Journal of Monetary Economics 129 (2022) S35–S51

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

News and uncertainty about COVID-19: Survey evidence and short-run economic impact☆

Alexander M. Dietrich[a], Keith Kuester[b], Gernot J. Müller[c], Raphael Schoenle[d,*]

[a] University of Tuebingen, Nauklerstraße 50, 72074 Tuebingen, Germany
[b] University of Bonn and CEPR, Adenauerallee 24–42, 53113 Bonn, Germany
[c] University of Tuebingen, CEPR and CESifo, Nauklerstraße 50, 72074 Tuebingen, Germany
[d] Brandeis University, CEPR and CESifo, 415 South Street, 02454 Waltham, MA, USA

ARTICLE INFO

Article history:
Received 30 November 2021
Revised 5 February 2022
Accepted 7 February 2022
Available online 9 February 2022

JEL classification:
C83
E32
E52

Keywords:
Consumer expectations
Survey
Large shock
Uncertainty
Monetary policy

ABSTRACT

A tailor-made survey documents consumers' perceptions of the US economy's response to a large shock: the advent of the COVID-19 pandemic. The survey ran at a daily frequency between March 2020 and July 2021. Consumer's perceptions regarding output and inflation react rapidly. Uncertainty is pervasive. A business-cycle model calibrated to the consumers' views provides an interpretation. The rise in household uncertainty accounts for two-thirds of the fall in output. Different perceptions about monetary policy can explain why consumers and professional forecasters agree on the recessionary impact, but have sharply divergent views about inflation.

1. Introduction

Expectations are central to economic decision making and so is the expectation formation process. The full information rational expectations hypothesis (FIRE) provides a natural benchmark according to which people adjust their expectations adequately and immediately in the face of new information. Survey evidence, by contrast, suggests that expectations tend to adjust only sluggishly to macroeconomic shocks. This holds not only for professional forecasters but also for policymakers, firms, and households (Coibion and Gorodnichenko, 2012).[1] In normal times, the response of households tends to be even more sluggish than the response of professional forecasts (Carroll, 2003; Carroll et al., 2020). There is, however, evidence

☆ We thank the audiences at several seminars, including the second joint European Central Bank-Federal Reserve Bank of New York conference on expectations surveys, for comments that helped improve an earlier version of the paper. The views stated in this paper are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System. Kuester gratefully acknowledges support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC2126/1 390838866.

* Corresponding author.

E-mail addresses: alexander.dietrich@uni-tuebingen.de (A.M. Dietrich), keith.kuester@uni-bonn.de (K. Kuester), gernot.mueller@uni-tuebingen.de (G.J. Müller), schoenle@brandeis.edu (R. Schoenle).

[1] However, there is also evidence of overreaction to news at the level of individual forecasters (Bordalo et al., 2020; Broer and Kohlhas, 2021). For firms, Born et al. (2021) find overreaction of expectations to firm-level news and underreaction to aggregate news.

https://doi.org/10.1016/j.jmoneco.2022.02.004

that expectations adjust more quickly in times of high uncertainty, in response to large shocks, and as media coverage intensifies (Baker et al., 2020b; Coibion and Gorodnichenko, 2015; Larsen et al., 2021). All of these conditions are met in the context of the COVID-19 pandemic: it offers a natural experiment to study the expectation formation process in some detail.

In order to do so, we exploit a special resource: a daily survey of consumer expectations that we have been running since the start of the pandemic. The survey asks a representative sample of consumers in real time—that is, starting with the onset of the pandemic—how they expect the COVID-19 shock to affect income and inflation over a 12-month horizon. We find that consumer expectations respond very rapidly and that uncertainty about the economic effects of the shock is pervasive—and much more so than what comparable measures for professional forecasters suggest. Our survey is unique in that it directly elicits the shift in consumers' conditional expectations that the COVID-19 shock brings about, in real time. Translated to a model context, the survey elicits consumers' views about the impulse response of the economy to the pandemic. In the second part of the paper, thus, we use the survey responses as "identified moments" (Nakamura and Steinsson, 2018) to which we calibrate a quantitative business-cycle model. We study how the response of consumer expectations shapes the adjustment of the economy to the pandemic. In the model, uncertainty rises because the effective lower bound becomes binding, amplifying fluctuations, and because there are shocks to uncertainty about demand preferences. The rise in consumer uncertainty accounts for two-thirds of the fall in GDP during the pandemic.

The paper starts out by documenting the high-frequency response of consumer expectations to the COVID-19 pandemic. For this purpose we rely on an online survey that we initiated on March 10, 2020. At that time the pandemic had only just started to arrive in the US, giving rise to about 1000 infections in the entire country. The survey ran each day and is representative of the US population according to age, gender, region, income and education. In our analysis we consider data up to July 11, 2021, approximately 60,000 responses in total. Four observations summarize the main results of the survey. Observation 1: consumer expectations respond quickly and strongly to the pandemic. In terms of magnitude, the income loss that consumers expect is consistent with what professional forecasts at that time imply: a 7% fall in output over the course of 12 months. Consumer expectations overshoot on impact. Observation 2: the uncertainty about the output loss reflected in *consumers'* responses is very large. It exceeds by an order of magnitude the uncertainty implied by the disagreement across *professional* forecasters.

Observation 3: respondents expect the pandemic to be inflationary. Over the 12 months that follow the shock, the average response sees inflation rise by about 5 percentage points. This stagflationary view of the recession emerges also when we consider responses for output and inflation at the individual level. Indeed, the majority of respondents expect the pandemic to raise prices while lowering GDP. In this, consumers' inflationary views stand in sharp contrast to the views held by professional forecasters. Observation 4: the pandemic strongly raises consumer uncertainty about future inflation. Once more, the impact on uncertainty is much more pronounced for consumers than for professional forecasters.

The second part of the paper rationalizes these observations based on a FIRE model. Here we make only a first pass at the data, but FIRE seems justified from an *ex ante point of view* in light of the distinct features of the pandemic—the massive increase in uncertainty, the large shock, and the intensive news coverage. The model is a simplified version of the representative-household New Keynesian framework with demand uncertainty shocks developed by Basu and Bundick (2017).[2] Next to such demand uncertainty shocks, business cycles arise in the model because of level shocks to demand preferences and to supply. The latter set of shocks also features news about future supply conditions, since anticipation of future supply disruptions seems essential for pandemics. We solve the model controlling for the effective lower bound on interest rates.

In the model, no single shock in isolation is sufficient to generate the identified moments. Hence, we devise a COVID-19 scenario based on several large shocks. So as to replicate the response of uncertainty in the survey, for example, the COVID-19 scenario sees an increase in the volatility of demand shocks by 17.5 standard deviations. These shocks reduce the natural rate of interest by 15 percentage points (annualized) and make the effective lower bound bind. Next, comparably large adverse news shocks about productivity are required to match households' stagflationary views. The representative-household FIRE approach, thus, provides a nuanced interpretation of the survey facts.

We continue to proceed under the working hypothesis that the model provides a reasonable first pass at the data, for the episode at hand. We use the model in two ways. First, we analyze the contribution of specific shocks. We find that the uncertainty shock (to demand preferences) is the main driver of the expected output loss: without this shock, output falls by a mere 2% (rather than by 7%). The same shock hardly affects inflation, however. Average inflation is driven mostly by the adverse news shock: because supply or, more specifically, total factor productivity is expected to decline in the near future, firms raise prices at the beginning of the pandemic, in anticipation of rising marginal costs.

Second, we quantify the role of the monetary response, consumers' perceptions of which the survey itself does not elicit. The baseline, calibrated to match the *consumers' view on outcomes*, relies on a conventional interest rate feedback rule. In a counterfactual, instead, we let the nominal interest rate track the natural rate of interest, to the extent that the effective lower bound allows. While the shocks are identical under both policies, the outcomes differ. In the counterfactual, the outcomes look much more similar to the *professional forecasters' view of the pandemic*: the pandemic's effect on uncertainty about output and inflation falls by half. In addition, the inflation response switches sign. The reason for this is that tracking

[2] Fernández-Villaverde and Guerrón-Quintana (2020) provide a survey of the related literature.

the natural rate of interest prevents two effects that are inflationary: it avoids the future accommodation built into the baseline policy and the uncertainty about marginal costs (Fernández-Villaverde et al., 2015).

Thus, the same model—with the same size and timing of shocks—can replicate both the consumers' and the professional forecasters' views on the impact of the pandemic, the only difference being the perceived monetary policy response. This suggests an important policy implication: communicating effectively with *the broader public* (and not only professional forecasters) about monetary policy and the state of the economy (as captured by the natural rate) could itself dampen economic uncertainty and the fallout after large unexpected shocks.

There is a different reading of the model-based exercise, of course: namely, that it points toward important gaps in modeling and survey methodology that future work should address. The shocks that are needed to make the model replicate the survey evidence are very large. Depending on one's view, this may adequately reflect the depth of the recession or cast doubt on the representative-household FIRE model. By its very nature this approach also cannot account for the *heterogeneity* of households and of households' expectations, both of which could make activity more exposed to shocks. Another weakness is that we cannot account for both household and firm expectations about inflation at the same time. Modeling such heterogeneity would seem important. At the same time, there is also the issue—given the current state of survey methodology—of whether consumers' self-declared expectations, notably about inflation, may be taken at face value, as we do throughout our analysis.

A number of studies use survey data to study the impact of the pandemic with a focus on inflation expectations of firms or consumers (Armantier et al., 2020; Binder, 2020; Candia et al., 2020; Meyer et al., 2021). Christelis et al. (2020) document a decline in consumption in response to the pandemic, based on survey evidence from Europe. Others have focused on implications of lockdown policies or the stock market's reaction for household expectations (Coibion et al., 2020b; Hanspal et al., 2021; Miescu and Rossi, 2021).[3] Relative to these papers, our survey makes three contributions. First, we identify expectations conditional on an exceptionally large shock. Second, we do so in real time at high frequency. Complementary work by Andre et al. (2021) also studies expectations conditional on shocks, but they consider hypothetical shocks rather than the exceptional event that is the focus of our paper. Relative to other existing surveys, such as the Survey of Consumer Expectations and its analysis in regard to COVID-19, for example, by Armantier et al. (2020), our analysis shows that conditional and unconditional expectations can differ substantially. Third, we use the survey responses as identified moments in order to calibrate a business-cycle model. This, in turn, allows us to analyze the role of expectations in transmitting the shock. Our analysis also relates to work by Bloom (2009) and many others who have stressed the role of uncertainty as a potential source and amplification channel of the business cycle, a view recently supported by direct survey evidence (Coibion et al., 2021).

