3), 1387–1446.

—, — and WEBER, M. (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, **130** (6), 1537–1584.

COVAL, J. D. and MOSKOWITZ, T. J. (2001). The geography of investment: Informed trading and asset prices. *Journal of Political Economy*, **109** (4), 811–841.

CRAVINO, J. and LEVCHENKO, A. A. (2017). Multinational firms and international business cycle transmission. *Quarterly Journal of Economics*, **132** (2), 921–962.

DAVIS, L. and KAHN, M. E. (2008). The effect of immigration on natives' well-being: Evidence from the london bombings. *Working Paper*, nBER Working Paper No. 14394.

EISENSEE, T. and STRÖMBERG, D. (2007). News droughts, news floods, and U.S. disaster relief. *Quarterly Journal of Economics*, **122** (2), 693–728.

EUSEPI, S. and PRESTON, B. (2011). Expectations, learning, and business cycle fluctuations. *American Economic Review*, **101** (6), 2844–2872.

FEDERLE, J., MEIER, A., MÜLLER, G. J., MUTSCHLER, W. and SCHULARICK, M. (2024). *The Price of War*. Tech. rep., CEPR Discussion Paper No. 18834.

FORBES, K. J. (2019). Inflation dynamics: Dead, dormant, or determined abroad? *Brookings Papers on Economic Activity*, pp. 257–319, fall 2019.

GABAIX, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, **79** (3), 733–772.

GOLUB, B. and JACKSON, M. O. (2010). Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, **2** (1), 112–149.

LUCAS, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, **4** (2), 103–124.

MAĆKOWIAK, B. and WIEDERHOLT, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, **99** (3), 769–803.

— and WIEDERHOLT, M. (2015). Business cycle dynamics under rational inattention. *Review of Economic Studies*, **82** (4), 1502–1532.

MALLOY, C. J. (2005). The geography of equity analysis. *Journal of Finance*, **60** (2), 719–755.

MANKIW, N. G. and REIS, R. (2002). Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve. *Quarterly Journal of Economics*, **117** (4), 1295–1328.

—, — and WOLFERS, J. (2003). Disagreement about inflation expectations. *NBER Macroeconomics Annual*, **18**, 209–248.

MORRIS, S. and SHIN, H. S. (2002). Social value of public information. *American Economic Review*, **92** (5), 1521–1534.

MULLAINATHAN, S. and SHLEIFER, A. (2005). The market for news. *American Economic Review*, **95** (4),

1031–1053.

PASTEN, E., SCHOENLE, R. and WEBER, M. (2020). The propagation of monetary policy shocks in a heterogeneous production economy. *Journal of Monetary Economics*, **116**, 1–22.

PATTON, A. J. and TIMMERMANN, A. (2010). Why do forecasters disagree? lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics*, **57** (7), 803–820.

SHILLER, R. J. (2017). Narrative economics. *American Economic Review*, **107** (4), 967–1004.

SIMS, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, **50** (3), 665–690.

SUNDBERG, R. and MELANDER, E. (2013). Introducing the UCDP georeferenced event dataset. *Journal of Peace Research*, **50** (4), 523–532.

Online Appendix for:
"Information Networks and Inflation Expectations: Evidence from Professional Forecasters and Conflicts"

Francesco Grigoli and Bertrand Gruss

October 6, 2026

A-I Data

This appendix lists the data sources, variable definitions, and aggregation rules used in the empirical analysis. The analysis combines four layers of data described in turn below.

Forecast data The Continuous Consensus Forecasts (CCF) from Consensus Economics, 2012–2026, cover 48 countries and 611 forecasting institutions. Each observation is at the country–forecaster–date level and includes inflation (CPI) forecasts at three annual horizons (current year, 1 year ahead, 2 years ahead) and eight quarterly horizons (current quarter through 7 quarters ahead). The forecaster–country coverage matrix—which forecasters submit forecasts for which countries, and when—is the central object for identifying the information channel.

Conflict events We identify candidate conflict episodes from the Uppsala Conflict Data Program (UCDP) Georeferenced Event Dataset (Sundberg and Melander, 2013), which records geocoded violent events at the day level for state-based, non-state, and one-sided violence. The episode-selection criterion is described in Section 2.

Bilateral economic linkages We assemble matrices of bilateral exposure at the country-pair-year level across three dimensions, each providing one of the exposure measures in Table 5:

- *Energy:* The IEA World Energy Balances, providing bilateral physical energy flows. We construct country j's import dependence on country i as i's share in j's total primary energy imports. For the components decomposition in Table 6 we separate oil and natural gas. For Israel-anchored episodes (Gaza 2014, Hamas 2023, Iran 2024, Hezbollah 2024, and the February 2026 escalation), we add Iran to the exporter set to reflect the supply-side adversary in those conflicts; the energy panel used in the dispersion decomposition (Table 8) and the NK calibration (Table 10) uses this Iran-extended construction.
- *Trade:* The BACI database from CEPII, which reconciles UN Comtrade importer- and exporter-reported flows and provides harmonized bilateral trade at the HS6 commodity level. For each country pair we compute the share of country j's total goods imports sourced from country i. For the components decomposition in Table 6 we further classify all HS6 lines into intermediate, capital, and consumption goods using the Comtrade HS6-to-BEC Rev. 5 concordance, allowing us to distinguish supply-chain from final-demand exposure.

- *Supply chains:* The OECD Trade in Value Added (TiVA) database, based on the OECD Inter-Country Input-Output (ICIO) tables. We compute, for each country pair, the foreign value-added share originating in i embodied in j's gross output.

Media coverage GDELT 1.0 Events provides daily, geocoded global media coverage of conflict events by country of publication. We use the (episode, importer) panel of mentions of the conflict country as a direct empirical test of the bilateral-narrative-bias assumption in the model, regressing log mentions on the log bilateral trade share across episode windows.

A-II Media coverage and bilateral exposure

Section 2 describes the conflict episodes as sharing a common public narrative, one that connects the conflict to inflation in non-conflict economies through bilateral energy and intermediate-goods exposure. The model of Section 4 formalizes this as a narrative that overstates the cross-country spillover from a conflict in country i to a non-conflict country j in proportion to bilateral exposure x_{ij} (Equation 5). The descriptive claim and the modeling assumption imply the same reduced-form prediction, that media coverage of the (conflict country, non-conflict country) pair scales with bilateral exposure across the cross section. This appendix tests that prediction directly in GDELT data.

For each episode and each non-conflict country, we count total mentions of the (conflict country, non-conflict country) pair in the 30 calendar days after the cutoff, using the GDELT 1.0 Events daily files (which report the number of articles per geocoded event and identify the two actor countries). We keep events in which one actor is a conflict country and the other is a non-conflict country, and aggregate article counts to the (episode, non-conflict country) level. We then merge the mentions panel with the bilateral total-trade share—the only exposure measure defined for every (importer, exporter) pair—and regress log mentions on the (log) trade share with episode fixed effects.[27]

Table A1 reports the pooled results. In the log–log specification, a one-log-point increase in the bilateral trade share raises mentions by 0.56 log points, so a country with twice the bilateral trade share to the conflict country receives roughly fifty percent more media coverage of the link. The level-share specification has the same positive sign but is only marginally significant, reflecting the strong right-skew of the trade-share distribution. The relationship is not an artifact of pooling. Estimated episode by episode, the log–log slope is positive in all eight episodes, ranging from 0.20 to 1.41, and is individually significant at conventional levels in five of them. The cross-sectional scaling of media coverage with bilateral exposure is therefore present in the data, which supports both the characterization of the episodes in Section 2 and the narrative-bias term in the model.

