Fiscal Policy and Household Over-Optimism

Alessia Papini*

July 2026

Abstract

This paper studies whether fiscal policy transmits to the economy through a sentiment channel. The shift in European fiscal policy after the pandemic changed how households across the euro area perceived future credit access. I study Italy, where the Draghi credibility event of February 2021 provides a sharp, datable domestic amplification of the common European shock. Using the ECB Consumer Expectations Survey, I construct individual-level credit-access forecast errors and estimate their response with the Forecasted Average Treatment estimator, designed for settings of universal treatment. I document systematic and persistent over-optimism: forecast errors increase by up to 7.8 percent of the survey scale, peaking several months after the announcement before gradually declining. Finally, I show that over-optimistic beliefs are associated with active credit-seeking behavior: a one-unit increase in the forecast error is associated with a 0.78 percentage-point higher probability of applying for credit in the subsequent quarter.

1 Introduction

Advanced economies are entering a prolonged period of fiscal constraint. Public debt has reached historic highs and is projected to rise further, narrowing the space available for discretionary government expenditure (IMF, 2026). In this environment, understanding whether fiscal policy can influence household expectations and behavior through channels beyond direct transfers becomes increasingly important. One such channel, the sentiment channel of fiscal policy, operates through the beliefs that policy action generates in households. Existing evidence provides limited causal evidence on whether fiscal policy systematically distorts

*Department of Economics, European University Institute. Email: alessia.papini@eui.eu. I am grateful to Giancarlo Corsetti, Alexander Ludwig, and Fabrizia Mealli for helpful comments and suggestions.

household beliefs, how persistent those distortions are, and whether they translate into real economic behavior. This paper addresses these questions directly.

I study how credit-access forecast errors respond to a fiscal policy shift. I measure credit access expectations using the ECB Consumer Expectations Survey, which asks households each month how they expect credit conditions to change over the next twelve months. The forecast error is defined as the difference between this forward-looking expectation and the realized credit access outcome twelve months later, a positive value indicating that the household was over-optimistic ex ante. The shift in European fiscal policy after the pandemic, coordinated through NextGenerationEU and the broader move from austerity to expansion, changed how households across the euro area perceived future credit access. In Italy, Draghi's appointment and Senate address in February 2021 turned this common shift into a sharp, datable domestic credibility event. Focusing on Italy as my case study, I estimate the effect of this credibility event on household credit-access forecast errors. The empirical setting, however, poses a fundamental identification challenge: because the announcement was observed simultaneously by all Italian households, there is no untreated comparison group. I address this universal-treatment problem using the Forecasted Average Treatment Effect estimator of Botosaru et al. (2026), which constructs individual-level counterfactuals from pre-treatment outcome histories without requiring untreated units, and corroborate the results using the Machine Learning Control Method of Cerqua et al. (2024). To the best of my knowledge, this is the first application of these estimators to the study of fiscal sentiment.

My main finding is that Italian households became systematically and persistently over-optimistic about future credit access after the February 2021 fiscal shift. The estimated effect is statistically insignificant in the announcement wave, consistent with partial treatment exposure, and becomes positive and significant from the following wave onward. The response builds gradually, peaks eight months after the announcement at $\widehat{FAT}_8 = 0.313$ (approximately 7.8 percent of the survey scale), and declines thereafter. An expectation-realization decomposition shows that while the realization-side contribution is shared across euro area countries, Italy's distinctiveness lies in the expectation component, consistent with a credibility channel tied to the Draghi appointment rather than a generic European shock. This hump-shaped dynamic is difficult to reconcile with rational expectations benchmarks and is instead more consistent with diagnostic expectations models in which agents overweight salient positive news (Bordalo et al., 2019). The MLCM corroboration delivers estimates of similar magnitude, indicating that the result is not driven by the particular counterfactual construction of the FAT estimator. Finally, the belief distortions are associated with observable behavior. A one-unit increase in the credit-access forecast error is associated with a 0.78 percentage-point higher probability of applying for credit in the subsequent quarter.

These results speak directly to a puzzle in the literature. Coibion et al. (2020) document that households tend to be poorly informed about fiscal policy and that even large announcements fail to move aggregate expectations. Their preferred explanation is that *"households do not believe that the policy responses [...] are effective: i.e., the multipliers they associate with the described policy responses are close to zero"* (Coibion et al., 2020, p. 2). This is not what I find. Italian households reacted strongly to the February 2021 fiscal regime shift, with belief distortions that are large, persistent. More importantly, belief distortions are associated with action: households with larger forecast errors became more likely to apply for credit afterward. A key methodological choice is to measure this response through forecast errors rather than expectation levels. Forecast errors are more informative because, under rational expectations, they should be orthogonal to predictable policy events (Born et al., 2023) — a systematic response therefore constitutes direct evidence of persistent belief distortion. This approach follows the most recent literature (Bellifemine et al., 2025) and allows me to distinguish genuine over-reaction from rational updating, while providing causal identification in a natural experiment rather than an RCT setting.

Related Literature. This paper contributes to three strands of the literature.

First, it speaks to the literature on departures from rational expectations and belief distortions (Coibion and Gorodnichenko, 2012; Broer and Kohlhas, 2024; Angeletos and Lian, 2022; Born et al., 2023; Bordalo et al., 2019). Existing evidence tends to find underreaction to aggregate shocks but overreaction to idiosyncratic ones. My results contribute to this debate by documenting systematic over-optimism — closer in nature to overreaction — in response to a salient aggregate fiscal announcement, and by using forecast errors rather than expectation levels as the outcome variable (Barrero, 2022; Ma et al., 2020).

Second, this paper connects to the literature on information, expectation formation, and the real effects of policy communication (Coibion and Gorodnichenko, 2015b; Coibion et al., 2020, 2021a,b, 2024; Weber, 2022). A central finding of this literature is that information provision shifts household beliefs and that these revisions translate into measurable changes in spending and borrowing behavior (Coibion et al., 2024). My contribution is to provide causal identification in a natural experiment setting rather than through an RCT, exploiting a universally observed policy event rather than experimentally assigned information.

Third, and most directly, this paper contributes to an emerging literature on unconventional fiscal policy and the sentiment channel of fiscal transmission (D'Acunto et al., 2016, 2018; Bachmann et al., 2021; Corrado et al., 2026; Melosi et al., 2025; Gallegos et al., 2025; Born et al., 2025; Bellifemine et al., 2025). D'Acunto et al. (2016) show that pre-announced VAT increases shift household inflation expectations and consumption plans,

operating through a channel analogous to forward guidance. Gallegos et al. (2025) provide cross-country evidence from the euro area that fiscal announcements shift household expectations, though they focus on the direction of belief revision rather than its accuracy. Melosi et al. (2025) study the signaling channel, showing that fiscal announcements convey information about the state of the economy beyond the direct transfer effect. Consistent with this literature, my results suggest that fiscal announcements affect household behavior not only through realized transfers, but also through systematic shifts in beliefs about future financial conditions and credit access. My paper is particularly close to Bellifemine et al. (2025), who show that positive government-demand shocks generate systematic over-optimism in Italian firms' sales expectations. I complement their analysis by providing analogous evidence for households rather than firms, focusing on credit-access expectations following a nationwide fiscal regime shift in Italy.

The paper is organized as follows. Section 2 describes the institutional background; Section 3 presents the data; Section 4 outlines the identification strategy; Section 5 presents the results and explores the mechanisms. The last section concludes.

2 Institutional Background

The European fiscal regime shift. The COVID-19 pandemic triggered the deepest recession in the euro area since the formation of the monetary union. Real GDP fell by 6.6% in 2020 (Eurostat, 2021), prompting a relaxation of the European fiscal framework. On 20 March 2020, the European Commission activated the general escape clause of the Stability and Growth Pact for the first time, allowing member states to depart temporarily from the Union's fiscal rules (European Commission, 2020; Council of the European Union, 2020).

The change in fiscal governance went beyond suspending existing constraints. In May 2020, France and Germany jointly proposed a €500 billion recovery fund financed through common European debt and distributed as grants, marking a departure from Germany's long-standing opposition to fiscal transfers and debt mutualization (Buti and Papaconstantinou, 2021; Macron and Merkel, 2020). Two months later, the European Council agreed on the €750 billion NextGenerationEU recovery instrument, financed through common borrowing and combining grants and loans to support member states' recovery (European Council, 2020; Dullien, 2020).

Rather than emphasizing fiscal consolidation, European institutions coordinated a large-scale expansionary response based on common borrowing and sizeable fiscal transfers. These decisions represented the first large-scale common borrowing in the history of European integration and were widely characterized as path-breaking (Schelkle, 2021; De la Porte and

Jensen, 2021). Buti and Papaconstantinou (2021) describe the policy response as a "*clear break with precedent*" (Buti and Papaconstantinou, 2021, p.4) in both the instruments deployed and the magnitude of the fiscal effort. Not all assessments share this interpretation. Abels (2024) argues that the pandemic measures are better understood as a "passive revolution" in which the EMU adapted elements of a fiscal-integrative agenda while preserving the underlying architecture of fiscally restrictive governance. For the purposes of this paper, what matters is the fiscal signal received by households in early 2021, a coordinated shift from consolidation toward expansion, rather than the permanence of the underlying institutional change.

A fiscal expansion of this magnitude could affect household credit-access expectations through several channels. In economies where banks hold substantial domestic sovereign debt, an improvement in the fiscal outlook compresses sovereign spreads, strengthens bank balance sheets, and eases funding conditions, which households may perceive as improving future access to credit (Bocola, 2016; Altavilla et al., 2017; Farhi and Tirole, 2018). Figure 1 plots mean credit-access forecast errors, defined formally in Section 3, for the six founding CES countries.

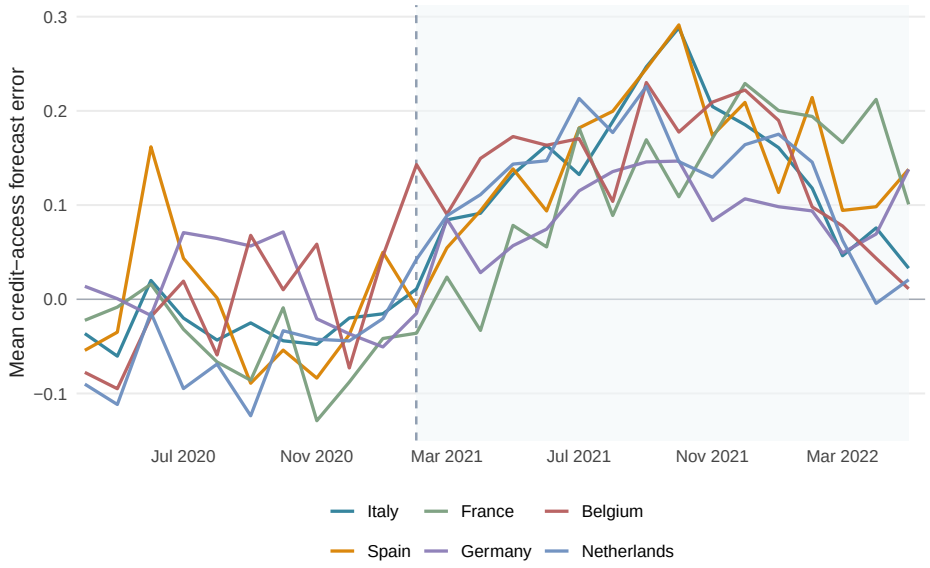


FIGURE 1: Mean credit-access forecast error by country. All six founding CES countries shift upward after early 2021, but Italy's response is larger and more persistent. The forecast error is defined in Section 3. The vertical line marks wave 14 (February 2021).

Credit-access forecast errors shift upward in all six countries after early 2021, consistent with a common European belief response to the fiscal regime shift. The movement is not unique to any single country: the post-pandemic fiscal expansion appears to have shifted household credit-access beliefs across the euro area.

The Italian credibility channel. The European fiscal shift described above was common in scope, but Italy is the natural setting in which to study how such a shift can be amplified by domestic credibility. Italy was one of the largest expected beneficiaries of NextGenerationEU. Its Recovery and Resilience Facility allocation amounted to approximately €191.5 billion, combining €68.9 billion in grants and €122.6 billion in loans, the largest national envelope in the programme (European Commission, 2021; Italian Ministry of Economy and Finance, 2021). The fiscal turn was therefore especially salient in Italy, both economically and politically.

Italy also entered 2021 with a more fragile fiscal-financial environment than most euro-area countries. Real GDP had fallen by 8.9% in 2020 (World Bank, 2020), the general government deficit reached 9.6% of GDP (Banca d'Italia, 2022), public debt stood at 150.8% of GDP, and the vaccination campaign was still in its early stages, with only approximately 3% of the population having received at least one dose by mid-February 2021 (OWID, 2024). The sovereign-bank nexus made domestic credit conditions particularly sensitive to changes in sovereign risk: in this environment, an improvement in fiscal credibility could plausibly affect household credit expectations through a bank-lending channel (Bocola, 2016; Altavilla et al., 2017; Farhi and Tirole, 2018).

It was in this context that Mario Draghi received the mandate to form a government in February 2021 and delivered his inaugural address to the Senate on 17 February. The speech marked a sharp break from the fiscal narrative that had prevailed since the European sovereign debt crisis of 2011–2013, during which Italy had pursued severe consolidation under the Monti government (Culpepper, 2014). Draghi explicitly reframed the relationship between public debt and fiscal responsibility: debt was not inherently problematic, he argued, but incurring it without spending and investing well was a wrong done to future generations (*"ogni spreco oggi è un torto che facciamo alle prossime generazioni"*) (Draghi, 2021). The speech signalled a shift away from austerity toward expansionary fiscal policy, backed by the resources of the NextGenerationEU programme. In this sense, Draghi's address aligned with the broader European narrative of a "*clear break with precedent*" (Buti and Papaconstantinou, 2021, p. 4), but gave it a sharper and more personally credible expression in Italy.[1] Google search intensity for terms related to Draghi's Senate address spiked around the speech date and was geographically widespread across Italian regions (Appendix A). The relevant treatment is not a single fiscal decree, but the broader fiscal credibility shift that the Draghi government represented for Italian households.

Financial markets reacted in a way consistent with this credibility interpretation. Figure 2

[1]The NextGenerationEU funds were not disbursed immediately; the signal effect of the speech operated through the credible commitment to spend, not through actual transfers at the time of the announcement.

plots daily sovereign spreads relative to Germany around the Draghi window. Between Draghi's mandate and the Senate speech, the BTP–Bund spread compressed by about 10.5 basis points, while the Spain–Bund spread rose and the France–Bund spread was essentially flat. This suggests that the announcement contained an Italy-specific credibility component, rather than only a common euro-area repricing.

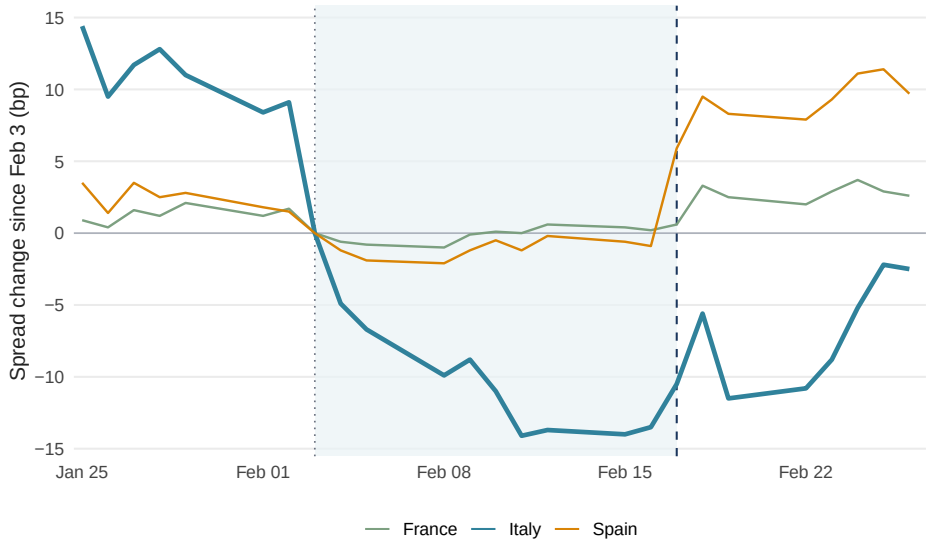


FIGURE 2: Daily sovereign-spread changes around the Draghi mandate-to-Senate-speech window. Spreads are measured as 10-year country yields minus the German 10-year yield. Each series is normalized to zero on 3 February 2021, the date of Draghi's mandate, so the y-axis reports changes in basis points relative to that date. The shaded area covers the period between Draghi's mandate and his Senate speech on 17 February 2021.

This market reaction matters because the outcome studied in this paper is household credit-access expectations. If households perceived the Draghi government as improving Italy's fiscal credibility and reducing sovereign risk, they may also have expected banks to face easier funding conditions and to supply credit more readily. This mechanism is specific to credit access: it links fiscal credibility to household beliefs through the sovereign-bank channel, rather than through a generic optimism channel.

Draghi's communication was followed rapidly by concrete implementing measures, as shown in Figure 3. In March 2021, D.L. 41/2021 introduced direct transfers to households and firms affected by the pandemic. D.L. 73/2021 of May 2021 extended the *Fondo Gasparrini* mortgage moratorium and expanded state-backed guarantees for new mortgage lending.[2]

[2]The *Fondo Gasparrini*, originally established in 2007 and expanded during the pandemic under D.L. 18/2020, was extended by the Draghi government through December 2021. The moratorium allowed households to suspend repayments on first-home loans for up to 18 months with the state covering half of accruing interest.

These measures reinforced the fiscal-expansion narrative and were closely related to households' borrowing conditions, the dimension captured by the forecast error constructed in Section 3.

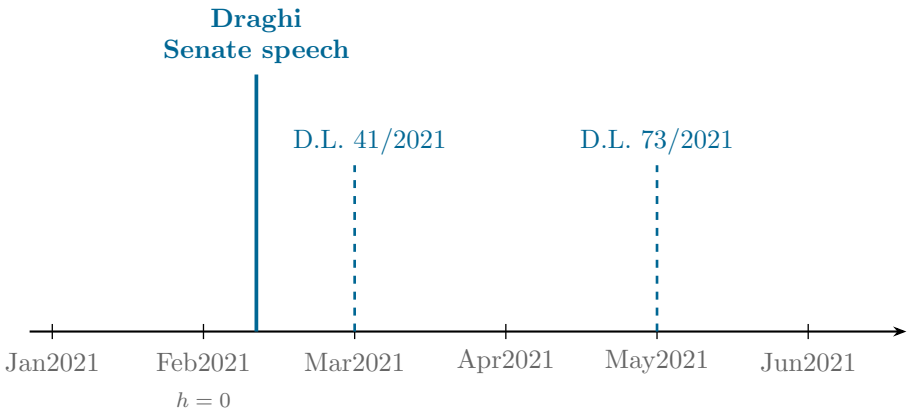


FIGURE 3: Timeline of the February 2021 fiscal regime shift in Italy. The solid line marks Draghi's Senate speech (17 February 2021). Dashed lines indicate the main implementing decrees.