The remainder of the paper is structured as follows. We introduce our survey in the next section and present the main results of the survey in Section 3. Section 4 introduces our business-cycle model, which allows us to develop a structural scenario for the expected impact of the COVID-19 shock. A final section concludes.

2. Survey design

The survey that we run is unique in two ways. First, our survey systematically introduces questions that elicit *conditional* expectations on prices, quantities and behavioral variables. Namely, we ask respondents to assess the impact of COVID-19 on their outlook for the economy. In doing so, our work presents the empirical counterpart of the hypothetical "vignettes" in Andre et al. (2021). These conditional expectations correspond closely to how shocks move expectations in the context of models. As such, the questions allow for a tighter identification of the impact of specific shocks than would eliciting conventional unconditional expectations. We find that conditional and unconditional expectations differ for GDP and personal household income but are similar for inflation. At the same time, disagreement is relatively similar across all variables (see Supplementary Material C.8).

Second, the high-frequency approach is a distinct feature of our survey. It is rooted in a daily sample of respondents, an approach that presents a large option value for policymaking in practice and real time. However, we do not exploit the high-frequency feature further.

2.1. Survey description and demographics

We contracted with Qualtrics Research Services to provide us with a survey of 60,003 nationally representative respondents for 16 months between March 10, 2020, and July 11, 2021. The survey was run with a daily sampling size of at least 100 respondents. Over the course of one month the number of survey responses (above 3000) compares favorably to that of existing consumer surveys. Balancing a more granular view on the expectations process with a larger, less noisy sample size, we mainly report 11-day moving averages below. The survey required all respondents to be US residents and speak English as their primary language. Other than this, our sample was taken to be representative of the US population.

In terms of demographics, respondents had to be male or female with 50% probability. Moreover, approximately one-third of respondents were targeted to be between 18 and 34 years of age, another third between ages 35 and 55, and a final

[3] Fetzer et al. (2021) assess the determinants of economic anxiety at the onset of the pandemic, based on survey evidence from a large set of countries.

Table 1
Survey Respondent Characteristics. *Notes:* The column "Survey" represents characteristics in our survey. The column "US population" reports the value for the US population, as obtained from the US Census Bureau (Household income: CPS ASEC, 2021; gender, education: ACS, 2019, age, race, region: National Population Estimate, 2019).

	Survey	US population		Survey	US population
Age			**Race**		
18–34	33.11%	29.8%	non-Hispanic white	72.75%	60.1%
35–55	33.82%	32.4%	non-Hispanic black	9.29%	12.5%
>55	33.07%	37.8%	Hispanic	10.08%	18.5%
			Asian or other	7.88%	8.9%
Gender					
female	49.92%	50.8%	**Household Income**		
male	49.69%	49.2%	less than 50k$	46.23%	37.8%
other	0.39%	-%	50k$ - 100k$	41.53%	28.6%
			more than 100k$	23.02%	33.6%
Region					
Midwest	20.64%	20.7%	**Education**		
Northeast	21.86%	17.3%	some college or less	50.62%	58.3%
South	39.54%	38.3%	bachelor's degree or more	49.38%	41.7%
West	17.96%	23.7%			
			N=60,003		

third older than age 55. We also required a distribution across US regions in proportion to population size, drawing 20% of our sample from the Midwest, 20% from the Northeast, 40% from the South and 20% from the West.

The survey includes filters to eliminate respondents who enter gibberish for at least one response or who complete the survey in less (more) than five (30) minutes. We also employ CAPTCHA tests to reduce the possibility that bots participate in the survey. Table 1 provides a breakdown of respondent characteristics and sampling targets.

Respondents match the US population demographics along key dimensions. To improve the fit further, we additionally compute a survey weight for each respondent. To do so, we apply iterative proportional fitting to create respondent weights after completion of the survey ("raking," see, for example, Bishop et al. (1975) or Idel (2016)). This allows us to calculate statistics that are *exactly* representative of the US population also according to ethnicity, income and education, that is, the variables in the right column of Table 1.

2.2. Survey questions

To elicit expectations for our variables of interest, we build on the Survey of Consumer Expectations (SCE) pioneered by the Federal Reserve Bank of New York. Whereas the SCE asks for an unconditional forecast, we directly elicit consumers' assessments of the "impact of the coronavirus" or changes in economic aggregates "because of the coronavirus." Otherwise, we stick to the wording of the SCE as closely as possible. While we keep the way of measuring inflation the same as in the SCE, we elicit responses for two different measures of income. On the one hand, we follow the SCE by asking for the "total income of all members of your household (including you)." On the other hand, we are interested in GDP as a measure of income, motivated by modeling purposes. Leading surveys like the University of Michigan's Survey of Consumers and the Federal Reserve Bank of New York's Survey of Consumer Expectations do not include questions on GDP. We elicit expectations for GDP, household income, and inflation at the 12-month horizon relative to today. Supplementary Material A summarizes the questions we ask. In the results section below, we discuss some possible limitations of asking consumers about the abstract concept of GDP and we show how well the consumer responses regarding GDP and personal household income forecast actual GDP realizations.

Following again the approach in the SCE, we first elicit point estimates and afterwards the probability that respondents assign to a particular outcome given a range of outcomes.

3. How consumer expectations responded to the pandemic: Survey evidence

This section presents the results of our survey, documenting the response of consumer expectations to the COVID-19 pandemic. We show how the COVID-19 shock moves the first and second moments of income and inflation expectations, and how the variables comove.

We summarize the survey results with four key observations. We state each of them and provide evidence to back up each observation. Regarding income expectations we make

Observation 1 (Income Expectations)**.** *Consumers adjust income expectations downward in response to the pandemic. The adjustment is stronger and faster than that of professional forecasters.*

Fig. 1 displays consumers' perceptions about how the pandemic would affect income over the next 12 months. Panel (a) refers to the impact on GDP. It shows four lines: expected impact of the pandemic in terms of GDP over the next 12 months, averaged across respondents daily (jagged black solid line), the 11-day moving average thereof (red dotted line), and a measure of the impact of the pandemic as viewed by the professional forecasters that contribute to the Blue Chip forecasts (blue dashed line).[4] In addition, panel (a) also reports how much *actual* GDP over the next 12 months has deviated from the pre-pandemic trend (triangles). Panel (b) refers to personal household income.

What is most striking, perhaps, is the speed with which consumer expectations react: by late March/early April 2020 the average expected GDP impact (across households) is close to -15%. The maximum effect in terms of the moving average is -18% and observed on April 01, 2020. For personal household income, we observe the maximum drop in conditional income expectations on March 24; it is -13%.

Importantly, we observe a strong reaction of expectations even though at this point in time the pandemic had barely arrived in the US. This fact is illustrated in panel (c) of Fig. 1. At the time when the expected impact on GDP is largest, the case rate is 47 (number of new infections during the last 7 days/100K people). The maximum case rate in our sample of approximately 500 is reached only much later, in December 2020/January 2021. Similarly, panel (d) shows the unemployment rate for our sample period. It, too, peaks much later than the response of household expectations.

What is striking, too, is that consumers' expectations initially react by more than professional forecasters'.[5] This is noteworthy because it seems to run counter to the received idea of sluggish responses of household expectations and a one-way information flow from professional forecasts to households (Carroll, 2003). Instead, household expectations and the Blue Chip forecasts converge to the middle ground: by May/June 2020 they are very well aligned at approximately a -7% impact on GDP, and remain surprisingly aligned all the way until the end of our sample period.

Our measures of income—GDP and household income—require different levels of abstraction of households. By national income accounting they are closely linked, GDP being equal to the survey's definition of personal household income plus taxes, deductions, depreciation and net foreign factor income. At 0.73, the time-series correlation in our survey of the two measures is high. *Ex post* the responses for expected household income growth turn out to be close to realized *GDP* changes (Figure 1, panel b). Direct GDP expectations do not align as closely (Figure 1, panel a). This is suggestive of a difference in forecasting ability that, below, leads us to rely on personal household income expectations for our model-based calibration targets for the impact on output.

The COVID-19 shock also triggered a massive increase in uncertainty. This increase has been documented based on a variety of indicators, including expectations of firms' sales growth as measured in business expectations surveys (Altig et al., 2020; Armantier et al., 2020; Baker et al., 2020a). Such an increase in uncertainty also shows up on the consumer side.[6] In particular, regarding our survey we make

Observation 2 (Income Uncertainty)**.** *Consumers' uncertainty about the impact of the pandemic in terms of GDP rises fast, faster and by much more than professional forecasters'.*

Specifically, Fig. 2 shows two measures of uncertainty regarding the impact of the pandemic in terms of GDP: disagreement (panel a) and subjective uncertainty (panel b). Disagreement is a widely used measure of uncertainty (Bloom, 2014). We measure it in panel (a) by the standard deviation of responses across consumers, on a daily basis (jagged solid black line) as well as the 11-day moving average (dotted red line). The relevant scale is the left axis. Against a different scale on the right, we show the corresponding measure based on the Blue Chip forecast (dashed blue line). Two patterns are particularly noteworthy. First, consumer disagreement in our data leads disagreement of professional forecasters, suggestive of a real-time information content of the daily consumer survey. Second, consumer disagreement rises by an order of magnitude more than disagreement of professional forecasters (recall that in panel (a) we measure disagreement against the left axis for our survey and against the right axis for the Blue Chip survey). This finding is, perhaps, not entirely unexpected, given that our survey's respondents, consumers, are considerably more heterogeneous than the respondents in the Blue Chip survey. Consumer uncertainty about the economic effects on income is high from the start of the survey/pandemic.

Later, in Section 4 we will use a representative-household model to take a first pass at the role that consumer uncertainty had in shaping the recession. By its nature, this model does not feature disagreement. Therefore, to calibrate the model, we rely on a second measure of household uncertainty about income: uncertainty about the expected impact of COVID-19 at the level of individual responses. We have this measure only for GDP. It relies on consumers' responses to a question that asks respondents to assign probabilities to specific outcomes. We then fit a beta distribution individually to the responses of each respondent and compute the standard deviation of the distribution, following the SCE methodology (Armantier et al., 2017). Panel (b) of Fig. 2 displays this measure of subjective uncertainty, averaged across respondents on a daily basis (jagged solid black line), and an 11-day moving average (dotted red line). For this measure, too, the rise in uncertainty is rapid. Relative to the later levels the initial increase in subjective uncertainty (measured by the standard deviation) is approximately 4 percentage points.