[27] We use the trade share rather than the energy share that enters the model's calibration because media coverage of a country pair reflects the overall salience of the bilateral relationship, not any single supply channel. The energy share is also far sparser.

Table A1: Media coverage of conflict-country pairs and bilateral exposure

	(1)	(2)
Bilateral trade share	3.817*	
	(2.250)	
Log bilateral trade share		0.561***
		(0.094)
Within-R^2	0.009	0.098
N	336	336

Notes: The table reports the relationship between bilateral exposure and media coverage. The dependent variable is the log number of article mentions of each (conflict country, non-conflict country) pair in the 30 calendar days after the episode cutoff. The regressor is the bilateral total-trade share of the non-conflict country's imports sourced from the conflict country, summed across conflict countries within multi-conflict episodes. Both regressions pool all eight episode windows with episode fixed effects: column (1) uses the raw trade share, column (2) its natural log. The within-R^2 is the share of the within-episode variation in log mentions explained by the regressor. Standard errors are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

A-III Robustness results

This appendix reports the results of the robustness tests.

Baseline estimates Table A2 reports the baseline estimates of Figure 4.

Table A2: Information spillover effects to non-conflict countries

	Forecast horizon						
	Q0	Q1	Q2	Q3	Q4	Q5	Q6
Network connection (β)	-0.295^{**}	-0.265^{**}	-0.317^{***}	-0.298^{***}	-0.096	-0.067	-0.207^{***}
	(0.126)	(0.119)	(0.114)	(0.078)	(0.069)	(0.109)	(0.074)
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	101	164	181	195	137	104	56
Countries	17	20	21	22	20	20	18
Forecasters	16	37	44	43	37	38	33
R^2	0.184	0.027	0.100	0.202	0.002	0.011	0.090

Notes: The table reports the estimated network coefficient isolating the information channel. The dependent variable uses an asymmetric pre/post window (pre 7 days, post 30 days). Qh denotes h quarters ahead. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Alternative samples of conflict episodes The 'restrictive' sample retains only the seven episodes that satisfy the on-territory battle-event threshold: Crimea 2014, Gaza 2014, Russia–Ukraine 2022, Hamas 2023, Iran 2024, Hezbollah 2024, and the Middle East escalation of February 2026. It excludes Saudi–Yemen 2015, which qualifies under the expeditionary threshold. The 'aggressive' sample relaxes each of the two thresholds and adds a third leg admitting wartime sovereign-authority shocks. The battle-event threshold is read on a monthly basis and is further lowered to 10 deaths in the cutoff month with no rolling-median requirement. The expeditionary threshold drops the 1,000-death magnitude floor. A third leg admits discrete sovereign-authority shocks during or near armed conflict. The aggressive sample contains fourteen episodes in total.[28]

Table A3 reports the baseline CPI inflation regression of Equation 2 on each of the three samples. The pattern of estimates is preserved across samples. Connected forecasters revise CPI inflation forecasts downward at quarterly horizons. The CPI inflation estimates at Q0–Q3 remain negative and significant in all three samples, although the magnitudes attenuate in the aggressive sample because the additional Israel–Gaza flare-ups and sovereign-authority shocks generated smaller forecaster revisions on average.

[28] We do not adopt the aggressive sample as the headline because each of the relaxations is hard to defend without external validation. The lowered battle-event thresholds have no obvious external benchmark; the expeditionary threshold without a magnitude floor admits any cross-border military operation regardless of scale; and the sovereign-authority leg conflates state-coordinated escalations of an active war with domestic political crises whose economic channel differs from the wartime-news mechanism the headline criterion targets.

Table A3: Alternative samples

	Forecast horizon						
	Q0	Q1	Q2	Q3	Q4	Q5	Q6
Restrictive (7)	-0.295^{**}	-0.265^{**}	-0.317^{***}	-0.298^{***}	-0.096	-0.067	-0.207^{***}
	(0.126)	(0.119)	(0.114)	(0.078)	(0.069)	(0.109)	(0.074)
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	97	158	175	193	137	104	56
Headline (8)	-0.295^{**}	-0.265^{**}	-0.317^{***}	-0.298^{***}	-0.096	-0.067	-0.207^{***}
	(0.126)	(0.119)	(0.114)	(0.078)	(0.069)	(0.109)	(0.074)
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	101	164	181	195	137	104	56
Aggressive (14)	-0.295^{**}	-0.252^{**}	-0.253^{**}	-0.171^{***}	-0.037	-0.016	-0.095^{*}
	(0.126)	(0.108)	(0.112)	(0.061)	(0.038)	(0.053)	(0.049)
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	108	205	278	300	235	207	133

Notes: The table reports the estimated network coefficient isolating the information channel, re-estimated on samples satisfying more restrictive or aggressive selection criteria. The dependent variable uses an asymmetric pre/post window (pre 7 days, post 30 days), with episode×country×date fixed effects. Qh denotes h quarters ahead. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Table A4 reports the heterogeneity-by-exposure regression of Equation 3 on the same three samples. The information-channel differential $\hat{\beta}_1$ is stable across episode samples. For the energy share, $\hat{\beta}_1$ ranges between -0.27 and -0.31 across all three samples, and remains statistically significant. For the value-added share, $\hat{\beta}_1$ is negative in every sample and statistically significant in the restrictive and headline samples. The total trade measure, which was borderline in the headline sample, gains significance in the aggressive sample. The economic conclusion—that the inflation-spillover dampening concentrates at high-exposure pairs and at the channels where the public narrative is most easily organized around bilateral linkages—is unchanged across the full range of episode-selection rules we consider.