I date the Italian treatment at wave 14 (February 2021). CES fieldwork runs from the 1st to the last day of each calendar month, so wave 14 covers 1–28 February 2021. Because the Draghi appointment and speech fall mid-fieldwork, wave 14 captures partial treatment exposure. I treat wave 14 as the announcement wave ($h = 0$) and wave 15 (March 2021) as the first fully treated wave; main results focus on wave 15 onward, with wave 14 reported as a robustness check.

3 Data

The ECB Consumer Expectations Survey. My primary data source is the ECB Consumer Expectations Survey (CES), a monthly online panel survey measuring household expectations and economic behavior across the euro area (Georgarakos and Kenny, 2022; Bańkowska et al., 2021). The survey was launched in April 2020 and initially covered six countries: Belgium, Germany, Spain, France, Italy, and the Netherlands. In 2022 the sample was extended to five additional countries - Ireland, Greece, Austria, Portugal, and Finland - bringing the total to eleven euro area economies.

The CES is designed as a rotating panel: respondents complete online questionnaires every month and remain in the survey for between 18 and 24 months after joining. This design provides a rich within-respondent time series while continuously refreshing the cross-section. For each country, the target sample size is approximately 2,000 respondents per wave. The sample is designed to be broadly representative of the adult population in each

country, though some limitations remain.[3] Italian respondents constitute the focus of the present analysis; the five other countries — Belgium, Germany, Spain, France, and the Netherlands — serve as the cross-country reference sample used to interpret the Italian response. Figure B1 in the Appendix reports the full wave-by-wave breakdown for Italian respondents.

Upon joining the panel, each respondent first completes a background questionnaire covering characteristics that change infrequently - including family situation, education level, housing tenure, and financial attitudes. Time-sensitive information is then collected at varying frequencies: monthly questionnaires capture expectations about inflation, labor market outcomes, credit access, and household finances, while quarterly questionnaires record spending behavior and credit applications.

Variable Construction. The analysis centers on a specific pair of survey questions that together allow me to construct individual-level forecast errors on credit access. Each month, respondents are asked: *"Looking ahead, do you think that 12 months from now it will generally be harder or easier for your household to obtain credit or loans (including credit cards, retail cards, car loans, student loans, and mortgages) than it is these days?"*, Answers are recorded on a five-point scale from 1 (*much harder*) to 5 (*much easier*). This question elicits respondent i's expectation at time t of the 12-month-ahead change in household credit access, which I denote $E_{i,t}(\Delta\text{CredAcc}_{i,t+12})$.

The rotating panel structure of the CES allows me to re-observe the same respondent twelve months later. At that point, respondents are asked: *"Compared with your household's situation 12 months ago, do you think it is generally harder or easier these days for your household to obtain credit or loans?"* answered on the same scale. This retrospective question elicits a subjective evaluation of the realized change in credit access, $\Delta\text{CredAcc}_{i,t+12}$.[4] Both questions refer to the same underlying quantity — the change in credit access over a 12-month horizon — measured first as a forward-looking expectation and then as a backward-looking assessment by the same individual.

[3]The ECB constructs survey weights calibrated to Eurostat population benchmarks on age, gender, and region through iterative proportional fitting (Georgarakos and Kenny, 2022). The online panel underrepresents respondents aged 71 and above and those with lower educational attainment. Baseline estimates in this paper are unweighted, following standard practice in the CES literature; survey-weighted estimates are reported in the Appendix D and confirm the main results.

[4]The use of subjective rather than administrative measures of credit access is standard in the expectations literature. Survey expectations have been shown to possess strong predictive power for household debt dynamics and borrowing behavior (Angelico, 2019). Matched survey and credit-registry evidence further shows that subjective expectations predict subsequent debt accumulation and defaults (D'Acunto et al., 2024). Importantly, since both the expectation and the realization are measured on the same subjective scale by the same respondent, any systematic reporting bias cancels in the construction of the forecast error.

I construct the credit access forecast error as:

$$\text{FE}_{i,t} = E_{i,t}(\Delta\text{CredAcc}_{i,t+12}) - \Delta\text{CredAcc}_{i,t+12} \tag{1}$$

Credit access matters because households adjust consumption and borrowing when they expect loans to become easier or harder to obtain, yet forecast errors on credit access have received little attention in the fiscal transmission literature relative to inflation or income expectations. Under this convention, $\text{FE}_{i,t} > 0$ indicates that the respondent expected credit conditions to improve by more than they actually did - that is, the respondent was over-optimistic ex ante. Conversely, $\text{FE}_{i,t} < 0$ reflects over-pessimism.[5]

Descriptive statistics. Before turning to the identification strategy, I describe the key features of the data used in estimation. Table 1 reports summary statistics for the full Italian CES sample and the analytical estimation sample. The two groups are broadly comparable across all observable dimensions: age, gender, education, labor market status, and household finances are similar in columns (A) and (B). This suggests that the rotating panel design, combined with the 12-month realization window required to construct $\text{FE}_{i,t}$, does not introduce systematic selection on observables. Reassuringly, pre-treatment means of the credit access expectation and its 12-month realization are nearly identical across the two groups (2.71 versus 2.72, and 2.73 versus 2.73, respectively), and the pre-treatment forecast error mean is -0.02 in both samples.

Figure 4 shows the distribution of $\text{FE}_{i,t}$ across respondent-wave observations in the pre-treatment period (waves 4-13). The distribution is highly concentrated at zero: 62.2% of observations record a forecast error of exactly zero, indicating that the respondent's expectation was precisely realized. This concentration partly reflects the coarse five-point response scale used for both expectations and realizations. The remaining mass is nearly symmetric around zero, with 19.5% of observations recording over-pessimism ($\text{FE}_{i,t} < 0$) and 18.2% recording over-optimism ($\text{FE}_{i,t} > 0$); skewness is 0.004. The mean pre-treatment forecast error of -0.02 is statistically indistinguishable from zero. Before February 2021, Italian households did not systematically over- or under-predict future credit access.

[5]Part of the forecasting literature defines forecast errors as realization minus expectation (Born et al., 2023; Bellifemine et al., 2025), so that positive values indicate under-prediction. All results in this paper use the convention in Equation (1); signs of estimates from that literature should be reversed for comparison.

TABLE 1: Summary statistics. The analytical sample remains broadly representative of the full Italian CES population despite attrition induced by the forecast-error construction.

Group	Variable	(A) Full Italian CES sample	(B) Analytical sample
Sample size	Respondents	16,767	2,154
Demographics	Age category (mean; 1 = 18–34, 4 = 71+)	2.1	2.1
	Female (%)	49.9	56.0
	At least upper-secondary education (%)	90.6	91.9
	Tertiary education (%)	47.2	44.4
Labor market	Employed (%)	67.7	65.3
	Self-employed / unpaid family worker (%)	20.9	21.3
	Retired (%)	11.4	13.3
Household finances	Household income quintile (mean)	3.0	2.9
	Income in top two quintiles (%)	39.9	34.1
	Owner-occupier (%)	78.2	76.6
	Owner-occupier with mortgage (%)	26.9	26.6
Expectations	Credit forecast error, pre-treatment mean	−0.02	−0.02
	Credit access expectation, pre-treatment mean	2.71	2.72
	Credit access realisation $t+12$, pre-treatment mean	2.73	2.73

Notes: Shares are reported in percent. Respondent-level covariates use each respondent's first non-missing value in the processed Italy panel. Pre-treatment expectation and forecast-error means use respondent-wave observations from CES waves 4–13. The analytical sample contains respondents with at least one non-missing pre-treatment forecast error. Treatment is February 2021 (CES wave 14). Column (A) reports statistics for all Italian CES respondents; column (B) restricts to the analytical estimation sample.

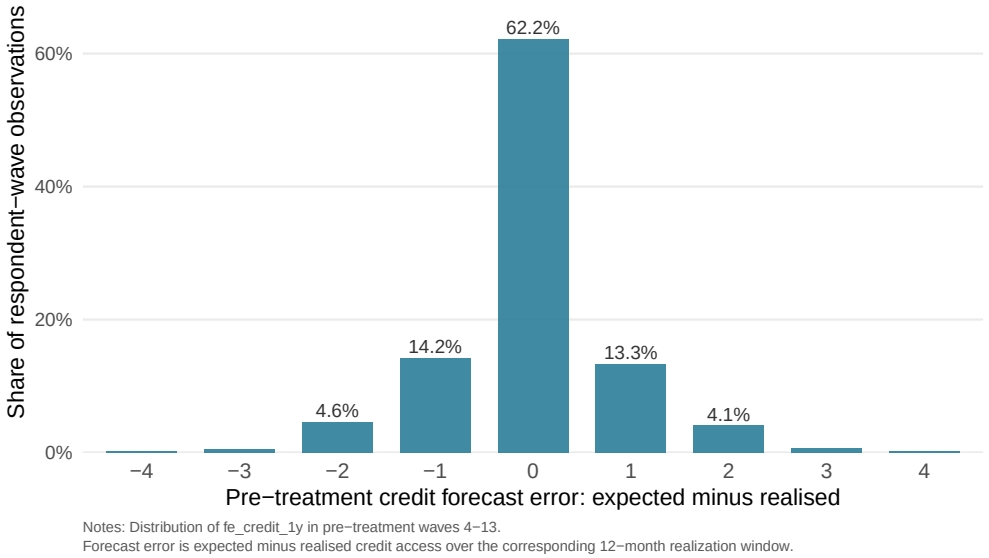


FIGURE 4: Pre-treatment distribution of credit access forecast errors.

4 Identification

The goal of this paper is to identify the causal effect of the February 2021 fiscal regime shift on Italian household credit access forecast errors. Standard panel identification strategies are not applicable in this setting. The treatment, the Draghi credibility event described in Section 2, was simultaneously observed by all Italian households: before the speech, no agent could anticipate the precise content and timing of the regime shift; after the Senate speech of 17 February 2021, all Italian respondents in the CES were exposed to the same fiscal narrative change. This setting poses a universal-treatment identification problem: after February 2021, every Italian household in the CES is treated. The salience and nationwide diffusion of the event are documented in Section 2 and Appendix A. Cross-country evidence from the other five founding CES countries (Section 2) is used not as a clean control group, but to interpret whether the Italian response reflects a common European movement or an Italy-specific amplification.

The object of interest is the causal effect of the fiscal regime shift on household beliefs, as measured by household credit-access forecast errors[6]:

$$ATT_h = \frac{1}{n} \sum_i E\left[y_{i,\tau+h}(1) - y_{i,\tau+h}(0)\right] \tag{2}$$

where τ denotes the treatment date (February 2021, wave 14), $y_{i,\tau+h}(1)$ is the potential

[6]Following the notation from (Botosaru et al., 2026, p.7).

outcome of respondent i at horizon h under treatment, and $y_{i,\tau+h}(0)$ is the counterfactual potential outcome that would have been observed in the absence of the regime shift. The outcome $y_{i,t}$ is the credit access forecast error $\text{FE}_{i,t}$ defined in Equation (1). Since $y_{i,\tau+h}(0)$ is never observed for any i, identification requires predicting what households would have expected without the fiscal shift. The FAT estimator of Botosaru et al. (2026) is useful here because it constructs counterfactual outcomes without untreated households.

The FAT estimator. The Forecasted Average Treatment Effect (FAT) estimator of Botosaru et al. (2026) identifies ATT_h by replacing the unobserved counterfactual outcome $y_{i,\tau+h}(0)$ with a forecast constructed from each unit's own pre-treatment history. The estimator is:

$$\widehat{FAT}_h := \frac{1}{n}\sum_{i=1}^{n}[y_{i,\tau+h} - \widehat{y}_{i,\tau+h}(0)] \tag{3}$$

where $\widehat{y}_{i,\tau+h}(0)$ denotes the predicted untreated outcome for unit i at horizon h.

The counterfactual forecast is obtained by fitting a unit-specific polynomial trend to the pre-treatment observations $\{y_{i,t}\}_{t<\tau}$ and extrapolating it into the post-treatment period. I set the polynomial order to $q = 0$, so the counterfactual for each unit equals its pre-treatment mean. This choice is appropriate because the pre-treatment series exhibits no visible deterministic trend (Figure 5), while the short panel dimension makes more flexible specifications noisy. I therefore prefer the unconditional FAT (UC-FAT) over the parametric variant (PM-FAT), which requires estimating unit-level autoregressive dynamics from a limited number of pre-treatment observations.

Identification of ATT_h through Equation (3) relies on three assumptions, which I discuss below.

Identifying Assumptions. Identification through the FAT estimator relies on three assumptions.

No anticipation. Treated households must not adjust expectations prior to treatment realization. This assumption is likely to hold in the present setting: formally accepted the mandate to form a government only on 13 February 2021, four days before his Senate address, and the content of the speech was not pre-announced. Given the monthly frequency of the CES survey, there is limited scope for anticipatory adjustment within the pre-treatment window.

Data-generating process for untreated outcomes. The FAT estimator requires assumptions on the process governing the counterfactual outcome $y_{it}(0)$ in the absence of treatment. Following Botosaru et al. (2026), the counterfactual is assumed to be the sum of

a mean-stationary component (which nests stationary ARMA processes with heterogeneous parameters) and a stochastic trend component (a random walk), with either or both components possibly absent and no deterministic time trend. Under this DGP, any weighted average of pre-treatment outcomes with weights summing to one produces individually unbiased forecasts of the counterfactual (Botosaru et al., 2026, Theorem 1, p. 10). The $q = 0$ specification, which sets each unit's counterfactual equal to its pre-treatment mean, belongs to this class. If the counterfactual process contained a deterministic trend, a higher polynomial order would be required. The substantive assumption is therefore that credit-access forecast errors, absent the fiscal regime shift, would not have trended deterministically over the sample period. Figure 5 shows that the mean pre-treatment forecast error remains approximately flat across waves 4–13, with no visible drift. Placebo exercises reinforce this: assigning fictitious treatment dates two and three waves before the actual treatment yields estimates close to zero and statistically insignificant ($\widehat{FAT} = -0.012, p = 0.568$ and $\widehat{FAT} = -0.003, p = 0.864$, respectively). Results under $q = 1$ are reported in Appendix D as a specification check and confirm the baseline.

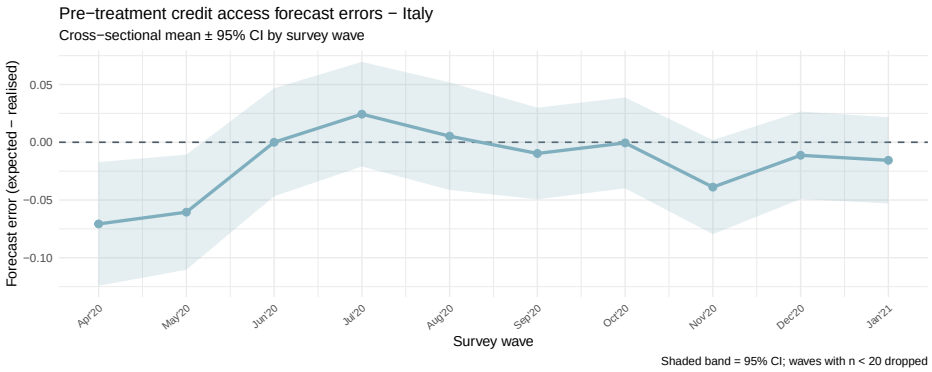


FIGURE 5: Mean credit access forecast error in the pre-treatment period (waves 4–13). The flat trajectory supports the $q = 0$ specification of the FAT counterfactual.

Absence of common unforecastable shocks. Consistency of the FAT estimator requires that no other major aggregate shock coinciding with the treatment period independently shift household credit-access expectations. Because the design relies on each unit's own pre-treatment history rather than on untreated comparison units, any such shock would be absorbed into the estimated treatment effect. This assumption warrants careful discussion. As documented in Section 2, the fiscal regime shift had a European dimension: credit-access forecast errors shifted upward across all six founding CES countries, not only in Italy. The FAT estimate therefore captures Italy's total response to the post-2021 fiscal environment, including any component shared with the rest of Europe. The cross-country decomposition in Section 2 disciplines the interpretation: the realization component is com-

mon across countries, while Italy's distinctiveness lies in the expectation leg, consistent with the credibility channel rather than with a generic European shock. Beyond the European fiscal shift, the most relevant alternative explanations are contemporaneous monetary-policy developments and pandemic-related expectation dynamics. Section 5 presents exercises designed to assess these threats, including controls for identified monetary-policy shocks and placebo tests on outcomes that the fiscal credibility event should not affect.

5 Results

Main Results. Figure 6 shows a sharp and persistent increase in household credit-access forecast errors following the February 2021 fiscal regime shift. Before the treatment wave, the observed series fluctuates around zero with no discernible trend, consistent with the stationarity assumption underlying the $q = 0$ specification. After wave 14, the observed series diverges persistently upward from the counterfactual: Italian households became systematically over-optimistic about future credit access following the fiscal regime shift. The shaded area between the observed series and the counterfactual represents the $\widehat{FAT}_h$ estimate at each horizon. Forecast errors continue rising for eight months before declining.

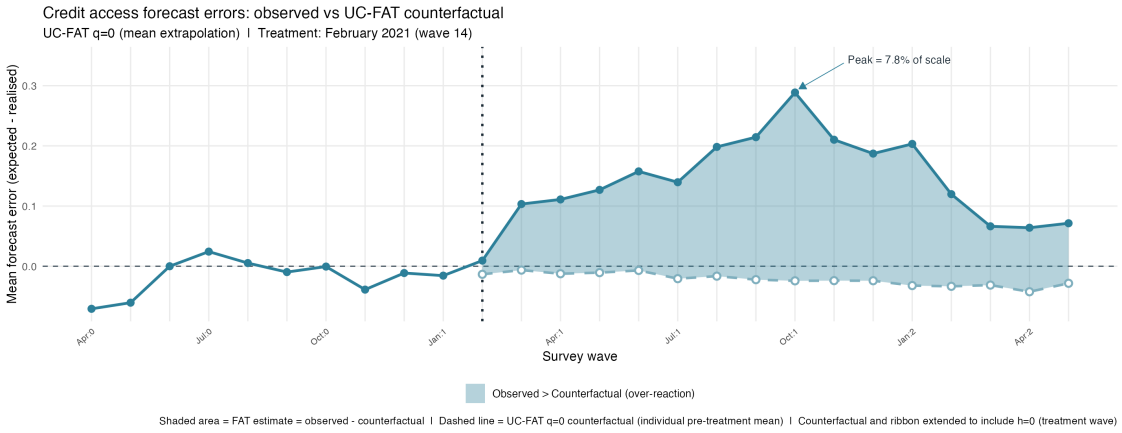


FIGURE 6: Observed mean credit access forecast error and UC-FAT counterfactual ($q = 0$). The shaded area is the $\widehat{FAT}_h$ estimate at each post-treatment horizon. Vertical dotted line marks the treatment wave (wave 14, February 2021).