[4] The figure plots the Blue Chip GDP forecast net of a pre-pandemic trend (see Supplementary Material B for details on the computation).

[5] This may have been an overreaction, suggestive of non-rational expectations (Bordalo et al., 2020; Broer and Kohlhas, 2021), a possibility we do not pursue further in this paper.

[6] Our analysis is complementary to the work by Andre et al. (2021), who use experimental "vignettes" to also get at conditional expectations. In the context of uncertainty, our results are in line with their findings that there is more disagreement among households than among experts. However, our work differs in that we ask about the effect of an actual shock rather than a hypothetical one.

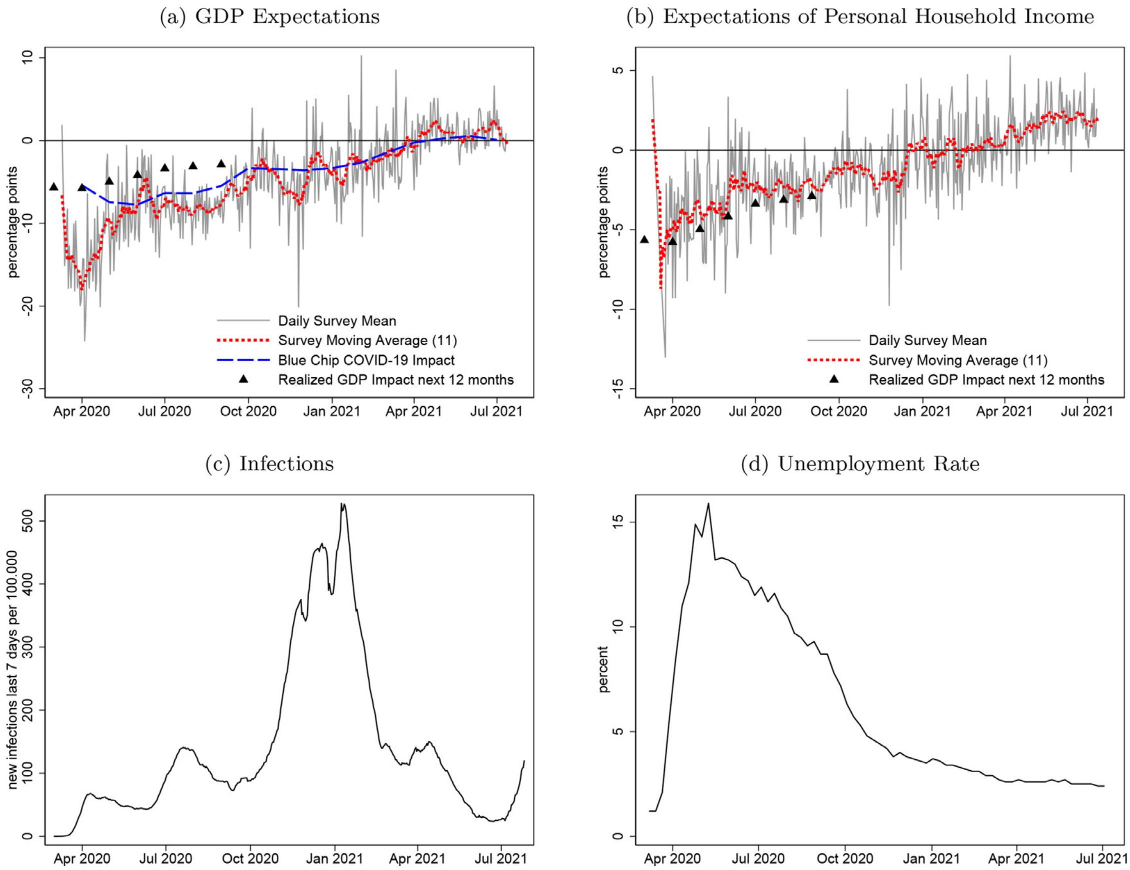

Fig. 1. Response of Income Expectations to COVID-19 Shock. *Notes:* The jagged black solid lines in panels (a) and (b) show the daily mean of survey responses (weighted using survey weights and Huber-robust weights), red dotted lines are an 11-day moving average. The blue dashed line in panel (a) represents Blue Chip forecasts: the average deviation of GDP from a pre-pandemic trend over the next 12 months. Blue Chip forecasts are a resource of Wolters Kluwer Legal and Regulatory Solutions US. Black triangles correspond to the realized GDP deviation from the pre-pandemic trend over the next 12 months; see Supplementary Material B for details; panel (c) shows new COVID-19 infections within the last 7 days per 100K people; panel (d): weekly unemployment claims as percent of workforce. For data sources see Supplementary Material B.3. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The Supplementary Material revisits the conceptual concern of how well households understand the concept of GDP (Figure C.9). It compares the disagreement in GDP expectations to the disagreement in expectations about household income. Both measures of income uncertainty show a massive increase amid substantial comovement. The increase in GDP uncertainty is larger, however. This difference—plus the larger increase in disagreement for consumers relative to professional forecasters—may potentially indicate some difficulty among consumers in forming GDP expectations.

In addition to income expectations, our survey asks respondents about the likely impact of the pandemic on inflation. We summarize the results with

Observation 3 (Inflation Expectations)**.** *On average, consumers expect inflation to rise strongly in response to the COVID-19 shock, in contrast to professional forecasters, who expect a deflationary effect. Moreover, most consumers expect an inflationary impact, independently of whether they expect economic activity to contract or rise in response to the pandemic.*

Fig. 3 shows the expected effect of the pandemic on inflation. Panel (a) is organized in the same way as panel (a) of Fig. 1. Panel (b) shows the break-even inflation rate for our sample, as a measure of expectations of financial market participants. We show the uncertainty measures for inflation in panels (c) and (d), analogous to Fig. 2. We observe that consumers see the pandemic as a cause of inflationary pressure; see panel (a): by March/April 2020, households on average expect an inflationary impact of COVID-19 on an order of magnitude of 7–8%. The number declines somewhat over the summer of 2020, but the expected impact on inflation remains high throughout the year. This result stands in sharp contrast to the Blue Chip survey and financial market-based expectations. Here the expected impact of the COVID-19 pandemic on inflation

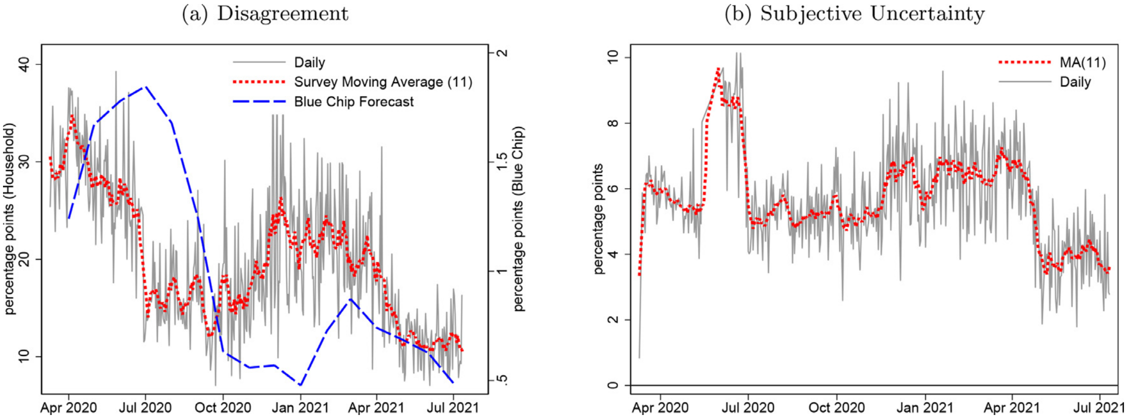

Fig. 2. Uncertainty about the GDP Impact of COVID-19. *Notes:* The left panel shows the uncertainty about the economic impact of COVID-19 on GDP, for consumers (jagged black solid and dotted red line) and Blue Chip forecasters (blue dashed line), measured by the standard deviation across responses in percentage points ("disagreement"). Panel (b) shows average subjective uncertainty across respondents, computed as the standard deviation of a beta distribution fitted on the probability distribution solicited in the consumer survey at the level of individual responses. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
Sign of COVID-19's Impact on Expectations. *Notes:* The table displays the reported sign of the expected GDP and inflation impact of COVID-19 over the next 12 months by respondents (N=60,003). "0" means that participants indicated a positive/negative impact in the first part of the question, but subsequently filled in a numerical value of 0.

		Inflation			
		negative	0	positive	
GDP	negative	16.18%	0.46%	44.16%	60.81%
	0	0.19%	0.71%	0.42%	1.32%
	positive	9.20%	0.28%	28.39%	37.87%
		25.57%	1.46%	72.97%	

is negative early in the sample.[7] Note that our evidence on inflation expectations also differs from Armantier et al. (2020). Based on the SCE they find that consumer inflation expectations move little.

An advantage of our survey is that it allows us to elicit the *joint* distribution of respondents' views about the effect of the pandemic on inflation and output. Table 2 cross-tabulates individual responses for inflation and output. Three-quarters of respondents see a positive impact of the pandemic on inflation. This is so regardless of what effect the respondents expect the pandemic to have on output. Of the households that see a recessionary impact of COVID-19 (60.81% of respondents), three-quarters, too, anticipate an inflationary impact. Overall, the dominant view is that the effect of the pandemic is stagflationary (44.16% of respondents). Still, whether households see COVID-19 as recessionary does not seem to matter for their anticipation of the inflationary impact. This pattern is consistent with the notion that a pandemic is an adverse event and a widespread view that "inflation is bad for the economy," no matter what (Candia et al., 2020). In line with this conclusion, and the framing of our survey as being COVID-related, indeed, the survey's unconditional inflation expectations are similar to the conditional expectations; see Figure C.8 in the Supplementary Material.

