What do forecasters say they don't know?

Elicited macroeconomic uncertainty in the US and euro-area Surveys of Professional Forecasters, 1968–2026 — working draft r5

Max Ghenis

2026-09-19

Abstract

Professional forecasters state their own uncertainty four times a year. Since 1968 the US Survey of Professional Forecasters has collected probability distributions — not just point forecasts — for output growth and inflation, and the ECB's survey has done the same for the euro area since 1999. We pool every reported histogram by the law of total variance, decomposing total uncertainty into the average forecaster's stated uncertainty and the disagreement between forecasters, for every variable, horizon, and survey round — 3,695 round-target distributions. Five findings organize the panorama. Stated growth uncertainty is stable: its pandemic spike decayed by 2026 on both continents, while euro-area inflation uncertainty remains 34% above its 2010–19 average. Disagreement — the only uncertainty visible in point-forecast data — is the minority share everywhere: a median 14.7–15.3% of total variance for growth, at most 30% for any variable. Stated uncertainty stops rising with horizon by year two. Past year three in the US and year five in the euro area, no survey elicits a distribution at all: the ten-year questions collect point forecasts, and their spread measures how much forecasters agree, which says nothing about how sure any of them is. A ±1σ band around the US consensus covered 66.7% of next-year growth outcomes over 33 years, and the pooled 90% interval 78.8%; euro-area bands covered less; both surveys miss in clusters at regime shifts, in the tail the regime moved toward. Scored properly, the densities beat climatology everywhere and roughly tie a Gaussian built from the consensus and its own error history (-0.9% to +8.7% CRPS skill). Density elicitation stops three to five years out. Published scenarios of AI's growth effects differ in outcome, horizon and conditioning (Shenk 2026; Cunningham 2025), and most concern horizons no survey density reaches; we show what asking would take. A live interactive tracks every series.

1 Introduction

In the first quarter of 2026, the forecasters in the US Survey of Professional Forecasters reported probability distributions for next year's real GDP growth that pool to a median of 2.0%, and the distributions say how sure they were: the middle half of the pooled distribution spans 1.4 points. Asked about average growth over the next ten years, the same panel gave a median of 2.1% and nothing else. The survey collects no distribution for the decade. The ten-year answers sit within 0.2 points of each other, which shows that the forecasters agree, and agreement is a different quantity from confidence. The histograms stop three years out.

That gap matters more now than at any point in the surveys' history. Published work on what artificial intelligence will do to growth shares neither an outcome nor a horizon. Acemoglu (2024) puts the gain in total factor productivity at under 0.1 percentage points a year over a decade; Briggs and Kodnani (2023) put the boost to annual labor productivity growth near 1.5 points; Aghion and Bunel (2024) get 0.07 to 1.24 points of productivity growth from one framework under different readings of the evidence; Davidson (2021) and Erdil and Besiroglu (2023) study the conditions under which gross world product could grow above 30% a year, given advanced AI and at horizons running to 2100. Compilations set

these side by side, almost entirely as point estimates (Shenk 2026; Cunningham 2025). A productivity increment, a conditional world-growth regime and an unconditional forecast of US GDP are different quantities, and no single range spans them. The one large elicitation, run by the Forecasting Research Institute with the Federal Reserve Bank of Chicago, collected medians and two quantiles, once (Karger et al. 2026). The instruments that would discipline this debate — surveys that make professionals state distributions and then score them — exist, run quarterly, and have five decades of track record. They stop at year three in the US and year five in the euro area.

This paper assembles that track record in full. We compute one construction — the law-of-total-variance decomposition of pooled survey histograms (Wallis 2005; Boero et al. 2008) — for every density variable, every horizon, and every round of the two long-running professional surveys, 3,695 round-target distributions in all, and evaluate every one against realized outcomes where a verified realization concept exists (3,070 forecast-outcome pairs). The distinction we track back to 1968 is Zarnowitz and Lambros's (1987): *disagreement* — the spread between forecasters' means — is not *uncertainty* — the spread each forecaster reports around their own mean.

> **i** Draft note
>
> Contribution list and positioning paragraph still to come; nearest antecedents: D'Amico and Orphanides (2008) (US long series, inflation, ends 2008), Clements et al. (2022) (both surveys, in prose), Allayioti et al. (2024) (ECB 25-year retrospective, one survey).