Table A4: Heterogeneity by bilateral exposure, episode-sample robustness

	$\hat{\beta}_0$ (Connected)	$\hat{\beta}_1$ (Connected × HighExp)	N
Energy share			
Restrictive (7)	-0.161^{***}	-0.305^{***}	158
	(0.050)	(0.065)	
Headline (8)	-0.161^{***}	-0.305^{***}	164
	(0.050)	(0.065)	
Aggressive (14)	-0.167^{***}	-0.271^{***}	253
	(0.047)	(0.061)	
Trade share			
Restrictive (7)	-0.163^{***}	-0.150	229
	(0.040)	(0.111)	
Headline (8)	-0.163^{***}	-0.150	235
	(0.040)	(0.111)	
Aggressive (14)	-0.125^{***}	-0.349^{**}	368
	(0.031)	(0.133)	
Value-added share			
Restrictive (7)	-0.116^{**}	-0.130^{*}	229
	(0.054)	(0.076)	
Headline (8)	-0.116^{**}	-0.130^{*}	235
	(0.054)	(0.076)	
Aggressive (14)	-0.148^{***}	-0.063	368
	(0.047)	(0.071)	
Episode × country × week FE	Yes	Yes	

Notes: The table reports the estimated network coefficient and its differential at high-exposure pairs, re-estimated on samples satisfying more restrictive or aggressive selection criteria. The dependent variable is the (forecaster, country, date) average revision across quarterly horizons. A pair is classified as high-exposure by the within-episode median split for energy and value-added, and by a fixed 2 percent threshold for trade. The regression includes episode×country×week fixed effects. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Single-post construction As a further check on the multi-post estimates in Figure 4, we collapse each forecaster–country pair to a single observation. For each forecaster f and country j, the pre-cutoff baseline is the closest forecast in $[t^* - 7, t^* - 1]$ and the post-cutoff observation is the closest forecast in $[t^* + 1, t^* + 30]$. Hence, a pair contributes to the regression only if both observations exist within these windows. The fixed effects are episode×country (rather than episode×country×date in the main specification, since each pair contributes a single observation).

The dampening pattern documented in the main text survives. $\hat{\beta}$ is negative and significant at multiple quarterly horizons. Sample sizes are roughly two-thirds smaller than in the multi-post specification because each pair is collapsed to a single observation, and standard errors widen accordingly, but the sign and approximate magnitude of $\hat{\beta}$ track the main estimates closely.

Table A5: Information network and forecast revisions, single forecast

	Forecast horizon						
	Q0	Q1	Q2	Q3	Q4	Q5	Q6
Network connection (β)	-0.260^{**}	-0.421^{*}	-0.384^{**}	-0.273^{**}	-0.047	-0.036	-0.180^{**}
	(0.110)	(0.212)	(0.182)	(0.129)	(0.071)	(0.089)	(0.073)
Episode × country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	52	96	111	121	89	78	50
Countries	17	20	21	22	20	20	18
Forecasters	16	37	44	43	37	38	33
R^2	0.218	0.101	0.120	0.085	0.003	0.002	0.059

Notes: The table reports the estimated network coefficient isolating the information channel, collapsing each forecaster–country pair to a single observation. The dependent variable uses an asymmetric pre/post window (pre 7 days, post 30 days) Qh denotes h quarters ahead. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

Forecaster coverage and connectedness strictness Table A6 re-estimates Equation 2 under two families of restrictions, reporting the coefficient at the peak forecast horizon. The first tightens the definition of connectedness from at least one forecast for the conflict country in the lookback to at least five, and then at least twenty. Forecasters with some but insufficient coverage are dropped rather than recoded as unconnected, so that the comparison is always between sustained coverage and none at all. The second restricts the estimation sample to forecasters covering at least five, and then at least ten, countries over the same window, which holds coverage breadth roughly constant across the connected and unconnected groups.

Table A6: Forecaster coverage and connectedness strictness

		Minimum submissions		Minimum countries	
	Baseline	5	20	5	10
	(1)	(2)	(3)	(4)	(5)
Network connection (β)	-0.317^{***}	-0.334^{***}	-0.427^{***}	-0.345^{***}	-0.358^{***}
	(0.114)	(0.118)	(0.107)	(0.109)	(0.081)
Episode × country × date FE	Yes	Yes	Yes	Yes	Yes
Observations	181	179	173	153	143
Countries	21	20	20	21	20
Forecasters	44	43	42	22	18
R^2	0.100	0.104	0.146	0.236	0.341

Notes: The table reports the estimated network coefficient at the peak forecast horizon, two quarters ahead. The dependent variable uses an asymmetric pre/post window (pre 7 days, post 30 days). Column (1) is the baseline specification. Columns (2) and (3) tighten the connectedness definition, requiring a forecaster to have submitted at least five and at least twenty forecasts for the conflict country during the two-year lookback ending the day before the cutoff, rather than at least one. Columns (4) and (5) restrict the estimation sample to forecasters covering at least five and at least ten distinct countries over the same window, which holds coverage breadth roughly constant across the connected and unconnected groups. Standard errors clustered at the forecaster level are in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

A-IV Model derivations

This appendix collects the formal derivations behind the Predictions, Propositions, and Corollaries stated in Section 4. We proceed in five blocks: (i) the Bayesian filter for the two forecaster types, (ii) the three predictions for individual forecasters, (iii) the mapping between the model objects and the regressions of Section 3, (iv) the consensus aggregator, the network-mediation theorem, and the macroeconomic comparative statics, and (v) the two extensions flagged at the end of Section 4.

A-IV.1 Bayesian filter

We collect the normality and independence assumptions used throughout. The fundamental shock is $\xi_e \sim \mathcal{N}(0, \sigma_\xi^2)$. The signal innovations $\{\eta_t^{\text{pub}}\}_{t\geq 1}$ and $\{\eta_t^{\text{priv}}\}_{t\geq 1}$ are mutually independent, independent of ξ_e, and i.i.d. across t within stream, with $\eta_t^{\text{pub}} \sim \mathcal{N}(0, \sigma^2_{\text{pub}})$ and $\eta_t^{\text{priv}} \sim \mathcal{N}(0, \sigma^2_{\text{priv}})$. The bilateral exposure x_{ij}, the spillover loading $\gamma_{ij} = \bar{\gamma} + \lambda x_{ij}$, the narrative bias slope b, and the connectedness indicator $z_{f,e}$ are all observable to the econometrician and predetermined relative to the realization of ξ_e. The realized spillover for country j in episode e is $\Delta y_{j,e} = \gamma_{ij}\xi_e + u_{j,e}$ with $u_{j,e} \sim \mathcal{N}(0, \sigma_u^2)$ independent of ξ_e and the signal innovations.

Reduced-form likelihoods Stack the t public observations into the vector $\mathbf{s}_{1:t}^{\text{pub}} = (s_1^{\text{pub}}, \ldots, s_t^{\text{pub}})^\top$. The τ-th entry equals $s_\tau^{\text{pub}} = \xi_e + bx_{ij} + \eta_\tau^{\text{pub}}$ for $\tau = 1, \ldots, t$, so the likelihood is

$$\mathbf{s}_{1:t}^{\text{pub}} \mid \xi_e \;\sim\; \mathcal{N}\big((\xi_e + bx_{ij})\mathbf{1}_t,\, \sigma^2_{\text{pub}}\, \mathbf{I}_t\big),$$

where $\mathbf{1}_t$ is the t-vector of ones and $\mathbf{I}_t$ the identity. The cumulative public mean

$$\bar{s}_t^{\text{pub}} \;=\; \tfrac{1}{t}\sum_{\tau=1}^{t} s_\tau^{\text{pub}} \;\sim\; \mathcal{N}\big(\xi_e + bx_{ij},\, \sigma^2_{\text{pub}}/t\big)$$

is a sufficient statistic for ξ_e given $\mathbf{s}_{1:t}^{\text{pub}}$ (Gaussian likelihoods admit one-dimensional sufficient statistics). The same is true for the private stream, with cumulative mean $\bar{s}_t^{\text{priv}} \sim \mathcal{N}\big(\xi_e,\, \sigma^2_{\text{priv}}/t\big)$.