The question naturally arises: what does the survey convey about consumers' perceptions about the economic mechanisms at work during the pandemic? On the one hand, consumers may see an overall demand shock, consistent with several model-based accounts of the economic impact of the pandemic (for example, Bayer et al., 2020; Fornaro and Wolf, 2020; Guerrieri et al., 2020). On the other hand, the data on conditional inflation expectations suggest that consumers might interpret the conditional effect of COVID-19 as having a supply component, too. There is a large ongoing research agenda that studies the deeper mechanisms of inflation expectations formation. It is relevant for the interpretation of our data. Behavioral interpretations may also support a conditional supply-side view: salience of shopping experiences drives inflation expectations as documented in Bryan and Venkatu (2001) and D'Acunto et al. (2021). Acute, salient shortages of certain goods may drive supply-shock perceptions on the side of consumers. These inflation expectations may also simply reflect personal consumption bundles (Cavallo et al., 2017) and, in particular, high-inflation items in these bundles rather than goods experiencing deflation during the evolution of COVID-19 (Cavallo, 2020). Dietrich et al. (2022) show that several salient product categories seem to drive aggregate inflation expectations. Carroll (2003) and Larsen et al. (2021) show how media coverage may affect consumers' inflation expectations, while, for example, Candia et al. (2020) and Coibion et al. (2022) study the role of policy communication. Importantly, our model-based analysis in Section 4 will point to yet another

[7] There is survey evidence suggesting that firms, too, expected a negative impact of the pandemic on inflation (Balleer et al., 2020; Meyer et al., 2021).

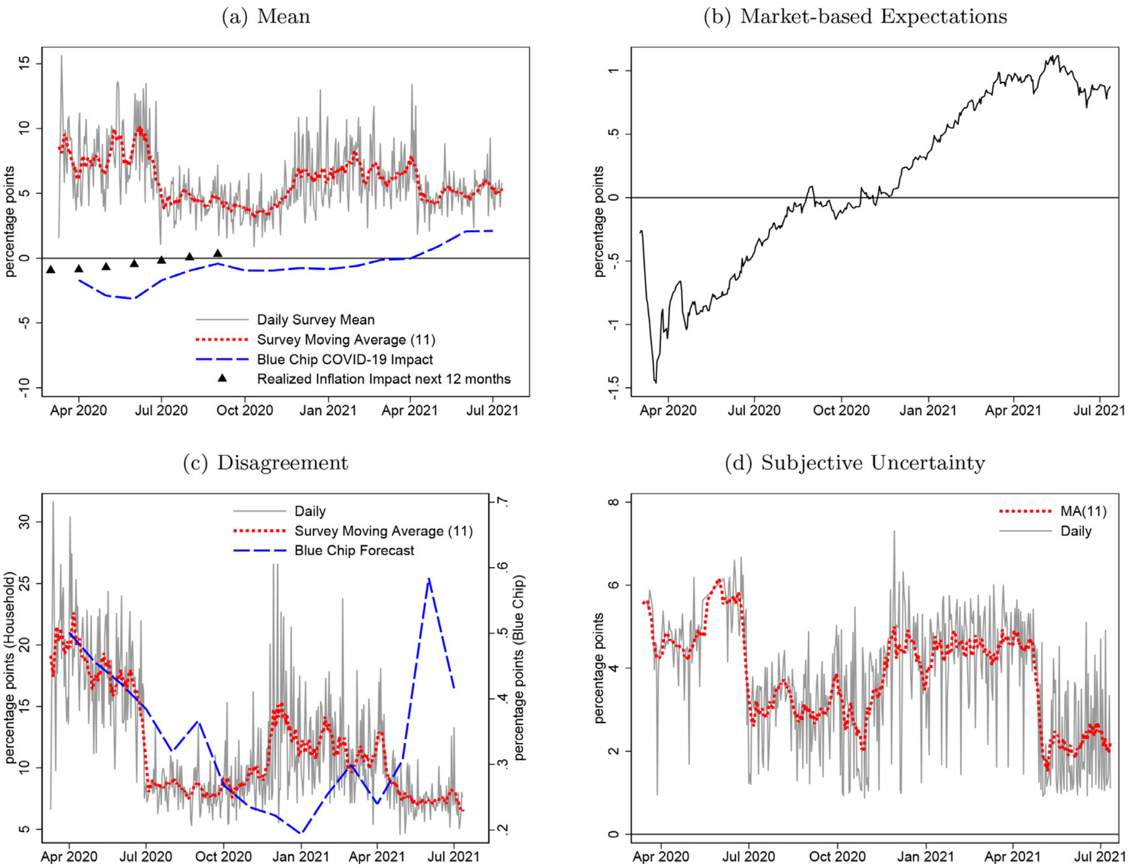

Fig. 3. Expected Impact of COVID-19 on Inflation. *Notes:* Expected impact of COVID-19 on average inflation for the next 12 months. Panel (a): means. Panel (b): 5 year break-even inflation expectations, net of mean of 2019. See Supplementary Material B.3 for data sources. Panel (c): disagreement. Panel (d): subjective uncertainty of consumers. See Figs. 1 and 2 for further notes.

explanation for the stagflationary views of consumers and their disagreement with professional forecasters: household views of the Fed's response to the pandemic may differ from the view of professional forecasters.

Last, we consider the extent of uncertainty about the inflationary impact of the pandemic and make the following

Observation 4 (Inflation Uncertainty)**.** *Consumer uncertainty about the impact of the pandemic on inflation is large and much larger than the uncertainty of professional forecasters.*

To see this, consider the bottom panels of Fig. 3, which are in line with Armantier et al. (2020). The pandemic's impact on consumers' uncertainty about inflation is similar to the patterns of uncertainty reflected in GDP expectations.[8] This finding holds whether we consider a measure of disagreement (panel c in the figure, measured against the left axis) or subjective uncertainty (panel d). Inflation uncertainty starts out high at the beginning of the sample. Then it drops in the summer of 2020, only to rise again at the end of 2020 going into April 2021, and then calms down. Uncertainty computed both from survey responses and from Blue Chip inflation forecasts (blue dashed line in panel a) shows similar patterns, though consumers' uncertainty once more is larger by an order of magnitude. As with the uncertainty about income losses, consumers' uncertainty regarding the inflationary impact of the pandemic is also pronounced if we turn to the measure of subjective uncertainty.[9]

[8] Our survey documents that households disagree about the impact of the COVID-19 shock on inflation. That households hold heterogeneous inflation expectations more generally is well documented, for example, by Mankiw et al. (2004).

[9] The survey also includes questions on savings and purchasing behavior and plans in response to COVID-19, the expected duration of the pandemic, and whether respondents have hoarded food and medical supplies. Economic expectations elicited within our survey vary in a meaningful way with the behavioral adjustments and financial decisions of survey participants. We also document demographic and socio-economic heterogeneity in expectations. Supplementary Material C provides these findings.

4. The expected impact of the COVID-19 shock: A Structural perspective

The survey responses paint a particular picture of consumers' expected short-run macroeconomic impact of the COVID-19 pandemic. What remains to be understood are the potential mechanisms that are behind the survey responses and their implications. Toward this end, we now put forward a business-cycle model for which we devise a specific COVID-19 scenario.

The model assumes rational expectations and full information (FIRE) even though evidence suggests that this assumption is generally too restrictive (Coibion and Gorodnichenko, 2012; 2015). Against this background, we understand our modelling exercise as a first pass in accounting for the response of consumer expectations to the COVID-19 shock. This seems reasonable because of three distinct features of this very special episode. First, the increase in uncertainty was massive and Coibion and Gorodnichenko (2015) find that information rigidities decline precisely in times of increased macroeconomic uncertainty. Specifically, their estimate of information rigidities for the volatile 1970s and early 1980s suggests that the distance to FIRE was smaller. Second, the media focus on COVID-19 was also exceptional. This matters since Larsen et al. (2021) show that information rigidities decline *for consumers* when the news coverage on specific topics intensifies. Third, the shock was exceptionally large, thus arguably capturing people's mindset.

That said, we stress two limitations of our analysis upfront. First, our results are specific to the episode under consideration and do not necessarily apply in all other contexts. Second, we assume that all agents in the model have the same expectations. Specifically, we model firms' expectations as being the same as consumers' expectations, even though our evidence concerns consumer expectations only. In light of these limitations, the following analysis offers an exploration of what it takes for the representative-agent FIRE model to be able to account for our survey evidence.

4.1. The model

Consider the following infinite-horizon model, where time t is discrete and runs forever. Expectations are rational and information is complete. The model is a slight modification of the framework in Basu and Bundick (2017), BB for short. It is a fairly conventional New Keynesian business-cycle model augmented by shocks to the level and volatility of demand and by news shocks to productivity. Relative to BB, we abstract from investment dynamics for clarity. In what follows we provide a compact description of the model that follows the exposition in BB (and their notation) closely.

There is a representative household that has Epstein-Zin preferences over current and future consumption, C_t, and hours worked, N_t. The household faces competitive labor, goods, and financial markets. Let $\sigma > 0$ mark the household's risk aversion and $\psi > 0$ its intertemporal elasticity of substitution and let $\theta_V := (1-\sigma)(1-1/\psi)^{-1}$. Letting V_t mark the value of the household's lifetime utility in period t,[10] the household's problem is given by

$$V_t = \max\left[a_t\left(C_t^{\eta}(1-N_t)^{1-\eta}\right)^{(1-\sigma)/\theta_V} + \beta\left(\mathbb{E}_t V_{t+1}^{1-\sigma}\right)^{1/\theta_V}\right]^{\theta_V/(1-\sigma)}$$

$$\text{s.t. } C_t + \frac{P_t^E}{P_t}S_{t+1} + \frac{1}{R_t^R}B_{t+1} = \frac{W_t}{P_t}N_t + \left(\frac{D_t^E + P_t^E}{P_t}\right)S_t + B_t. \tag{1}$$

$\mathbb{E}_t$ is the expectation operator, a_t is a preference shifter ("demand shock") and $\eta \in (0,1)$. The household purchases consumption at nominal price P_t per unit. In addition, the household can buy infinitely lived shares S_{t+1} at price P_t^E or a real one-period pure discount bond bearing real gross interest R_t^R. The household funds these expenditures through labor income (with W_t marking the nominal wage rate) and past savings. S_t are share holdings going into the period, B_t are bond holdings, and D_t^E are the dividends that shares pay at the beginning of the period.

The final good, Y_t, is a conventional Dixit-Stiglitz aggregator that consists of a bundle of intermediate goods $Y_t(i)$ with $i \in [0,1]$. Intermediate goods producers thus operate under monopolistic competition and solve

$$\max \mathbb{E}_t \sum_{s=0}^{\infty} M_{t,t+s}\frac{D_{t+s}(i)}{P_{t+s}}, \tag{2}$$

subject to the production function

$$\left[\frac{P_t(i)}{P_t}\right]^{-\theta_\mu} Y_t = K^{\alpha}[Z_t N_t(i)]^{1-\alpha} - \Phi,$$

where

$$\frac{D_t(i)}{P_t} = \left[\frac{P_t(i)}{P_t}\right]^{1-\theta_\mu} Y_t - \frac{W_t}{P_t}N_t(i) - \frac{\phi_P}{2}\left[\frac{P_t(i)}{\overline{\Pi}P_{t-1}} - 1\right]^2 Y_t.$$

Above, $M_{t,t+s}$ is the stochastic discount factor arising from the household problem. It prices in period t claims in $t+s$. $\theta_\mu > 1$ is the elasticity of demand, $\alpha \in [0,1)$, Φ is a fixed cost measured in terms of goods used in the production process, and

[10] Here and in the following, to preserve on notation, we use t interchangeably as an indicator of time, a summary measure of the information set in period t, or to mark states in period t.