Table 2 reports the formal estimates. The treatment effect at $h = 0$ (wave 14, February 2021) is positive but not statistically significant ($\widehat{FAT}_0 = 0.023$, $p = 0.277$), consistent with the partial treatment exposure documented in Section 2: not all respondents in wave 14 had been surveyed after the 17 February Senate address. From $h = 1$ onward, the effect is positive and statistically significant at the 1% level throughout. The estimate builds gradually, from $\widehat{FAT}_1 = 0.110$ at $h = 1$ (March 2021) to a peak of $\widehat{FAT}_8 = 0.313$ at $h = 8$ (October 2021),

before declining at longer horizons. Given that the forecast-error measure ranges from -4 to $+4$, the peak effect corresponds to roughly $+7.8\%$ of the directional support of the survey scale. As documented in Section 2, part of this response reflects a broader European shift in how households perceived future credit access. The FAT estimate captures Italy's total response, which the cross-country decomposition shows to be amplified relative to comparator countries through the expectation channel. I focus the main discussion on horizons $h = 1$ to $h = 8$, where identification is cleanest and the sample size is largest. Estimates at $h \geq 9$ are statistically significant but based on a shrinking panel, as the rotating design and the 12-month realization window progressively reduce the number of contributing units.

The $q = 1$ specification, which extrapolates unit-level linear trends rather than pre-treatment means, yields qualitatively similar short-run estimates but substantially wider confidence intervals at longer horizons. Given the short panel dimension, higher-order specifications are noisier and more sensitive to trend misspecification, so I focus on the more stable $q = 0$ estimator throughout.[7]

MLCM corroboration. To assess whether the results depend on the specific counterfactual construction underlying the FAT estimator, I complement the analysis with the Machine Learning Control Method (MLCM) of Cerqua et al. (2024), which predicts respondent-level untreated outcomes using supervised machine-learning forecasts estimated on pre-treatment data.[8]

Relative to FAT, the main advantage of the MLCM in this application is that counterfactual outcomes can be constructed using stable respondent characteristics rather than individual pre-treatment histories alone. This is particularly useful in the rotating-panel structure of the CES, where reliance on individual pre-treatment trajectories mechanically reduces the estimation sample at longer horizons. Because the MLCM predictor set consists primarily of fixed background characteristics, the resulting counterfactual remains conceptually close to the stable pre-treatment extrapolation underlying the FAT estimator. This also allows later entrants into the survey to contribute to post-treatment estimation, avoiding the sample shrinkage that affects the FAT at longer horizons. The preferred specification selects a Random Forest predictor based on pre-treatment out-of-sample cross-validation performance.

The MLCM estimates closely mirror the FAT results. The estimated effect is near zero in the treatment month, becomes positive and statistically significant from $h = 1$ onward, and peaks around 8 months after the fiscal regime shift. At $h = 1$, the MLCM estimate is 0.118,

[7] Figure D1 in the Appendix compares the $q = 0$ and $q = 1$ estimates directly.

[8] The MLCM implementation uses only predetermined respondent characteristics measured before treatment. Details on covariate selection, model training, and cross-validation are reported in Appendix E.

TABLE 2: FAT estimates: effect on credit access forecast errors

		$q = 0$ (mean)		$q = 1$ (linear)	
h	Wave	$\widehat{FAT}_h$	N	$\widehat{FAT}_h$	N
0	Feb'21	0.023	1,813	0.045*	1,778
		(0.021)		(0.025)	
1	Mar'21	0.110***	1,819	0.156***	1,780
		(0.022)		(0.029)	
2	Apr'21	0.123***	1,793	0.172***	1,759
		(0.022)		(0.033)	
3	May'21	0.138***	1,647	0.180***	1,618
		(0.023)		(0.036)	
4	Jun'21	0.165***	1,542	0.225***	1,511
		(0.023)		(0.047)	
5	Jul'21	0.160***	1,354	0.181***	1,327
		(0.027)		(0.058)	
6	Aug'21	0.215***	1,286	0.320***	1,258
		(0.026)		(0.068)	
7	Sep'21	0.237***	1,203	0.352***	1,174
		(0.027)		(0.079)	
8	Oct'21	0.313***	1,112	0.441***	1,079
		(0.028)		(0.092)	
9	Nov'21	0.234***	823	0.407***	789
		(0.035)		(0.130)	
10	Dec'21	0.211***	796	0.473***	763
		(0.032)		(0.116)	
11	Jan'22	0.235***	743	0.480***	731
		(0.033)		(0.127)	
12	Feb'22	0.153***	693	0.411***	687
		(0.035)		(0.110)	
13	Mar'22	0.098***	648	0.256**	636
		(0.038)		(0.117)	
14	Apr'22	0.106***	500	0.263*	496
		(0.041)		(0.137)	
15	May'22	0.100**	392	0.082	390
		(0.049)		(0.155)	

Notes: $\widehat{FAT}_h$ is the Forecasted Average Treatment effect at horizon h waves after the treatment date (wave 14, February 2021). Cross-sectional standard errors in parentheses, following Lemma 1 of Botosaru et al. (2026). N denotes the number of units contributing to the estimate at horizon h; the decline at longer horizons reflects the rotating panel design and the 12-month realization window required to construct the forecast error. $q = 0$: counterfactual equals the pre-treatment sample mean (UC-FAT). $q = 1$: counterfactual extrapolated from a unit-level linear pre-treatment trend.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

compared with $\widehat{FAT}_1 = 0.110$; the MLCM peak is 0.302 at $h = 8$, compared with the FAT peak of 0.313.[9] Figure E1 shows that the dynamic response obtained from MLCM closely tracks the baseline FAT estimates despite relying on a fundamentally different counterfactual construction. This consistency suggests that the main findings are not driven by the specific forecasting structure imposed by the FAT estimator.

Assessing Alternative Explanations. The main identification concern is that the post-February 2021 increase in credit-access forecast errors reflects broader macroeconomic developments overlapping with the fiscal regime shift rather than the fiscal announcement itself. I assess the leading alternative explanations through a range of placebo and validation exercises reported in Appendix C.

First, contemporaneous monetary-policy developments do not appear to explain the results: neither ECB monetary-policy surprises nor alternative measures based on sovereign spreads and PEPP purchases systematically predict pre-treatment credit forecast errors (Appendix C.1). Second, the dynamics are difficult to reconcile with a general pandemic-reopening interpretation. The estimated response builds gradually and peaks several months after the treatment date — a timing inconsistent with pandemic fading, which predates the buildup. Cross-country comparisons further show no systematic relationship between pandemic severity, vaccination dynamics, and the size of forecast-error revisions (Appendix C.3). Third, placebo treatment dates centered on earlier policy events yield estimates close to zero, suggesting that the timing of the effect is specific to the February 2021 fiscal regime shift. Finally, applying the FAT design to alternative forecast-error domains, including inflation and spending expectations, produces different dynamics from those of credit-access forecast errors. This does not rule out a broader improvement in sentiment after the European fiscal shift, but it suggests that the credit-access response is not simply the mechanical reflection of a one-factor optimism shock.

These exercises suggest that the February 2021 fiscal shift contributed to the increase in credit-access forecast errors, while also clarifying that the effect occurred in a broader European macro-financial environment.

Robustness checks. The main findings are robust to a range of alternative specifications and sample restrictions. Survey-weighted estimates, restricting the sample to respondents with longer pre-treatment histories ($R_i \geq 6$), and excluding mortgage holders all leave the estimated effects between $h = 1$ and $h = 8$ largely unchanged. Excluding mortgage holders

[9] The slight difference at the peak horizon reflects in part the larger and more stable MLCM estimation sample, which does not shrink with the rotating panel design.

attenuates the estimates only at longer horizons ($h \geq 11$), consistent with the gradual expiration of the Fondo Gasparrini mortgage moratorium. Results are also robust to alternative FAT specifications, including the PM-FAT variant, and to alternative MLCM implementations. Full robustness results are reported in Appendix D and Appendix E.

6 Mechanism and Behavioral Relevance

Decomposing the Italian forecast-error response. The headline forecast-error response combines two distinct components. By construction, the forecast error is given by the difference between expectations and ex-post realization, as framed in Equation 1. A positive forecast error can therefore arise because households revise expectations upward, because realized credit access later deteriorates, or both.

To distinguish these channels, I re-estimate the FAT decomposition separately for the expectation and realization components. Figure 7 reports the resulting paths. At the peak horizon, $h = 8$, the total FAT estimate is 0.313. The expectation component accounts for 0.132 (42% of the total), while the realization contribution is 0.182 (58%). Thus, the headline forecast-error response is not driven exclusively by contemporaneous belief updating: later realized credit conditions also contribute to the gap.

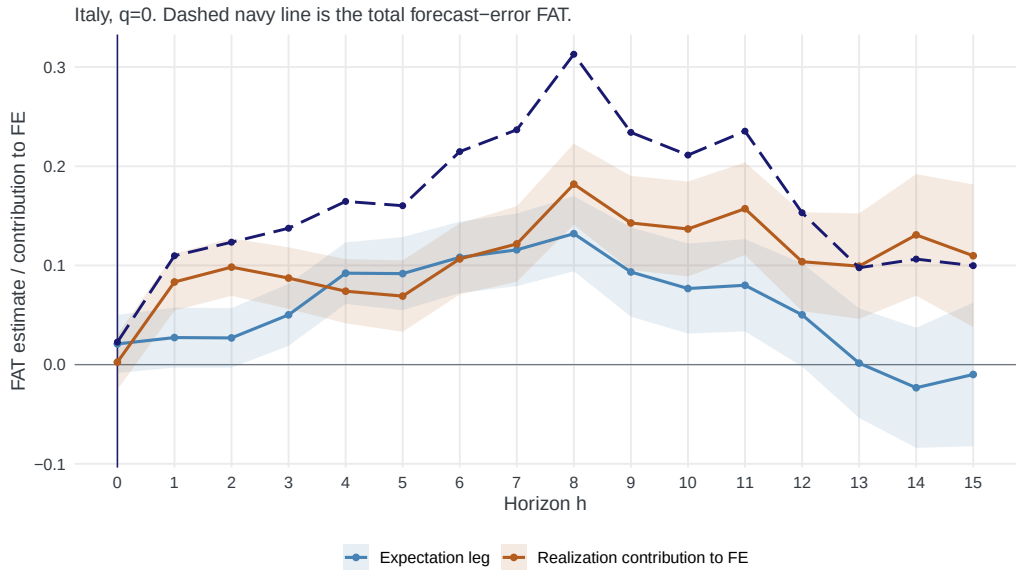


FIGURE 7: Decomposition of the Italian credit-access forecast-error response. The figure reports FAT estimates for the total forecast error, the expectation component, and the realization contribution. The treatment wave is wave 14, February 2021.

This realization-side component is important because the relevant realization windows overlap with a turbulent macro-financial period in 2022, including the energy crisis and

the beginning of monetary tightening. Appendix F shows that comparable realization-side contributions appear in the other CES countries as well, suggesting that this part of the forecast-error response contains a common European component.

The expectation leg, however, remains central. First, the expectation component is positive and statistically significant at the peak horizon. Second, Italy's expectation response is larger than that of the comparison countries at $h = 8$, while the realization component is more broadly shared. Third, when I subtract the average donor-country realization-side FAT from the Italian realization component, the remaining forecast-error FAT is 0.079, still positive and statistically significant ($p = 0.005$).

The dynamic shape of the expectation component is also informative. It builds gradually after the February 2021 event, peaks several months later, and then declines. This hump-shaped profile is consistent with diagnostic expectations (Bordalo et al., 2019): households initially overweight salient fiscal news and subsequently correct as new information arrives. Coibion–Gorodnichenko regressions (Coibion and Gorodnichenko, 2015a), which test whether forecast revisions predict subsequent forecast errors, provide complementary evidence: credit-access beliefs display overreaction-like predictability both before and after the treatment ($\hat{\lambda}^{post} = -0.494$, $\hat{\lambda}^{pre} = -0.483$), suggesting that this expectation domain is structurally prone to belief distortions.[10]

Who responded most. The MLCM delivers respondent-level treatment-effect estimates, which I summarize using regression trees following Cerqua et al. (2024). The tree recursively partitions respondents by observed characteristics to maximize explained variation in estimated effects. I interpret the results descriptively: split variables are measured contemporaneously, so groups identify respondent profiles associated with larger or smaller responses, not causal moderators.

The most informative partition emerges at the peak horizon, $h = 8$. The primary split occurs on trust in the European Commission: respondents with lower institutional trust display substantially larger estimated responses, with effects reaching approximately 0.46 among high-income households within this group.[11] One possible interpretation is that households with lower prior trust in European institutions experienced a larger informational update

[10] The CG slopes do not independently reproduce the hump shape of the FAT event study. The overreaction-like relation is present throughout the sample period, with only a modest post-treatment amplification. Details are reported in Appendix G.

[11] Within the lower-trust branch, the tree further splits by income quintile: the largest effects concentrate in the top two quintiles. Among lower-income respondents in this branch, a further split by gender emerges, with women displaying larger estimated effects than men. Within the higher-trust branch, estimated responses are generally smaller, though renters and mortgage holders exhibit relatively stronger effects than outright homeowners.

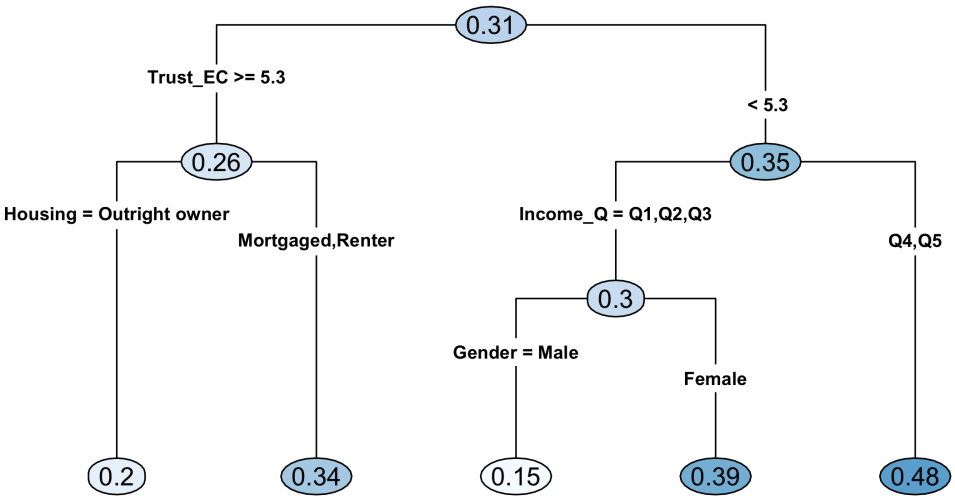


FIGURE 8: Regression-tree partition of estimated MLCM treatment effects at horizon $h = 8$.

following Draghi's Senate address and the broader NextGenerationEU fiscal program. These households revised borrowing expectations upward more strongly precisely because prior beliefs were initially more skeptical. Within the higher-trust branch, estimated responses are considerably smaller across all household types. Overall, the heterogeneity patterns suggest that the aggregate response was concentrated among respondent groups for whom changes in future borrowing conditions were particularly salient, consistent with the diagnostic expectations framework of Bordalo et al. (2019).

Beliefs and Credit-Seeking Behavior. The preceding evidence shows that households became more optimistic about future credit access after the fiscal-credibility shift. A natural question is whether these belief distortions are associated with subsequent economic behavior. I study this by matching each post-treatment credit-access forecast-error observation to the respondent's subsequent quarterly credit-application modules. I then estimate descriptive linear probability models and probit models of credit applications on the individual forecast error.

For the linear probability model, I estimate

$$A_{i,t,h_q} = \alpha + \beta\, \mathrm{FE}_{i,t} + X'_{i,t}\gamma + \lambda_{h_q} + \varepsilon_{i,h_q}, \tag{4}$$

where A_{i,h_q} is an indicator equal to one if respondent i applied for credit in the matched quarterly application module h_q, $\mathrm{FE}_{i,t}$ is the credit-access forecast error measured at monthly wave t, $X_{i,t}$ denotes the control vector, and λ_{h_q} are application-wave fixed effects. The coefficient β is reported in percentage points per one-unit increase in the forecast error.

For the probit specification, I estimate

$$\Pr(A_{i,h_q} = 1 \mid \mathrm{FE}_{i,t}, X_{i,t}, \lambda_{h_q}) = \Phi\left(\alpha + \beta\,\mathrm{FE}_{i,t} + X'_{i,t}\gamma + \lambda_{h_q}\right), \tag{5}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. I report average marginal effects, computed as the average change in the predicted probability associated with a one-unit increase in $\mathrm{FE}_{i,t}$. Table 3 presents the results.

TABLE 3: Credit-access forecast errors and subsequent credit applications

	Linear Probability Model			Probit (AME)		
	(1) Baseline	(2) Controls	(3) Resp. FE	(4) Baseline	(5) Controls	(6) Resp. FE
Panel A: $h_q = 0$						
$FE_{i,t}$	0.95***	0.78***	−0.25	0.90***	0.75***	−0.60
	(0.30)	(0.30)	(0.33)	(0.29)	(0.29)	(0.91)
Panel B: $h_q = 1$						
$FE_{i,t}$	0.96***	0.78**	−0.17	0.90***	0.76**	−0.27
	(0.31)	(0.31)	(0.33)	(0.31)	(0.30)	(0.87)
Panel C: $h_q = 2$						
$FE_{i,t}$	1.01***	0.73**	0.40	0.95***	0.72**	0.88
	(0.32)	(0.32)	(0.35)	(0.32)	(0.31)	(1.01)
Panel D: $h_q = 3$						
$FE_{i,t}$	0.76**	0.55	0.07	0.71**	0.55*	0.12
	(0.33)	(0.34)	(0.42)	(0.33)	(0.32)	(1.21)
Panel E: $h_q = 4$						
$FE_{i,t}$	0.76**	0.57	−0.30	0.71*	0.57	−0.65
	(0.37)	(0.36)	(0.43)	(0.37)	(0.35)	(1.21)
App.-wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	No	Yes	No	No	Yes	No
Respondent FE	No	No	Yes	No	No	Yes
Life-event controls	No	No	Yes	No	No	Yes

Notes: Outcome is an indicator equal to one if the respondent applied for any credit product in the matched quarterly application module. Coefficients are in percentage points per one-unit increase in the credit-access forecast error. Standard errors in parentheses, clustered at the respondent level. Standard errors are derived from the 95% confidence intervals reported by the estimation procedure. Demographic controls include age group, education, employment status, housing tenure, income quintile, and generalized trust. Probit estimates are reported as average marginal effects.
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

The baseline specification controls for application-wave fixed effects. I then add predetermined demographic controls: age group, education, employment status, housing tenure, income quintile, and generalized trust. Finally, I estimate a within-person specification with respondent fixed effects and time-varying life-event controls, which absorbs all stable respondent characteristics and identifies the association only from changes within the same household over time. Standard errors are clustered at the respondent level. The probit results are reported as average marginal effects, so all estimates can be interpreted in percentage points per one-unit increase in the credit-access forecast error.