Posterior for unconnected forecasters An unconnected forecaster f ($z_{f,e} = 0$) observes only $\mathbf{s}_{1:t}^{\text{pub}}$, equivalently the sufficient statistic $\bar{s}_t^{\text{pub}}$. The data-generating process for the public signal has mean $\xi_e + bx_{ij}$, but the modeling assumption is that the unconnected forecaster does not know about the bias b and treats $\bar{s}_t^{\text{pub}}$ as an unbiased Gaussian signal of ξ_e. Under this perceived likelihood and the prior $\xi_e \sim \mathcal{N}(0, \sigma_\xi^2)$, the standard Gaussian conjugate-update formula gives

$$\mathbb{E}[\xi_e \mid \mathcal{I}_t^U] \;=\; \frac{\sigma_\xi^{-2}\cdot 0 \;+\; t\sigma^{-2}_{\text{pub}}\cdot \bar{s}_t^{\text{pub}}}{\sigma_\xi^{-2} + t\sigma^{-2}_{\text{pub}}} \;=\; w_t^U\, \bar{s}_t^{\text{pub}}, \tag{16}$$

where the Kalman weight is

$$w_t^U \;=\; \frac{t\sigma_{\text{pub}}^{-2}}{\sigma_\xi^{-2} + t\sigma_{\text{pub}}^{-2}} \;\in\; [0,1). \tag{17}$$

This matches Equation 7 in the main text. The unconnected forecaster's posterior is not asymptotically unbiased: because $\bar{s}_t^{\text{pub}}$ has true mean $\xi_e + bx_{ij}$, its conditional expectation given ξ_e is $\mathbb{E}[\hat{\xi}_t^U \mid \xi_e] = w_t^U(\xi_e + bx_{ij})$. The trajectory starts below ξ_e at small t, when prior shrinkage toward zero dominates, crosses ξ_e at the working-day count at which $w_t^U = \xi_e/(\xi_e + bx_{ij})$, and converges to $\xi_e + bx_{ij}$ as $t \to \infty$, when $w_t^U \to 1$. The build-up of the unconnected forecaster's posterior over working days is the mechanism behind Prediction 1, matching the positive slope documented in the RDD-in-time of Subsection 3.1. A sophisticated unconnected forecaster who knew b would instead subtract bx_{ij} from $\bar{s}_t^{\text{pub}}$ before updating—the same correction the connected forecaster applies in Equation 18 below—and the cross-forecaster wedge documented in the data would not arise from public-narrative exposure. The naive-narrative convention adopted here is therefore the source of the empirical pattern the paper rationalizes. Subsection A-IV.6 below reports the alternative case in which both forecaster types know b, and shows that all three predictions go through with the same signs; only the magnitudes of the gaps change.

Posterior for connected forecasters A connected forecaster ($z_{f,e} = 1$) observes both streams. The joint sufficient statistic is the pair $(\bar{s}_t^{\text{pub}}, \bar{s}_t^{\text{priv}})$. Stacking and applying the same conjugate-update formula—now with three precision contributions, one each from the prior, the public stream, and the private stream—gives the posterior mean

$$\mathbb{E}[\xi_e \mid \mathcal{I}_t^C] \;=\; \frac{\sigma_\xi^{-2}\cdot 0 \;+\; t\sigma_{\text{pub}}^{-2}\cdot(\bar{s}_t^{\text{pub}} - bx_{ij}) \;+\; t\sigma_{\text{priv}}^{-2}\cdot \bar{s}_t^{\text{priv}}}{\sigma_\xi^{-2} + t\sigma_{\text{pub}}^{-2} + t\sigma_{\text{priv}}^{-2}}. \tag{18}$$

Connected forecasters know about the public-narrative bias because their access to country i's granular data lets them detect and correct it; the term $-bx_{ij}$ inside the public-stream contribution captures this. Equivalently, after correction, the connected forecaster sees two unbiased Gaussian signals of ξ_e with precisions $t\sigma_{\text{pub}}^{-2}$ and $t\sigma_{\text{priv}}^{-2}$. Defining the connected Kalman weights

$$w_t^{C,\text{pub}} \;=\; \frac{t\sigma_{\text{pub}}^{-2}}{\sigma_\xi^{-2} + t\sigma_{\text{pub}}^{-2} + t\sigma_{\text{priv}}^{-2}}, \tag{19}$$

$$w_t^{C,\text{priv}} \;=\; \frac{t\sigma_{\text{priv}}^{-2}}{\sigma_\xi^{-2} + t\sigma_{\text{pub}}^{-2} + t\sigma_{\text{priv}}^{-2}}, \tag{20}$$

Equation 18 simplifies to

$$\mathbb{E}[\xi_e \mid \mathcal{I}_t^C] \;=\; w_t^{C,\text{pub}}\big(\bar{s}_t^{\text{pub}} - bx_{ij}\big) \;+\; w_t^{C,\text{priv}}\,\bar{s}_t^{\text{priv}}, \tag{21}$$

which matches Equation 8 in the main text once the bias correction is folded into the public stream. The text aliases $\bar{s}_t^{\text{pub}} - bx_{ij} \to \bar{s}_t^{\text{pub}}$ for the connected forecaster, since the connected forecaster removes the bias before using the signal.

Properties of the Kalman weights The following four properties of Equation 17, Equation 19, and Equation 20 are used repeatedly below; all follow by direct inspection.

(i) Monotone in t. For every $t \geq 1$, w_t^U is strictly increasing in t, and similarly for $w_t^{C,\text{pub}} + w_t^{C,\text{priv}}$. *(ii) Long-run limits.* As $t \to \infty$, $w_t^U \to 1$ and $w_t^{C,\text{pub}} + w_t^{C,\text{priv}} \to 1$. *(iii) Public-weight gap.* For every $t \geq 1$ and every $\sigma^2_{\text{priv}} \in (0, \sigma^2_{\text{pub}})$,

$$w_t^U - w_t^{C,\text{pub}} \;=\; \frac{t\sigma_{\text{pub}}^{-2}\left(t\sigma_{\text{priv}}^{-2}\right)}{\left(\sigma_\xi^{-2} + t\sigma_{\text{pub}}^{-2}\right)\left(\sigma_\xi^{-2} + t\sigma_{\text{pub}}^{-2} + t\sigma_{\text{priv}}^{-2}\right)} \;>\; 0, \tag{22}$$

with the gap growing in t and in the precision ratio $\sigma^2_{\text{pub}}/\sigma^2_{\text{priv}}$. *(iv) Limits.* $\lim_{t\to\infty} w_t^{C,\text{pub}} \equiv w_\infty^{C,\text{pub}} = \sigma^2_{\text{priv}}/(\sigma^2_{\text{pub}} + \sigma^2_{\text{priv}}) \in (0,1)$, so the public weight at the connected forecaster is bounded strictly below 1 for all t.