2 The elicitation landscape

Two surveys regularly ask professional forecasters for probability distributions over macroeconomic outcomes. The US SPF has collected histograms since 1968Q4 — for nominal GNP growth at first, real GNP from 1981Q3, real GDP from 1992Q1 — for output growth and the GDP price index throughout, joined by fourth-quarter-over-fourth-quarter core CPI and core PCE inflation in 2007Q1 and the unemployment rate in 2009Q2 (Federal Reserve Bank of Philadelphia 2026). Density horizons ran to the next calendar year until 2009Q2 and to three years out since. The ten-year variables — real GDP growth, CPI, PCE — are point forecasts only. The ECB SPF has run quarterly since 1999Q1: densities for HICP inflation, real GDP growth, and unemployment (core HICP since 2017Q1) at current-year through year-after-next calendar targets, rolling one- and two-year horizons, and a longer-term calendar year four years ahead in Q1/Q2 rounds and five ahead in Q3/Q4 (García 2003; Allayioti et al. 2024).

Elsewhere, elicited densities stop early or do not exist. The Bank of England's Survey of External Forecasters collects UK inflation and GDP densities (Boero et al. 2008); the NY Fed's consumer survey elicits household densities to three years; Consensus Economics elicits by-range probabilities at one year and publishes its 5–10-year forecasts as points; Livingston (since 1946) and Blue Chip are point products; the FOMC's projections come with fan charts built from historical errors, not elicited beliefs. The two long-horizon precedents are one-off quantile elicitations (Christensen et al. 2018; Karger et al. 2026). No survey anywhere regularly elicits probability distributions over ten-year-ahead growth.

3 Measures

Each forecaster attaches probabilities to bins. The questionnaires print closed one-decimal ranges — "2.0 to 2.9", then "3.0 to 3.9" — and growth is continuous, so we read each range as running to the next range's lower label: "2.0 to 2.9" covers [2, 3). Lower labels stay where the questionnaire prints them. We spread each bin's mass uniformly inside the bin, give open tails the width of the adjacent bin, keep responses whose probabilities lie between 0 and 100 and sum to within two points of 100, and renormalize. Moments, quantiles, and scores all use this one distribution, so a respondent's variance carries the within-bin term $w_b^2/12$, and disagreement is the population variance across respondent

means — the pooled mixture's variance then equals the two parts exactly. For forecaster i, the histogram yields a mean m_i and a variance v_i. The law of total variance splits the pooled mixture (Wallis 2005):

$$\text{total variance} = \mathbb{E}[v_i] + \text{Var}\ (m_i).$$

The first term is the average uncertainty forecasters state individually; the second is disagreement, the only term visible in point-forecast data (Zarnowitz and Lambros 1987; Lahiri and Sheng 2010). We report the square roots, the between share, and quantiles of the pooled mixture. The interquartile range of the mixture doubles as a robustness measure: bin-scheme changes move midpoint-based SDs mechanically (Shoja and Soofi 2017; Knüppel and Pavlova 2026), and the 2020Q2 US bin widening is the worked example (Table 1).

Table 1: The wide-bin era, averaging every round and horizon in each window equally, and the clean 2024Q1-to-Q2 scheme boundary at the next-year horizon. Moments move further than the pooled IQR, which is the signature of coarser bins rather than of belief alone.

Variable	Window	Rounds	Within SD	Total SD	Pooled IQR	Total/IQR
US real GDP	pre (2016Q2–2020Q1)	16	1.142	1.238	1.440	0.859
US real GDP	wide (2020Q2–2024Q1)	16	1.848	2.046	1.988	1.029
US real GDP	post (2024Q2–2026Q3)	10	1.291	1.360	1.466	0.928
US real GDP	2024Q1	1	1.737	1.862	1.551	1.200
US real GDP	2024Q2	1	1.353	1.453	1.483	0.979
US unemployment	pre (2016Q2–2020Q1)	16	0.667	0.767	1.077	0.712
US unemployment	wide (2020Q2–2024Q1)	16	0.996	1.194	1.451	0.823
US unemployment	post (2024Q2–2026Q3)	10	0.605	0.658	0.742	0.886
US unemployment	2024Q1	1	1.012	1.091	1.208	0.903
US unemployment	2024Q2	1	0.656	0.705	0.789	0.893

Targets are calendar years, so a "next year" forecast made in Q1 and one made in Q4 differ in horizon by three quarters (Ganics et al. 2024). We treat round-target pairs explicitly and show the within-year shrinkage — current-year uncertainty falls 23–35% from Q1 to Q4, depending on the variable (Table 2) — rather than averaging over it (Clements et al. 2025).

