Temperature surprise shocks

Natoli, Filippo

Bank of Italy

29 March 2022

Temperature surprise shocks

Filippo Natoli

Bank of Italy

March 29, 2022

Abstract

Using daily county-level data since 1970, we construct a series of temperature shocks for the United States that capture the average *surprise* effect of heat and cold events experienced in each season, net of climate trends and adaptation. Temperature surprise shocks in the global warming era have been a balanced mix of heat and cold surprises and reduced in size in recent times, in contrast to common belief. Estimates made with local projections show a negative impact on the US economy at business cycle frequency via both consumption and investment, while the effect on prices is more muted and varies over time. The central bank does react to the shocks by adjusting its economic projections and cutting interest rates, with effects spreading out through the yield curve.

Keywords: climate change, temperatures, surprise shocks, business cycle, monetary policy

JEL Classification Numbers: C32, E32, E52, Q54

Email address: `filippo.natoli@bancaditalia.it.` (Filippo Natoli)

Bank of Italy, Directorate General for Economics, Statistics and Research. I thank Lars Peter Hansen, Johannes Stroebel, Piergiorgio Alessandri, Fabrizio Ferriani, Valerio Nispi Landi and seminar participants at Banca d'Italia for their helpful comments and suggestions. I also thank Luca Missori for making textual analysis on FOMC meeting transcripts. I certify that I have the right to deposit the contribution with MPRA. The views expressed in this paper are those of the author and do not necessarily reflect those of the Bank of Italy. All the remaining errors are mine.

1. Introduction

The analysis of the economic effects of climate change has gained traction in recent years, with a renewed focus on very high and low temperatures. In the macroeconomic literature, the impact of temperatures has commonly been investigated empirically under the assumption that, in the short run, temperatures vary randomly in a given place. However, findings in climate research cast doubt on the plausibility of this assumption, as global warming is considered to be impacting not only temperature trends, but also their local variability over time, increasing the incidence of extreme values. If this is the case, deviations of temperatures from their long-run averages may not be considered as exogenous shocks, because part of this variation depends on climatic developments, making it predictable.

We take up this issue and propose a new way to construct temperature shocks. For this purpose, we resort to a concept that is common in the empirical macro literature but still absent from climate research, which is that of surprise shocks. Based on that notion, pointing to agents' reaction to a surprising event, we construct a time series of temperature shocks for the United States by isolating the component of temperature variations that cannot be reasonably anticipated. We do this at business cycle frequency and by aggregating local (county-level) surprises, as average country-year dynamics may hide relevant variations at seasonal frequency and sub-national level. We make our shock series in two steps. First, using daily data since 1970, we construct a series of surprises in each county, where surprises are defined as the number of hot and cold days in excess to those agents expect in each quarter, based on their experience in the preceding years for the same period. Temperature beliefs, defined in terms of a distribution of experienced temperatures, are constructed over multi-year spans reflecting the documented learning behavior regarding climate-related phenomena, and updated yearly to identify surprises. Second, county-level surprises are averaged based on counties' weather exposure and weight on the national economy, pinning down US-wide unanticipated surprise shocks. The obtained series nets out any secular variation in the temperature distribution, including trends in temperature level and volatility, as well as any form of adaptation of agents to climate change.

We build US temperature surprise shocks quarterly starting from the beginning of the global warming era (the 1970s) to 2019. Shocks have been relatively small in the last 15 years, suggesting that adjustments in the shape of the temperature distribution have been stronger in the early phase of climatic changes. This is not in contrast with the intensification of climate-related manifestations (including temperatures oscillations) over time, as large weather events with respect to the pre-global-warming period may have become increasingly frequent, but also less surprising than in the past. We use these shocks to explore the effects of temperatures on the US economy, studying the impacts on different expenditure components of output, on consumer prices and interest rates. In a narratively identified local projection framework, we find that temperature surprise shocks negatively affect both consumption and investment, but weigh more on the latter. On aggregate, a one-standard deviation shock causes a contraction in real GDP by 0.34% after 2 years; by contrast, the impact on prices appears to be more muted, with significant effects on the CPI index showing up only discontinuously as if demand- and supply-side effects almost balance out. The detrimental economic effect induces an expansionary monetary policy reaction, with a fall in short- and long-term nominal interest rates. These findings survive to a large set of robustness exercises on the design of the temperature shock – experimenting different learning times, thresholds for extreme values and counties' weighting schemes – and on regression specifications. Moreover, the effects appear to be heterogeneous across seasons – largest in summer – and between right-tail (heat) shocks and left-tail (cold) shocks, with the latter generating slightly stronger responses.

In order to investigate the central bank's reaction further, we rely on Fed's staff economic expectations provided to the Federal Open Market Committee ahead of their meetings (Greebook forecasts), as well as on textual analysis on the transcripts of each meeting since the 1970s. Following the shock, the Federal Reserve is found to revise its forecasts down for growth, slightly up for prices in the short run and down in the longer term. Moreover, shocks have an impact on policy discussion within the Committee: while climate change has been rarely debated, the Fed did mention temperatures quite often and the incidence of this wording increases after temperature shocks occurr.

Our contribution to the literature is twofold. First, we provide a new source of exogenous macroeconomic variation, related to climate change, which affects the economy at aggregate level and business cycle frequency. Having a measure of unanticipated climate-related shock can help tackle a wider range of questions than those faced in this paper, and can be the basis of further climate research. Moreover, our shock is easy to be replicated for other countries or at wider/narrower geographic dimension, and can serve as reference to construct other weather-related shocks under the same logic. Second, we contribute to the literature by quantifying the aggregate short-term implications of temperature oscillations for the United States, including the extent of monetary policy reaction. In this respect, looking at temperature shocks in isolation is key, as temperatures impact the economy through a set of transmission channels (including demand-side ones) that can shape output, prices and interest rates differently than other weather-related extreme events. While our shock, being temporary in nature, cannot speak about longer run effects of climate change, it can nonetheless inform the debate on the needed policy response, in particular regarding the engagement of central banks.

The paper is organized as follows. Section 1.1 reviews the theoretical and empirical contributions related to our study. Section 2 illustrates the possible transmission mechanisms of temperature variations to the economy. Section 3 discusses the key issues surrounding the identification of an exogenous temperature shocks and describes how the shock is constructed. Section 4 presents the data used in the analysis and works out the shock series. Section 5 proposes an empirical application to estimate the domestic effect of temperatures on the US economy using the previously constructed shock. Section 6 provides robustness checks, and Section 7 concludes.

1.1. Related literature

The effects of climate change on the global economy are found to be multifaceted. Increasing temperatures substantially reduce activity and growth, especially in hot and poor countries (Dell et al., 2012; Burke et al., 2015; Acevedo et al., 2020); see Dell et al., 2014 for a review of the literature). Advanced economies appear not to be immune either, and evidence of an impact of temperatures also in the United States, initially mixed, is also growing. Indeed, US output is negatively affected in a wide range of industries (Hsiang et al., 2017; Colacito et al., 2019), with income per capita losses being concentrated during business days (Deryugina and Hsiang, 2014).[1] From a wider perspective, extreme weather events are found to have increased their negative impact on the US economy in the most recent decades, according to Kim et al. (2021). However, challenges in isolating the economic effects of temperatures remain, also because their implications differ substantially in different seasons and sectors (Addoum et al., 2021).[2] We contribute to this literature by shedding light on the impact of temperatures shocks on the US economy, exploring their effects beyond aggregate output including consumption, investment, consumer prices and bond yields.

From a methodological point of view, this paper contributes to the strand that identifies short-term impacts of climate change, in particular temperature shocks. A common empirical approach is to retrieve shocks in a recursively-identified structural VAR, with average temperatures ordered as the most exogenous variable (Donadelli et al., 2017); the same logic has been applied to extrapolate shocks from other weather series.[3] Other studies employ temperature anomalies, constructed as deviations from the

[1] A significantly negative effect is also found on agriculture (Fisher et al., 2012; Burke et al., 2015), despite earlier evidence suggested no impact (Deschênes and Greenstone, 2007).

[2] While pointing to a detrimental effect on firm activity, micro evidence is still mixed. Some papers do not find any consequence of high temperatures on firm sales and labor productivity (Addoum et al., 2020), while others document negative effects on upstream firms in the supply chain (Custódio et al., 2021; Pankratz and Schiller, 2021) and a reduction in the number of employees and firm establishments in the medium run (Jin et al., 2021).