$\phi_p > 0$ indexes price adjustment costs. Capital is in fixed supply and does not depreciate. A bar on top of a variable marks the variable's steady-state value. So, for example, the presence of $\overline{\Pi}$ above reflects that prices are indexed to steady-state inflation, where inflation is given by $\Pi_t := P_t/P_{t-1}$.

As in BB, each intermediate goods firm is assumed to issue real bonds in proportion to the capital stock, $B_t(i) = \nu K$, with $\nu \in [0, 1)$. Total cash flows of the firm are divided between dividends to equity holders and interest paid to bond holders, so that dividends are given by $\frac{D_t^E(i)}{P_t} = \frac{D_t(i)}{P_t} - \nu K\left(1 - \frac{1}{R_t^R}\right)$. The financing structure of the firm is without consequence; it only serves to introduce the return to equity as an observable variable.

The monetary policy instrument is the gross nominal interest rate R_t on one-period risk-free nominal bonds that are in zero net supply. Let R_t^{tar} denote the target interest rate and $\underline{R}$ the effective lower bound on gross interest rates, such that $R_t = \max[R_t^{\text{tar}}, \underline{R}]$. For the target interest rate itself we assume a conventional Taylor rule:

$$\log\left(R_t^{\text{tar}}/\overline{R}\right) = \left[\rho_\Pi \cdot \log(\Pi_t/\overline{\Pi}) + \rho_y \cdot \log(Y_t/Y_t^n)\right], \tag{3}$$

where $\rho_\Pi > 1$ and $\rho_y \geq 0$ determine the responses to inflation and the output gap, respectively. The output gap, Y_t^n, is defined as the gap between actual output and its natural level. The equilibrium condition for nominal bonds is given via the conventional consumption Euler equation

$$1 = \mathbb{E}_t\{M_{t,t+1}R_t/\Pi_{t+1}\}.$$

Let $\epsilon_t^{(\cdot)}$'s mark iid, zero-mean, standard-normal innovations. Following BB, there are both first-moment shocks to demand and "uncertainty" shocks to demand, namely:

$$a_t = (1-\rho_a) + \rho_a a_{t-1} + \sigma_{t-1}^a \epsilon_t^a,$$

$$\sigma_t^a = (1-\rho_{\sigma^a})\sigma^a + \rho_{\sigma^a}\sigma_{t-1}^a + \sigma^{\sigma^a}\epsilon_t^{\sigma^a}.$$

Productivity is a convolute of two components. The first is a front-loaded productivity component as in BB. The second component allows a gradual build-up of productivity (the news shock). Namely, productivity follows

$$\log(Z_t) = \log(A_t) + \log(X_t),$$

with

$$\log\left(A_t/\overline{Z}\right) = \rho_A \log\left(A_{t-1}/\overline{Z}\right) + \sigma^A\epsilon_t^A,$$

and

$$\log(X_t) = \rho_{X,1}\log(X_{t-1}) + \rho_{X,2}\log(X_{t-2}) + \sigma^X\epsilon_t^X.$$

We include a news component X_t about future productivity since we consider the news important for tracing some of the features of the COVID-19 crisis, the anticipation inherent in the survey, in particular. In each case, the shock processes' parameters are restricted such that all the shocks are stationary.

In equilibrium, all intermediate goods firms choose the same price. Hence, they all have the same level of production, the same demand for inputs, and the same financing structure. Goods market clearing implies $Y_t(i) = Y_t$ and

$$Y_t = C_t + \phi_p/2\left[\Pi_t/\overline{\Pi} - 1\right]^2 Y_t.$$

Labor-market clearing implies $N_t(i) = N_t$. Next, the bond and equity markets clear, so that $D_t^E(i) = D_t^E$ and $B_t(i) = B_t$.

4.2. Calibration

We calibrate the model to perform a quantitative analysis at two levels. First, we set parameters to make sure the model performs reasonably well in capturing regular business-cycle dynamics in "normal times." In the next section, instead, we will devise a specific shock scenario to target the identified moments in the survey, that is, the response of consumer expectations to the pandemic.

Table 3 reports the values we assign to all parameters of the model. Here we generally follow BB. In fact, most of the parameters come directly from their paper. Here we discuss only those parameters that do not. We fix the capital stock K at a value of 2.5 times annual GDP (steady-state GDP itself being fixed at unity). We calibrate the fixed costs of production Φ such that the dividend/GDP ratio is 1%, in line with the calibration in BB. As to price rigidities, we choose a value of $\phi_p = 400$. This value delivers a slope of the Phillips curve that is commensurate with a Calvo rigidity of about 0.867, thus, bringing the calibration into line with conventional estimates of the slope of the Phillips curve (e.g., Gali and Gertler, 1999).

The parameters we choose for the shocks are as in BB. There is only one exception: the parameters pertaining to the news-type shock to productivity (X_t), a shock that BB do not have. We choose parameters $\rho_{X,1}$ and $\rho_{X,2}$ such that a negative news shock means that productivity falls on impact, that half a year later productivity has fallen by a further 75% of the

Table 3
Parameters – Calibration. *Notes:* Parameters for the baseline calibration, see the main text for details.

param.	value	source/target	param.	value	source/target
Preferences			*Monetary policy*		
β	0.994	Basu and Bundick (2017) (BB).	ρ_Π	1.5	conventional value, as in BB.
η	0.326	Frisch elasticity of 2, BB.	ρ_y	0.5/4	conventional value.
ψ	0.95	BB.	$\overline{\Pi}$	1.0057	mean inflation rate 2% p.a.
σ	80	BB.	*Shocks*		
Production			ρ_a	0.935	BB.
α	1/3	BB.	σ^a	0.0026	BB.
K	10	capital stock 2.5 times annual GDP.	ρ_{σ^a}	0.742	BB.
θ_μ	6	BB.	σ^{σ^a}	0.0025	BB.
Φ	0.584	dividend/GDP ratio of 1%, BB.	$\overline{Z}$	2.206	Targets $\overline{Y} = 1$.
ν	0.85	BB.	ρ_A	0.987	BB.
ϕ_p	400	slope of Phillips curve, see text.	σ^A	0.0013	BB.
			$\rho_{X,1}$	1.5	judgmental, see text.
			$\rho_{X,2}$	-0.6	judgmental, see text.
			σ^X	0.001	judgmental, see text.

initial impact, and that, thereafter, it rapidly normalizes, implying that three years after an innovation to the news component, the effect of the shock has essentially vanished. In addition, we choose a standard deviation for the innovation, σ^X. The value we choose implies that in normal times the news shock only has little effect on economic activity, in line with results by Schmitt-Grohé and Uribe (2012) for the news component in neutral productivity shocks.

Table D.1 in Supplementary Material D shows one implication of the calibration, namely, second moments. The table compares the unconditional second moments of the model to the US business cycle, as measured by HP-filtered data. The model-based moments are computed without any adjustments for a potential zero lower bound, using third-order perturbation. In line with this, we also report the data counterparts only for the period before the lower bound became binding (1984Q1 to 2008Q3). Overall, the model appears to paint a reasonable picture of the standard business cycle.

In all the simulations that follow, we allow the conditional mean dynamics of the nominal interest rate to fall at most 1.5 percentage points (annualized) below the steady state. We do so to mimic the room for interest-rate cuts that the Fed had going into the recession (from November 2019 through February 2020, the effective federal funds rate stood at roughly 1.5 percentage points). Accounting for the lower bound in the context of uncertainty raises some challenges in terms of computation. Supplementary Material D.2 explains how we solve the model numerically.

4.3. Mapping the survey responses into the model

To map the survey responses into the model, we devise a COVID-19 scenario based on a range of shocks. In the survey we do not ask respondents about specific macro shocks as, for instance, Andre et al. (2021) do. Rather, we ask respondents about the impact of the pandemic. Here we therefore specify a combination of shocks that is meant to rationalize the conditional expectations we obtain from the survey; that is, we target moments identified by the survey. The set of shocks includes a demand uncertainty shock, TFP shocks, and a level shock to demand preferences. The aim of the *scenario* we develop is to replicate the main patterns of the survey responses, taking a representative-agent perspective.

A key feature that emerges is that the required shocks are large, reflecting the extent of the effects manifest in the survey responses. An uncertainty shock to demand helps us to target the patterns of **observation 2** and **observation 4**, namely, the rise in uncertainty about output and inflation that consumers express; recall Fig. 2 and the bottom row of Fig. 3. To get anywhere near the numbers reported in the survey, we assume a 17.5-standard-deviations pandemic rise in the volatility of the demand shock (σ_t^a).

Next, a fall in TFP now or later is essential for mimicking consumers' stagflationary view of the recession, **observation 3**. We allow for a 5-standard-deviations shock to the persistent component of TFP (A_t) and a 15-standard-deviations fall in the news component of TFP (X_t). The split between the persistent component A_t and the rather transitory news component X_t arises from fact that household expectations of the recession are of limited persistence; recall Fig. 1. Last, we have a 15-standard-deviations shock to the level of demand preferences (a_t), so as to replicate the drop in GDP that consumers expect, **observation 1** of the survey; refer again to Fig. 1.[11]

All of the COVID-19 scenario simulations assume that monetary policy is expected to stabilize economic activity at its pre-pandemic (no-shock) level. That is, monetary policy is expected to follow rule (3) as in normal times, but with ρ_y multiplying a conventional measure of the output gap, $\log(Y_t/\overline{Y})$, rather than the gap between output and flex-price output.[12]

[11] In our view, there are two possible ways to think of our COVID-19 scenario. First, the pandemic was just exceptional, and hence, it requires exceptionally large shocks to rationalize the response of expectations to the pandemic. Second, the size of the shocks testifies to the limitations of the representative-agent FIRE model and strengthens the case to move beyond that framework.

[12] This choice may appear natural in light of the March 15, 2020 FOMC statement, which reads: *The Committee expects to maintain this target range [of 0 to 0.25 percent for the federal funds rate] until it is confident that the economy has weathered recent events and is on track to achieve its maximum employment*

The model does not feature cost-push shocks. At the same time, consumer survey expectations are stagflationary on average. In the context of the model, stagflation will result only if monetary policy is overly accommodative in some period (relative to the natural interest rate).