Table 2: Mean current-year total SD by survey quarter, 2010-19 rounds. The fixed calendar event arrives whatever the round, so the fall from the Q1 to the Q4 round tracks three more quarters of data on the target year; the panel's composition also changes between rounds.

Variable	Q1	Q2	Q3	Q4	Q1 to Q4
US real GDP	1.123	1.010	0.893	0.727	−35.3%
US GDP prices	0.743	0.702	0.665	0.555	−25.4%
US unemployment	0.602	0.548	0.479	0.394	−34.6%
US core CPI	0.622	0.587	0.530	0.461	−25.9%
US core PCE	0.605	0.561	0.542	0.468	−22.7%
ECB HICP	0.601	0.562	0.494	0.422	−29.8%
ECB core HICP	0.527	0.494	0.470	0.406	−22.8%
ECB real GDP	0.661	0.614	0.554	0.464	−29.8%
ECB unemployment	0.639	0.595	0.549	0.493	−22.9%

Coverage rates test the density at two thresholds; proper scores evaluate all of it (Gneiting and Raftery 2007; Boero et al. 2011). The pooled CDF is piecewise linear, so the continuous ranked probability score has an exact closed form, and pinball losses at the published quantile levels decompose it by tail. We score every calibrated round against two no-lookahead benchmarks — expanding-window climatology and a Gaussian centered on the consensus with historical-error RMSE — and against the average individual forecaster. One structural note: the pooled CDF is the mean of individual CDFs and the CRPS is convex in the CDF, so the pool beats the average individual on every observation by construction; the empirical content is the size of that gain (9.1% on average, 3.4–18.4% by cell) and where it thins.

4 Five stylized facts

> **i** Draft note
>
> Each fact still wants a figure. Every inline value and table below computes at render from `outputs/panorama_*.csv`, including the balanced-sample term structure and the post-2021 ECB comparison.

Stated growth uncertainty is stable — and its pandemic spike decayed on both continents, while euro-area inflation uncertainty has not come home. The pooled SD of next-year US growth was 1.27 points in 1992 and averaged 1.29 through 2020; it peaked at 2.18 in 2021 and stands at 1.41 in 2026 (decade means: 1.28 over 2010–19, 1.80 over 2020–24). On the matched construction — Q1 rounds, next calendar year — euro-area growth uncertainty peaked at 1.73 in 2021, below the US peak, stands at 0.78 in 2026, and sat below the US level in 28 of the 28 years both surveys cover. Euro-area inflation uncertainty did not: the one-year series remains 48% above its pre-2021 level in 2026 (Table 3).

Table 3: Total SD of the headline series: US Q1 next-calendar-year rounds, and for the ECB the annual mean of the quarterly rolling-one-year rounds. US real GDP and GDP prices use the 1992 onward real-GDP sample.

Series	Years	Through 2020	2010-19	2020-24	Peak	2026
US real GDP	1992-2026	1.287	1.281	1.804	2.185 (2021)	1.410
US GDP prices	1992-2026	0.992	0.832	0.772	1.300 (2008)	0.667
US unemployment	2010-2026	0.762	0.763	1.014	1.227 (2021)	0.615
US core CPI	2007-2026	0.701	0.695	0.808	0.981 (2022)	0.651
US core PCE	2007-2026	0.673	0.667	0.748	1.002 (2022)	0.628
ECB HICP	1999-2026	0.624	0.690	0.973	1.223 (2023)	0.922
ECB core HICP	2017-2026	0.603	0.573	0.846	1.077 (2023)	0.721
ECB real GDP	1999-2026	0.768	0.742	1.595	2.626 (2020)	0.800
ECB unemployment	1999-2026	0.656	0.713	0.946	1.282 (2020)	0.710

Disagreement is the minority share everywhere, but not uniformly. The median between share is about a sixth for growth (US 15%, euro area 15%) and euro-area prices (13%), and a quarter to a third for unemployment and the US inflation densities (core CPI 30%, core PCE 29%; Table 4). Disagreement here is the dispersion of respondents' density-implied means; no reported point forecast enters it. That dispersion understates stated uncertainty by a median factor of 1.8 to 2.8 across the whole panorama (Rich and Tracy 2010; Giordani and Söderlind 2003).