A-IV.2 Predictions for individual forecasters

The forecaster's expected spillover for country j in episode e is $\mathbb{E}^f[\Delta y_{j,e} \mid \mathcal{I}_t^f] = \gamma_{ij}\,\mathbb{E}[\xi_e \mid \mathcal{I}_t^f]$, since $u_{j,e}$ is mean-zero and independent. Substituting Equation 16 and Equation 21 gives

$$\begin{aligned}\mathbb{E}^U[\Delta y_{j,e} \mid \mathcal{I}_t^U] \;&=\; \gamma_{ij}\, w_t^U\left(\bar{s}_t^{\text{pub}} - bx_{ij}\right) \;+\; \gamma_{ij}\, w_t^U \cdot bx_{ij} \\ &=\; \gamma_{ij}\, w_t^U\, \bar{s}_t^{\text{pub}}.\end{aligned} \tag{23}$$

It should be noted that Equation 23 uses the unconnected forecaster's naive treatment: she does not know b and therefore does not subtract bx_{ij} from s_t^{pub} before forming her posterior. Algebraically, this is equivalent to adopting Equation 16 but conditioning the expectation on the same $\bar{s}_t^{\text{pub}}$ that an external observer sees. Equivalently:

$$\mathbb{E}^U[\Delta y_{j,e} \mid \mathcal{I}_t^U] \;=\; \gamma_{ij}\, w_t^U\, \bar{s}_t^{\text{pub}} \;=\; \gamma_{ij}\, w_t^U\,(\xi_e + bx_{ij} + \bar{\eta}_t^{\text{pub}}), \tag{24}$$

with $\bar{\eta}_t^{\text{pub}} = t^{-1}\sum_{\tau\leq t}\eta_\tau^{\text{pub}} \sim \mathcal{N}(0, \sigma^2_{\text{pub}}/t)$. Similarly, for the connected forecaster (who removes the bias):

$$\mathbb{E}^C[\Delta y_{j,e} \mid \mathcal{I}_t^C] \;=\; \gamma_{ij}\left[w_t^{C,\text{pub}}(\bar{s}_t^{\text{pub}} - bx_{ij}) + w_t^{C,\text{priv}}\bar{s}_t^{\text{priv}}\right]. \tag{25}$$

Prediction 1 (Gradual revision) Equation 23 and Equation 25 both have level zero at $t = 0$: when no signal has arrived, the posterior equals the prior mean (zero), so the expected spillover is zero. For

$t \geq 1$, taking expectations over the signal noise $(\bar{\eta}_t^{\text{pub}}, \bar{\eta}_t^{\text{priv}})$ conditional on the realized fundamental,

$$\mathbb{E}\big[\mathbb{E}^U[\Delta y_{j,e} \mid \mathcal{I}_t^U] \,\big|\, \xi_e\big] \;=\; \gamma_{ij}\, w_t^U\, (\xi_e + b x_{ij}), \tag{26}$$

$$\mathbb{E}\big[\mathbb{E}^C[\Delta y_{j,e} \mid \mathcal{I}_t^C] \,\big|\, \xi_e\big] \;=\; \gamma_{ij}\, (w_t^{C,\text{pub}} + w_t^{C,\text{priv}})\, \xi_e. \tag{27}$$

For both types, expected spillovers are increasing in t (because the Kalman weights are monotone in t, by Property (i) above) and equal zero at $t = 0$. Translating to working-day calendar time, this delivers the no-jump-positive-slope path documented in Table 3. □

Prediction 2 (Information channel) Subtracting Equation 25 from Equation 23,

$$\begin{aligned}\beta_t \;&\equiv\; \mathbb{E}^C[\Delta y_{j,e} \mid \mathcal{I}_t] - \mathbb{E}^U[\Delta y_{j,e} \mid \mathcal{I}_t] \\ &=\; \gamma_{ij}\big\{w_t^{C,\text{pub}}(\bar{s}_t^{\text{pub}} - b x_{ij}) + w_t^{C,\text{priv}} \bar{s}_t^{\text{priv}} - w_t^U \bar{s}_t^{\text{pub}}\big\}.\end{aligned} \tag{28}$$

Taking the expectation over the signal noise conditional on ξ_e (so $\bar{s}_t^{\text{pub}} \to \xi_e + b x_{ij}$ and $\bar{s}_t^{\text{priv}} \to \xi_e$),

$$\begin{aligned}\mathbb{E}[\beta_t \mid \xi_e] \;&=\; \gamma_{ij}\big\{w_t^{C,\text{pub}} \xi_e + w_t^{C,\text{priv}} \xi_e - w_t^U(\xi_e + b x_{ij})\big\} \\ &=\; \gamma_{ij}\big\{(w_t^{C,\text{pub}} + w_t^{C,\text{priv}} - w_t^U)\xi_e - w_t^U\, b x_{ij}\big\}.\end{aligned} \tag{29}$$

Taking the further expectation over $\xi_e \sim \mathcal{N}(0, \sigma_\xi^2)$ kills the first term:

$$\mathbb{E}[\beta_t] \;=\; -\,\gamma_{ij}\, w_t^U\, b\, x_{ij} \;\leq\; 0. \tag{30}$$

This is the main text's Equation 9. The dampening is non-positive, equal to zero only when $b = 0$ or $x_{ij} = 0$, and monotonically increasing in magnitude as t grows (Property (i)) and as the precision ratio $\sigma^2_{\text{pub}}/\sigma^2_{\text{priv}}$ grows. As $t \to \infty$, $w_t^U \to 1$ and $|\mathbb{E}[\beta_t]|$ converges to the finite long-run limit $\gamma_{ij}\, b\, x_{ij}$. The connected-vs-unconnected wedge therefore builds monotonically over working days toward a finite steady-state value. □

Prediction 3 (Heterogeneity by exposure) Differentiate Equation 30 with respect to x_{ij}, treating the spillover loading $\gamma_{ij} = \bar{\gamma} + \lambda x_{ij}$ as a function of x_{ij} (so $\partial\gamma_{ij}/\partial x_{ij} = \lambda > 0$):

$$\begin{aligned}\beta_{1,t} \;&\equiv\; \frac{\partial \mathbb{E}[\beta_t]}{\partial x_{ij}} \\ &=\; -\,b\, w_t^U\, \big(\gamma_{ij} + x_{ij}\, \lambda\big) \\ &=\; -\,b\, w_t^U\, \big(\bar{\gamma} + 2\lambda x_{ij}\big) \;<\; 0.\end{aligned} \tag{31}$$

Two forces compose: the bias term $b x_{ij}$ in the public narrative grows in x_{ij} (the linear factor outside the differentiation), and the spillover loading γ_{ij} grows in x_{ij} as well (the linear factor inside). Both forces pull the connected-vs-unconnected gap in the same negative direction. The text writes this as Equation 10, with the same sign and the same scaling in $\bar{\gamma} + 2\lambda x_{ij}$. □

Why the heterogeneity loads on energy and intermediates but not on capital or consumption goods Equation 31 is proportional to b, the public-narrative bias slope. Where the public narrative is most easily organized around a quantitative bilateral linkage, b is large; where it is not, b is small. Empirically, the energy and intermediate-goods channels are repeatedly told as quantitative narratives in macroeconomic media coverage ("country j imports 50 percent of its gas from country i," "country j's industries depend on country i's critical inputs"), while capital- and consumption-goods exposures are too granular to organize into a headline. The components decomposition in Table 6 thereby identifies which margins are the salient bilateral narratives the model's b is meant to capture.