[3] Gallic and Vermandel (2020) uses a de-trended version of the soil moisture index, a measure capturing the combined effect of temperatures and precipitations, in a VAR model; Cashin et al. (2017) includes deviation of the Southern Oscillation Index (SOI) from their historical averages in a Global VAR framework.

averages of a pre-global-warming reference period, as a direct measure of temperature shocks (Makkonen et al., 2021, among others). In the widely used fixed-effects panel settings, shocks to temperatures are identified as the deviation of current temperature levels from their long run mean. We argue that all the methods that do not take into account the continued evolution of the temperature distribution suffer from an endogeneity issue, and propose a new way to compute shocks capturing the surprise component embedded in temperature episodes. Moreover, most research employs country-level temperatures built as simple averages of sub-national values, disregarding the fact that extreme events occurring in different areas of the country can produce profoundly different aggregate effects. We start our analysis from daily county data and test different weighting schemes to evaluate the robustness of our results. Last, we take into account seasonal patterns in line with the relevance of seasonality found in Colacito et al. (2019) and Addoum et al. (2021).

As our notion of temperature shock is based on variations in the shape of the whole temperature distribution, we also connect to papers studying the effects of actual or expected temperature volatility (Kotz et al., 2021; Donadelli et al., 2021; Alessandri and Mumtaz, 2021). While these papers isolate second-moment shocks, which might have the characteristics of forward-looking climatic risk, our shock is defined over the level of temperatures currently experienced by agents. As such, it should be interpreted as an unanticipated variation, more than as a news over future temperature oscillations.

The idea of macroeconomic temperature surprises is inherited from the empirical macro literature, in which the identification of shocks typically rely on extrapolating an exogenous component from policy announcements or decisions (see Nakamura and Steinsson, 2018 for a review of the literature). For example, the notion of surprise shocks related to monetary policy refers to the contemporaneous surprise component of monetary policy announcements (Gurkaynak et al., 2005; Gertler and Karadi, 2015; Miranda-Agrippino and Ricco, 2021 among others); regarding fiscal policy, government spending surprises can reflect the unanticipated component of public spending decisions (Forni and Gambetti, 2016, among others). As for policy surprises, what matters here is defining agents' ex-ante expectations in a proper way: in our setting, what agents expect in a given year and season cannot be any different from temperatures directly experienced during the same season in the most recent years, i.e. those agents are used to.

Last, our paper links to those investigating the effects of climate change on consumer prices and the reaction of monetary policy. With respect to output effects, the implications for prices are less clear in the literature, sometimes providing opposite evidence (Mukherjee and Ouattara, 2021; Faccia et al., 2021). This can be due to the fact that temperatures might have relevant demand-side – other than supply-side – effects, which might offset the final impact on consumer prices at some particular horizons. We find that the effects of temperatures on prices are more muted than those on output, but skewed towards a price fall in the medium run. This price response is broadly in line with that found in Faccia et al. (2021) for a panel of countries, for which the impact is initially positive and it becomes negative in the longer term: our result of a barely significant response is coherent with that found in that paper for advanced economies.

On monetary policy, the literature has mainly focused on how changes in the monetary stance might have implications for climate and how central banks might cooperate to stimulate the low-carbon transition (Ferrari and Nispi Landi, 2020; Papoutsi et al., 2021; Hansen, 2021, among others), almost disregarding the effects of actual weather occurrences on its conduct. We fill this gap by documenting that monetary policy responds to the economic damage caused by temperature shocks by cutting short-term rates, with effects passing through the whole term structure of government bond yields.

2. Transmission channels

While climate change is a long run phenomenon, the literature has widely documented that temperatures are able to affect economic outcomes also in the short run. These effects unfold through different transmission mechanisms. The one that received most attention so far is the heat stress channel. Extreme temperatures might cause illnesses (e.g., heat stroke), reducing hours worked and time allocated to outdoor leisure (Graff Zivin and Neidell, 2014), and generate a fall in individual productivity related to heat-exposed working tasks (Cachon et al., 2012; Somanathan et al., 2021). As temperature extremes are not destructive events, supply effects are mainly on the labor side, while direct effects on capital are much more limited. Through the same channel, temperature oscillations might also hit the economy on the demand side. Indeed, some papers have documented significant behavioral effects of temperatures, according to which extreme temperatures would discourage open air activities: perception of waiting time worsens in hot days (Baker and Cameron, 1996) and social interactions with strangers are felt as more unpleasant (Griffit and Veitch, 1971). There is initial evidence that these effects can have economic implications, putting downward pressure on consumer spending through a decrease in shop retail sales (Starr-McCluer, 2000; Roth Tran, 2022).

Another transmission mechanism of temperature extremes to the economy is through the cost channel, for which the literature is scant. In the United States, based on climate scenarios demand for electricity will rise in the next decades, with a substantial increase in system costs (McFarland, 2015). In the shorter run, firms may suffer larger expenses to cool/heat work spaces or production processes: this can hit firms' performance and higher costs may pass through higher consumer prices. Along the same lines, if extreme temperatures damage agricultural output (crop yield channel), this may entail an increase in the food component of consumer prices, as shown in Faccia et al. (2021).

A third channel of transmission is related to uncertainty over future climatic developments. Extreme temperature oscillations, if wide or frequent enough, can influence decision making by raising attention towards the future repercussions of climate change (wake-up call effect): this is documented in Choi et al. (2020), who link temperature episodes to higher financial investment in green than brown assets. From a firm point of view, in times of very hot or cold periods forward-looking entrepreneurs might decide to undertake adaptation strategies to cope with increasingly frequent temperature extremes, thereby reducing more productive investments.

All in all, the economic effects of temperature oscillations can be diverse, potentially impacting both the consumption and investment components of output, as well as final consumer prices. In what follows, we explain how our shock is constructed, starting from methodological issues and insights, and test its effects on GDP and its sub-components, on prices and interest rates.

3. Constructing the shock

3.1. Key points

Using temperatures to construct a climate-related macro shock is the natural choice, as local temperatures have a direct link with the current global warming trend.[4] However, the advantages of using temperatures go beyond that. For example, while all other climate-related phenomena are rare, temperatures manifest everyday, allowing to compute statistics at high frequency. Moreover, they are collected simultaneously in different places and are easily comparable. We use daily temperatures in each US county to construct a nation-wide shock at the lower (quarterly) frequency.

The shock design is based on the following arguments. First, it rests on the idea that very high ad low temperatures are intrinsically undesirable and, as documented in the literature, generate mostly detrimental effects on the economy. Indeed, while the direction of the impact can be different by season and depend on the type of business exposed, studies recognize that, overall, both extremes are bad for the economy. At the hearth of this outcome lies, for example, the U-shaped relationship between temperatures and mortality documented for the United States by Deschenes (2014) and Barreca et al. (2016), among others. A second key point is that, as science suggests, temperatures turning very hot or cold in a short period of time can be considered as exogenous to current and recent past economic

[4]Other weather events might also act as shocks to the economy, but whether and how they occur because of human-induced climate change is less clear. On the attribution of extreme events to climate change, see Schiermeier (2018).

activity, as feedback effects to local temperatures from human-generated CO2 emissions unravel only in the longer term.

Third, and most important, we argue that while agents can be able to find a workaround to isolated extreme temperature episodes, e.g., by rescheduling working tasks or outdoor activities, exposure to very hot or cold temperatures could become impossible to avoid if these events are frequent enough within a short time span, with negative effects on human and firm activity accumulating over time. This might be so because there are limits to adaptation in the short run, and time constraints (notably in business) that impede to postpone any important activity. For households, although available at very short horizons, daily weather forecasts may in principle help to cope with undesirable temperatures: however, evidence shows that an anticipatory behavior – including taking protective actions – is limited at best (Morss et al., 2010; Graff Zivin and Neidell, 2014).

All these arguments suggest that a shock to temperatures can be inferred by looking at the incidence of multiple extreme temperature episodes within a specific period of time. We choose to work at quarterly frequency as quarters roughly coincide with calendar seasons. Because many economic activities (not only agriculture) follow seasonal patterns, there could be some degree of substitution over time of outdoor tasks – due to extreme temperatures – within seasons, but much less so across seasons, so quarters are ideal time spans to measure temperature surprises. Moreover, quarterly frequency allows to empirically treat calendar year seasonality in temperatures: highs or lows in each seasons can be damaging even if they do not reach extreme levels in absolute terms. For example, the impact of an exceptionally cold summer, while not reaching winter lows, can be nonetheless material for tourism. If this is so, to be surprising extreme realizations need to be evaluated with respect to what the *distribution* of temperatures is expected to be for that period.