Table 4: Distribution of the between-forecaster share of total variance across round-targets, in percent. The last column is the median ratio of pooled total SD to disagreement SD: the factor by which the dispersion of density-implied means understates stated uncertainty.

Series	Round-targets	1st decile	Median	9th decile	At or below 20%	Total/disagreement
US real GDP	418	8.2	15.3	27.0	74.2	2.56
US GDP prices	278	14.8	23.9	43.1	30.6	2.04
US unemployment	280	12.9	22.9	39.8	39.3	2.09
US core CPI	158	16.4	29.7	53.3	25.9	1.83
US core PCE	158	17.6	29.2	48.1	17.7	1.85
ECB HICP	633	6.7	13.1	28.6	74.4	2.77
ECB core HICP	234	6.4	14.0	36.6	70.9	2.67
ECB real GDP	633	7.7	14.7	31.3	69.8	2.61
ECB unemployment	633	10.4	24.7	47.9	35.2	2.01

Stated uncertainty rises with horizon only until year two — then plateaus, and for growth it inverts. US growth within-SD climbs from the current year to two years out (1.41), then gains 0.02 at three (1.43). In the euro area, total growth uncertainty is mildly higher at the one-year rolling horizon (0.87) than at the four-to-five-year horizon (0.83). Only unemployment keeps steepening (Table 5, Table 6). Past year three the US record stops. The ten-year questions collect point forecasts; their interquartile range (0.2 points in 2026) measures agreement among forecasters, and no density stands behind it.

Table 5: US balanced Q1 term structure, `within SD/total SD` by horizon in years, over the Q1 window in brackets. Each window keeps only the rounds carrying every horizon it offers, so no cell is an average over a different set of rounds than its neighbour.

Variable	Rounds	h=0	h=1	h=2	h=3
US real GDP (2010–26)	17	1.195/1.324	1.343/1.453	1.412/1.518	1.428/1.522
US unemployment (2010–26)	17	0.617/0.689	0.713/0.827	0.767/0.920	0.800/0.976
US GDP prices (1992–26)	35	0.722/0.855	0.817/0.952	–	–
US core CPI (2007–26)	20	0.520/0.668	0.595/0.727	–	–
US core PCE (2007–26)	20	0.512/0.639	0.572/0.691	–	–

Table 6: ECB round-balanced term structure, `within SD/total SD` by horizon class. Rounds enter only if they carry all three modern horizons, and duplicate early longer-term targets are averaged within the round first.

Variable	Rounds	Rolling 1y	Rolling 2y	Longer term
ECB HICP	105	0.622/0.708	0.682/0.745	0.713/0.759
ECB core HICP	39	0.638/0.741	0.698/0.780	0.749/0.815
ECB real GDP	105	0.733/0.872	0.772/0.863	0.775/0.826
ECB unemployment	105	0.608/0.702	0.687/0.825	0.812/1.026

Both surveys' bands miss in tails, in clusters, and the euro area's cover less. A ±1σ band around the consensus covered 22 of 33 next-year US growth outcomes (66.7%) over 1993–2025, and the pooled 90% interval covered 26 of 33 (78.8%). A histogram mixture has no universal nominal ±1σ coverage, so the first figure describes the record and only the second has a nominal rate to fall short of; with 33 dependent outcomes, neither supports a formal verdict on calibration. The misses cluster: six straight target years (1996–2001) around the late-1990s productivity boom, then 2008–09, then 2020–21. Euro-area one-sigma coverage runs 49–61% by variable (Table 7) (Kenny et al. 2014; 2015). Scored properly, the densities beat climatology everywhere (US growth +39% CRPS skill; Table 8) and roughly tie the consensus-plus-track-record Gaussian: skill runs from −0.9% to +8.7%, with ECB HICP −0.0%, US GDP prices −0.9% on the losing side and US growth at +6.3%. Most of what the elicited densities know, the point forecast and its error history already carry, and for several variables the sign of the comparison turns on how forecast rows are weighted (Clements 2018; Knüppel and Krüger 2022). The pinball decomposition puts the 2008–09 growth losses in the lower tail and the 2020–22 inflation losses in the upper tail by factors of 4.6 to 8.8 (Table 9).

Table 7: Interval coverage on every calibrated round-target, and on the comparable one-year rows. Forecasts of a shared outcome are correlated, so these are descriptive counts rather than tests.