A-IV.3 Mapping the model to the empirical regressions

Forecast revisions The data observe a forecaster's posterior mean as a real-time submitted forecast; the empirical revisions Δy in Equation 2 are differences between two such posteriors. Within a forecaster type f and at a single horizon h, the revision between the pre-cutoff baseline at $t = 0$ and a post-cutoff observation at $t \geq 1$ is

$$\Delta y^{h}_{j,f,t,e} \;=\; \mathbb{E}^{f}[\Delta y_{j,e} \mid \mathcal{I}^{f}_{t}] - \mathbb{E}^{f}[\Delta y_{j,e} \mid \mathcal{I}^{f}_{0}] \;=\; \mathbb{E}^{f}[\Delta y_{j,e} \mid \mathcal{I}^{f}_{t}], \tag{32}$$

since the prior mean (the baseline) is zero by assumption. Substituting Equation 24 and Equation 25 yields type-specific revisions that are exactly what we estimate in Equation 2.

Mapping to β in Equation 2 The main text's Equation 2 regresses the empirical revision on $\text{Connected}_{f,i(e)}$ within episode×country×date groups. The within-pair expected difference between connected and unconnected forecasters is exactly β_t in Equation 30. Averaging over the post-cutoff window $t \in [1, T]$ of the regression and over the cross-section of episodes and countries gives

$$\hat{\beta} \;\xrightarrow{p}\; -\,\bar{w}^{U}\,\overline{\gamma_{ij}\, b\, x_{ij}}, \tag{33}$$

where $\bar{w}^{U}$ is the time-averaged public-Kalman weight over the window and the overline denotes the cross-section average over (i, j, e) tuples. The model's prediction is therefore that $\hat{\beta} < 0$, with magnitude scaled by the average exposure-weighted bias—both observed in Table A2.

Mapping to β_1 in Equation 3 The main text's Equation 3 replaces the linear x_{ij} argument with a coarser $\text{HighExp}_{j,e}$ indicator (above the within-episode median or above the BACI 2 percent threshold). Coarsening turns the continuous linear comparative statics in Equation 31 into a step-function statement: $\hat{\beta}_0$ is the model's $\mathbb{E}[\beta_t]$ averaged over low-exposure pairs, and $\hat{\beta}_0 + \hat{\beta}_1$ is the same average over high-exposure pairs. Both are negative (and the second more so) by Equation 30 and Equation 31.

Mapping to δ in the RDD Equation 1 The RDD slope δ measures the rate at which post-cutoff forecasts walk upward per working day. In the model, δ corresponds to $\gamma_{ij}\, \partial \mathbb{E}[\xi_e \mid \mathcal{I}^{f}_{t}] / \partial t$. Differentiating

Equation 26 with respect to t and evaluating in the cross-section average yields a positive slope that decays in t (the Kalman weight is concave in t for $t \geq 1$). The pooled-across-eight-episodes Figure 2 is consistent with the positive post-cutoff slope predicted by the model; the slope decay implied by the concavity of the Kalman weight is harder to read off the figure.

A-IV.4 Aggregation and macroeconomic implications

Consensus aggregator Let $\rho_{j,e}$ denote the share of forecasters covering country j in episode e who are connected (i.e., satisfy $z_{f,e} = 1$). The unweighted cross-forecaster average within an episode×country×date group is

$$\begin{aligned} \mathbb{E}^{\text{cons}}_{e,t}[\Delta y_{j,e}] &= \frac{1}{|\mathcal{F}_{j,e}|} \sum_{f \in \mathcal{F}_{j,e}} \mathbb{E}^{f}[\Delta y_{j,e} \mid \mathcal{I}^{f}_{t}] \\ &= \rho_{j,e}\, \mathbb{E}^{C}[\Delta y_{j,e} \mid \mathcal{I}^{C}_{t}] + (1 - \rho_{j,e})\, \mathbb{E}^{U}[\Delta y_{j,e} \mid \mathcal{I}^{U}_{t}], \end{aligned} \tag{34}$$

where $\mathcal{F}_{j,e}$ is the set of forecasters who covered j in the post-cutoff window of episode e. Equation 34 is the main text's Equation 11.

Aggregate bias (Proposition 1, network mediation) Take expectations of Equation 34 over the signal noise conditional on ξ_e (so $\bar{s}^{\text{pub}}_t \to \xi_e + b x_{ij}$ and $\bar{s}^{\text{priv}}_t \to \xi_e$), then over $\xi_e \sim \mathcal{N}(0, \sigma^2_\xi)$. The unconnected contribution becomes

$$\mathbb{E}\big[\mathbb{E}^{U}[\Delta y_{j,e} \mid \mathcal{I}^{U}_{t}]\big] = \gamma_{ij}\, w^{U}_{t}\, \mathbb{E}[\xi_e + b x_{ij}] = \gamma_{ij}\, w^{U}_{t}\, b x_{ij},$$

since $\mathbb{E}[\xi_e] = 0$. The connected contribution becomes

$$\mathbb{E}\big[\mathbb{E}^{C}[\Delta y_{j,e} \mid \mathcal{I}^{C}_{t}]\big] = \gamma_{ij}\, (w^{C,\text{pub}}_{t} + w^{C,\text{priv}}_{t}) \cdot \mathbb{E}[\xi_e] = 0,$$

because the connected forecaster has corrected the bias, so the bias term drops out of his posterior. Combining,

$$\begin{aligned} \mathbb{E}\big[\mathbb{E}^{\text{cons}}_{e,t}[\Delta y_{j,e}]\big] &= \rho_{j,e} \cdot 0 + (1 - \rho_{j,e}) \cdot \gamma_{ij}\, w^{U}_{t}\, b x_{ij} \\ &= (1 - \rho_{j,e})\, \gamma_{ij}\, b\, x_{ij}\, w^{U}_{t}. \end{aligned} \tag{35}$$

Subtracting the rational-expectations benchmark $\mathbb{E}[\gamma_{ij}\xi_e] = 0$ and taking the long-t limit ($w^{U}_{\infty} \to 1$) yields

$$\mathbb{E}\big[\mathbb{E}^{\text{cons}}_{e,\infty}[\Delta y_{j,e}] - \gamma_{ij}\xi_e\big] = (1 - \rho_{j,e})\, \gamma_{ij}\, b\, x_{ij}\, w^{U}_{\infty}. \tag{36}$$

This is the main text's Equation 12. □

Network mediation: necessity of both networks Equation 36 factors as a product of (i) the spillover loading $\gamma_{ij} = \bar{\gamma} + \lambda x_{ij}$, (ii) the bilateral-exposure scalar bx_{ij}, and (iii) the connectedness-complement scalar $1 - \rho_{j,e}$. Setting $\rho_{j,e} = 1$ (every forecaster covering j also covers i) collapses the bias to zero regardless of x_{ij}. Setting $x_{ij} = 0$ (no bilateral linkage) collapses the bias to zero regardless of $\rho_{j,e}$. The two networks compose multiplicatively and each is necessary. □