In this perspective, climate change causing temperatures to drift upward and oscillations to become wider over time, need to be properly considered. Indeed, the scientific literature agrees on the fact that climate change manifests not only as an upward trend in temperatures, but also as a trend in temperature volatility, leading to an increasing frequency of very hot and cold days. As the entire distribution of seasonal temperatures modifies over time because of that, a reference set to proxy agents' beliefs on current temperatures must relate to past realizations that cannot go too far back in time. One reason for that, from an economic viewpoint, is that agents have memories and learn from their past experience: this form of experience-based Bayesian learning, which takes multiple years, has been documented to drive the dynamics of beliefs also related to global warming (Kelly et al., 2005; Deryugina, 2013; Moore, 2017; Kala, 2019; Choi et al., 2020; Pankratz and Schiller, 2021). In our shock computation, we assume that the reference distribution for each quarter rolls over time, in order to compare current values with updated temperature beliefs. This identification procedure marks a stark difference with all past approaches based on pre-global-warming anomalies or on deviations from historical averages, which implicitly assume that agents anchor their beliefs to some average values and do not update them over time. According to those metrics, observing a larger number of extreme realizations than in the pre-global-warming period would always imply a positive shock, even though agents can be, in fact, used to them: therefore, outdated beliefs may lead to overestimate the size of the current shock. In the following section we provide the formula to recover our shock at local level, as well as the procedure to aggregate it to country level.

3.2. A US wide climate shock

We build our US surprise shock in two steps. First, we construct temperature surprises separately for each county; then, we aggregate them to obtain the nation-level shock. We start by collecting average temperatures in each county at daily frequency. For each series, we group observations by quarter and compare the within-quarter distribution of temperatures to a reference distribution, which is made by pooling daily observations recorded in the same quarter of the past years. We argue that five years is a sufficiently long period for agents to figure out the shape of the underlying temperature distribution, so we construct the reference distribution based on that time span.[5] The reference distribution rolls over time, i.e. it is updated every year for each quarter. In each period, we compute the 10th and 90th percentiles, which are taken as upper and lower thresholds for current temperatures, i.e. the values beyond which actual observations are labeled as very high or low. In order to be perfectly aligned with expectations, we posit that the share of extreme days in current quarter must be the same of that distribution: differently, a larger number of extremes represent a positive surprise, while a lower number makes a negative one.

In formulas, let $T^i_{d,q,y}$ being average daily temperatures in day d, quarter q and year y, recorded in county i, for $i = 1, \ldots, K$ counties. Denote with $F^i_{q,y} = \{T^i_{d,q,y-j},\ \text{j} = 1, \ldots, 5\}$ the empirical cumulative distribution function of the reference temperature distribution for quarter q and year y, totalling $N_{q,y}$ days. The reference temperature values are

$$\text{ut}(q)^i_y = F^{-1}_{q,y}(0.9) \qquad \forall \quad q = \{\text{q1}, \text{q2}, \text{q3}, \text{q4}\} \tag{1}$$

$$\text{lt}(q)^i_y = F^{-1}_{q,y}(0.1) \qquad \forall \quad q = \{\text{q1}, \text{q2}, \text{q3}, \text{q4}\} \tag{2}$$

where $\text{ut}(q)^i_y$ and $\text{lt}(q)^i_y$ are yearly series for upper (lower) thresholds for county i, quarter q and year y. By combining threshold values together by quarter, we get quarterly series of upper and lower thresholds

[5] Pankratz and Schiller (2021) test learning periods of five, ten, and fifteen years length. As a robustness check, we construct an alternative version of the shock using a 10-year learning window as in Choi et al. (2020), see Section 6.

for each county, i.e.

$$\mathrm{ut}^i_{t,y} = \{\mathrm{ut(q)}^i_y,\ q = \mathrm{q}1, \ldots, \mathrm{q}4\}$$

$$\mathrm{lt}^i_{t,y} = \{\mathrm{lt(q)}^i_y,\ q = \mathrm{q}1, \ldots, \mathrm{q}4\}$$

The same can be done with the size of the reference distribution, yielding $\mathrm{N}_t = \{\mathrm{N}_{q,y},\ q = \mathrm{q}1, \ldots, \mathrm{q}4\}$. Note that we change notation for the quarter (from q to t) as q denotes quarters in yearly series, while t indicates quarterly frequency. Denoting with n^i_t the number of days in quarter t, county surprises are then evaluated as the number of beyond-threshold days in the current quarter *in excess* of those in the reference distribution. Provided that, by construction, the share of extreme days in the reference distribution is 20 percent, county surprises can be expressed as

$$\text{county\_surprise}^i_{t,y} = \sum_{j=1}^{n^i_t} \Big[\mathrm{I}(T^i_{d,t,y} < \mathrm{lt}^i_{t,y}) + \mathrm{I}(T^i_{d,t,y} > \mathrm{ut}^i_{t,y})\Big] - N_{t,y} \times 0.2 \tag{3}$$

where $I(x)$ is an indicator function that values 1 if x is true, 0 otherwise. County surprises are the difference between the number of extremely hot and cold days in quarter t and the number of extreme days in the reference distribution for that quarter. The underlying idea is that the reference distribution, which updates quarterly, represents the information set of economic agents, who directly experienced a range of temperatures for that season in the recent past. As they have no reason to foresee some change in the distribution with respect to the very recent past, agents expect that the same number of hot and cold days also occur in current quarter. In this setting, temperatures go beyond expectations if the number of extreme days exceed (or is below) what agents expect.

In order to make a US-wide surprise shock, we make a weighted average of county-level surprises occurred in the K counties by quarter:

$$\text{US\_shock}_{t,y} = \sum_{i=1}^{K} \Big(\text{county\_surprise}^i_{t,y} \ \times \ \mathrm{w}^i_{y-1}\Big) \tag{4}$$

where w are county-level weights, proxying counties' vulnerability to temperatures, which vary at annual frequency. Weights are lagged to capture the ex-ante exposure to temperatures.

Example. Imagine to compute the US-wide temperature shock for q2 (spring) of a generic year (numbers are fictitious). Suppose the US is divided in only two large territories, West (40% population) and East (60%). Assume that, considering spring time in the past 5y, the 10th percentile of the reference temperature distribution for the West – made of 90 days × 5 years – is 40 F, while the 90th percentile is 60 F. In spring this year, West had 8 days with average daily temperature below 40F, and 15 days above 60F. As norm in a single quarter should be 9 days per tail, the temperature surprise for the West is quantified as

$$\text{shock}_W = (8 - 9) + (15 - 9) = 5 \text{ days}$$

if surprise in East results, instead, to be 2 days, the US-wide shock is

$$\text{shock}_{US} = 5 \times 0.4 + 2 \times 0.6 = 3.2 \text{ days}$$

3.3. Insights

The constructed shock presents the following characteristics. First, it is designed to have an unambiguous economic interpretation: a positive one, meaning an unexpected increase in the occurrence of abnormal daily temperatures, should have a negative economic impact, while a negative one (milder temperatures than expected) should be beneficial. Second, it is season-specific, as surprises are measured within the same season. Therefore, the largest surprises during the year may not necessarily come in winter and summer when temperatures show record lows and highs, as what matters is how temperatures deviate from their seasonal norm. Third, as the shock is measured with respect to updated expectations, it nets out long-run phenomena such as the intensification of climatic trends – upward-drifting and increasing volatile temperatures – and adaptation to extreme temperatures: agents learning about the evolution of temperatures can continuously adjust their resilience to them. Last, the shock is two-sided and surprises coming from extremely hot and cold days are treated as equivalent. This can be considered as a neutral assumption, as taking a stance on which tails matter more in each season is not easy (Addoum et al., 2021). In a robustness exercise, we decompose the US-wide shock into heat and cold shocks and explore their effects separately using local projections, finding that the response of economic variables to them are comparable in size.

Overall, the shock measures the size of the variation in the distribution of temperatures in the short run. Being an aggregation of county-level surprises, it will be largest in absolute value the largest temperature variations will be in highly relevant US counties. In this respect, the way surprises are aggregated could also make the difference: for this purpose, the economic vulnerability of each place to temperature shocks, which can itself vary over time, should be properly considered.

4. Data

4.1. Temperature data

To construct our shock, we rely on two data sources. The first one is the gridded air temperature database for continental US taken from the Northeast Regional Climate Center. From that source we extract daily temperatures, i.e. the mean value of temperatures recorded during the 24 hours, averaged at county level. We consider data starting from Jan 1, 1970, i.e. the beginning of the human-induced global warming trend, and ending on Dec 31, 2019. To construct the weights to aggregate county-level surprises, we take annual data on US counties' economies from the Census Bureau, from 1969 to 2018. In our baseline formulation of the shock, we consider county-level population and construct weights as the county share over nation-wide values. The rationale for this choice is that the higher the population, the higher the incidence of agents that are exposed to extreme temperatures, so the higher the potential economic impact. As a robustness, we also consider land extension – used as an alternative to population in Colacito et al. (2019) – or the number of employed people and personal income to weight temperature surprises: while the first two are also proxies of counties' temperature exposure, personal income might instead proxy the vulnerability of the local economy to temperature shocks. Merging the two datasets yields time series of temperatures and weights for 3053 counties. This sample is highly representative of the US economy, covering more than 98% of the country's population, jobs and personal income as of 2019.