Series	Sample	n	Inside +/-1 SD	Inside pooled 90%	Outcome years
US real GDP	all horizons	382	288 (75.4%)	328 (85.9%)	1992-2025
US real GDP	Q1 next year	33	22 (66.7%)	26 (78.8%)	1993-2025
US GDP prices	all horizons	268	221 (82.5%)	246 (91.8%)	1992-2025
US GDP prices	Q1 next year	33	23 (69.7%)	30 (90.9%)	1993-2025
US unemployment	all horizons	228	141 (61.8%)	198 (86.8%)	2009-2024
US unemployment	Q1 next year	14	8 (57.1%)	13 (92.9%)	2011-2024
US core CPI	all horizons	140	114 (81.4%)	124 (88.6%)	2007-2024
US core CPI	Q1 next year	17	14 (82.4%)	15 (88.2%)	2008-2024
US core PCE	all horizons	148	113 (76.4%)	129 (87.2%)	2007-2025
US core PCE	Q1 next year	18	12 (66.7%)	16 (88.9%)	2008-2025
ECB HICP	all horizons	575	342 (59.5%)	418 (72.7%)	1999-2025
ECB HICP	Q1 next year	26	14 (53.8%)	18 (69.2%)	2000-2025
ECB HICP	rolling 1y	105	56 (53.3%)	77 (73.3%)	1999-2025
ECB core HICP	all horizons	176	107 (60.8%)	129 (73.3%)	2017-2025
ECB core HICP	Q1 next year	8	5 (62.5%)	6 (75.0%)	2018-2025
ECB core HICP	rolling 1y	33	22 (66.7%)	24 (72.7%)	2017-2025
ECB real GDP	all horizons	579	284 (49.1%)	382 (66.0%)	1999-2026
ECB real GDP	Q1 next year	26	11 (42.3%)	14 (53.8%)	2000-2025
ECB real GDP	rolling 1y	107	55 (51.4%)	73 (68.2%)	1999-2026
ECB unemployment	all horizons	574	327 (57.0%)	474 (82.6%)	2000-2026
ECB unemployment	Q1 next year	26	13 (50.0%)	21 (80.8%)	2000-2025
ECB unemployment	rolling 1y	106	70 (66.0%)	93 (87.7%)	2000-2026

Table 8: CRPS skill against two no-lookahead benchmarks, as one minus the ratio of mean losses on matched eligible rows, with row-level mean skill as the sensitivity in brackets. The last column is the pooled distribution's gain over the average individual forecaster.

Series	n clim.	Climatology skill	n Gauss.	Gaussian skill	Pooling gain
US real GDP	382	38.5% [41.2%]	312	6.3% [5.5%]	10.9%
US GDP prices	268	46.2% [53.6%]	240	−0.9% [−6.4%]	14.5%
US unemployment	228	45.2% [22.5%]	160	7.0% [−2.6%]	12.7%
US core CPI	140	49.7% [57.6%]	112	3.5% [−4.1%]	15.0%
US core PCE	148	48.9% [55.9%]	120	3.0% [−2.6%]	14.7%
ECB HICP	471	17.1% [16.6%]	474	−0.0% [0.6%]	5.8%
ECB core HICP	176	33.7% [20.6%]	80	8.7% [8.4%]	7.4%
ECB real GDP	489	19.9% [16.6%]	480	1.2% [17.2%]	5.7%
ECB unemployment	415	57.5% [27.3%]	473	0.6% [3.8%]	12.5%

Table 9: Mean pinball loss at the 05 and 95 quantile levels inside each clustered-event window, pooling every calibrated horizon. A window is assigned by the realization's event year, not the forecast round.