Figure 9 reports the results for any credit application. A one-unit larger credit-access forecast error is associated with a 0.78 percentage point higher probability of applying for credit in the next quarterly module, controlling for predetermined respondent characteristics and application-wave fixed effects. The estimate is statistically significant at the 5% level. The baseline estimates are positive across the matched quarterly horizons. With demographic controls, a one-unit larger credit-access forecast error is associated with a 0.78 percentage point higher probability of applying for credit at $h_q = 0$, 0.78 percentage points at $h_q = 1$, and 0.73 percentage points at $h_q = 2$. The corresponding probit average marginal effects are very similar: 0.75, 0.76, and 0.72 percentage points. The estimates become smaller and statistically indistinguishable from zero in the respondent fixed-effect specification, suggesting that part of the association reflects persistent differences across households rather than purely within-person changes. However, this could be expected: the within-person specification asks whether the same household applies for more credit in months when their forecast error is larger, but the rotating panel gives very few within-person observations to identify this, so low power is the natural explanation alongside genuine cross-household heterogeneity.

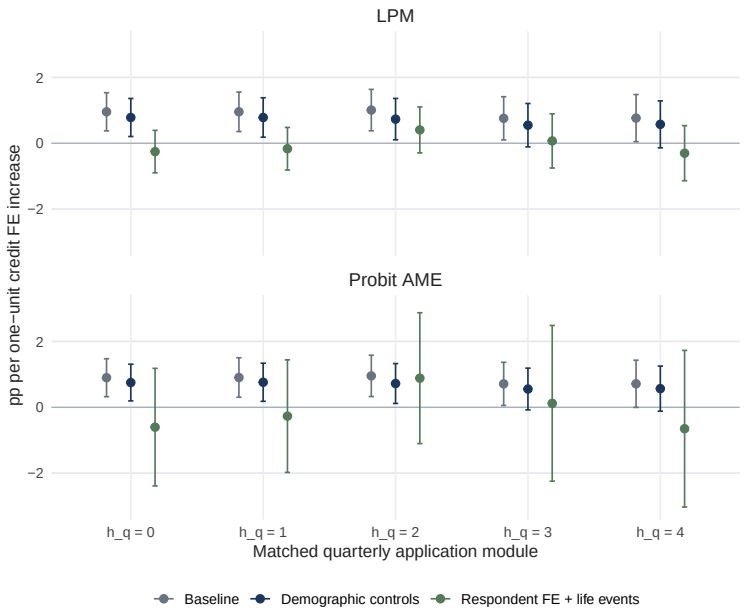


FIGURE 9: Credit-access forecast errors and subsequent credit applications. The figure reports linear probability model estimates and probit average marginal effects for any credit application across matched quarterly application horizons. Coefficients are expressed in percentage points per one-unit increase in the credit-access forecast error. Confidence intervals are 95%, with standard errors clustered at the respondent level.

Product-level estimates, reported in Appendix H, show that the positive association is most visible for consumer loans and, to a lesser extent, vehicle loans and credit-card or overdraft applications. For consumer loans, the controlled linear probability estimates are 0.56 percentage points at $h_q = 0$, 0.95 percentage points at $h_q = 1$, and 0.69 percentage points at $h_q = 2$. These results suggest that the belief distortions documented in the paper are not confined to survey responses: they are positively associated with subsequent credit-market activity, especially for consumer credit.

To understand the magnitude, I combine the peak FAT estimate for credit-access forecast errors, $\widehat{FAT}_8 = 0.313$, with the preferred controlled credit-application slope. For any credit application, the controlled estimate at $h_q = 1$ is 0.78 percentage points, implying an illustrative increase of about 0.24 percentage points in the probability of applying for any credit product at the peak of the belief response. For consumer loans specifically, where the association is strongest, the corresponding slope is 0.95 percentage points, implying an increase of about 0.30 percentage points. These calculations are back-of-the-envelope scaling exercises, not causal mediation estimates, since the fiscal-credibility shift may affect credit applications through channels beyond household credit expectations.

7 Conclusion

This paper shows that the shift in European fiscal policy after the pandemic made households across the euro area persistently over-optimistic about future credit access. Italy, where a domestic credibility event amplified the common European response, provides the case for causal analysis. Using household-level survey data and the Forecasted Average Treatment estimator designed for settings of universal treatment, I document that the effect built gradually over time before declining. The expectation-realization decomposition shows that Italy's distinctiveness lies in the expectation component: while the realization-side contribution is shared across euro area countries, Italy's expectation response is the largest, consistent with a credibility channel tied to the Draghi appointment. The hump-shaped dynamics are difficult to reconcile with standard rational-expectations benchmarks and are instead more consistent with models of diagnostic expectations and overreaction to salient news. I further show that the documented belief distortions are associated with subsequent credit-market activity, particularly for consumer loans, suggesting that the sentiment channel operates not only through expectations but also through economically meaningful behavior.

Taken together, the evidence suggests that fiscal policy communication can shape how households perceive future borrowing conditions, especially when backed by a credible domestic signal within a broader shift in the fiscal environment. Several questions remain open, including the development of a structural model of fiscally induced overreaction, a more formal analysis of the transmission from distorted beliefs to real economic behavior, and the extent to which the magnitude of the sentiment channel depends on the credibility of the announcing policymaker and the institutional environment of the target population.

References

J. Abels. Does the current crisis mark the end of the eu's austerity era? competing political projects in european fiscal governance. Comparative European Politics, 22(2): 192–211, 2024. doi: 10.1057/s41295-023-00346-4. URL https://doi.org/10.1057/s41295-023-00346-4.

C. Altavilla, M. Pagano, and S. Simonelli. Bank exposures and sovereign stress transmission. Review of Finance, 21(6):2103–2139, 2017.

C. Altavilla, L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa. Measuring euro area monetary policy. Journal of Monetary Economics, 108:162–179, 2019.

G.-M. Angeletos and C. Lian. Confidence and the propagation of demand shocks. The Review of Economic Studies, 89(3):1085–1119, 2022.

C. Angelico. Household expectations and the credit cycle. Available at SSRN 3491685, 2019.

R. Bachmann, B. Born, O. Goldfayn-Frank, G. Kocharkov, R. Luetticke, and M. Weber. A temporary vat cut as unconventional fiscal policy. Technical report, National Bureau of Economic Research, 2021.

Banca d'Italia. The 2021 annual report at a glance, 2022. URL https://www.bancaditalia.it/pubblicazioni/relazione-annuale/2021/sintesi/index.html. Published May 2022. Accessed May 2026.

K. Bańkowska, A. Borlescu, E. Charalambakis, A. Dias da Silva, D. Di Laurea, M. Dossche, D. Georgarakos, J. Honkkila, N. Kennedy, G. Kenny, et al. Ecb consumer expectations survey: an overview and first evaluation. ECB Occasional Paper, (2021/287), 2021.

J. M. Barrero. The micro and macro of managerial beliefs. Journal of Financial Economics, 143(2):640–667, 2022.

G. Bellifemine, T. Couturier, and A. Tozzo. The sentiment channel of fiscal policy. 2025. Working Paper.

L. Bocola. The pass-through of sovereign risk. Journal of Political Economy, 124(4):879–926, 2016.

P. Bordalo, N. Gennaioli, R. L. Porta, and A. Shleifer. Diagnostic expectations and stock returns. The Journal of Finance, 74(6):2839–2874, 2019.

B. Born, Z. Enders, and G. J. Müller. On fire, news, and expectations. In The Routledge Handbook of Economic Expectations in Historical Perspective, pages 178–194. Routledge, 2023.

B. Born, N. Lamersdorf, J.-L. Schuster, and S. Steffen. From tweets to transactions: High-frequency inflation expectations, consumption, and stock returns. Technical report, CESifo Working Paper, 2025.

I. Botosaru, R. Giacomini, and M. Weidner. Forecasted treatment effects, 2026. URL https://arxiv.org/abs/2309.05639.

T. Broer and A. N. Kohlhas. Forecaster (mis-) behavior. Review of Economics and Statistics, 106(5):1334–1351, 2024.

M. Buti and G. Papaconstantinou. Does the EU ever learn? – From the financial crisis to the Covid-19 crisis – Looking back to move forward. European University Institute, 2021. doi: doi/10.2870/565110.

A. Cerqua, M. Letta, and F. Menchetti. Causal inference and policy evaluation without a control group, 2024. URL https://arxiv.org/abs/2312.05858.

O. Coibion and Y. Gorodnichenko. What can survey forecasts tell us about information rigidities? Journal of political economy, 120(1):116–159, 2012.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review, 105(8):2644–2678, 2015a. doi: 10.1257/aer.20110306.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review, 105(8):2644–2678, 2015b.

O. Coibion, Y. Gorodnichenko, and M. Weber. Does policy communication during covid work? Technical report, National Bureau of Economic Research, 2020.

O. Coibion, Y. Gorodnichenko, and D. Koustas. Consumption inequality and the frequency of purchases. American Economic Journal: Macroeconomics, 13(4):449–482, 2021a.

O. Coibion, Y. Gorodnichenko, and M. Weber. Fiscal policy and households' inflation expectations: Evidence from a randomized control trial. Technical report, National Bureau of Economic Research, 2021b.

O. Coibion, D. Georgarakos, Y. Gorodnichenko, G. Kenny, and M. Weber. The effect of macroeconomic uncertainty on household spending. American economic review, 114(3): 645–677, 2024.

L. Corrado, D. Fantozzi, and S. Giglioli. Quasi-fiscal policies in times of crisis: A high-frequency data analysis. European Economic Review, page 105265, 2026.

Council of the European Union. Statement of eu ministers of finance on the stability and growth pact in light of the covid-19 crisis, Mar. 2020. URL https://www.consilium.europa.eu/en/press/press-releases/2020/03/23/statement-of-eu-ministers-of-finance-on-the-stability-and-growth-pact-in-light-of-the Press release. Accessed: 2026-06-28.

P. D. Culpepper. The political economy of unmediated democracy: Italian austerity under mario monti. West European Politics, 37(6):1264–1281, 2014.

F. D'Acunto, M. Weber, and X. Yin. Subjective income expectations and household debt choices. University of Chicago, Becker Friedman Institute for Economics Working Paper, 84, 2024.

C. De la Porte and M. D. Jensen. The next generation EU: An analysis of the dimensions of conflict behind the deal. Social Policy & Administration, 55(2):388–402, 2021.

M. Draghi. Comunicazioni del presidente del consiglio dei ministri Mario Draghi al Senato della Repubblica. Senate of the Republic, Rome, February 2021. URL https://www.publicpolicy.net/wp-content/uploads/2021/02/Mario-Draghi-1.pdf. Inaugural address requesting confidence vote.

S. Dullien. Next generation eu: A large common response to the covid-19 crisis. Intereconomics, 55(4):211–213, 2020. doi: 10.1007/s10272-020-0900-6. URL https://www.intereconomics.eu/contents/year/2020/number/4/article/next-generation-eu-a-large-common-response-to-the-covid-19-crisis.html.

F. D'Acunto, D. Hoang, and M. Weber. The effect of unconventional fiscal policy on consumption expenditure. Technical report, National Bureau of Economic Research, 2016.

F. D'Acunto, D. Hoang, and M. Weber. Unconventional fiscal policy. In AEA Papers and Proceedings, volume 108, pages 519–523. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2018.

European Commission. Communication from the commission to the council on the activation of the general escape clause of the stability and growth pact. COM(2020) 123 final, Mar. 2020. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52020DC0123. Accessed: 2026-06-28.

European Commission. NextGenerationEU: European Commission endorses Italy's €191.5 billion recovery and resilience plan. European Commission Press Corner, June 2021. URL https://ec.europa.eu/commission/presscorner/detail/en/ip_21_3126.

European Council. Special meeting of the european council (17, 18, 19, 20 and 21 july 2020) – conclusions. EUCO 10/20, July 2020. URL https://www.consilium.europa.eu/en/policies/eu-recovery-plan/. Accessed: 2026-06-28.

Eurostat. Gdp down by 0.7% and employment up by 0.3% in the euro area. News Release 33/2021, European Commission, Mar. 2021. URL https://ec.europa.eu/eurostat/documents/portlet_file_entry/2995521/2-09032021-AP-EN.pdf/2cf0fd87-a11d-a0eb-ca36-2092f1574f80. Accessed: 2026-06-28.

E. Farhi and J. Tirole. Deadly embrace: Sovereign and financial balance sheets doom loops. Review of Economic Studies, 85(3):1781–1823, 2018.

J.-E. Gallegos, E. García-Miralles, I. Kataryniuk, and S. P. Rodríguez. Fiscal announcements and households' beliefs: Evidence from the euro area. ECB Working Paper Series 3139, European Central Bank, 2025. URL https://www.ecb.europa.eu/pub/research/working-papers/html/index.en.html. ECB Working Paper No. 3139.

D. Georgarakos and G. Kenny. Household spending and fiscal support during the covid-19 pandemic: Insights from a new consumer survey. Journal of monetary economics, 129: S1–S14, 2022.

IMF. Fiscal policy under pressure: High debt, rising risks. Fiscal monitor, International Monetary Fund, Washington, DC, April 2026. URL https://www.imf.org/-/media/files/publications/fiscal-monitor/2026/april/english/text.pdf.

Italian Ministry of Economy and Finance. The national recovery and resilience plan (nrrp), May 2021. URL https://www.mef.gov.it/en/focus/The-National-Recovery-and-Resilience-Plan-NRRP/. Accessed: 2026-06-28.

Y. Ma, T. Ropele, D. Sraer, and D. Thesmar. A quantitative analysis of distortions in managerial forecasts. Technical report, National Bureau of Economic Research, 2020.

E. Macron and A. Merkel. French-German initiative for the European recovery from the coronavirus crisis, May 2020. URL https://www.elysee.fr/en/emmanuel-macron/2020/05/18/french-german-initiative-for-the-european-recovery-from-the-coronavirus-crisis. Joint declaration, 18 May 2020.

L. Melosi, H. Morita, A. Rogantini Picco, and F. Zanetti. Fiscal announcements and the signaling channel of fiscal policy. 2025. ECB Working Paper No. 3110.

OWID. Share of people who received at least one dose of COVID-19 vaccine — Italy. Our World in Data, 2024. URL https://ourworldindata.org/explorers/covid?Metric=People+vaccinated&Interval=Cumulative&Relative+to+population=true&country=~ITA. Data from Official data collated by Our World in Data (2024); World Health Organisation (2025). Downloaded May 2026.

W. Schelkle. Fiscal integration in an experimental union: How path-breaking was the EU's response to the COVID-19 pandemic? JCMS: Journal of Common Market Studies, 59 (S1):44–55, 2021. doi: 10.1111/jcms.13246.

M. Weber. Subjective inflation expectations of households. Business Economics, 57(4): 217–221, 2022.

World Bank. Gdp growth (annual %) — Italy. World Development Indicators, 2020. URL https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=IT. Accessed May 2026.

Appendix

A Salience and Regional Attention

This appendix provides supporting evidence on the public salience of the February 2021 Draghi event. Figure A1 shows that Google search intensity for speech-specific terms rose sharply around Draghi's Senate address. Figure A2 shows that search activity for "Mario Draghi" and "discorso Draghi" was geographically widespread across Italian regions in February 2021. Taken together, these patterns support the interpretation that the event was nationally salient rather than confined to a narrow political or financial audience.

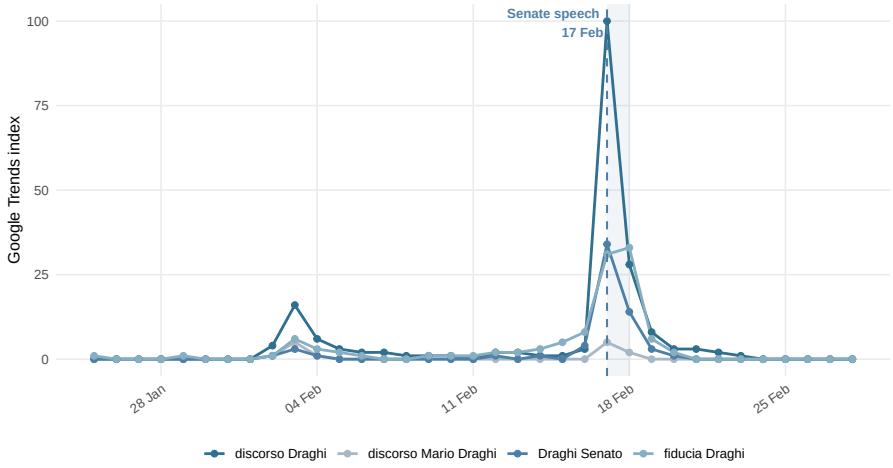


FIGURE A1: Google search intensity for speech-specific terms around Draghi's Senate address. Search intensity is normalized by Google Trends within the selected query and time window.

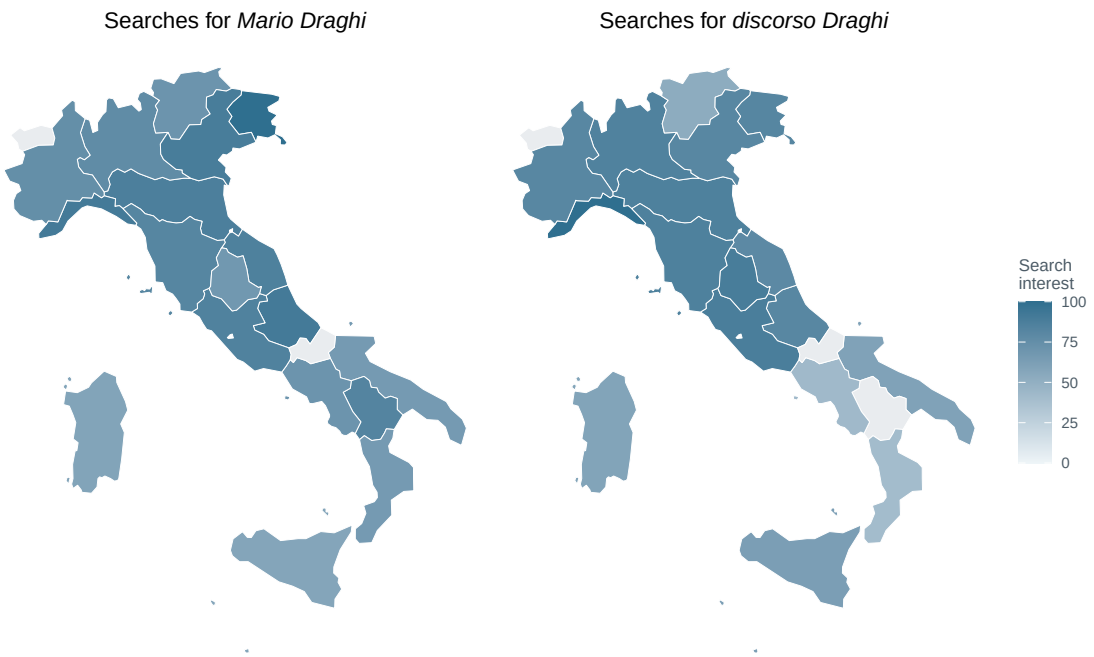


FIGURE A2: Regional Google Trends intensity for searches related to "Mario Draghi" and "discorso Draghi" in February 2021. Darker regions indicate higher relative search intensity within the Google Trends normalization.