4.4. The baseline

Fig. 4 shows how the COVID-19 shock defined above affects consumers' perceptions of the distribution of future economic activity and inflation at the time of impact (*in spring 2020*). The left column shows the pandemic's effect on the distribution of future output; the right column shows the pandemic's effect on inflation. The first row shows the average perceived effect along with 95% coverage bands. The bottom row shows the effect of the pandemic on the standard deviation as another measure of uncertainty. What is directly comparable to the survey is the impact response (in period 0 on the x-axis) in the graphs. The time dimension shown here serves as corroborating evidence.

Compare the mean response of (consumer) expectations of GDP to COVID-19 to **observation 1** in Section 3. On impact, consumers in the model scenario expect COVID-19 to make output over the course of the next 12 months fall by 7%. This is consistent with the survey responses except for the very strong overshooting of initial perceptions in the survey. The mean response for inflation (top-right panel) should be compared to **observation 3** in the survey. In the model-based scenario, as in the survey, consumers expect COVID-19 to make inflation rise. At the same time, the sheer extent of the stagflationary impact of COVID-19 that consumers expect may be hard to replicate with the model, unless we have even larger shocks or allow for cost-push shocks. Recall that respondents in the survey report that they expect COVID-19 to raise inflation by more than 5 percentage points.

The panels in the bottom row show the implied standard deviation of the conditional distribution of output and inflation as induced by the advent of COVID-19. The standard deviations reported here are consistent with the strong rise in subjective uncertainty that consumers report in the survey. COVID-19 makes the standard deviation of output as perceived by consumers in our model scenario rise by 3.5 percentage points; compare **observation 2**. Similarly, compare the standard deviation of inflation (bottom-right panel) directly to **observation 4** for inflation and, in particular, to the bottom-right panel of Fig. 3.

4.5. The role of news and uncertainty

There are three salient observations in the survey. First, household expectations on average are stagflationary. Second, households expect as deep a recession as professional forecasters. Third, households are much more uncertain about the impact of COVID-19 than professional forecasters. The current section gets at what may be behind the stagflation view and how household uncertainty is related to the average depth of the recession. The next section, then, provides one potential reason why households' expectations may differ from professional forecasters'.

The top row of Fig. 5 shows the role that news about future productivity (X_t) plays in shaping the model-based recession. The panels plot the perceived impact of COVID-19 in the baseline scenario (red dashed lines) against a counterfactual that is identical except that the news component is mute (blue dashed-dotted line). The right panel shows the response of inflation: the negative news shock to productivity is essential for explaining the stagflationary response, **observation 3** (inflationary beliefs). The reason is simple. In the scenarios here, the central bank leans against lower future productivity, keeping the real rate below the natural rate of interest—unless it is constrained by the effective lower bound. This policy response raises future marginal costs. Forward-looking price setters respond by raising prices already at the onset of the pandemic (red dashed line and corresponding bounds). Absent the news shock, instead, future marginal costs do not rise, and inflation falls on impact (blue dashed dotted line). Note that this means that—whenever the effective lower bound is binding—the real rate of interest is higher without the news shock. This in turn explains why the response of output is of comparable magnitude both with and without an initial negative news shock (see the left panel of the first row).

The bottom row of Fig. 5 shows the role that the shock to demand uncertainty plays in the model-based recession. The baseline is identical to the top row (and is shown as a red dashed line again). The blue dashed-dotted line now shows the response of the economy if the shock to demand uncertainty is mute. In interpreting this, it is important to note that—unless monetary policy is constrained—in the context of the model, monetary policy could perfectly absorb the shock to demand uncertainty (Basu and Bundick, 2017). The uncertainty bands, thus, reflect the shock's propagation. Quite clearly, looking through the lens of the model, the shock to demand uncertainty is the primary driver of **observation 1** (the deep recession). The rise in demand uncertainty means that the natural rate of interest falls sharply in the baseline, by 15 pp. annualized (the bottom left panel of Fig. 6 shows the response of the natural rate). Most of this is due to the shock to demand uncertainty.[13] This means that the effective lower bound on interest rates becomes binding, and monetary policy cannot accommodate this shock. A deep recession ensues. Absent the shock to demand uncertainty, output falls only about a third as much as in the baseline (bottom row, left panel). Note that the demand uncertainty shock is also the primary driver of the rise in consumer

and price stability goals. This action will help support economic activity, strong labor market conditions, and inflation returning to the Committee's symmetric 2 percent objective.

[13] Figure D.11 in the Supplementary Material shows the impact of individual shocks. It shows that the demand uncertainty shock is the main driver of the fall in the natural rate. Intuitively, the natural rate falls in response to consumers' increased demand for precautionary savings.

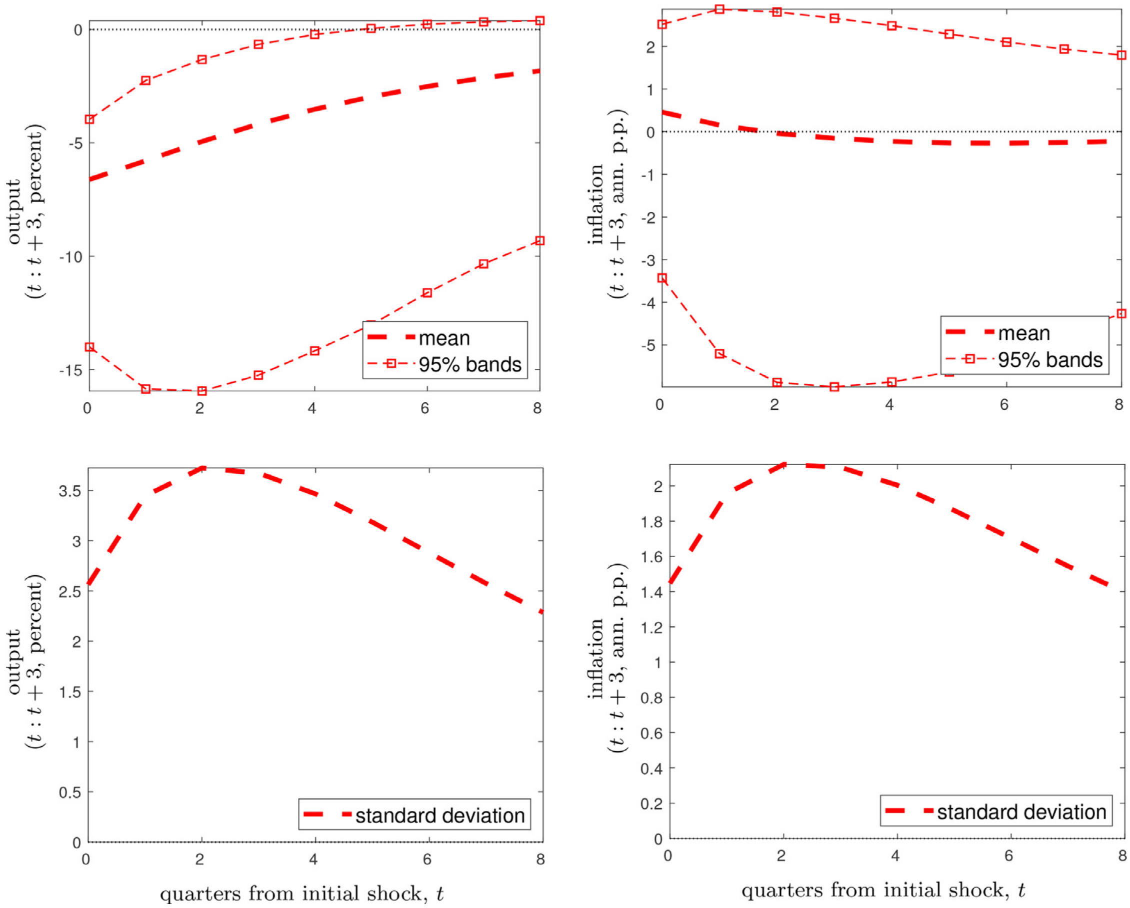

Fig. 4. COVID-19 Baseline Scenario in Simulated Model. *Notes:* The left (right) column shows the effect on output (inflation) in the coming 12 months (as expected in the impact period). Top row: expected effect and ± standard deviation bands of the effect. Bottom row: standard deviation of the shock's effect. Supplementary Material D reports the expected response of other model variables.

uncertainty itself, **observations 2 and 4**. Absent the direct effect on uncertainty, the pandemic shock would hardly affect the standard deviation of (that is, the uncertainty about) output and inflation, as can be seen by observing that the bands almost coincide with the effect of the shock on the means.

In sum, the above suggests that heightened consumer uncertainty about demand itself may have been an important factor behind the depth of the recession. This raises the policy-relevant question: through what means exactly could consumer uncertainty have been reduced? This is clearly outside the scope of the current analysis. That said, the role of central bank communication has recently figured prominently in both research and the policy discussion (see, e.g., Coibion et al., 2020a). It is also important here to bear in mind that—in the context of the model—expectations about the response of monetary policy have a role in shaping the uncertainty that consumers face, a role to which we turn next.

4.6. Road to reconciliation? the role of monetary policy

Our consumer survey did not ask consumers directly about their perceptions of how monetary policy would respond to the pandemic. The model-based analysis of the survey responses took a particular stand, though. It assumed that monetary policy does not correct for the sharp movements in the natural rate of interest. This resulted in dynamics consistent with the households' view of the impact of the pandemic. This section, instead, looks at a policy counterfactual in which monetary policy does adjust the level of the nominal rate for movements in the natural rate of interest—and in which households know that it does. The results are as follows. Keeping the shocks as before, the response of output is as in the baseline. However, inflation now falls, and uncertainty as measured by the standard deviations is smaller. All of these are characteristics that the projections of professional forecasts had (recall Figs. 1 and 2). This suggests that different perceptions about policy are potentially part of the disagreement between professional forecasters and consumers.