Outcome years	Series	n	Pinball 05	Pinball 95	Larger tail
1996–2001	ECB HICP	34	0.057	0.041	lower (05)
1996–2001	ECB real GDP	36	0.098	0.208	upper (95)
1996–2001	ECB unemployment	29	0.127	0.069	lower (05)
1996–2001	US real GDP	48	0.151	0.231	upper (95)
1996–2001	US GDP prices	48	0.046	0.103	upper (95)
2008–09	ECB HICP	44	0.376	0.265	lower (05)
2008–09	ECB real GDP	44	2.293	0.217	lower (05)
2008–09	ECB unemployment	44	0.103	0.459	upper (95)
2008–09	US core CPI	16	0.042	0.060	upper (95)
2008–09	US core PCE	16	0.044	0.077	upper (95)
2008–09	US real GDP	16	0.581	0.199	lower (05)
2008–09	US GDP prices	16	0.076	0.134	upper (95)
2008–09	US unemployment	3	0.035	0.038	upper (95)
2020–22	ECB HICP	72	0.245	1.837	upper (95)
2020–22	ECB core HICP	66	0.102	0.472	upper (95)
2020–22	ECB real GDP	72	1.839	0.901	lower (05)
2020–22	ECB unemployment	72	0.052	0.115	upper (95)
2020–22	US core CPI	24	0.145	1.141	upper (95)
2020–22	US core PCE	24	0.133	0.921	upper (95)
2020–22	US real GDP	48	0.587	0.574	lower (05)
2020–22	US GDP prices	24	0.152	1.335	upper (95)
2020–22	US unemployment	48	0.114	0.693	upper (95)

Table 10: Probability integral transform of the pooled distribution at the realization. Under ideal calibration each decile would hold about 10% of rows, so mass in the two extreme deciles is the underdispersion signal.

Series	n	Mean PIT	Lowest decile	Highest decile	Both extremes
US real GDP	382	0.533	7.6	12.8	20.4
US GDP prices	268	0.447	7.1	6.0	13.1
US unemployment	228	0.338	23.7	6.6	30.3
US core CPI	140	0.531	1.4	12.1	13.6
US core PCE	148	0.467	6.1	10.8	16.9
ECB HICP	575	0.547	16.3	19.1	35.5
ECB core HICP	176	0.568	6.2	26.1	32.4
ECB real GDP	579	0.494	25.0	20.4	45.4
ECB unemployment	574	0.494	12.9	18.8	31.7

The euro area asks the question the US doesn't — and after 2021 the long answer moved less than the short one, but it moved. The ECB longer-term density's total SD rose about 29% after 2021, against 54% at one year; its consensus moved from 1.9% to 2.0% while the one-year consensus moved from 1.6% to 2.3% (Górnicka and Meyler 2022; Vincent-Humphreys et al. 2019). Stated long-run uncertainty peaked at 1.00 in 2023 and stands at 0.92 in 2026, still well above its pre-2021 level, at the horizon where the US elicits nothing (Table 11).

Table 11: ECB one-year and longer-term uncertainty and consensus, through 2020 against 2021-26. Both horizons are read from the same round-balanced panel, so neither change is averaged over a different set of rounds.

Variable	Horizon	Total SD	Change	Consensus	Series high	2026
ECB HICP	Longer term	0.714 → 0.923	+29.4%	1.87 → 2.05	0.997 (2023)	0.916
ECB HICP	Rolling 1y	0.633 → 0.973	+53.7%	1.63 → 2.34	1.223 (2023)	0.922
ECB core HICP	Longer term	0.746 → 0.863	+15.7%	1.65 → 2.00	0.943 (2023)	0.810
ECB core HICP	Rolling 1y	0.603 → 0.837	+38.7%	1.28 → 2.28	1.077 (2023)	0.721
ECB real GDP	Longer term	0.792 → 0.948	+19.7%	1.92 → 1.45	1.071 (2020)	0.930
ECB real GDP	Rolling 1y	0.784 → 1.189	+51.7%	1.45 → 1.62	2.626 (2020)	0.800
ECB unemployment	Longer term	1.017 → 1.055	+3.7%	8.23 → 6.71	1.244 (2012)	0.956
ECB unemployment	Rolling 1y	0.667 → 0.825	+23.7%	9.32 → 6.95	1.282 (2020)	0.710

5 The elicitation gap

i Draft note

Section drafted at beat level. The argument: put the published AI-growth scenarios (Shenk 2026; Cunningham 2025) — which differ in outcome, horizon and conditioning, so no single range spans them — and their anchors (Acemoglu 2024; Briggs and Kodnani 2023; Aghion and Bunel 2024; Davidson 2021; Erdil et al. 2025; Korinek and Suh 2024) against the inventory of section 2; Karger et al. (2026) as the one-off flagship; Chow et al. (2026) extracting long-horizon growth beliefs from asset prices precisely because no survey elicits them. The observation the paper owns: the instruments that would discipline this debate exist, run quarterly, have 50 years of calibration record — and stop at year three (US) or five (euro area).