Cross-country dispersion (Corollary 1) Conditional on ξ_e, the cross-country variance of consensus inflation expectations is

$$\begin{aligned}\mathrm{Var}_j\big(\mathbb{E}^{\text{cons}}_{e,\infty}[\Delta y_{j,e}] \,\big|\, \xi_e\big) &= \mathrm{Var}_j\big(\gamma_{ij}\xi_e + (1-\rho_{j,e})\gamma_{ij} b x_{ij}\, w^U_\infty\big) \\ &= \xi_e^2\,\mathrm{Var}_j(\gamma_{ij}) \;+\; (b\,w^U_\infty)^2\,\xi_e^2\,\mathrm{Var}_j\big[(1-\rho_{j,e})\gamma_{ij}x_{ij}\big] \\ &\quad +\; 2\,b\,w^U_\infty\,\xi_e^2\,\mathrm{Cov}_j\big[\gamma_{ij},(1-\rho_{j,e})\gamma_{ij}x_{ij}\big]. \end{aligned} \tag{37}$$

The text reports the first two terms in Equation 13; the cross-covariance term is positive when $\rho_{j,e}$ is unrelated to γ_{ij} (countries with high γ_{ij} are not systematically the most-covered) and is in any case dominated by the first two terms when the cross-section of $\rho_{j,e}$ is wide. Both first two terms are non-negative; the second is the network-bias contribution and is the term the model identifies as "new" relative to a rational-expectations benchmark. □

Phillips-curve closure (Proposition 2) Take a country-level New Keynesian Phillips curve at horizon t:

$$\pi_{j,t} \;=\; \kappa\, mc_{j,t} \;+\; \theta\,\mathbb{E}^{\text{cons}}_{e,t}[\pi_{j,t+1}], \tag{38}$$

where $\pi_{j,t}$ and $mc_{j,t}$ are the log-deviations of inflation and real marginal cost from their steady-state values in country j at date t, $\kappa > 0$ is the standard New-Keynesian slope, and $\theta \in (0,1)$ is the subjective discount factor. The wedge derivation below isolates the deviation generated by the conflict shock relative to the same steady state, so all magnitudes reported below (basis points of inflation) are differences in this deviation between the consensus and full-information solutions; the common steady-state component cancels. Iterating Equation 38 forward, with $s \geq 0$ indexing future periods,

$$\pi_{j,t} \;=\; \kappa \sum_{s=0}^{\infty} \theta^s\, \mathbb{E}^{\text{cons}}_{e,t}[mc_{j,t+s}].$$

For comparison, the full-information benchmark replaces $\mathbb{E}^{\text{cons}}_e$ with the true expectation $\mathbb{E}$, yielding $\pi^{\text{full info}}_{j,t} = \kappa \sum_{s\geq 0} \theta^s \mathbb{E}[mc_{j,t+s}]$. The wedge between realized inflation and the full-information benchmark is therefore

$$\pi_{j,t} - \pi^{\text{full info}}_{j,t} \;=\; \kappa \sum_{s=0}^{\infty} \theta^s \left(\mathbb{E}^{\text{cons}}_{e,t}[mc_{j,t+s}] - \mathbb{E}[mc_{j,t+s}]\right).$$

For a one-period spillover, $mc_{j,t+1} = c\,\Delta y_{j,e}$ for some constant $c > 0$ (real marginal cost is shifted by the conflict spillover), so the leading term in the wedge is the bias in $\mathbb{E}^{\text{cons}}_{e,t}[\Delta y_{j,e}] - \mathbb{E}[\Delta y_{j,e}]$. Substituting Equation 35 and discounting one period, and absorbing the product $\kappa\, c$ into the calibrated wedge constant,

$$\pi_{j,t} - \pi^{\text{full info}}_{j,t} \;\approx\; \theta\,(1-\rho_{j,e})\,\gamma_{ij}\,b\,x_{ij}\,w^{U}_{\infty} \;+\; O(\theta^2). \tag{39}$$

This is the main text's Equation 14. The higher-order θ^2 correction collects the contributions from $s \geq 2$, which are damped by additional powers of the discount factor and by the public-data revelation of ξ_e over time. □

Public-data substitutability (Corollary 2) Define the aggregate bias as $B(\rho, \sigma^{-2}_{\text{pub}}) = (1-\rho)\,\gamma_{ij}\,b\,x_{ij}\,w^{U}_{\infty}(\sigma^{-2}_{\text{pub}})$, where $w^{U}_{\infty} = \sigma^{-2}_{\text{pub}}/(\sigma^{-2}_{\xi} + \sigma^{-2}_{\text{pub}})$ is increasing in σ^{-2}_{pub} in the long-t limit. Total differentiation along the iso-bias contour $dB = 0$:

$$\frac{\partial B}{\partial \rho}\,d\rho \;+\; \frac{\partial B}{\partial \sigma^{-2}_{\text{pub}}}\,d\sigma^{-2}_{\text{pub}} \;=\; 0.$$

Substituting,

$$\frac{\partial B}{\partial \rho} \;=\; -\gamma_{ij}\,b\,x_{ij}\,w^{U}_{\infty} \;<\; 0, \tag{40}$$

$$\frac{\partial B}{\partial \sigma^{-2}_{\text{pub}}} \;=\; (1-\rho)\,\gamma_{ij}\,b\,x_{ij}\,\frac{\partial w^{U}_{\infty}}{\partial \sigma^{-2}_{\text{pub}}} \;>\; 0. \tag{41}$$

Solving,

$$\left.\frac{d\rho}{d\sigma^{-2}_{\text{pub}}}\right|_{B \text{ fixed}} \;=\; -\,\frac{\partial B/\partial \sigma^{-2}_{\text{pub}}}{\partial B/\partial \rho} \;=\; \frac{(1-\rho)}{w^{U}_{\infty}}\,\frac{\partial w^{U}_{\infty}}{\partial \sigma^{-2}_{\text{pub}}} \;>\; 0, \tag{42}$$

which is the main text's Corollary 2. The substitutability is exact and computable from the network primitives and the precision of public data. □

Proof of Corollary 3 (policy wedge) The central bank minimizes $\frac{1}{2}\,\mathbb{E}[\pi^2_{j,t} + \vartheta\, y^2_{j,t}]$ subject to $\pi_{j,t} = \kappa\, y_{j,t} + \theta\,\mathbb{E}_t[\pi_{j,t+1}] + u_{j,t}$, taking the forward expectation as given (discretion). The first-order condition in $y_{j,t}$ is $\kappa\,\pi_{j,t} + \vartheta\, y_{j,t} = 0$, i.e. the targeting rule $y_{j,t} = -(\kappa/\vartheta)\,\pi_{j,t}$. Substituting the targeting rule into the Phillips curve and solving,

$$y_{j,t} \;=\; -\,\frac{\kappa}{\vartheta + \kappa^2}\Big(\theta\,\mathbb{E}_t[\pi_{j,t+1}] + u_{j,t}\Big). \tag{43}$$