B Data

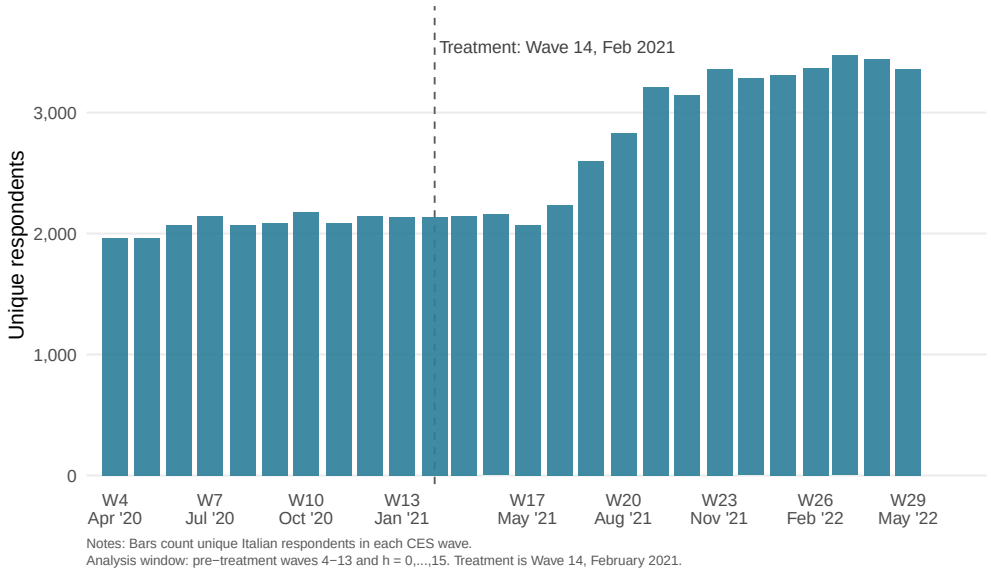


FIGURE B1: Italian CES respondents per wave. Bars report the number of unique Italian respondents in each CES wave. The figure covers the analysis window: pre-treatment waves 4–13 and treatment/post-treatment horizons $h = 0, \ldots, 15$. The treatment wave is wave 14, February 2021.

C Identification Checks

The exercises below are designed to assess the plausibility of the identifying assumptions discussed in Section 4, with particular emphasis on alternative aggregate explanations overlapping with the February 2021 fiscal regime shift.

C.1 Monetary Policy

This subsection tests whether ECB monetary policy can explain the credit forecast-error dynamics before the February 2021 treatment wave. All regressions use only pre-treatment waves 4–13, so the treatment wave itself is excluded. Figure C1 uses the EA-MPD database (Altavilla et al., 2019). I aggregate high-frequency ECB monetary policy surprises to the CES wave-month level and regress Italian mean credit access forecast errors on short-rate surprises, long-rate and QE surprises, Italian sovereign yield changes, and forward-guidance surprises. None of these channels systematically predicts credit forecast errors in the pre-treatment window. Figure C2 implements a complementary check using BTP-Bund spreads and PEPP purchase volumes; it yields the same flat result.

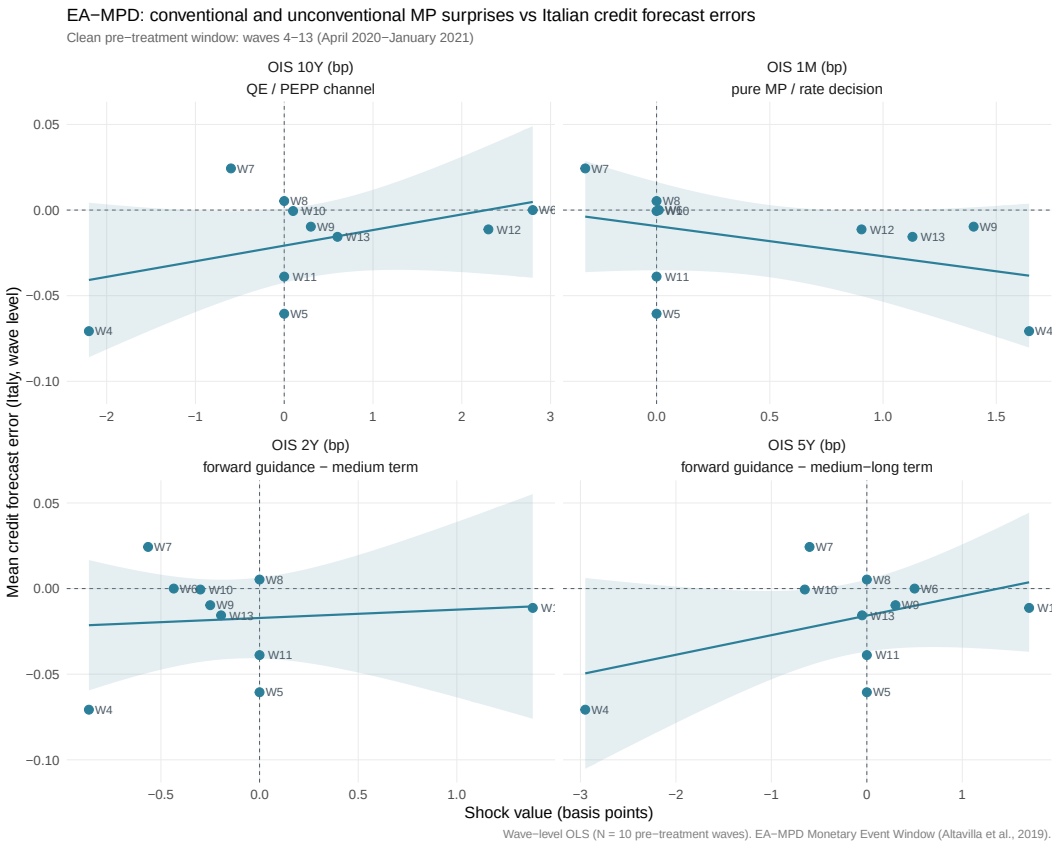


FIGURE C1: EA-MPD structural monetary policy surprises vs. Italian mean credit forecast errors, waves 4–13. Four ECB transmission channels are tested: pure rate decision (OIS 1M), QE/PEPP channel (OIS 10Y), and forward guidance at medium and medium-long term (OIS 2Y, OIS 5Y). Wave-level OLS with $N = 10$ waves. Flat slopes across channels indicate that ECB monetary surprises did not systematically predict Italian credit forecast errors before the February 2021 treatment wave. Source: Altavilla et al. (2019).

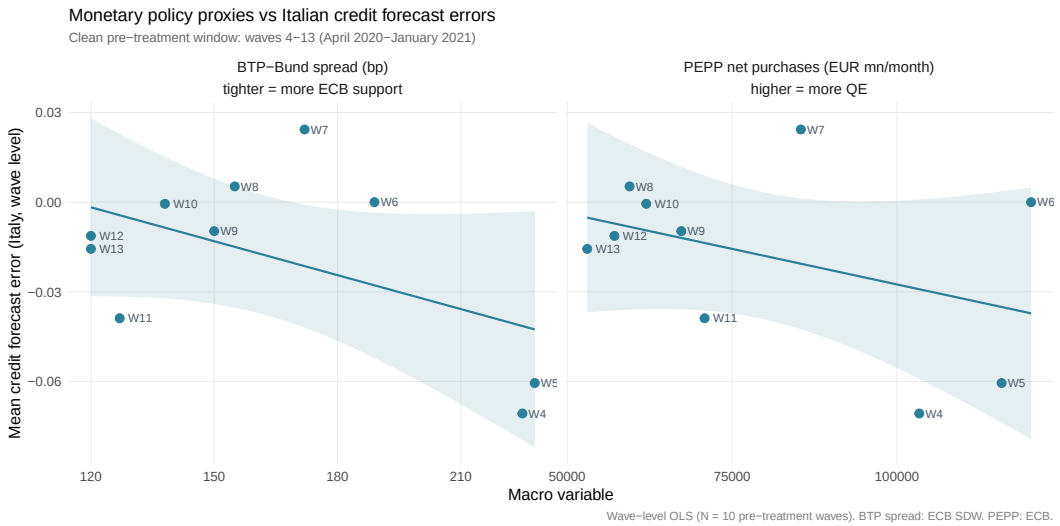


FIGURE C2: Monetary policy proxies (BTP-Bund spread and PEPP net purchases) vs. Italian mean credit forecast errors, waves 4–13. Wave-level OLS with $N = 10$ waves. Flat slopes indicate that sovereign spread compression and ECB asset purchases were not systematically transmitted through household credit access forecast errors before the February 2021 treatment wave. Sources: ECB Statistical Data Warehouse; ECB.

C.2 Earlier Policy Events

Figures C3 and C4 apply the same FAT design to fictitious treatment dates assigned to prior policy events. The placebo coefficient is intentionally evaluated at $h = 1$ because the baseline fiscal effect also appears one wave after the February 2021 treatment month; if those earlier events were the true driver, the estimator should detect effects at those placebo horizons. It does not.

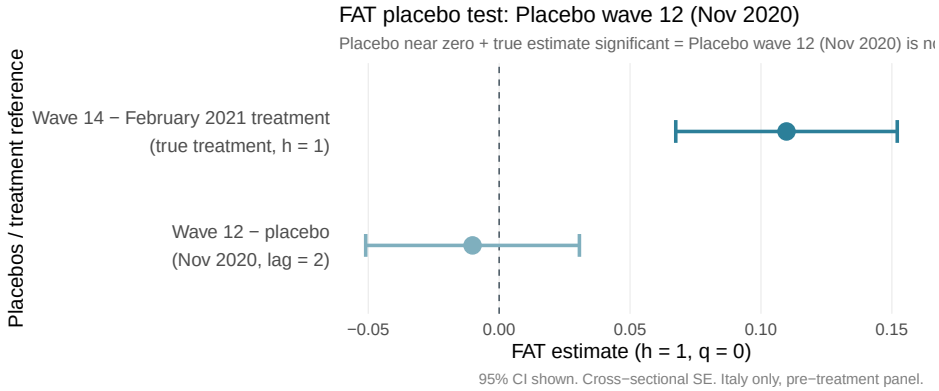


FIGURE C3: FAT placebo test with fictitious treatment dated at wave 12. The plotted coefficient is the placebo response at $h = 1$, one survey wave after the placebo date, matching the timing convention used for the baseline wave-14 treatment. The placebo estimate is close to zero and statistically insignificant, while the true wave-14 treatment estimate at $h = 1$ is positive and significant. 95% confidence intervals shown. Cross-sectional standard errors follow Lemma 1 of Botosaru et al. (2026). Italy only; pre-treatment panel used for placebo construction.

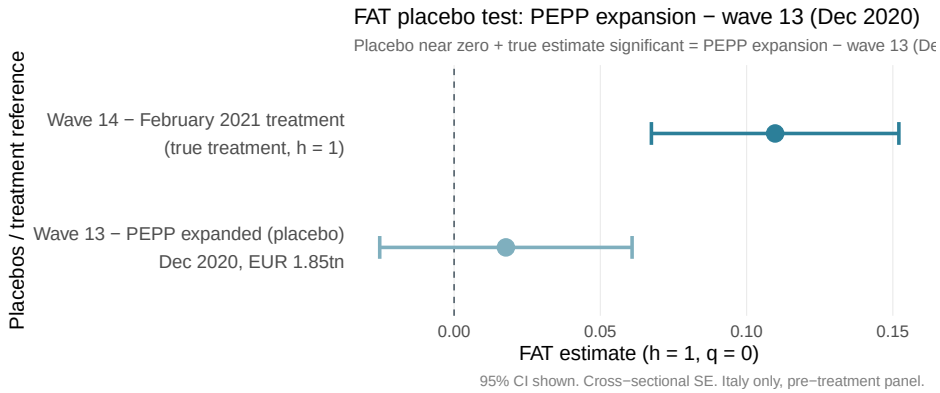


FIGURE C4: FAT placebo test with fictitious treatment dated at wave 13, corresponding to the PEPP expansion placebo. The plotted coefficient is the placebo response at $h = 1$, one survey wave after the placebo date, matching the timing convention used for the baseline wave-14 treatment. The placebo estimate is close to zero and statistically insignificant, while the true wave-14 treatment estimate at $h = 1$ is positive and significant. 95% confidence intervals shown. Cross-sectional standard errors. Italy only.

C.3 Pandemic and Reopening Dynamics

Figure C5 asks whether countries with more severe March 2021 COVID waves experienced larger credit forecast-error jumps. Figure C6 asks whether countries with faster vaccination rollouts or higher June 2021 coverage experienced larger jumps. Both tests use the same jump definition: the mean in waves 14–20 minus the mean in waves 4–13. The cross-country patterns provide little support for either explanation.

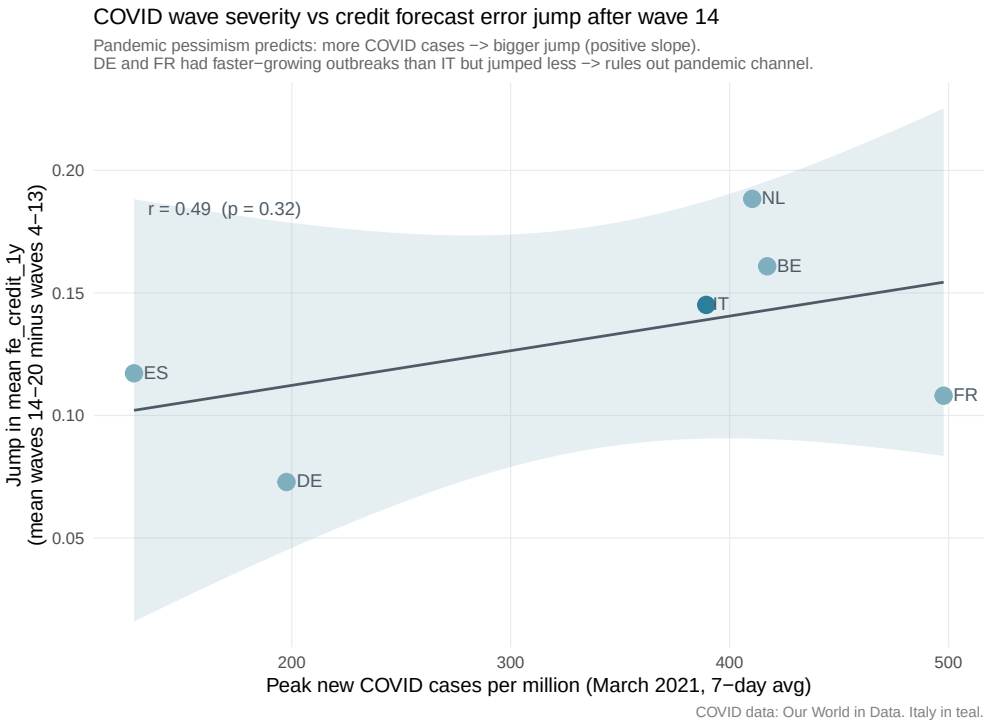


FIGURE C5: COVID wave severity (peak new cases per million, March 2021, 7-day average) vs. the jump in mean credit access forecast errors after wave 14, across six euro area CES countries. The jump is defined as the mean in waves 14–20 minus the mean in waves 4–13. Italy is shown in teal. The correlation is positive but statistically insignificant ($r = 0.49$, $p = 0.32$); Germany and France had faster-growing outbreaks than Italy but smaller forecast-error jumps, inconsistent with a pandemic-severity channel. COVID data: OWID (2024).

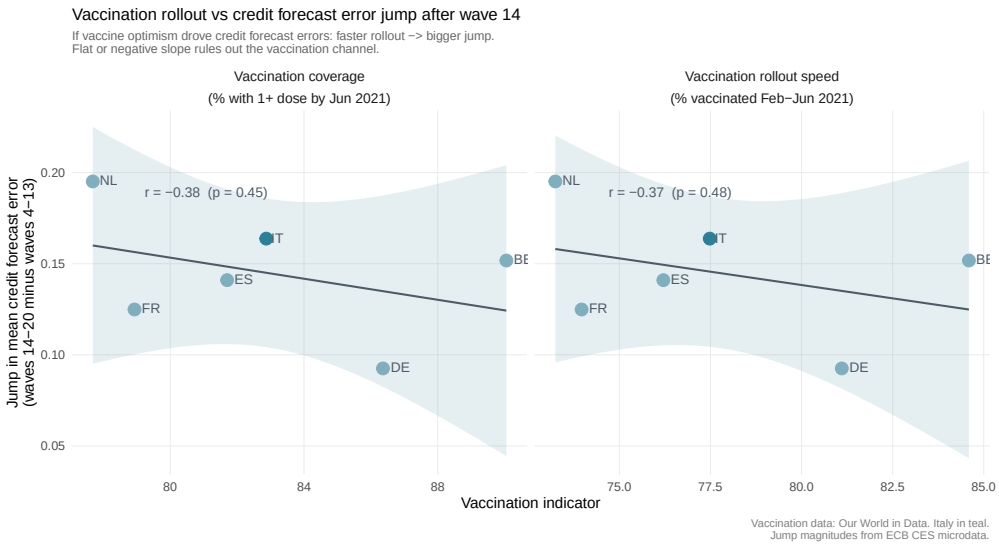


FIGURE C6: Vaccination rollout vs. the jump in mean credit access forecast errors after wave 14, across six euro area CES countries. The jump is defined as the mean in waves 14–20 minus the mean in waves 4–13. Italy is shown in teal. Both correlations are negative and statistically insignificant ($r = -0.16$ for rollout speed and $r = -0.19$ for June 2021 coverage, $p > 0.70$), inconsistent with vaccine optimism as the driver of the post-treatment increase. Vaccination data: OWID (2024).

C.4 Outcome Placebos

This subsection assesses whether the baseline credit-access result reflects a broad shift in household optimism rather than a credit-specific response. If the February 2021 fiscal shift simply made households more optimistic about all future outcomes, the same FAT design should produce similar dynamics across other forecast-error domains.

Figure C7 applies the FAT estimator to credit-access, inflation, and spending-growth forecast errors. The figure reports both pre-treatment deviations and post-treatment FAT estimates. After wave 14, the credit-access forecast error rises gradually and persistently, while the inflation and spending forecast-error paths display different dynamics. This suggests that the main result is not simply the mechanical consequence of common survey timing, the twelve-month realization window, or a generalized optimism shock.

Figure C8 summarizes the same comparison at the peak horizon, $h = 8$, after standardizing each outcome by its own pre-treatment standard deviation. This makes the magnitudes easier to compare across variables measured on different scales (inflation and spending forecast errors are in percentage-point units). The credit-access response is positive and precisely estimated, while the alternative forecast-error outcomes are either much smaller or move in a different direction. Taken together, the two figures support the interpretation that the February 2021 response is concentrated in the credit-access domain.