4.2. County-level statistics

Figure 1 represents county-level statistics over geographic maps of the United States. Panel 1 shows the growth rate of temperatures in the years 2010s with respect to the 1960s. On average, temperature grew by 1.7%, but climatic trends have been quite diverse, with some counties experiencing temperatures rising by more than 8%, while others in which temperatures decreased by 5%. Overall, northeastern counties and the west coast experienced the steepest temperature increase. The picture looks different when evaluating the size of county-level surprises occurred over time, which depend on the evolution of the entire temperature distribution: surprise shocks in northeastern counties have been small on average, while they have been very large in the south (Panel 2). Overall, what matters is the combination of shocks and weights: even small surprises can be important if occurred in highly-populated counties. Panel 3 shows counties' population in 2019. At that time, the 20 most populated counties, covering about 20 percent of US population, were mostly in California, Texas, New York and Florida.

In addition to having no one-to-one correspondence with local temperature trends, surprises do not necessarily pair with a high incidence of temperature extremes. Figure 2 shows the histogram of the correlation coefficient, by county, of the surprise series with the count of days with temperatures above 80F or below 20F in each quarter. Within country correlation has a mode between 0.1 and 0.2 and is rarely above 0.5, confirming that surprises and extremes are generally uncorrelated. To confirm that we are correctly capturing unexpected events even at the very short horizon, we conduct an illustrative exercise for the Washington DC county (for ease of computation) to evaluate how temperature predictability is related to realized temperature levels. Results, summarized in Appendix B suggest that days that are labeled as surprisingly hot or cold under our procedure are also associated with largest ex-post forecast errors at one- and two-day horizons.

To have an idea of how surprises at quarterly frequency evolve over time, we compare those occurred in two of the richest counties in the US, New York County and Los Angeles County, see Figure 3. The two surprise series look substantially different – correlation is approximately null – showing peaks of positive surprises in different periods: 1989q4 and 2003q1 in NY and 1978q4, 2001q1 and 2015q1 in LA. Also, those peaks reflect episodes that are different in nature: all the aforementioned events in New York were mainly due to exceptionally cold falls and winters while, in Los Angeles, heat- and cold-related surprises have been more balanced (e.g., temperatures in winter 2015, on account of a very warm March, mainly surprised on the upside).

Figure 1: **Within country comparison**

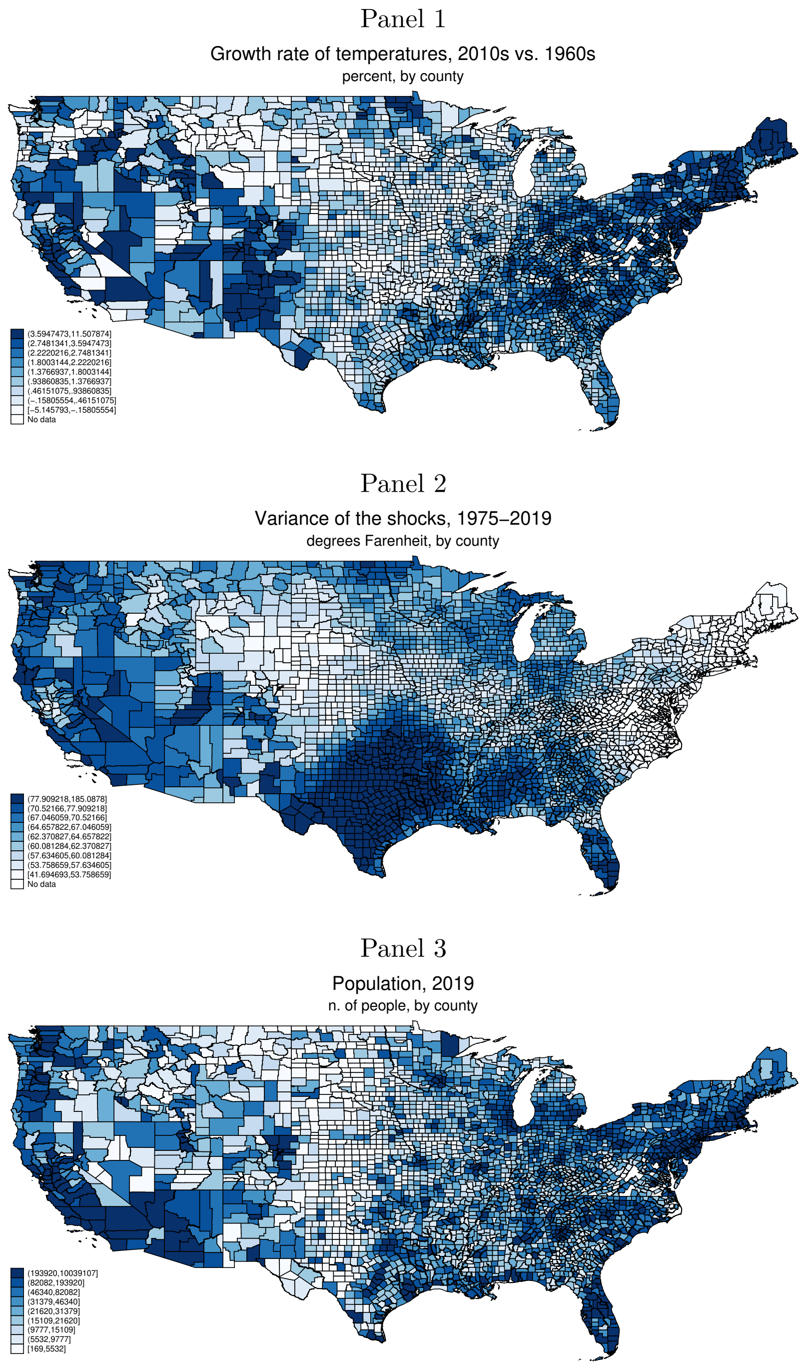

Notes: The figure shows ...

Figure 2: **Surprises vs incidence of extremes**

Time series correlation, by county

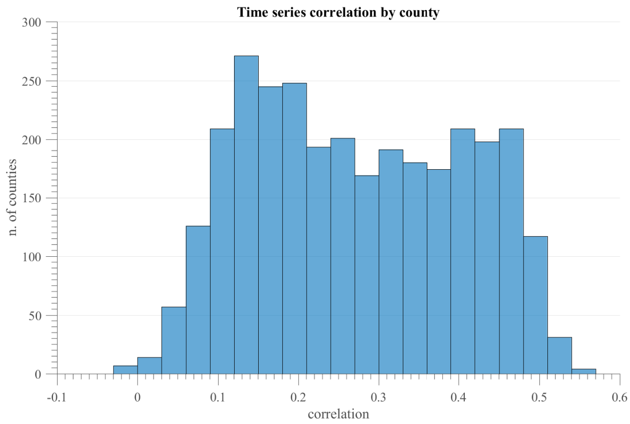

Notes: The histogram reports the correlation of the series of county-level surprises with that of temperature extremes, proxied by the number of days above 80F and below 20F in each quarter.

Figure 3: **Surprises in NY vs. LA**

New York County, NY | Los Angeles County, CA

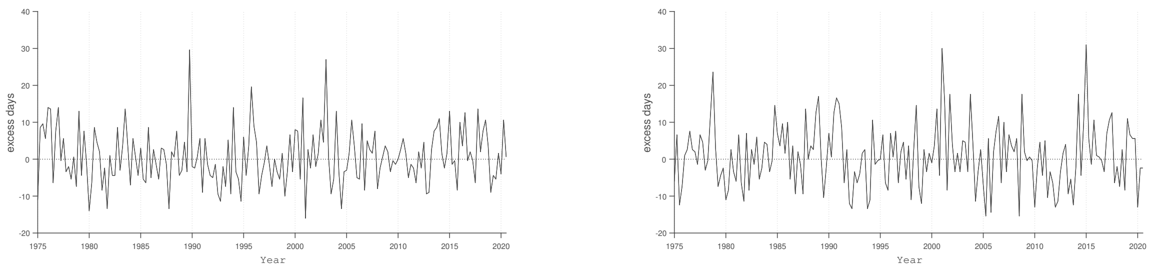

Notes: NY County (FIPS code 36061) and LA County (FIPS code 06037).

4.3. The shock series

The time series of the US-wide shock is displayed in Figure 4. The shock is expressed as the total number of surprisingly hot and cold days per quarter, and reaches a minimum at -9 and a maximum at 14. The shock embeds the basic econometric requirements to be used in empirical analysis: it is zero-mean and serially uncorrelated (see Appendix A). From an economic point of view, it features the three characteristics, explained in Ramey (2016), which make it suitable for macroeconomic applications: it is exogenous with respect to current and lagged outcome variables, it can be considered as uncorrelated with other exogenous economic shocks and it represents unanticipated movements in an exogenous variable.