6 What better elicitation would look like

> **i** Draft note
>
> Design beats, not advocacy: extend one existing survey block to a ten-year growth density (the ECB's longer-term question shows a panel will answer it, and its 2023 special survey asked whether ten-year expectations would differ (Allayioti et al. 2024)); bins wide enough to be honest about tails (the 2020 widening as precedent (European Central Bank 2020)); ship the question with the scoring rule that disciplines it — CRPS with pinball components — because the record shows stated densities break at regime shifts and the AI question is a regime-shift question; pair each long-horizon density with the respondent's AI-adoption assumption so the disagreement between scenarios becomes decomposable.

7 Conclusion

> **i** Draft note
>
> One page: the 2026 pair again — for next year, a median and 1.4 points of stated spread; for the decade, a median and no distribution — with the whole panorama behind it.

Data and limitations

All measures compute from the public microdata of the Philadelphia Fed SPF (Federal Reserve Bank of Philadelphia 2026) and the ECB SPF, with bin schemes taken from each survey's documentation for every era. Realizations are latest-vintage (Stark 2010); real-time-vintage evaluation is future work. Forecasts of a shared outcome are correlated, so miss counts and score comparisons are descriptive. Benchmark windows are expanding with a ten-observation minimum; euro-area annual histories are short. The reading of the bin labels is a choice, and `outputs/reconstruction_sensitivity.csv` records what it moves. Against the printed labels taken literally, our reading raises the median round-target mean by 0.05 points (5th to 95th percentile 0.04 to 0.05) and the median pooled SD by 1.0%, and it leaves the comparisons between periods, horizons and surveys as reported. The panel is not uniform. Where a round's mass sits in an open tail, that tail's closure width moves with the reading: 0.7% of means fall instead of rising, and the largest SD difference is 12%, in the euro-area current-year growth distribution of 2009Q1. The one headline the choice moves is the US ±1σ coverage count: 22 of 33 under our reading and 23 of 33 under either alternative, because the 2011 outcome sits on the band's edge. The pipeline, tests, and every series behind the live interactive are open source at https://github.com/MaxGhenis/expectations.

References

Bibliography

Acemoglu, Daron. 2024. *The Simple Macroeconomics of AI.* Working Paper No. 32487.

Aghion, Philippe, and Simon Bunel. 2024. "AI and Growth: Where Do We Stand?.". https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf.

Allayioti, Anastasia, Rodolfo Arioli, Colm Bates, et al. 2024. *A Look Back at 25 Years of the ECB SPF.* Occasional Paper Series No. 364.

Boero, Gianna, Jeremy Smith, and Kenneth F. Wallis. 2008. "Uncertainty and Disagreement in Economic Prediction: The Bank of England Survey of External Forecasters." *Economic Journal* 118 (530): 1107–27.

Boero, Gianna, Jeremy Smith, and Kenneth F. Wallis. 2011. "Scoring Rules and Survey Density Forecasts." *International Journal of Forecasting* 27 (2): 379–93.

Briggs, Joseph, and Devesh Kodnani. 2023. "The Potentially Large Effects of Artificial Intelligence on Economic Growth.".

Chow, Trevor, Basil Halperin, and J. Zachary Mazlish. 2026. "Transformative AI, Existential Risk, And Real Interest Rates.". https://basilhalperin.com/papers/agi_emh.pdf.

Christensen, Peter, Kenneth Gillingham, and William Nordhaus. 2018. "Uncertainty in Forecasts of Long-Run Economic Growth." *Proceedings of the National Academy of Sciences* 115 (21): 5409–14.

Clements, Michael P. 2018. "Are Macroeconomic Density Forecasts Informative?." *International Journal of Forecasting* 34 (2): 181–98.

Clements, Michael P., Robert W. Rich, and Joseph Tracy. 2025. "An Investigation into the Uncertainty Revision Process of Professional Forecasters." *Journal of Economic Dynamics and Control* 173.

Clements, Michael P., Robert W. Rich, and Joseph S. Tracy. 2022. *Surveys of Professionals*. Working Paper Nos. 22–13.

Cunningham, Tom. 2025. "Forecasts of AI & Economic Growth.". https://tecunningham.github.io/posts/2025-10-19-forecasts-of-AI-growth.html.

D'Amico, Stefania, and Athanasios Orphanides. 2008. *Uncertainty and Disagreement in Economic Forecasting*. Finance and Economics Discussion Series Nos. 2008–56.

Davidson, Tom. 2021. "Could Advanced AI Drive Explosive Economic Growth?.". https://coefficientgiving.org/research/could-advanced-ai-drive-explosive-economic-growth/.