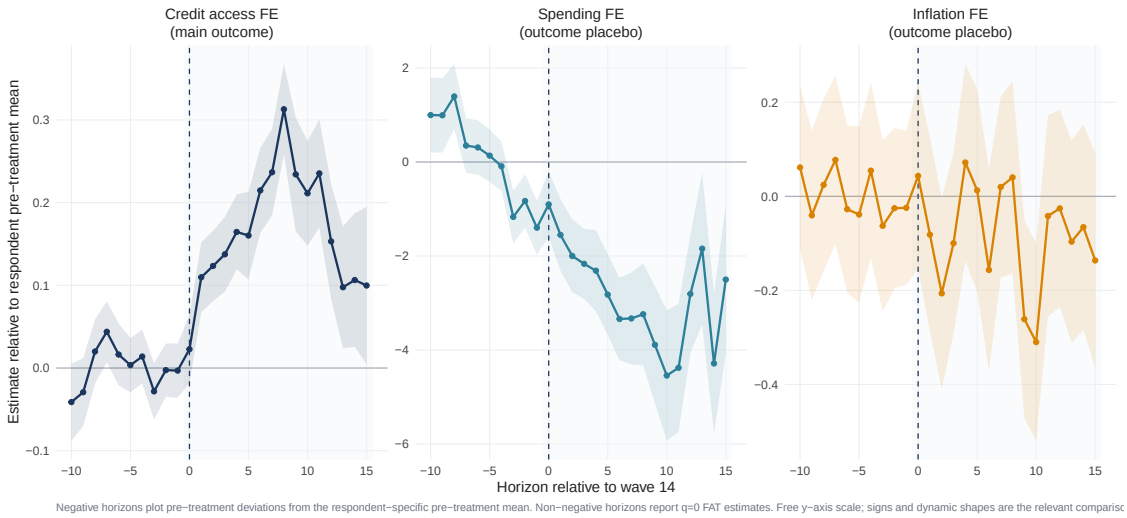


FIGURE C7: Outcome placebo test with pre-treatment deviations and post-treatment FAT estimates. Negative horizons plot pre-treatment deviations from the respondent-specific pre-treatment mean; non-negative horizons report FAT estimates for credit-access forecast errors, inflation forecast errors, and spending-growth forecast errors. Treatment wave 14, February 2021; $q = 0$; Italy only. 95% confidence intervals shown. Outcomes are measured on different scales, so signs and dynamic shapes are the relevant comparison.

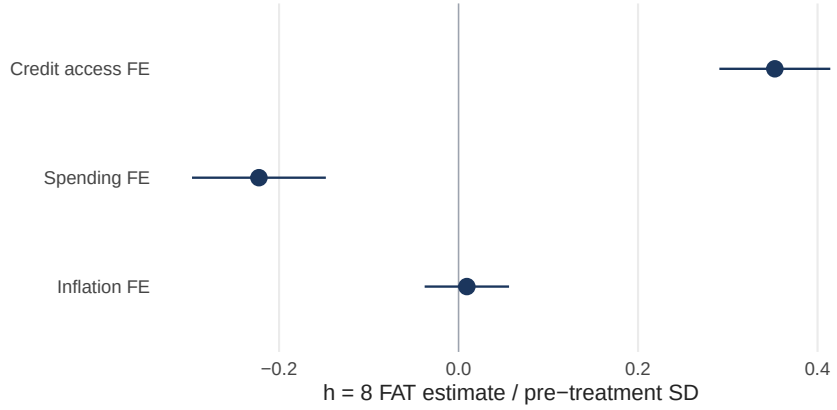


FIGURE C8: Outcome-gradient comparison at $h = 8$. Points report FAT estimates standardized by each outcome's pre-treatment standard deviation. Horizontal bars show 95% confidence intervals.

D FAT Robustness

This appendix reports robustness checks for the FAT estimates. The baseline specification uses the UC-FAT estimator with treatment at wave 14, corresponding to February 2021, and defines horizons as $h = 0, \ldots, 15$. The preferred specification sets $q = 0$, so each respondent's counterfactual is given by their own pre-treatment mean credit-access forecast error.

The checks below assess sensitivity to the polynomial order, inference, pre-treatment window, respondent pre-treatment history, sample composition, survey weights, and persistence assumptions.

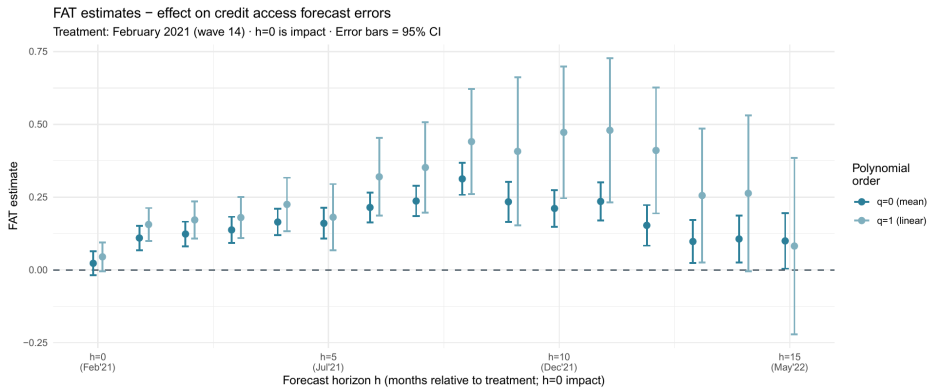


FIGURE D1: FAT estimates for $q = 0$ and $q = 1$. The preferred specification is $q = 0$, which uses each respondent's pre-treatment mean as the counterfactual. The $q = 1$ specification allows for an individual linear pre-treatment trend. Treatment occurs at wave 14, February 2021, and horizons run from $h = 0$ to $h = 15$.

D.1 Placebo Treatment Dates

Figure D2 reports placebo estimates using fictitious treatment dates in the pre-treatment period, together with the post-treatment FAT estimates. These placebo exercises test whether the estimator mechanically detects effects before the actual February 2021 treatment. The placebo estimates are close to zero relative to the post-treatment response.

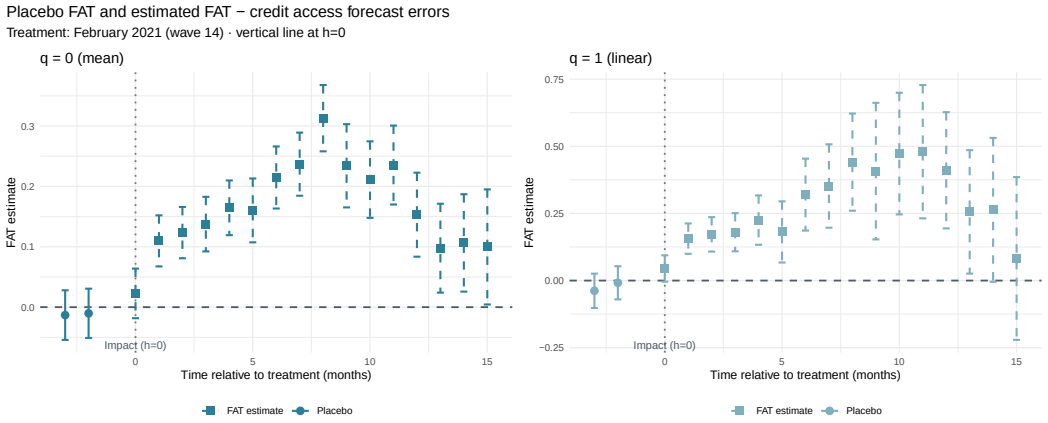


FIGURE D2: Placebo and post-treatment FAT estimates. Negative values on the horizontal axis denote placebo treatment dates before February 2021; non-negative values denote post-treatment horizons. The absence of sizeable placebo effects supports the interpretation that the estimated response is specific to the February 2021 treatment timing.

D.2 Respondent-Clustered Inference

Figure D3 compares the baseline confidence intervals with respondent-clustered confidence intervals. This check accounts for the fact that the same respondent may contribute observations at multiple horizons in the rotating CES panel.

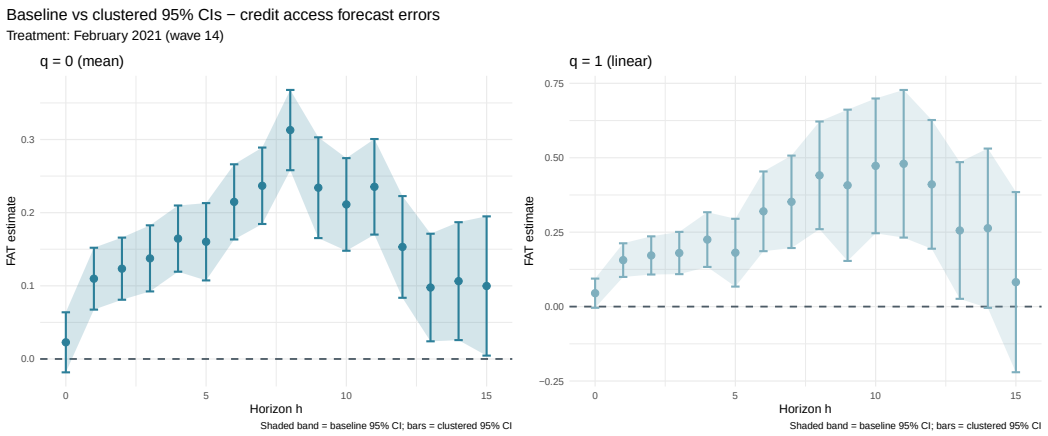


FIGURE D3: Baseline and respondent-clustered confidence intervals for the FAT estimates. Shaded bands show the baseline 95% confidence intervals; error bars show respondent-clustered 95% confidence intervals. The comparison shows that inference is not driven by ignoring repeated observations from the same respondent across horizons.

D.3 Restricted Pre-Treatment Window

Figures D4 and D5 restrict the pre-treatment estimation window to waves 6–13, excluding April and May 2020. This addresses the concern that the earliest pandemic months may represent an unstable regime for credit expectations. The resulting estimates remain close to the baseline profile.

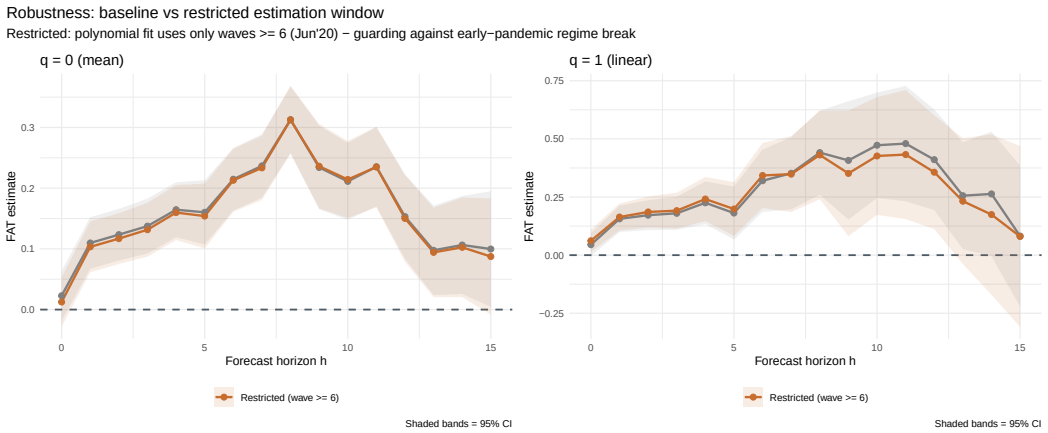


FIGURE D4: Robustness to restricting the pre-treatment estimation window. The restricted specification uses only waves 6–13 to construct the counterfactual, excluding the earliest pandemic waves. The comparison shows that the main FAT profile is not driven by April–May 2020 observations.

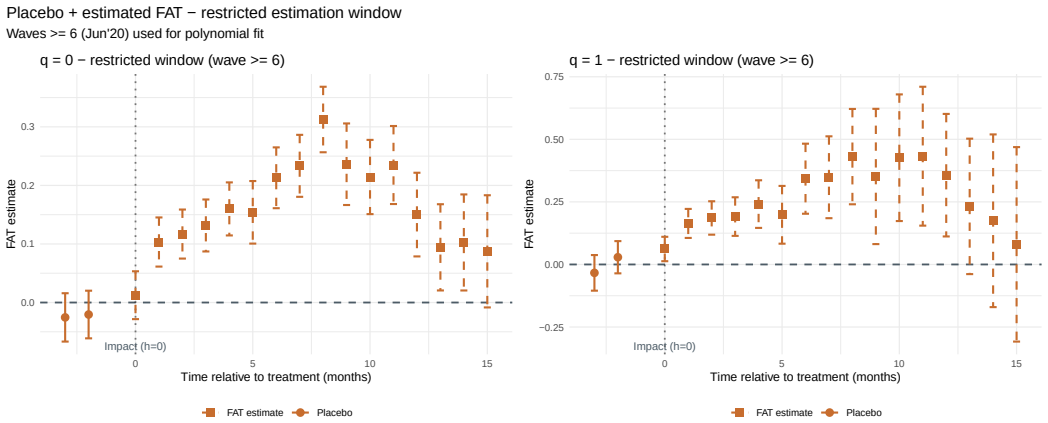


FIGURE D5: Placebo and post-treatment FAT estimates using the restricted pre-treatment window. The figure repeats the combined placebo/event-study exercise after constructing the counterfactual from waves 6–13 only.

D.4 Sample-Composition Robustness

Figure D6 restricts the sample to respondents with at least six valid pre-treatment observations of the credit-access forecast error. This checks whether the baseline estimates are driven by respondents with very short pre-treatment histories. Figure D7 excludes respondents who ever report having a mortgage on their main residence, addressing the possibility that mortgage-specific policy interventions affect the outcome. Figure D8 re-estimates the FAT using CES survey weights to assess whether sample composition drives the results.

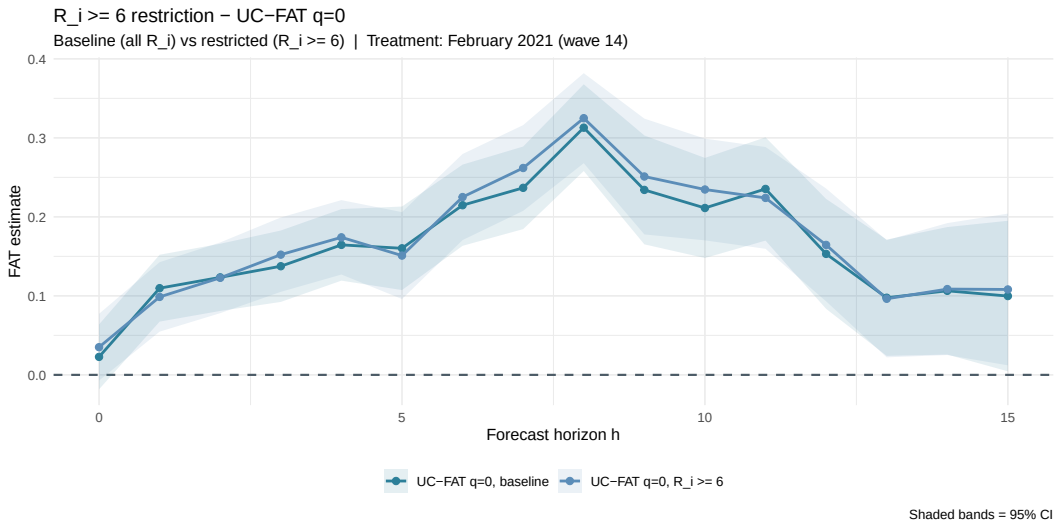


FIGURE D6: Robustness to requiring at least six pre-treatment forecast-error observations per respondent. The restricted sample keeps only respondents with $R_i \geq 6$. The similarity with the baseline indicates that the main result is not driven by respondents with sparse pre-treatment histories.

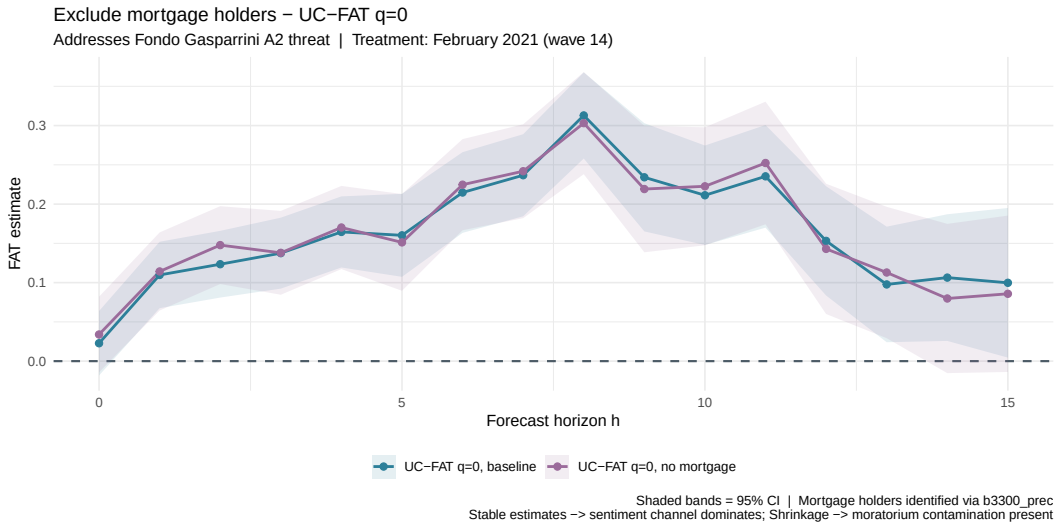


FIGURE D7: Robustness to excluding mortgage holders. Respondents who ever report having a mortgage on their main residence are removed from the sample. The check addresses the concern that mortgage-specific policies or moratoria could contaminate credit-access forecast errors.

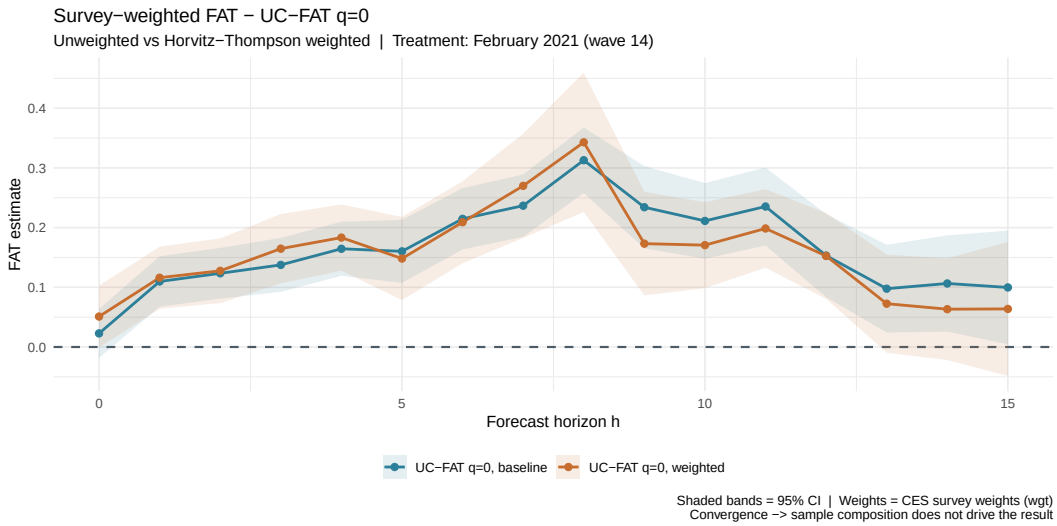


FIGURE D8: Survey-weighted FAT estimates. The weighted specification uses CES survey weights and compares the resulting estimates with the unweighted baseline. The stability of the estimates indicates that sample composition is unlikely to drive the result.