Figure 4: **US-wide temperature shock**

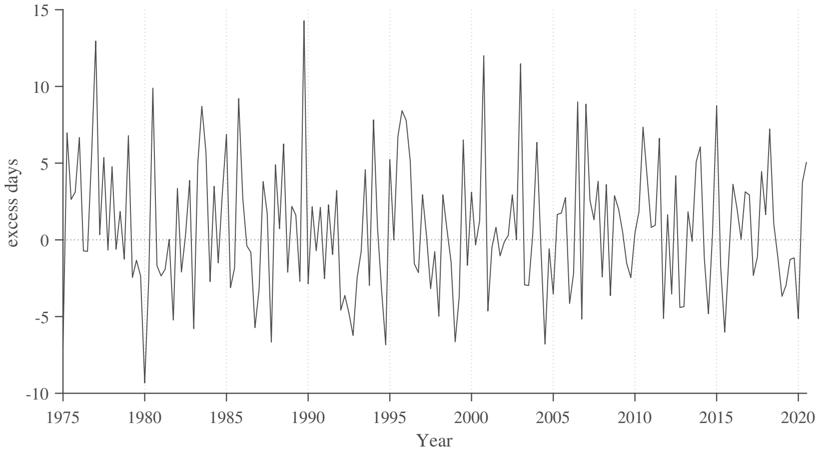

Looking at the plotted series, one thing that catches the eye is that temperature shocks have not been particularly large in the last 15 years with respect to the earlier period. Indeed, the volatility of the shock series computed on a rolling 10-year window slightly decreases throughout the sample, from a peak of 5 to about 3.5 days. It suggests that adjustments in the shape of the temperature distribution have been largest − inducing greater surprises − in the early phase of global warming than in recent times. This evidence reverses the common wisdom that climate change is generating shocks of increasing size. In fact, our shock dynamics is not in contrast with the intensification of climate-related weather events (including temperatures extremes), as climate indicators such as the Actuaries Climate Index show.[6] As we take out predictable climatic trends to capture the unexpected component of temperature fluctuations, it is well possible that large weather episodes with respect to the pre-global-warming period have become increasingly frequent, but also less surprising than in the past. This higher predictability speaks to the evidence of an upward trend in annual temperature volatility found for a panel of countries in Alessandri and Mumtaz (2021), over which the US economy may have adapted over time.

The largest positive values of our shock are recorded in 1977q1, 1989q4, 2000q4 and 2003q1. As the New York Times described, 1977 was "a year of weather extremes", which spurred new climate research in the United States and worldwide (Sterba, 1977). Analyses made at that time by the National Weather Service reported that winter in 1977, especially in the eastern part of the country, was one of the coldest

[6]The Actuaries Climate Index is a composite indicator of the frequency of a set of climate-related natural events including extreme temperatures, precipitation, wind, drought and sea level rise (American Academy of Actuaries, Canadian Institute of Actuaries, Casualty Actuarial Society and Society of Actuaries, 2020). The components are constructed as anomalies, i.e. in difference with respect to a fixed reference period in the past (1961-1990).

in the 20th century, with temperatures close to the record lows of 1917-18. The first quarter of 1989 also stood apart for its exceptionally low temperatures. Surges of Arctic air in December was, according to the National Weather Service, "a historic event, with many locations establishing monthly or all-time record lows" (National Weather Service, 1989). While the fall in temperatures was strongest in the northwest, the southeastern coast was also particularly damaged, with losses caused by damaged crops in Florida and broken water pipes in Texas causing failures at manufacturing plants. In that area, the 1989 cold wave also generated the largest snowstorm in recent history. The fourth quarter of year 2000 saw temperatures passing from one extreme to the other. The late summer period was characterized by a prolonged heat wave which, particularly, in southern states, resulted in record high temperatures. Then, after a warm October, temperatures plunged in November, standing as the second coldest since records began in 1895, according to NOAA's National Climate Report. This turnaround in temperatures caused power outages and triggered other weather-related extreme events, such as tornadoes. The fourth peak in the series is that of winter of 2003. That period will be mainly remembered in northern states such as Minnesota and Michigan. While it was not among the coldest on record, it stood out for its persisting cold, which started in December and went through mid-March. In particular, Michigan recorded the second longest continuous streak of very cold temperatures (76% of days in January and February) on record, just one notch below the exceptionally cold winter of 1977.

All the four aforementioned peaks were mainly due to abnormally cold periods. However, heat-related surprises have also been frequent in the history of US temperatures. To have a flavor on how both extremes contribute in shaping US-wide shocks, we inspect the incidence of surprisingly hot and cold days in the shock series within each quarter. Figure 5 displays positive US shocks in a bar plot where observations are labeled as orange or blue depending on whether surprisingly hot or cold days prevailed at each point in time. Abnormally hot and cold quarters have been almost equally frequent, with the former being slightly more (55% of times). Moreover, their incidence have been mostly balanced throughout the sample, with no apparent clusters in the most recent period. This finding also dispels the myth that, in a global warming era, abnormally hot days largely predominate over cold ones. In Section 6, we break down the US-wide shock into heat and cold shocks and estimate the economic effect of each component separately, finding comparable albeit not equal effects.

Figure 5: **Prevailing heat vs. cold surprising days**

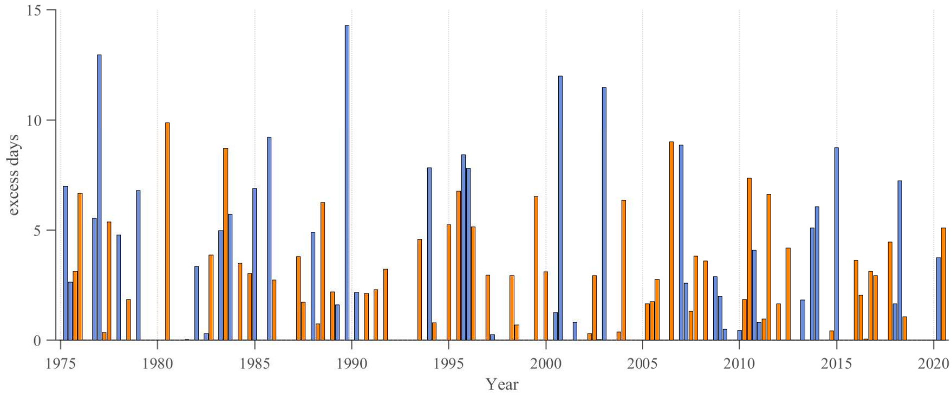

Notes: Labeling of positive shocks based on the number of surprisingly hot and cold days within the quarter. Orange bars represent quarters in which surprisingly hot days predominate, while blue bars denote those in which there are relatively more surprisingly cold than hot days.

5. Application: the impact on the US economy and monetary policy

In this section we employ our US-wide shock to evaluate the effect of temperature surprises on the US economy. In the following, we describe our approach and comment on the main findings.

5.1. Impulse-response analysis

We estimate the response of US domestic variables to the previously constructed temperature shock using the local projections framework of Jorda (2005). Impulse response functions (IRFs) are obtained from the following linear regressions:

$$y_{t+s} = \alpha_s + \beta_s \text{ US\_shock}_t + \psi_s(L)\mathbf{X}_{t-1} + u_{t+s} \qquad s = 0, 1, 2, \ldots, H \tag{5}$$

where t are quarters, y is the target variable and $\mathbf{X}_t$ is a vector of controls.[7] Estimation is made separately for each time horizon s and for each dependent variable. IRFs are defined by the sequence $\{\beta_s\}_{s=0}^{H}$, and inference is performed with Newey-West standard errors.

5.2. Target variables

In our baseline estimates, we test the effects of temperature shocks on the following set of domestic dependent variables: GDP, private consumption and investment (all in real terms); the CPI index, 3-

[7]Note that, with respect to previous notation, we here suppress subscript y to indicate years.

month interest rate and the 10-year Treasury yield. All variables except the 3-month rate and the 10-year yield are expressed in natural logarithm. Among control variables, we include linear and quadratic time trends, seasonal dummies, four lags of the shock and of the aforementioned variables and the 2-year less 3-month yield to control for the impact on the short-term portion of the yield curve. A summary of all data used in this paper is reported in Appendix C. Impulse responses are estimated up to 16-quarter horizon, and displayed with 68% and 90% confidence bands. The estimates are carried out on the sample going from 1975 Q1 to 2019 Q4.

5.3. Results

Impulse response functions from one a standard deviation shock are displayed in Figure 6. In response to a positive temperature surprise, real GDP in the United States significantly decreases. As it is common using local projections, the impact slowly builds up reaching a trough after 2 years. Both private consumption and investment shrink, with investment resulting much more impacted – response is five times larger at the trough. The CPI index falls too, but its variation is not significant. In line with the detrimental economic effects, short-term interest rates also decline, suggesting an expansionary monetary policy response. This effect pass through the long-end of the government yield curve, as the 10-year yield also decreases. A slight above-zero overshooting of the responses is visible four years after the shock, but it never looks significant.