Equation 43 is linear in the forward expectation, so replacing the full-information expectation $\mathbb{E}^{*}_{t}[\pi_{j,t+1}]$ with the observed consensus shifts the chosen gap by $-\kappa\theta/(\vartheta + \kappa^2)$ times the expectation bias. By Equation 12 that bias is $(1-\rho_{j,e})\gamma_{ij}\,b\,x_{ij}\,w^{U}_{\infty}$, and by Equation 14 the inflation wedge is θ times it,

$W_{j,e} = \theta\,(1-\rho_{j,e})\gamma_{ij}\,b\,x_{ij}\,w_\infty^U$. Writing the shift in terms of $W_{j,e}$ therefore cancels the discount factor:

$$\hat{y}_{j,t} - y^*_{j,t} \;=\; -\frac{\kappa\,\theta}{\vartheta+\kappa^2}\cdot\frac{W_{j,e}}{\theta} \;=\; -\frac{\kappa}{\vartheta+\kappa^2}\,W_{j,e},$$

which is Equation 15. It is negative because $W_{j,e} \geq 0$: the observed consensus overstates inflationary pressure, so the bank contracts by more than the fundamentals require.

For the loss, note that $y^*_{j,t}$ minimizes $L(y) = \frac{1}{2}(\pi^2 + \vartheta y^2)$ evaluated along the true Phillips curve $\pi = \kappa y + \theta\mathbb{E}^* + u$. Since L is quadratic in y with $L'' = \vartheta + \kappa^2$, a second-order expansion around the optimum is exact:

$$\Delta L_{j,e} \;=\; \tfrac{1}{2}\,(\vartheta+\kappa^2)\,\big(\hat{y}_{j,t} - y^*_{j,t}\big)^2 \;=\; \tfrac{1}{2}\,\frac{\kappa^2}{\vartheta+\kappa^2}\,W_{j,e}^2.$$

The excess loss is therefore quadratic in the wedge while the policy error is linear in it. Because $W_{j,e}$ is increasing in $(1-\rho_{j,e})$ and in x_{ij}, both expressions inherit those comparative statics. □

A-IV.5 Extensions

Repeated episodes and anchoring Consider a sequence of episodes $e = 1, 2, \ldots$ with conflict shocks ξ_e drawn i.i.d. from a stationary distribution. Let $\bar{B} = \mathbb{E}_e[(1-\rho_{j,e})\gamma_{i(e),j}\,b\,x_{i(e),j}\,w_\infty^U]$ denote the average aggregate bias per episode, where the expectation is over the random episode draw $i(e)$ and over the cross-section of j. After E episodes, the cumulative wedge in long-run inflation expectations satisfies

$$\mathbb{E}\big[\mathbb{E}^{\text{cons}}_{j,\infty}[\pi_j] - \pi_j^{\text{full info}}\big] \;\sim\; E\cdot\bar{B},$$

which diverges as $E \to \infty$ unless $\bar{B} = 0$. The sufficient condition for anchoring is therefore that $\mathbb{E}_e[(1-\rho_{j,e})x_{i(e),j}] = 0$, which holds if and only if for every (importer j, exporter i) pair with $x_{ij} > 0$, the forecaster network supplies positive coverage in expectation: $\mathbb{E}_e[1-\rho_{j,e} \mid i(e) = i] < 1$. In words, the forecaster network must not leave any high-exposure pair entirely uncovered across the episode distribution. □

Asymmetric narrative bias Allow the bias slope to depend on the sign of the shock: $b = b_+$ if $\xi_e \geq 0$ and $b = b_-$ if $\xi_e < 0$, with $b_+ > b_- \geq 0$ to capture media-attention asymmetries that emphasize upside inflation news. The asymmetry is documented in the media-economics literature (Mullainathan and Shleifer, 2005; Eisensee and Strömberg, 2007). We do not test it directly in our own data because the eight conflict episodes in our headline sample are all upside-inflation events—supply-side disruptions of energy or intermediate inputs—so b_+ and b_- cannot be separately identified from the cross-episode variation we have. We treat the asymmetry as a microfoundational extension rather than as a primitive estimated from our sample. Replacing b in Equation 16 and Equation 18 with $b_\pm$ and following the steps in §D-IV, the long-run aggregate bias becomes

$$\mathbb{E}\big[\mathbb{E}^{\text{cons}}_{e,\infty}[\Delta y_{j,e}]\big] \;=\; (1-\rho_{j,e})\,\gamma_{ij}\,x_{ij}\,w_\infty^U\,\mathbb{E}_{\xi_e}\big[b_+\mathbf{1}\{\xi_e \geq 0\} + b_-\mathbf{1}\{\xi_e < 0\}\big]. \tag{44}$$

The expectation in brackets equals $\frac{1}{2}(b_+ + b_-)$ under the symmetric prior $\xi_e \sim \mathcal{N}(0, \sigma_\xi^2)$, but the conditional bias on positive ξ_e is b_+, larger than the conditional bias on negative ξ_e. The unconditional aggregate bias therefore tilts upward whenever positive realizations are more salient than negative ones. Equation 44 provides a microfoundation for the systematically upward direction of the post-cutoff revisions documented in Subsection 3.1: the asymmetric narrative interacts with sparse forecaster connectedness to produce a structurally upward-biased consensus. □

A-IV.6 The naive-naive case

For completeness, we report the alternative case in which both connected and unconnected forecasters treat s_t^{pub} as unbiased (i.e., neither knows the bias slope b). In that case, the connected forecaster's posterior is Equation 18 with $b = 0$ in the public-stream contribution:

$$\mathbb{E}[\xi_e \mid \mathcal{I}_t^C] \;=\; w_t^{C,\text{pub}}\, \bar{s}_t^{\text{pub}} + w_t^{C,\text{priv}}\, \bar{s}_t^{\text{priv}}.$$

The information channel becomes

$$\mathbb{E}[\beta_t] \;=\; -\,\gamma_{ij}\,(w_t^{U} - w_t^{C,\text{pub}})\, b\, x_{ij}, \tag{45}$$

an alternative to the main text's Equation 9, which adopts the bias-correcting case in Equation 30. The sign and the qualitative comparative statics are unchanged. The two cases (connected forecaster removes the bias vs. does not) produce algebraically distinct expressions but identical predictions for sign, comparative statics, and the network-mediation theorem; they differ only in the quantitative magnitude of the dampening, which is larger in Equation 30 (the bias-correcting case, adopted in the main text) than in Equation 45 (the naive case) by a factor of $w_t^U / (w_t^U - w_t^{C,\text{pub}}) > 1$.

The two cases bracket the empirical estimates: the bias-correcting case predicts a larger dampening, the naive-naive case a smaller one, and both predict negative $\hat{\beta}$ that grows in $b\, x_{ij}$. The headline empirical patterns reported in Section 3 are consistent with either microfounding; distinguishing them would require auxiliary identifying assumptions on b, which we leave to future work.