D.5 Persistence-Based Counterfactual

As an additional specification check, I compare the preferred UC-FAT estimator with a persistence-based PM-FAT specification that uses an AR(1)-type counterfactual. This exercise evaluates whether the results are sensitive to allowing persistence in individual forecast errors rather than using the respondent-specific pre-treatment mean. The placebo version of the PM-FAT exercise checks whether the AR(1) counterfactual detects spurious effects before treatment, while the distribution of estimated persistence parameters documents how much individual-level persistence the PM-FAT specification actually introduces.

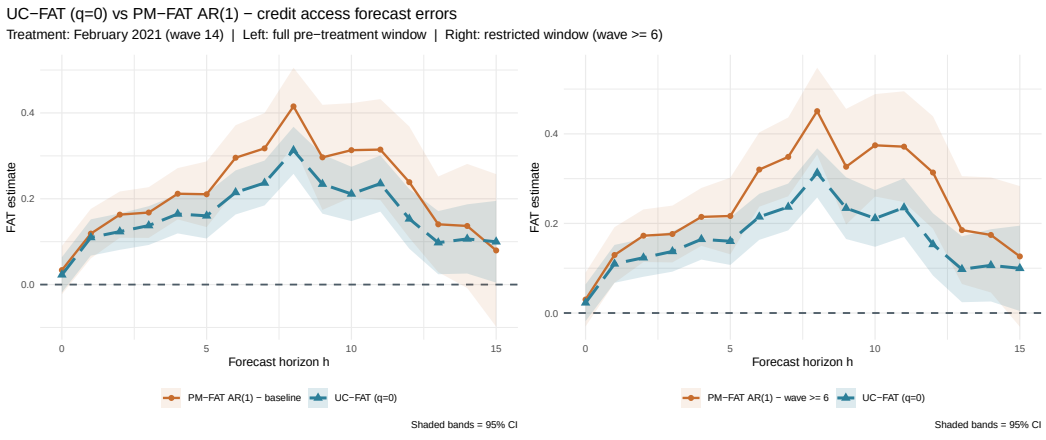


FIGURE D9: UC-FAT versus PM-FAT estimates. The figure compares the preferred unconditional FAT specification with a persistence-based parametric FAT specification.

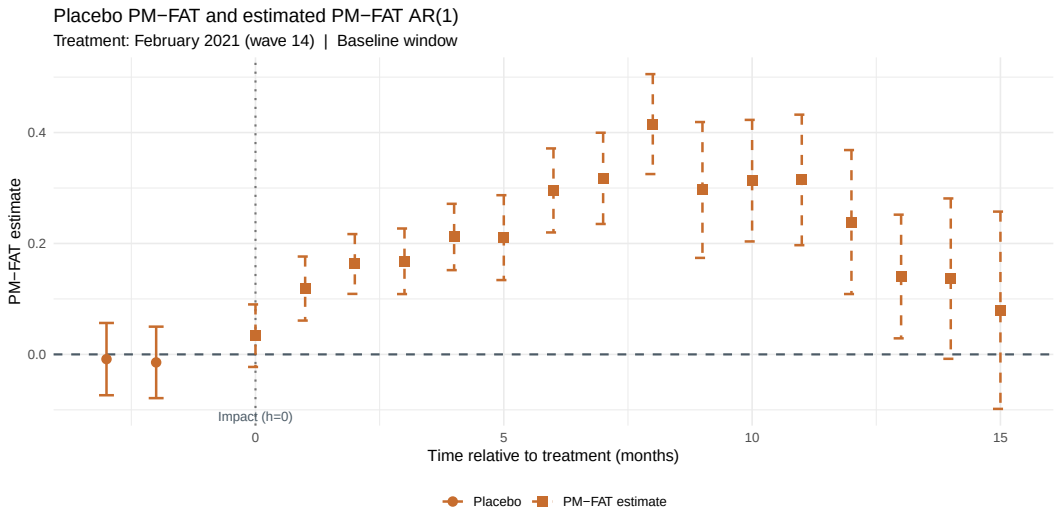


FIGURE D10: PM-FAT placebo and post-treatment estimates. The figure repeats the placebo and event-study exercise using the persistence-based PM-FAT counterfactual.

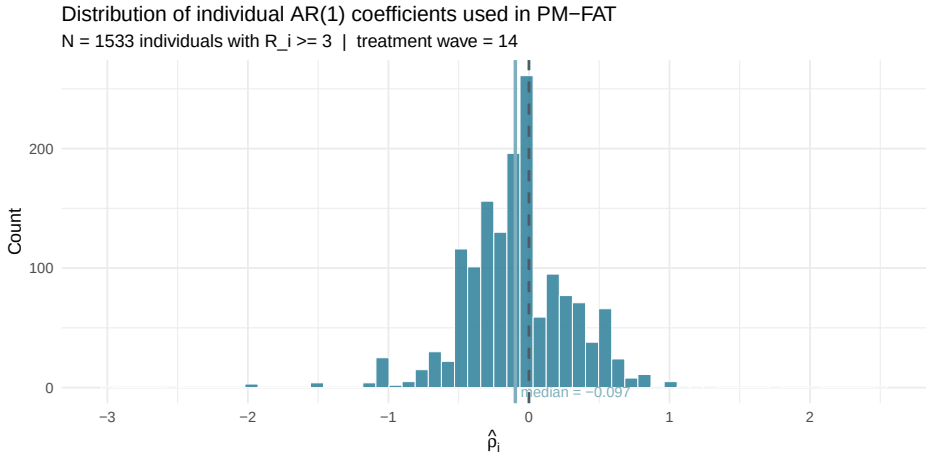


FIGURE D11: Distribution of estimated persistence parameters in the PM-FAT specification. The figure documents the respondent-level persistence estimates used to construct the PM-FAT counterfactual.

Overall, the FAT response is robust across alternative polynomial orders, placebo treatment dates, respondent-clustered inference, restricted pre-treatment windows, minimum pre-treatment observation requirements, exclusion of mortgage holders, survey weighting, and persistence-based counterfactuals. The robustness checks support the interpretation that the post-February 2021 increase in credit-access forecast errors is not driven by pre-treatment drift, sparse respondent histories, sample composition, mortgage-specific policies, or the particular counterfactual specification.

E MLCM Implementation and Validation

This appendix provides additional details on the implementation of the Machine Learning Control Method (MLCM), which I use as a corroboration exercise in the main text. Following Cerqua et al. (2024), the MLCM constructs respondent-level untreated counterfactual outcomes by training supervised machine-learning models on pre-treatment data and forecasting post-treatment outcomes out of sample.

Covariates and sample construction. The baseline specification uses a conservative set of predetermined respondent characteristics: gender, household size, number of children, partner status, education, employment status, housing tenure, generalized trust, financial decision-making responsibility, planning horizon, and age group. Income quintiles and institutional-trust variables are excluded from the baseline counterfactual model because they are time-varying survey measures that could themselves respond to the post-treatment environment. Age is retained because it evolves mechanically and cannot be affected by the fiscal regime shift.

The outcome is the credit-access forecast error, constructed as expected credit access at wave t minus realized credit access over the corresponding twelve-month realization window. As in the FAT analysis, the treatment wave is wave 14, February 2021, and horizons are indexed as $h = 0, \ldots, 15$.

The MLCM sample includes Italian CES respondents with non-missing forecast-error realizations and available covariates in the relevant estimation window. Unlike the FAT estimator, MLCM does not require each respondent to contribute a long individual pre-treatment outcome history. This is useful in the CES rotating panel: the effective post-treatment sample remains large at each horizon, ranging from 1,557 to 2,159 respondents.

Model selection. I estimate four candidate prediction models on the pre-treatment sample: Least Absolute Shrinkage and Selection Operator (LASSO), Partial Least Squares (PLS), Random Forest, and Stochastic Gradient Boosting (SGB). Model selection is based on out-of-sample pre-treatment predictive performance using the panel cross-validation procedure in Cerqua et al. (2024).

Random Forest is selected as the preferred specification. The preferred specification uses ten covariates and has an out-of-sample RMSE of 0.834. Table E1 reports the best cross-validated RMSE across the variable-count grid used in the model-selection step. The predictive performance is very similar across grid choices, confirming that individual-level credit-access forecast errors are difficult to predict from respondent characteristics alone.

TABLE E1: MLCM panel cross-validation performance

Number of covariates	Best RMSE
8	0.838
10	0.834
12	0.840

Notes: RMSEs are computed from out-of-sample pre-treatment predictions. The preferred specification is Random Forest with ten covariates, which minimizes prediction error. The final Random Forest tuning parameter is $mtry = 2$.

Treatment-effect construction. Let $\widehat{\text{FE}}_{i,\tau+h}(0)$ denote the predicted untreated forecast error for respondent i at horizon h, obtained from the fitted Random Forest model. The MLCM estimate at horizon h is computed as:

$$\widehat{ATT}_h^{MLCM} = \frac{1}{N_h} \sum_i \left(\text{FE}_{i,\tau+h} - \widehat{\text{FE}}_{i,\tau+h}(0) \right),$$

where τ corresponds to CES wave 14, February 2021.

Results. Figure E1 reports the dynamic MLCM estimates. The estimated effect is small and statistically indistinguishable from zero in the treatment month. It becomes positive and statistically significant from $h = 1$ onward, rises through late 2021, and peaks at $h = 8$. The average post-treatment MLCM estimate is 0.168, with a 95% confidence interval of 0.153–0.183. The peak estimate at $h = 8$ is 0.302, with a 95% confidence interval of 0.267–0.340.

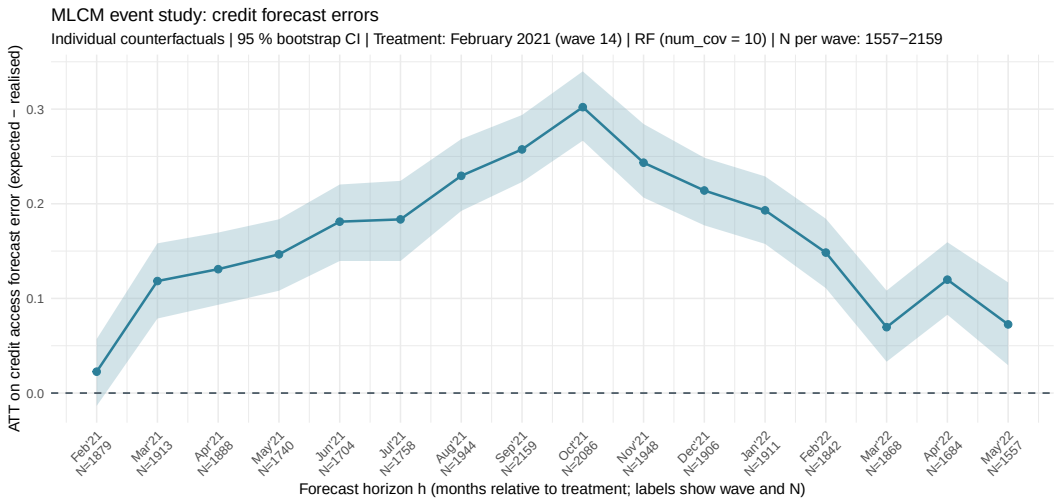


FIGURE E1: Dynamic MLCM estimates of the fiscal regime shift on credit-access forecast errors. The treatment wave is wave 14, February 2021. Shaded bands report 95% bootstrap confidence intervals. Axis labels report the calendar wave and the number of respondents contributing to each horizon.

Robustness checks. I assess the robustness of the MLCM corroboration exercise along four dimensions. First, I re-estimate the model using only background characteristics, excluding age from the covariate set. Second, I restrict the training window to later pre-treatment waves, using either waves 9–13 or waves 7–13, to verify that the results are not driven by the early pandemic months. Third, I combine the two restrictions, using background covariates only and training on waves 9–13. Finally, I run a placebo exercise that assigns the treatment to wave 12 and evaluates the one-wave-ahead placebo response at $h = 1$.

Figure E2 summarizes the robustness checks. The post-treatment average effect remains very close to the baseline estimate across all alternative specifications: 0.165 using background covariates only, 0.170 when training on waves 9–13, 0.161 when training on waves 7–13, and 0.166 when combining background-only covariates with the waves 9–13 training window. In contrast, the wave-12 placebo estimate at $h = 1$ is essentially zero, -0.002, with a confidence interval that includes zero. These checks indicate that the MLCM corroboration is not sensitive to the covariate set, the pre-treatment training window, or a nearby placebo treatment date.

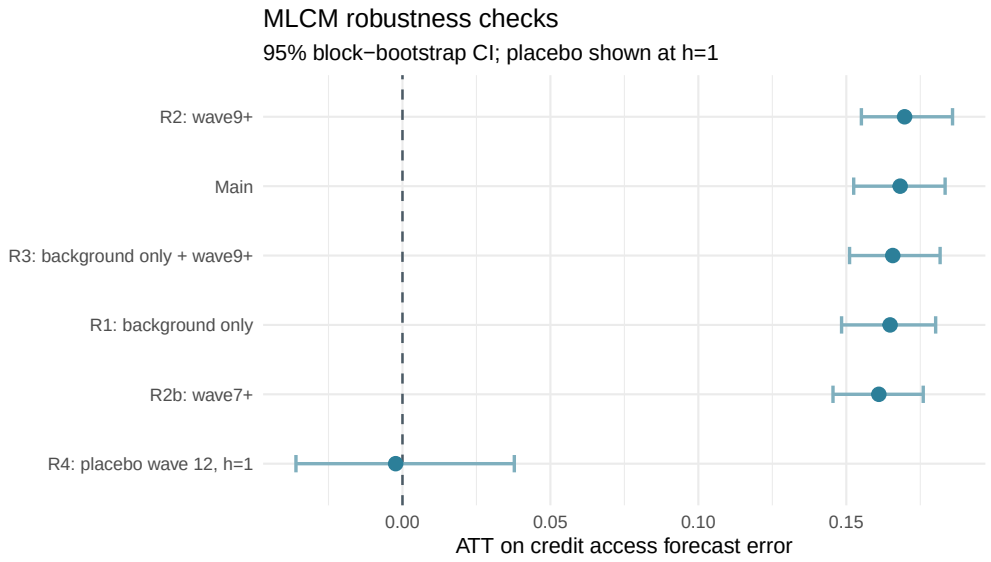


FIGURE E2: MLCM robustness checks. The figure reports the average post-treatment MLCM estimate under alternative covariate sets and pre-treatment training windows. The placebo specification assigns treatment to wave 12 and reports the one-wave-ahead placebo estimate at $h = 1$. Horizontal bars report 95% bootstrap confidence intervals.

TABLE E2: MLCM robustness summary

Specification	Estimate	95% CI lower	95% CI upper
Baseline	0.168	0.153	0.183
Background covariates only	0.165	0.148	0.180
Training waves 9–13	0.170	0.155	0.186
Training waves 7–13	0.161	0.146	0.176
Background only, training waves 9–13	0.166	0.151	0.182
Placebo wave 12, $h = 1$	-0.002	-0.036	0.038

Notes: The baseline and robustness estimates report the average post-treatment MLCM effect. The placebo row reports the one-wave-ahead placebo response after assigning treatment to wave 12. Confidence intervals are bootstrap confidence intervals.

Interpretation and limitations. The MLCM exercise is best interpreted as a corroboration analysis rather than the primary identification strategy. Unlike the FAT estimator, which relies only on respondents' own pre-treatment outcome dynamics, MLCM also assumes that the relationship between respondent characteristics and untreated potential outcomes remains stable enough for the pre-treatment prediction rule to provide a valid counterfactual forecast. The close agreement between FAT and MLCM estimates across all horizons supports the robustness of the main findings to the choice of counterfactual construction.

F Cross-Country Decomposition of Forecast Errors

This appendix reports the cross-country decomposition of credit-access forecast-error responses into expectation and realization components. The exercise is used to assess whether the realization-side contribution documented for Italy is specific to the Italian treatment window or instead reflects a broader European macro-financial component.

Figure F1 plots the expectation and realization contributions by country over the post-treatment horizons. The realization contribution is positive in several countries around the same horizons, consistent with common macro-financial shocks affecting the later realization windows. By contrast, Italy stands out more clearly on the expectation side, where the response is larger and more persistent.

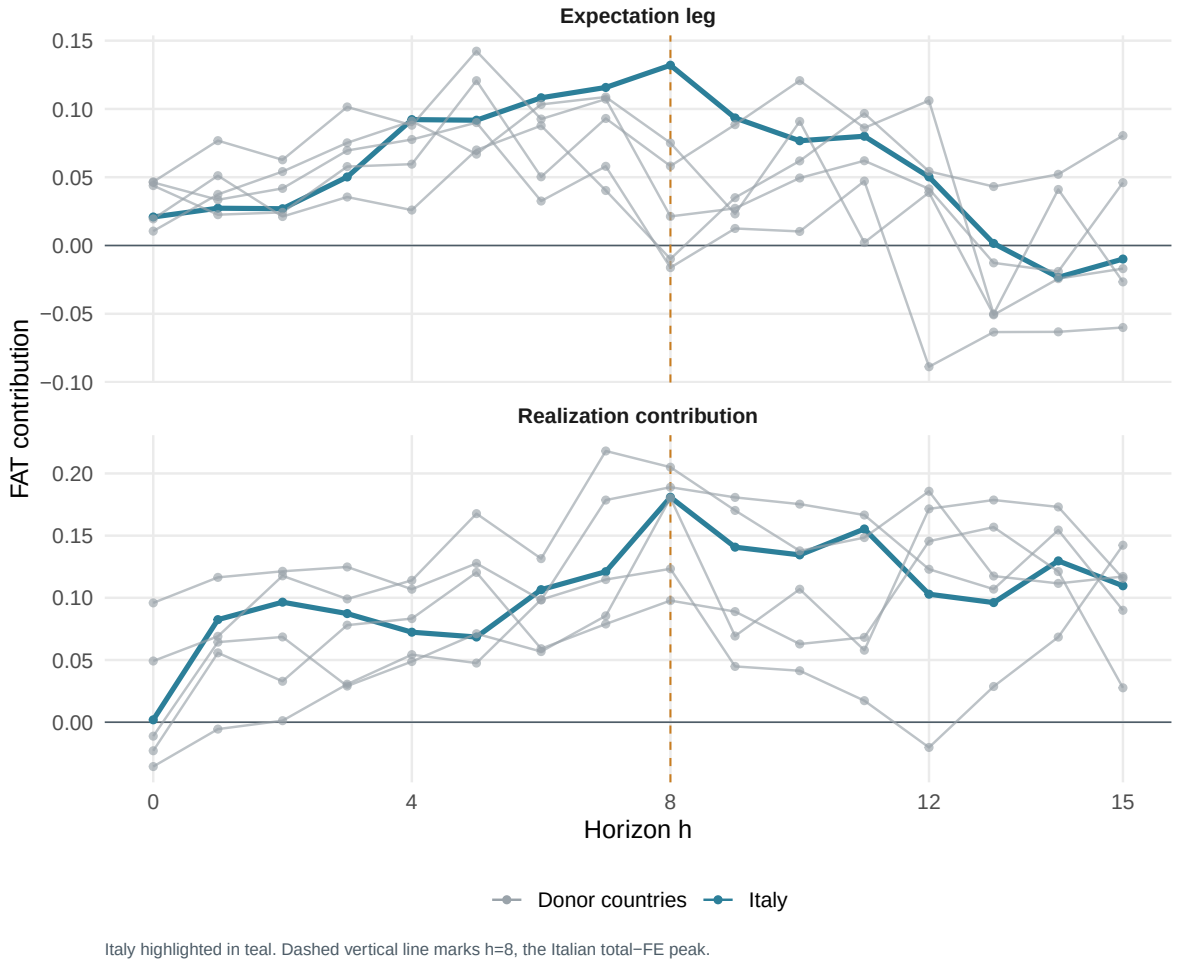


FIGURE F1: Cross-country decomposition of credit-access forecast-error responses. The figure reports FAT-style decompositions into the expectation component and the realization contribution for the six CES countries. Horizons are indexed relative to wave 14, February 2021.