In terms of size, the economic effects can be quantified as follows: a positive shock in a single quarter by one standard deviation, which is equal to 4.3 days (5% of the days in the quarter), entails a fall in real GDP by 0.34% in about two years. This effect speaks to the impact found in Deryugina and Hsiang (2014), where a single hot day generates a reduction by 0.065% in annual income.[8] Moreover, it implies a decrease in real private consumption by 0.35% and in real investment by 1.7%. Regarding interest rates, it generates a fall in the 3-month rate by 25 basis points, and in the 10-year yield by 18 bps at the same horizon. This effect stands out in the literature, as the response of US interest rates to extremely hot and cold temperatures has recently been found as insignificant (Kim et al., 2021). In fact, the two approaches capture different aspects of climate change, as we identify an unexpected surprise shock on the entire temperature distribution taking out climatic trends.

[8] 4.3 days would cause a reduction in annual income by 0.28%, very close to our 0.34%.

Figure 6: **Baseline estimates**

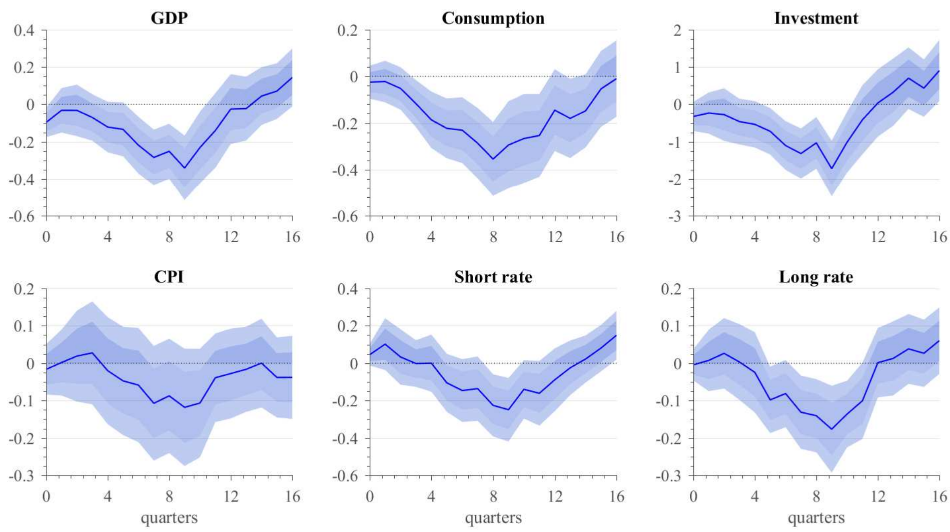

Notes: 68% and 90% confidence bands.

5.4. Shock transmission over time

The impacts found in the previous estimates should be interpreted as the average economic effects of the shock throughout the sample period. However, as climate manifestations are the more and more visible and knowledge about them becomes more widespread, the economic impact to temperature surprises may have also changed since the early global warming phase. To gauge such potential time variation, we repeat local projections on sub-periods, focusing on the effects on GDP, CPI and the 3-month rate only. We set a trailing sample of 124 quarters (31 years) and perform new estimates by rolling it by one quarter at a time. In this way, we get 57 sets of estimates, which begin from the sample period 1975q1-2005q4 and end with the sample 1989q1-2019q4. To summarize the impulse responses obtained at each iteration, we display for both variables the IRFs at the horizon displaying the largest negative impact (trough) in the baseline estimates; we also display the response at the trough reached at each iteration to capture possible changes in the dynamics of the effects.[9]

Figure 7 shows the results for GDP (Panel a and b), CPI (Panel c and d) and the short rate (Panel e and f), where pointwise median responses are displayed together with their 68% and 90% confidence bands. The estimates suggest two interesting facts. First, the elasticities of GDP and CPI to temperature shocks have slightly reduced in size over time: after reaching -0.47% at the trough, the response of GDP almost halved to a maximum (negative) impact of -0.25% (panel b); the response of CPI passed from being -0.27% to about -0.09%, before widening again to -0.16% (panel d). This reduced impact in both

[9]IRFs computed on different samples may reach the same trough but at different horizons.

variables reflects the fact that the increasing adoption of heating and cooling systems in houses and firm establishments and, more generally, the adaptation measures put in place to safeguard from temperature oscillations have made agents more resilient to shocks than they were 50 years ago. Second, the effects on GDP and CPI are substantially different between them, as the negative impact on GDP remained always significant throughout the whole period, while that of CPI is not significant during part of it. This may suggest that demand- and supply-side effects balance out in some part of the sample, muting price effects. Moreover, the response of CPI appears to be less significant in Panel (c) than in Panel (d), suggesting that the trough effect may have moved ahead (i.e, the pass-through to lower prices have become slower) over time. This is coherent with the idea that demand-side effects of temperatures are lagged and affect prices with a delay (Faccia et al., 2021). Panel (e) and (f) show the response of the short-term rate. In line with the slight reduction in the impact on GDP and prices, also the effect on interest rates diminishes throughout the sample. Overall, interest rates fall in response to the shock, suggesting an expansionary monetary policy response to sustain the economy. However, this response seems to be absent in the last part of the sample, where the variation in the short rate is barely or even not significant. This happens despite the effects on GDP and CPI continue to be material.

Figure 7: **Maximum negative impacts on GDP, CPI and 3-month rate, rolling estimates**

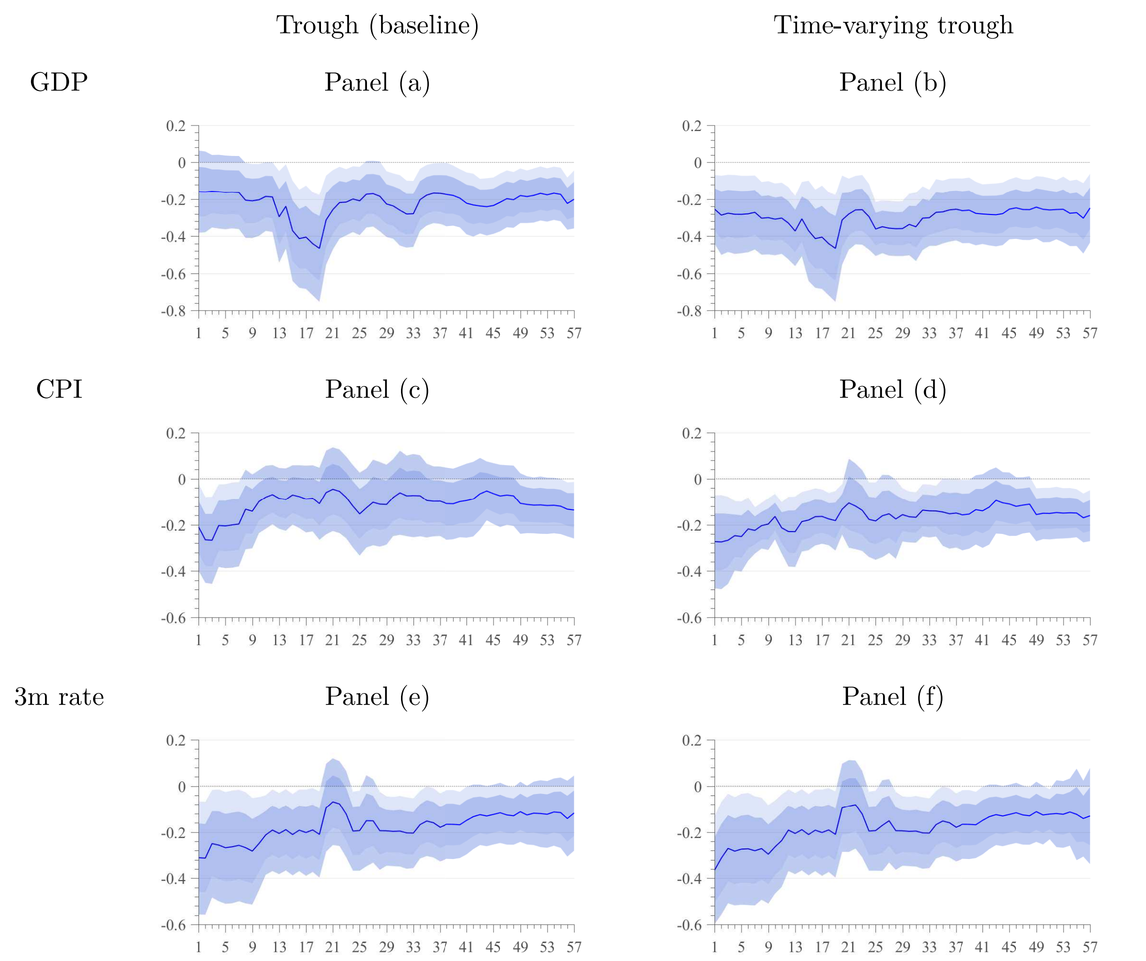

Notes: .