Fig. 6 shows the effect of the COVID shock on the evolution of the economy under both the baseline policy rule (red dashed lines as before) and under the policy alternative. The policy alternative has the interest-rate response perfectly

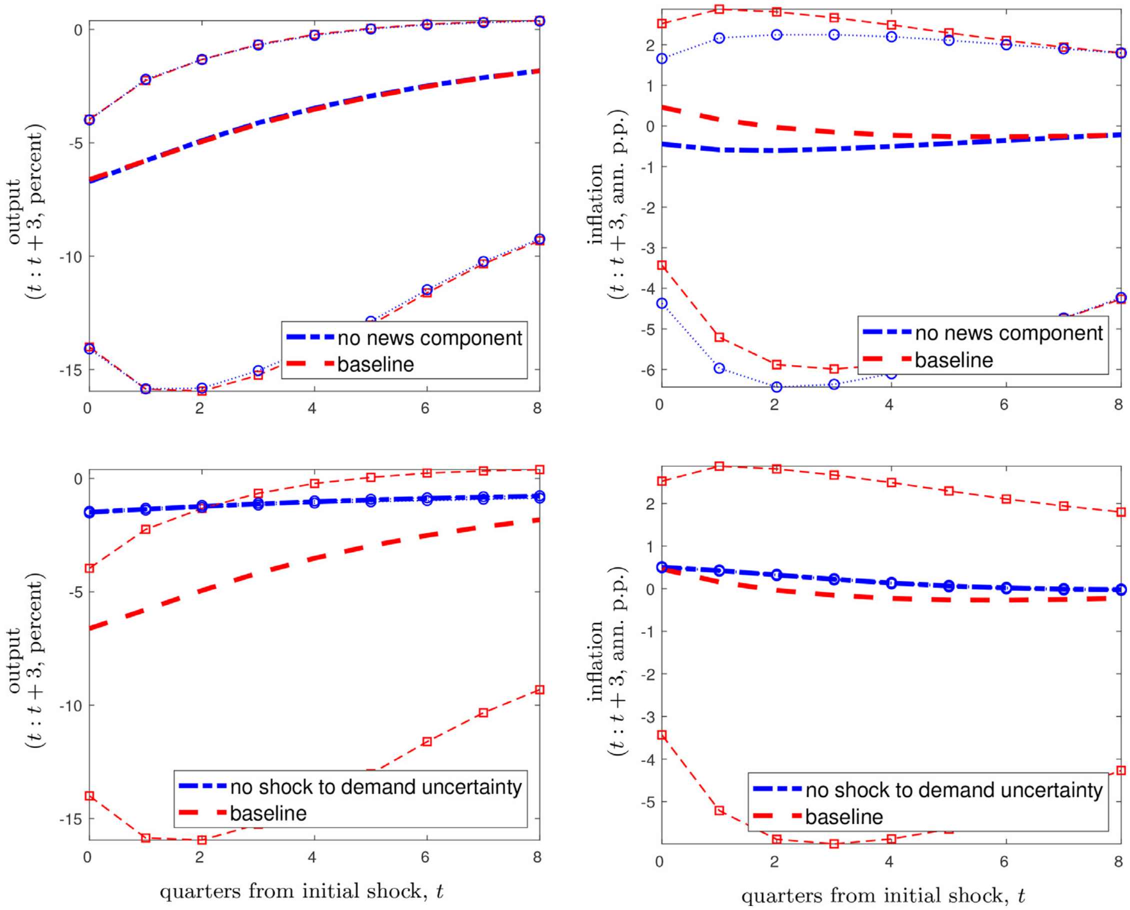

Fig. 5. The Role of News and Uncertainty. *Notes:* Same as top row of Fig. 4, but contrasting two alternative scenarios. Top row: contrasting the baseline response (red dashed lines) with a scenario in which there is no shock to the news component of TFP (blue dashed dotted). Bottom row: contrasting the baseline response (red dashed lines) with a scenario in which there is no shock to demand uncertainty (blue dashed dotted). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

indexed to the natural rate of interest (blue lines), provided monetary policy can achieve this in light of the effective lower bound.

In detail, under the alternative policy, the target interest rate is now governed by

$$\log\left(R_t^{\text{tar}}/R_t^n\right) = \rho_\Pi \cdot \log(\Pi_t/\overline{\Pi}) + \rho_y \cdot \log(Y_t/\overline{Y}), \quad (4)$$

so that—all else equal—the interest rate tracks the actual natural rate of interest one-to-one.

The top row of Fig. 6 shows how the change in policy affects output and inflation, and uncertainty about the two. The most important result is that the alternative policy notably reduces the uncertainty bands for all variables. The bands are between one-third and one-half as wide as under the baseline policy. It is important to note that the shocks are identical in both scenarios shown here. What differs is only the policy response: the policy response and perceptions thereof matter.

The difference in the responses of uncertainty is most easily explained for inflation. Absent the lower bound, the alternative policy would stabilize inflation almost perfectly. This means that any shock would hardly affect uncertainty about inflation. With the lower bound, however, such tracking is not perfect so that some uncertainty remains. The bottom-right panel of Fig. 6 illustrates the constraints imposed by the lower bound. It reports for each period for what share of the simulations the economy is at the effective lower bound. Since the natural rate falls markedly, tracking the natural rate directly also implies considerably more accommodative monetary policy.

The degree of accommodation that monetary policy provides (if it can) also explains why the bands for output are notably narrower under the alternative policy. The fact that there is a sizable recession at all (so that the demand uncertainty shock is not fully absorbed) comes from monetary policy being constrained in the short run, even under the alternative policy rule. Hence, in the initial periods, the recession is similar across the two policy alternatives. Importantly, however, tracking the natural rate of interest anchors expectations about inflation and output more firmly.

Another observation from the top-right panel is that the natural rate policy reduces not only the *downward* risk to inflation but also the *upward* risk to inflation. The reason for this has to do with firms' price-setting behavior. A rise in

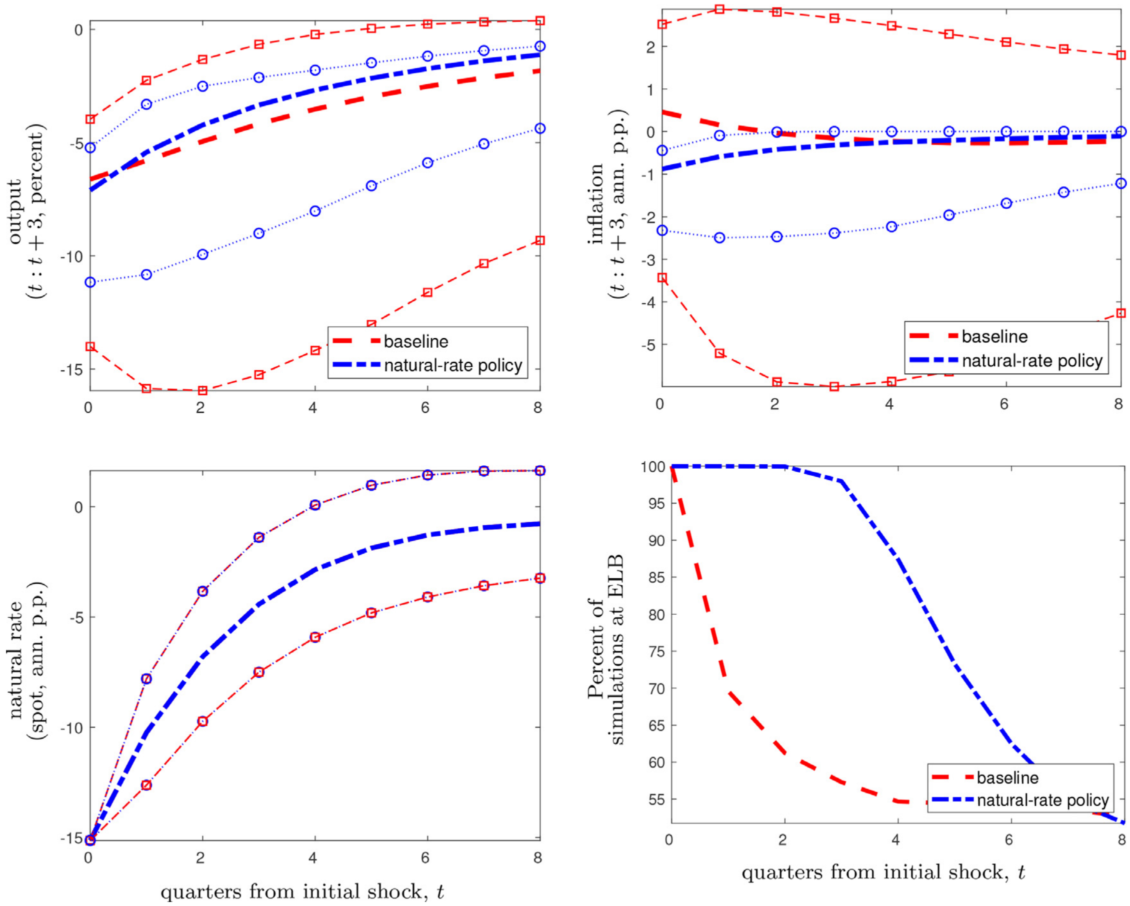

Fig. 6. The COVID-19 Effect: The Role of the Perceived Response of Monetary Policy. *Notes:* Same scenario as in Fig. 4, but contrasting the baseline COVID-19 effect (red lines) with a scenario in which the target interest rate is adjusted perfectly for movements in the natural rate of interest (blue lines), provided this is possible in light of the effective lower bound. Top row: output and inflation. Bottom row, left panel: natural rate of interest, the evolution of which is identical under both scenarios. Bottom row, right panel: Share of simulations for which economy is at the ELB, by period. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

demand uncertainty induces firms to charge higher prices for precautionary reasons (so as to insure against having to face a lot of demand for its products when marginal costs are high), an effect explored in detail in Fernández-Villaverde et al. (2015). A central bank that puts more consistent focus on inflation prevents this precautionary pricing: firms forgo raising prices precautionarily in the early periods of the recession. Whereas in the baseline inflation rose on impact, thus, under the natural rate policy, inflation falls; see the top row right panel of Fig. 6.

In sum, under the natural rate policy, for the same set of shocks as in our baseline, a deep recession arises that is accompanied by a fall in prices and one in which uncertainty about output and inflation rises much more moderately than in our baseline (and than in our consumer survey). Along all of these dimensions, the dynamics of the perceived impact of COVID-19 under the natural rate policy seem to be rather more consistent with the expectations of *professional* forecasters shown in Figs. 1 and 2 (and of firms; see Footnote 7 above). Different implicit assumptions about the monetary policy response may help explain why consumers declared that they were much more uncertain about the impact of COVID-19 than professional forecasters.