Figure F2 summarizes the same decomposition at selected horizons. The bar plots show

that realization-side contributions are broadly present across countries, while the Italian expectation component is especially large at the peak horizon.



FIGURE F2: Expectation and realization contributions at selected horizons across countries. Bars report the contribution of each component to the total credit-access forecast-error response. The figure highlights that realization-side movements are common across countries, while Italy's expectation component is particularly pronounced around the peak horizon.

As an additional summary, Figure F3 reports the full horizon profile as heatmaps. This representation makes clear that the realization-side contribution is not uniquely Italian, whereas the timing and persistence of the expectation component are more distinctive in Italy.

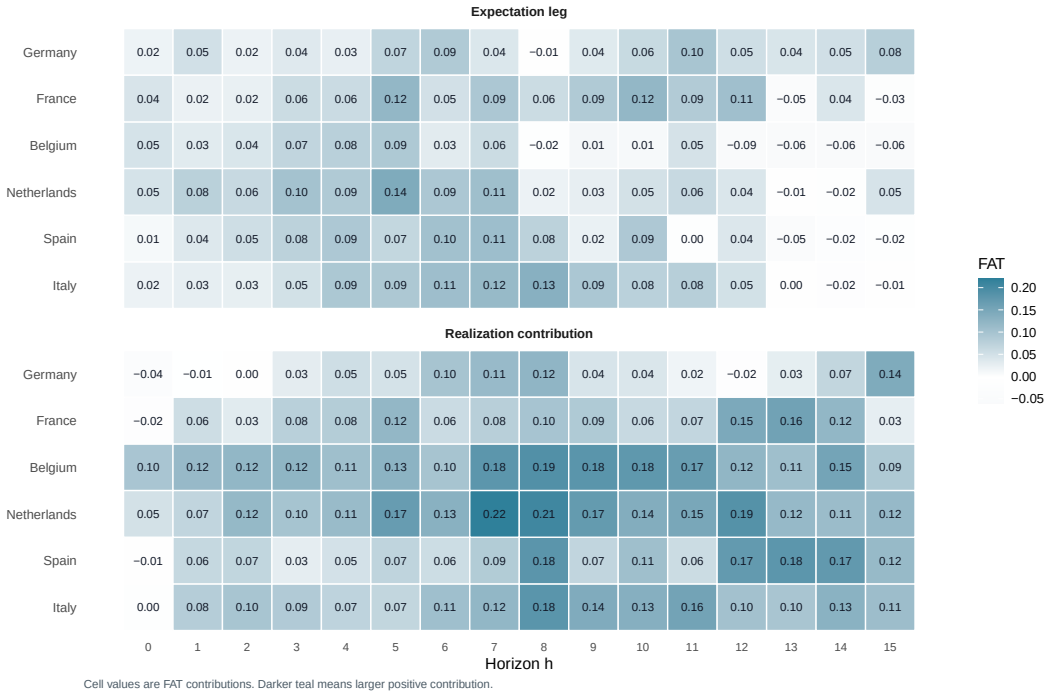


FIGURE F3: Heatmap summary of expectation and realization contributions across countries and horizons. Darker values indicate larger positive contributions to the total credit-access forecast-error response.

G Forecast Revisions and FIRE

This appendix reports Coibion–Gorodnichenko-style regressions (Coibion and Gorodnichenko, 2015a) that test whether forecast revisions predict subsequent forecast errors. Under full information rational expectations, forecast revisions should not systematically predict subsequent forecast errors. In this setting, I estimate:

$$FE_{i,t}^{CG} = \alpha_i + \delta_t + \lambda \left(E_{i,t} CA_{i,t+12} - E_{i,t-1} CA_{i,t+11} \right) + \varepsilon_{i,t},$$

where $FE_{i,t}^{CG}$ is signed so that negative values of λ indicate overreaction: upward forecast revisions predict lower subsequent realizations relative to expectations. The regression uses respondent and wave fixed effects, with standard errors clustered at the respondent level. Because revisions require consecutive survey waves, the sample is restricted to respondent-wave observations with valid consecutive credit-access expectations.

Table G1 reports the main pre/post estimates. The slope is negative and statistically significant both before and after the February 2021 treatment wave. The post-treatment slope is slightly more negative than the pre-treatment slope, but the main takeaway is that overreaction-like predictability is present throughout the sample rather than appearing only

after the treatment.

TABLE G1: Coibion–Gorodnichenko regressions for credit-access expectations

Specification	$\hat{\lambda}$	Std. error	Observations
Pre-treatment waves 5–13	-0.483	0.008	13,573
Post-treatment waves 14–29	-0.494	0.006	27,505
Pooled pre slope	-0.469	0.008	41,186
Post minus pre	-0.025	0.010	41,186

Notes: The dependent variable is the credit-access forecast error under the Coibion–Gorodnichenko sign convention. The regressor is the one-wave forecast revision in expected future credit access. Pre- and post-treatment split regressions include respondent and wave fixed effects; standard errors are clustered by respondent. The pooled specification interacts the revision term with a post-treatment indicator.

Figure G1 visualizes the pre- and post-treatment slopes. Both estimates are negative and tightly estimated. The modest difference between the two slopes indicates that the February 2021 event did not create the revision-error relationship from scratch; rather, it operated in a belief domain already characterized by overreaction-like forecast dynamics.

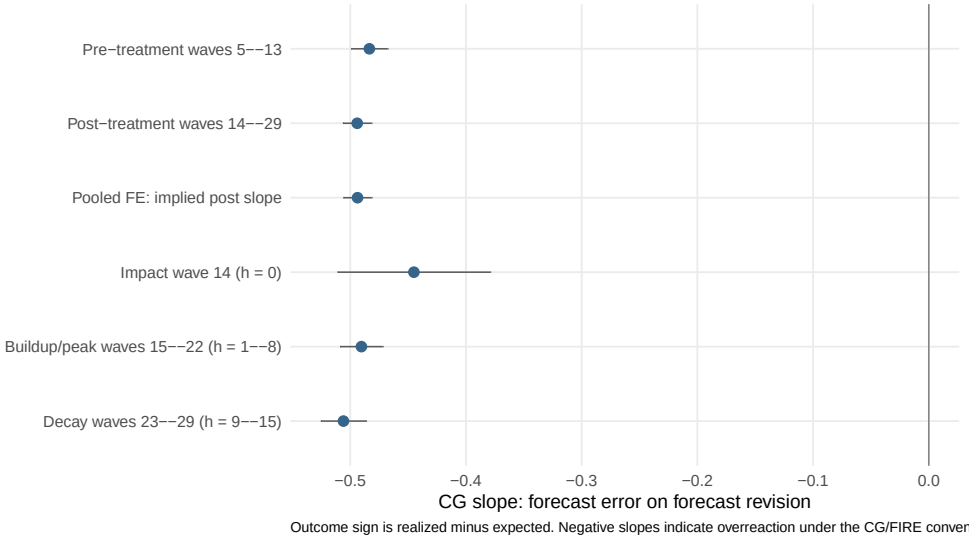


FIGURE G1: Pre- and post-treatment Coibion–Gorodnichenko slopes for credit-access expectations. Negative slopes indicate overreaction-like predictability under the sign convention used in the appendix.

Figure G2 reports the same regression over dynamic windows: the pre-treatment period, the impact wave, the buildup/peak window, and the decay window. The estimates remain negative across all windows. This confirms that the CG regression supports overreaction-like

belief dynamics, but it does not independently reproduce the hump-shaped FAT event-study profile.

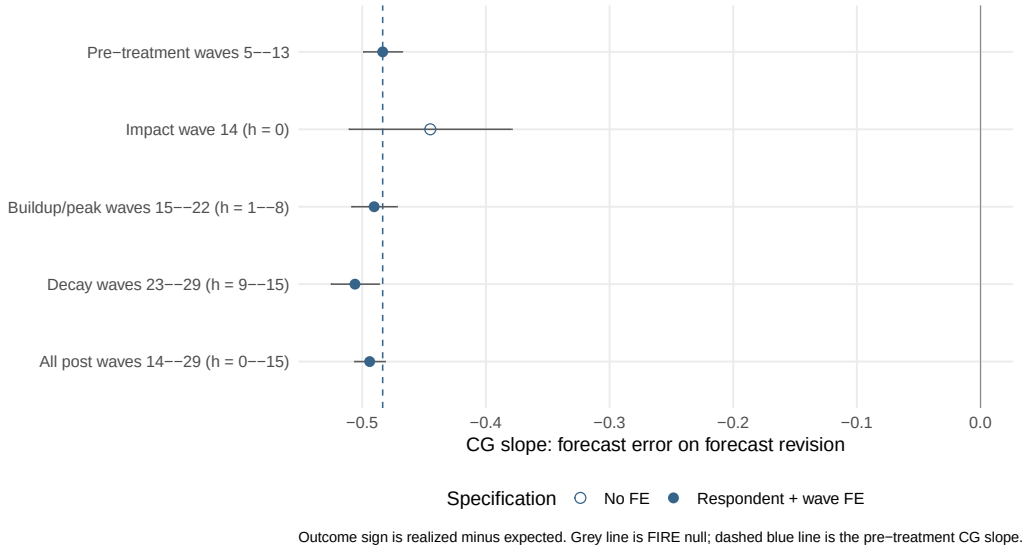


FIGURE G2: Coibion–Gorodnichenko slopes by dynamic window. The impact-wave estimate is cross-sectional and does not include fixed effects; the remaining windows include respondent and wave fixed effects.

Finally, Figure G3 reports the corresponding cross-country phase comparison. Overreaction-like slopes are visible in other CES countries as well, reinforcing the interpretation that this feature is not uniquely Italian. The CG evidence is therefore best read as supporting the belief-formation mechanism, while the Italy-specific part of the paper comes from the credibility channel and the expectation-leg decomposition.

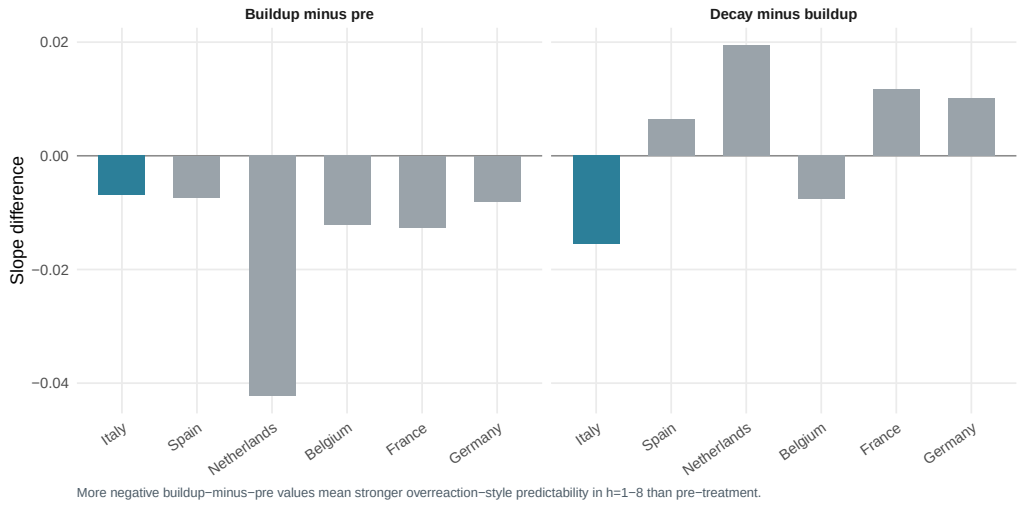


FIGURE G3: Cross-country changes in Coibion–Gorodnichenko slopes across phases. The figure compares how the revision-error relation changes from the pre-treatment period to the buildup/peak window and from the buildup/peak window to the decay window across CES countries.

H Credit Applications and Behavioral Relevance

This appendix reports supporting product-level evidence for the relationship between credit-access forecast errors and subsequent credit-application behavior. The analysis is descriptive and should not be interpreted as formal causal mediation. The main text reports the aggregate association for any credit application using linear probability models and probit average marginal effects. Here, I show how the association varies across specific credit products.

For each post-treatment forecast-error observation, I match the respondent to subsequent quarterly CES credit-application modules. The product-level specification estimates:

$$A^{j}_{i,t,h_q} = \alpha^{j}_{h_q} + \beta^{j}_{h_q} FE^{credit}_{i,t} + X_i'\gamma + \lambda_{\text{wave}(t,h_q)} + \varepsilon^{j}_{i,t,h_q},$$

where A^{j}_{i,t,h_q} is an indicator equal to one if respondent i applies for credit product j in the matched quarterly module h_q. The vector X_i includes predetermined demographic controls, and $\lambda_{\text{wave}(t,h_q)}$ denotes application-wave fixed effects. Standard errors are clustered at the respondent level. Coefficients are reported in percentage points per one-unit increase in the credit-access forecast error.

Figure H1 reports product-level associations. The strongest positive associations appear for consumer loans, especially at the first matched quarterly horizons. Vehicle loans and credit-card or overdraft applications also show positive associations at selected horizons, while mortgage refinancing and education loans do not display the same pattern.

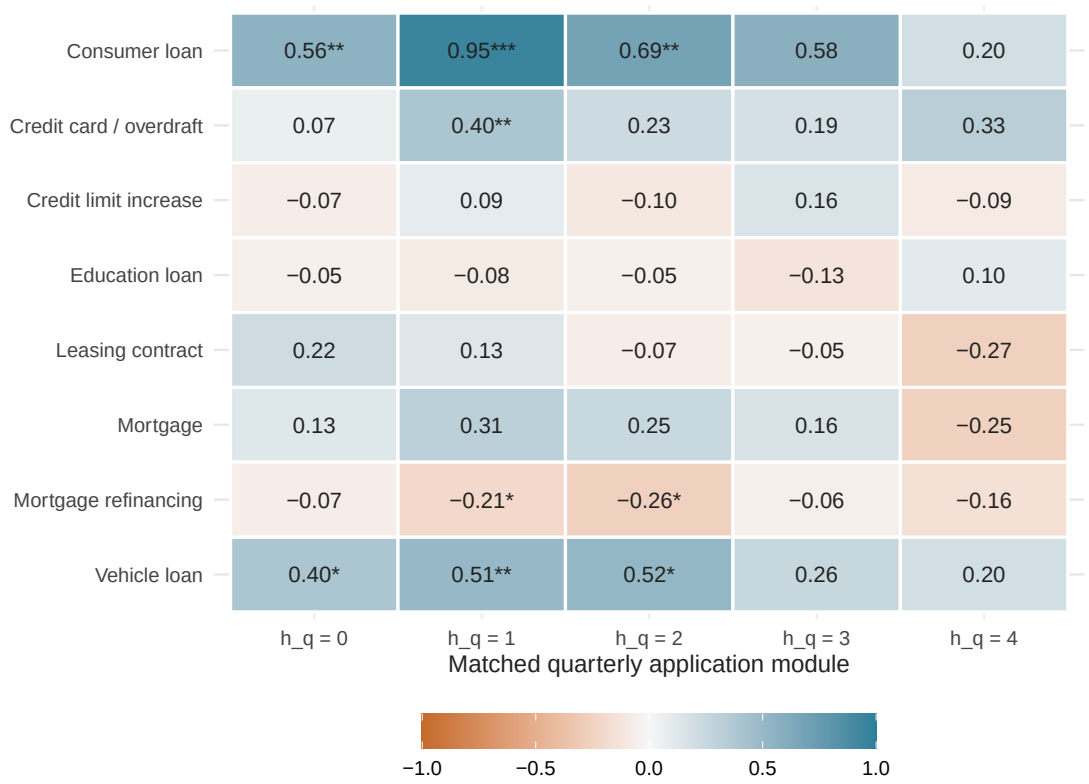


FIGURE H1: Product-level associations between credit-access forecast errors and subsequent credit applications. Cells report controlled linear probability model slopes in percentage points per one-unit increase in the credit-access forecast error. Specifications include demographic controls and application-wave fixed effects; standard errors are clustered at the respondent level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

For comparison, Figure H2 reports the baseline product-by-horizon association using only application-wave fixed effects. The product-level pattern is similar: consumer-loan applications display the most robust positive association with credit-access forecast errors across early matched horizons.

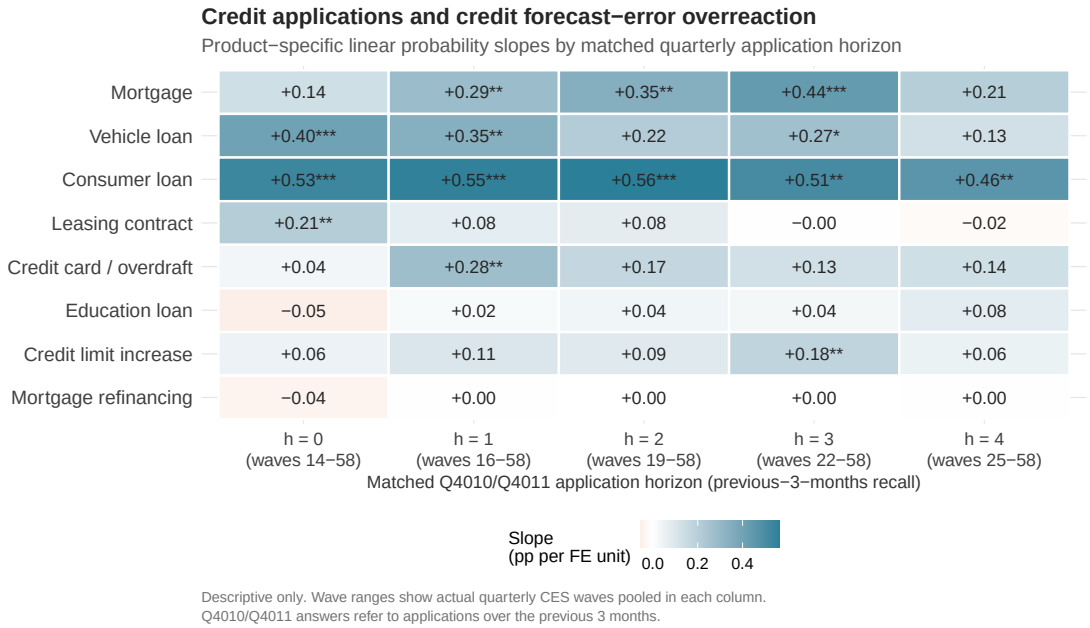


FIGURE H2: Baseline product-by-horizon association between credit-access forecast errors and subsequent credit applications. Cells report linear probability model slopes in percentage points per one-unit increase in the credit-access forecast error, controlling for application-wave fixed effects.