5.5. *Impact on Fed's economic forecasts*

The impact of temperature shocks on the short-term rate observed in Figure 6 and 7 suggests that, on average, the Fed does react to the induced economic impact by lowering policy rates. A natural question that arises is why this is so: do temperature shocks affect central bank's view on the economy for present and future quarters? Is that view aligned with the economic outcomes? To have insights on this, we explore whether Fed's economic forecasts are affected by temperature shocks. We consider Greenbook Forecasts for GDP and CPI, which are available to the Federal Open Market Committee prior to each meeting, and re-run local projections on those outcome variables. Greenbook Forecasts are available at quarterly level between 1980 and 2015. We investigate the effects of temperatures on current quarter nowcasts, as well as on the forecasts for the following two quarters.

Figure 8 displays the effects on Fed's forecasts, compared with the effects on actual GDP and CPI in the same sample period. The Fed reacts to the shock by adjusting both its nowcasts and short-term forecasts for both variables. Regarding GDP, it predicts a fall in current quarter that is in line with the actual fall, but it overestimates the effects on the subsequent quarters. On CPI, the impact in current quarter is predicted to be initially positive, then turning negative three years after the shock.

Figure 8: **Effects on Fed's Greenbook Forecasts**

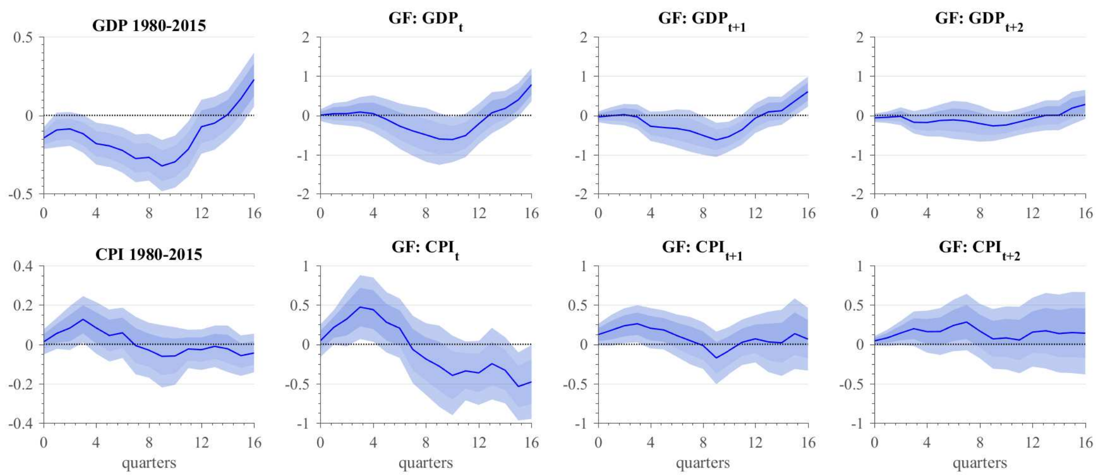

Notes: 68% and 90% confidence bands.

5.6. Attention to temperatures and climate

Does the Fed recognize climate and temperatures as a source of economic shocks? To investigate this issue, we collect the transcripts of all FOMC meetings historically available (i.e., from 1976 to 2015) and analyze wording therein. We define three sets of climate-related terms: a wording related to the climate change phenomenon (*Climate change wording*); a wording related to temperatures and strictly related phenomena (*Temperatures wording*); a wording related to climate-related non-temperature events (*Natural disasters wording*). For details on single words, couples and triples included in each list, see Appendix D. Figure 9 displays the word count over time for each set of words. The first panel reveals that the FOMC has explicitely mentioned the climate change phenomenon very rarely (only 7 times since 1976), but temperatures and natural disasters much more frequently. Between them, most attention has been devoted to temperature oscillations. Overall, evidence suggests that during the global warming era, climate change has rarely been a hot topic in official central bank's conversations, at least before 2016 (temperatures had mentioned, on average, just 4 times per quarter).[10]

[10]Nonetheless, this attitude could have changed in the last years as climate change debate increased strongly in popularity.

Figure 9: **Climate mentions by FOMC members**

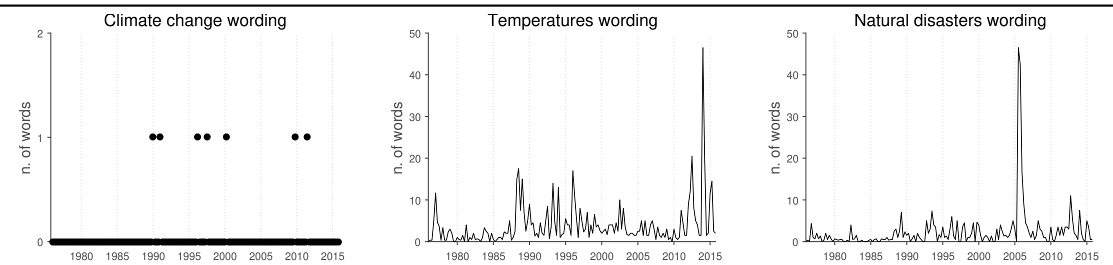

Notes: 68% and 90% confidence bands.

Do shocks stimulate more discussion on the topic within the FOMC? We re-run local projections on the count of words related to temperatures and natural disasters. According to Figure 10, it turns out that the Committee slightly increases mentions of temperatures after shocks. This increase is tiny – just 2 words against a one-standard deviation temperature shock. Differently, shocks to temperatures do not stimulate other climate-related debate such as that on natural disasters.

Figure 10: **Effect on Fed's climate-related wording**

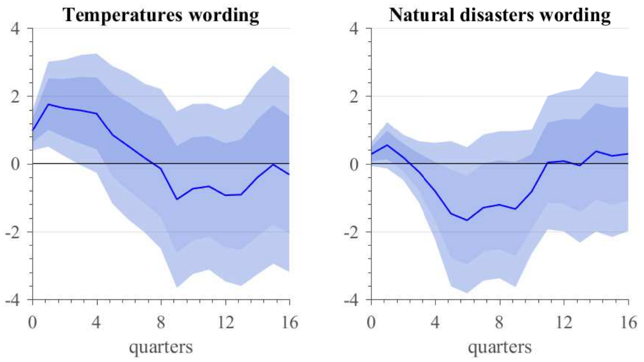

Notes: 68% and 90% confidence bands.

6. Robustness exercises

In this Section we perform a number of robustness exercises to confirm our baseline findings, based on different specifications of our temperature shock, of control variables and of the econometric specification.

6.1. Heat vs cold shocks

The baseline temperature surprise shock is constructed by summing left and right tail surprises in each quarter. In this section we dissect the economic effects of these two components by including them separately in local projection estimates. According to Equations 3 and 4, left- and right-tail shocks, which we call cold and heat surprise shocks, can be retrieved as

$$\text{US\_shock}_{t,y}^{b} = \sum_{i=1}^{K} \left(\text{county\_surprise}_{t,y}^{i,b} \ \times \ \text{w}_{y-1}^{i} \right) \qquad b = \{\text{cold}, \text{heat}\} \tag{6}$$

where

$$\text{county\_surprise}_{t,y}^{i,\text{cold}} = \sum_{j=1}^{n_t^i} \text{I}(T_{d,t,y}^{i} < \text{lt}_{t,y}^{i}) - N_{t,y} \times 0.1$$

$$\text{county\_surprise}_{t,y}^{i,\text{heat}} = \sum_{j=1}^{n_t^i} \text{I}(T_{d,t,y}^{i} > \text{ut}_{t,y}^{i}) - N_{t,y} \times 0.1$$

The two components, which sum up to the baseline shock in each quarter, are plotted in Figure Appendix E.2. We then employ those variables into linear local projections as

$$y_{t+s} = \alpha_s + \gamma_s \ \text{US\_shock}_t^{\text{cold}} + \delta_s \ \text{US\_shock}_t^{\text{heat}} + \psi_s(L)\mathbf{X}_{t-1} + u_{t+s} \qquad s = 0, 1, 2, \ldots, H \tag{7}$$

where γ_s and δ_s are impulse response functions to cold and heat shocks, respectively, and X is the set of control variables. Figures Appendix E.3 and Appendix E.4 show the response to the two shocks of all the variables considered in the baseline estimates. The economic impact of heat and cold shocks looks overall quite similar, although the response to cold shocks is somewhat larger (especially for GDP, consumption and investment) and slightly more significant. The difference is starker for the response of CPI, which is negative and even barely significant for cold shocks while it hovers around zero in case of heat shocks.