Erdil, Ege, and Tamay Besiroglu. 2023. "Explosive Growth from AI Automation: A Review of the Arguments.".

Erdil, Ege, Andrei Potlogea, Tamay Besiroglu, et al. 2025. "GATE: An Integrated Assessment Model for AI Automation.".

European Central Bank. 2020. "The ECB Survey of Professional Forecasters — Second Quarter of 2020.".

Federal Reserve Bank of Philadelphia. 2026. "Survey of Professional Forecasters: Documentation.".

Ganics, Gergely, Barbara Rossi, and Tatevik Sekhposyan. 2024. "From Fixed-Event to Fixed-Horizon Density Forecasts: Obtaining Measures of Multihorizon Uncertainty from Survey Density Forecasts." *Journal of Money, Credit and Banking* 56 (7): 1675–704.

García, Juan Angel. 2003. *An Introduction to the ECB's Survey of Professional Forecasters*. Occasional Paper Series No. 8.

Giordani, Paolo, and Paul Söderlind. 2003. "Inflation Forecast Uncertainty." *European Economic Review* 47 (6): 1037–59.

Gneiting, Tilmann, and Adrian E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, And Estimation." *Journal of the American Statistical Association* 102 (477): 359–78. https://doi.org/10.1198/016214506000001437.

Górnicka, Lucyna, and Aidan Meyler. 2022. "Does the Tail Wag the Dog? A Closer Look at Recent Movements in the Distributions of Professional Forecasters' Inflation Expectations." *ECB Economic Bulletin,*.

Karger, Ezra, Otto Kuusela, Jason Abaluck, et al. 2026. *Forecasting the Economic Effects of AI.* Working Paper No. 35046. https://forecastingresearch.org/research/economic-effects-of-ai.

Kenny, Geoff, Thomas Kostka, and Federico Masera. 2014. "How Informative Are the Subjective Density Forecasts of Macroeconomists?." *Journal of Forecasting* 33 (3): 163–85.

Kenny, Geoff, Thomas Kostka, and Federico Masera. 2015. "Density Characteristics and Density Forecast Performance: A Panel Analysis." *Empirical Economics* 48 (3): 1203–31.

Knüppel, Malte, and Fabian Krüger. 2022. "Forecast Uncertainty, Disagreement, And the Linear Pool." *Journal of Applied Econometrics* 37 (1): 23–41.

Knüppel, Malte, and Lora Pavlova. 2026. *Survey Design and Professional Forecasters: The Case of Uncertainty in the US SPF.* Discussion Paper Nos. 26–17.

Korinek, Anton, and Donghyun Suh. 2024. *Scenarios for the Transition to AGI.* Working Paper No. 32255.

Lahiri, Kajal, and Xuguang Simon Sheng. 2010. "Measuring Forecast Uncertainty by Disagreement: The Missing Link." *Journal of Applied Econometrics* 25 (4): 514–38.

Rich, Robert W., and Joseph S. Tracy. 2010. "The Relationships among Expected Inflation, Disagreement, And Uncertainty: Evidence from Matched Point and Density Forecasts." *Review of Economics and Statistics* 92 (1): 200–207.

Shenk, Anton. 2026. "The Quadrillion-Dollar Disagreement on AI and the Economy.". https://ai-frontiers.org/articles/the-quadrillion-dollar-disagreement-on-ai-and-the-economy.

Shoja, Mehdi, and Ehsan S. Soofi. 2017. "Uncertainty, Information, And Disagreement of Economic Forecasters." *Econometric Reviews* 36 (6–9): 796–817.

Stark, Tom. 2010. *Realistic Evaluation of Real-Time Forecasts in the Survey of Professional Forecasters.* Research Rap Special Report.

Vincent-Humphreys, Rupert de, Ivelina Dimitrova, Elisabeth Falck, Lukas Henkel, and Aidan Meyler. 2019. "Twenty Years of the ECB Survey of Professional Forecasters." *ECB Economic Bulletin,*.

Wallis, Kenneth F. 2005. "Combining Density and Interval Forecasts: A Modest Proposal." *Oxford Bulletin of Economics and Statistics* 67 (s1): 983–94.

Zarnowitz, Victor, and Louis A. Lambros. 1987. "Consensus and Uncertainty in Economic Prediction." *Journal of Political Economy* 95 (3): 591–621.