5. Conclusions

In this paper, we assess the response of consumer expectations to the pandemic. We do so at two levels. First, a real-time survey shows that consumer expectations respond strongly and swiftly to the COVID-19 shock. At the same time, consumers are highly uncertain about the size of its economic effects—and much more so than professional forecasters. Second, we show that it is possible to account for essential patterns embedded in the survey responses on the basis of a FIRE model. We also run counterfactuals to illustrate the importance of expectations for the unfolding of the crisis and of monetary policy. What appears to account for most of the pandemic's recessionary impact is the rise in households' uncertainty.

The model-based results rest on a representative-household FIRE framework. An important insight of our analysis is that it takes exceptionally large shocks to account for the survey evidence—perhaps underscoring the framework's limitations, not least its assumption of homogeneous expectations across consumers and between consumers and firms. Future modelling will hopefully be able to account for the heterogeneity of expectations. Our survey provides important constraints for such efforts. First, it shows that expectations can be very responsive to large, salient shocks. This suggests extending sticky-information settings with heterogeneous households, as in Auclert et al. (2020), to noisy-information environments. Second, the survey documents a strong response of the uncertainty that households express. The current heterogeneous-household literature provides algorithms that provide second-order approximations, for example, Gornemann et al. (2021). Accounting for the effects of changes in uncertainty likely requires extensions to a higher order still.

Even accounting for heterogeneity, however, our main policy conclusion should stand: if an event brings about a steep rise in household uncertainty, this rise in uncertainty will adversely affect the economy. Effective policy communication with the household sector could dampen the very rise in uncertainty. This suggests that such policy communication itself could be an important tool that helps limit the fallout from large adverse events.

Transparency declaration

This article is part of a Supplement funded by the European Central Bank.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2022.02.004

References

Altig, D., Baker, S., Barrero, J.M., Bloom, N., Bunn, P., Chen, S., Davis, S.J., Leather, J., Meyer, B., Mihaylov, E., Mizen, P., Parker, N., Renault, T., Smietanka, P., Thwaites, G., 2020. Economic uncertainty before and during the COVID-19 pandemic. J Public Econ 191.
Andre, P., Pizzinelli, C., Roth, C., Wohlfart, J., 2021. Subjective models of the macroeconomy: evidence from experts and representative samples. Review of Economic Studies. Forthcoming.
Armantier, O., Kosar, G., Pomerantz, R., Skandalis, D., Smith, K., Topa, G., van der Klaauw, W., 2020. How Economic Crises Affect Inflation Beliefs: Evidence from the COVID-19 Pandemic. Staff Report. Federal Reserve Bank of New York. 949.
Armantier, O., Topa, G., van der Klaauw, W., Zafar, B., 2017. An overview of the survey of consumer expectations. FRBNY Economic Policy Review 23 (2).
Auclert, A., Rognlie, M., Straub, L., 2020. Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model. Working Paper. National Bureau of Economic Research. 26647 doi: 10.3386/w26647.
Baker, S.R., Bloom, N., Davis, S.J., Terry, S.J., 2020. COVID-Induced Economic Uncertainty. Working Paper. National Bureau of Economic Research. 26983 doi: 10.3386/w26983.
Baker, S.R., McElroy, T.S., Sheng, X.S., 2020. Expectation formation following large, unexpected shocks. Review of Economics and Statistics 102 (2), 287–303. doi:10.1162/rest_a_00826.
Balleer, A., Link, S., Menkhoff, M., Zorn, P., 2020. Demand or Supply? Price Adjustment During the Covid-19 Pandemic, 31, pp. 59–102. Covid Economics: Vetted and Real-Time Papers.
Basu, S., Bundick, B., 2017. Uncertainty shocks in a model of effective demand. Econometrica 85 (3), 937–958. doi:10.3982/ECTA13960.
Bayer, C., Born, B., Luetticke, R., Müller, G. J., 2020. The Coronavirus Stimulus Package: How Large is the Transfer Multiplier?CEPR Discussion paper 14600.
Binder, C., 2020. Coronavirus fears and macroeconomic expectations. Review of Economics and Statistics 102 (4), 721–730. doi:10.1162/rest_a_00931.
Bishop, Y., Fienberg, S., Holland, P., 1975. Discrete multivariate analysis: theory and practice. MIT Press.
Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77 (3), 623–685. doi:10.3982/ECTA6248.
Bloom, N., 2014. Fluctuations in uncertainty. Journal of Economic Perspectives 28 (2), 153–176. doi:10.1257/jep.28.2.153.
Bordalo, P., Gennaioli, N., Ma, Y., Shleifer, A., 2020. Overreaction in macroeconomic expectations. American Economic Review 110 (9), 2748–2782. doi:10.1257/aer.20181219.
Born, B., Enders, Z., Menkhoff, M., Müller, G., Niemann, K., 2021. How Economic Expectations React to News: Evidence from German Firms.
Broer, T., Kohlhas, A., 2021. Forecaster (mis-)behavior. Review of Economics and Statistics. Forthcoming.
Bryan, M.F., Venkatu, G., 2001. The demographics of inflation opinion surveys. Economic Commentary 11, 2.
Candia, B., Coibion, O., Gorodnichenko, Y., 2020. Communication and the Beliefs of Economic Agents. Economic Policy Symposium (Jackson Hole, WY) Proceedings.
Carroll, C.D., 2003. Macroeconomic expectations of households and professional forecasters. Quarterly Journal of Economics 118 (1), 269–298. doi:10.1162/00335530360535207.
Carroll, C.D., Crawley, E., Slacalek, J., Tokuoka, K., White, M.N., 2020. Sticky expectations and consumption dynamics. American Economic Journal: Macroeconomics 12 (3), 40–76. doi:10.1257/mac.20180286.
Cavallo, A., 2020. Inflation with Covid Consumption Baskets. Working Paper. National Bureau of Economic Research. 27352 doi: 10.3386/w27352.
Cavallo, A., Cruces, G., Perez-Truglia, R., 2017. Inflation expectations, learning, and supermarket prices: evidence from survey experiments. American Economic Journal: Macroeconomics 9 (3), 1–35. doi:10.1257/mac.20150147.
Christelis, D., Georgarakos, D., Jappelli, T., Kenny, G., 2020. The Covid-19 Crisis and Consumption: Survey Evidence from Six EU Countries. Working Paper Series. European Central Bank. 2507.
Coibion, O., Georgarakos, D., Gorodnichenko, Y., Kenny, G., Weber, M., 2021. The Effect of Macroeconomic Uncertainty on Household Spending. Working Paper. National Bureau of Economic Research. 28625 doi: 10.3386/w28625.
Coibion, O., Gorodnichenko, Y., 2012. What can survey forecasts tell us about information rigidities? Journal of Political Economy 120 (1), 116–159.
Coibion, O., Gorodnichenko, Y., 2015. Information rigidity and the expectations formation process: a simple framework and new facts. American Economic Review 105 (8), 2644–2678. doi:10.1257/aer.20110306.
Coibion, O., Gorodnichenko, Y., Knotek, E.S., Schoenle, R., 2020. Average Inflation Targeting and Household Expectations. Working Paper. National Bureau of Economic Research. 27836 doi: 10.3386/w27836.
Coibion, O., Gorodnichenko, Y., Weber, M., 2020. The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending. Working Paper. National Bureau of Economic Research. 27141 doi: 10.3386/w27141.
Coibion, O., Gorodnichenko, Y., Weber, M., 2022. Monetary Policy Communications and their Effects on Household Inflation Expectations. Journal of Political Economy, forthcoming.

D'Acunto, F., Malmendier, U., Weber, M., 2021. Gender roles produce divergent economic expectations. Proceedings of the National Academy of Sciences 118 (21). doi:10.1073/pnas.2008534118.
Dietrich, A., Knotek, E. S., Myrseth, K. O. R., Rich, R. W., Schoenle, R. S., Weber, M., 2022. Inflation Expectations: Category Beliefs and Spending Plans.
Fernández-Villaverde, J., Guerrón-Quintana, P., Kuester, K., Rubio-Ramírez, J., 2015. Fiscal volatility shocks and economic activity. American Economic Review 105 (11), 3352–3384. doi:10.1257/aer.20121236.
Fernández-Villaverde, J., Guerrón-Quintana, P.A., 2020. Uncertainty shocks and business cycle research. Rev Econ Dyn 37, 118–146. doi:10.1016/j.red.2020.06.005.
Fetzer, T., Hensel, L., Hermle, J., Roth, C., 2021. Coronavirus perceptions and economic anxiety. Rev Econ Stat. Forthcoming.
Fornaro, L., Wolf, M., 2020. Covid-19 Coronavirus and Macroeconomic Policy. Working Paper. Barcelona Graduate School of Economics. 1168.
Gali, J., Gertler, M., 1999. Inflation dynamics: A Structural econometric analysis. J Monet Econ 44 (2), 195–222. doi:10.1016/S0304-3932(99)00023-9.
Gornemann, N., Kuester, K., Nakajima, M., 2021. Doves for the Rich, Hawks for the Poor? Distributional Consequences of Systematic Monetary Policy. Opportunity and Inclusive Growth Institute Working Papers. Federal Reserve Bank of Minneapolis doi:10.21034/iwp.50. 50.
Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I., 2020. Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? Working Paper. National Bureau of Economic Research. 26918 doi: 10.3386/w26918.
Hanspal, T., Weber, A., Wohlfart, J., 2021. Exposure to the COVID-19 stock market crash and its effect on household expectations. Review of Economics and Statistics. Forthcoming.
Idel, M., 2016. A review of matrix scaling and Sinkhorn's normal form for matrices and positive maps. arXiv: Rings and Algebras.
Larsen, V.H., Thorsrud, L.A., Zhulanova, J., 2021. News-driven inflation expectations and information rigidities. J Monet Econ 117, 507–520. doi:10.1016/j.jmoneco.2020.03.004.
Mankiw, N.G., Reis, R., Wolfers, J., 2004. Disagreement about Inflation Expectations. In: NBER Macroeconomics Annual 2003, Volume 18. National Bureau of Economic Research, Inc, pp. 209–270. NBER Chapters.
Meyer, B.H., Prescott, B., Sheng, X.S., 2021. The impact of the COVID-19 pandemic on business expectations. Int J Forecast. In press.
Miescu, M., Rossi, R., 2021. COVID-19-Induced shocks and uncertainty. Eur Econ Rev 139, 103893. doi:10.1016/j.euroecorev.2021.103893.
Nakamura, E., Steinsson, J., 2018. Identification in macroeconomics. The Journal of Economic Perspectives 32 (3), 59–86.
Schmitt-Grohé, S., Uribe, M., 2012. What'S news in business cycles. Econometrica 80 (6), 2733–2764.