6.2. Impacts by season

How does the impact of heat and cold shocks vary by season? For this purpose, we estimate the following variant of Equation 7 :

$$y_{t+s} = \alpha_s + \sum_{j=1}^{4} \left(\gamma_s^j \ \text{US\_shock}_t^{\text{cold}} \times D_t^j + \delta_s^j \ \text{US\_shock}_t^{\text{heat}} \times D_t^j \right) + \psi_s(L)\mathbf{X}_{t-1} + u_{t+s} \qquad s = 0, 1, 2, \ldots, H \tag{8}$$

[10]The set of controls is the same we employed in the baseline estimate. Four lags of the shocks are substituted here with four lags of both heat and cold shocks.

where D_t^j with $j = \{1, \ldots, 4\}$ are four dummies that equal 1 if current quarter is quarter j of the year, 0 otherwise. As we interpret quarters as calendar seasons, γ_s^j and δ_s^j represent impulse response functions to cold and heat shocks in winter (q1), spring (q2), summer (q3) and fall (q4). Figure Appendix E.5 and Figure Appendix E.6 represent IRFs for heat and cold shocks by season, which are displayed for GDP, CPI and short rate only to save space.[11] Overall, the two figures show that the impact of heat and cold shocks are generally more pronounced in summer than in other seasons. Regarding heat shocks, summer shocks have a stronger impact on GDP and the short rate than shocks in other seasons; the sign of the response of CPI vary substantially across seasons, appearing to be positive in winter and fall. Instead, cold temperature shocks have a detrimental economic effect not only in summer, but also in winter, while cold shocks in spring and fall seem to play a minor role.

6.3. Surprise thresholds

The temperature surprise shock is constructed in each quarter with respect to some threshold values, which are the 10th and 90th percentiles of the reference distribution. We propose alternative specifications of the shock in which we enlarge or restrict the definition of surprising days by setting threshold values first at the 5h and 95th percentiles, and then at the 25th and 75th percentiles of the reference distribution. Figure Appendix E.7 shows the median responses for our 6 variables of interest in these two cases, together with the baseline IRFs. Impulse responses of the two alternative specifications look very similar to the baseline, with shocks computed with larger tails having a slightly stronger impact. In that case, the effect on the CPI index is more persistent than in the baseline. Correlations between baseline US shock and the two alternatives are about 96%.

6.4. Weighting scheme

We here compute alternative formulation of the US-wide shock using counties' employment, personal income and land extension instead of population shares to aggregate county-level surprises. The median impulse responses obtained from these alternative shocks are displayed in Figure Appendix E.8. Using employment share and personal income share leave results unchanged with respect to the baseline. Using land weights change responses a bit more, although remaining broadly in line with baseline results.

6.5. Learning period

In the baseline specification we construct reference distributions by aggregating temperatures observed in the five years prior to the shock. Here, we assume that learning takes longer and construct those

[11] IRFs to heat and cold shocks by season of all the other variables in the baseline are available upon request.

distributions over 10-year rolling windows. The shock, plotted in Figure Appendix E.9, is very similar to the baseline version (95% correlation).

6.6. Control variables

Baseline local projections are constructed using a rich set of controls, including four lags of all of the six variables. We explore the effects of temperature shocks using a reduced set of controls by reducing lags of all control variables (except the shock) to two instead of four. Figure Appendix E.10 shows that estimates are robust to the choice of the lag structure, as impulse responses remain almost unchanged.

6.7. One-sided shock

The baseline shock is two-sided under the assumption that both extremes are bad for the economy. One may take a different stance and consider extreme days in one tail only. In this case, the underlying assumption is that temperatures in the other tail should be considered as beneficial for the economy as mild temperatures are. While this assumption looks quite strong, we construct an alternative, one-sided version of the shock by considering surprises only for the tail that is supposed to count more (in negative terms) in each season. Based on the literature (Colacito et al., 2019; Addoum et al., 2021), we define positive surprises as number of days in excess on the left tail (cold days) for all seasons but summer, for which we consider surprises as excess days in the right tail. This alternative formulation of the shock is reported in Figure Appendix E.11. Local projections carried out using this version of the shock, shown in Figure Appendix E.12 yield results that are in line with the baseline, with slightly smaller coefficients (e.g., the response of real GDP is -0.29%).

6.8. Asymmetric responses

Have shocks a symmetric impact on the economy? We investigate this issue by running state-dependent Local Projections, where the state variable is a dummy equal to 1 for positive shocks (a higher number of abnormal daily temperatures than expected), 0 otherwise (a lower number than expected). In the spirit of Iacoviello and Navarro (2019), we adopt a partial state-dependency specification:

$$y_{t+s} = \alpha_s + \beta_{1s}\ \text{US\_shock}_t \times d_t + \beta_{2s}\ \text{US\_shock}_t \times (1 - d_t) + \psi_s(L)\mathbf{X}_{t-1} + u_{t+s}$$

where d_t is the dummy that equals one in state 1 (positive shocks), 0 otherwise (negative shocks). Responses, displayed in Figure Appendix E.13, show that only positive temperature shocks (i.e., very hot and cold days going beyond expectations) affect GDP and the short rate, while negative ones (temperatures being milder than in the past) do not significant improve activity nor prices, with no effect on the short-term rate.

7. Conclusions

Climate change has multiple effects on the global economy. We propose a method to isolate exogenous temperature surprise shocks in one country based on the incidence of relatively high and low temperatures in excess to what agents expect in each season. Estimates made with local projections show that temperature surprise shocks in the United States have been significantly hurting the economy, and that the Federal Reserve has been reacting to the deteriorating environment by lowering short-term rates, slightly increasing its attention to temperature oscillations after the shocks. All in all, results show that temperatures are an autonomous source of macroeconomic variation, adding another piece of evidence in the debate on the needed policy response to climate change. The shock constructed in this paper can be replicated for other countries and at wider or narrower geographic dimension, and can serve as reference to build other weather-related shocks under the same logic.

References

Acevedo, S., Mrkaic, M., Novta, N., Pugacheva, E., Topalova, P., 2020. The effects of weather shocks on economic activity: What are the channels of impact? Journal of Macroeconomics 65, 103–207.

Addoum, J.M., Ng, D.T., Ortiz-Bobea, A., 2020. Temperature Shocks and Establishment Sales. The Review of Financial Studies 33, 1331–1366.

Addoum, J.M., Ng, D.T., Ortiz-Bobea, A., 2021. Temperature Shocks and Industry Earnings News. Mimeo, available at SSRN: https://ssrn.com/abstract=3480695.

Alessandri, P., Mumtaz, H., 2021. The macroeconomic cost of climate volatility. Papers 2108.01617. arXiv.org.

American Academy of Actuaries, Canadian Institute of Actuaries, Casualty Actuarial Society and Society of Actuaries, 2020. Actuaries climate index .

Baker, J., Cameron, M., 1996. The Effects of the Service Environment on Affect and Consumer Perception of Waiting Time: An Integrative Review and Research Propositions. Journal of the Academy of Marketing Science 24, 338–349.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., Shapiro, J.S., 2016. Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. Journal of Political Economy 124, 105–159.

Burke, M., Hsiang, S., Miguel, E., 2015. Global non-linear effect of temperature on economic production. Nature 527, 235–239.

Cachon, G.P., Gallino, S., Olivares, M., 2012. Severe weather and automobile assembly productivity. Columbia Business School Research Paper No. 12/37, Available at SSRN URL: http://dx.doi.org/10.2139/ssrn.2099798.

Cashin, P., Mohaddes, K., Raissi, M., 2017. Fair weather or foul? the macroeconomic effects of el nino. Journal of International Economics 106, 37–54.

Choi, D., Gao, Z., Jiang, W., 2020. Attention to Global Warming. The Review of Financial Studies 33, 1112–1145.

Colacito, R., Hoffmann, B., Phan, T., 2019. Temperature and growth: A panel analysis of the united states. Journal of Money, Credit and Banking 51, 313–368.

Custódio, C., Ferreira, M.A., Garcia-Appendini, E., Lam, A., 2021. Economic Impact of Climate Change. Mimeo, available at SSRN: https://ssrn.com/abstract=3724940.

Dell, M., Jones, B., Olken, B., 2012. Temperature shocks and economic growth: Evidence from the last half century. American Economic Journal: Macroeconomics 4, 66–95.

Dell, M., Jones, B.F., Olken, B.A., 2014. What do we learn from the weather? the new climate-economy literature. Journal of Economic Literature 52, 740–98. URL: https://www.aeaweb.org/articles?id=10.1257/jel.52.3.740, doi:10.1257/jel.52.3.740.

Deryugina, T., 2013. How do people update? The effects of local weather fluctuations on beliefs about global warming. Climatic Change 118, 397–416.

Deryugina, T., Hsiang, S.M., 2014. Does the Environment Still Matter? Daily Temperature and Income in the United States. NBER Working Papers 20750. National Bureau of Economic Research, Inc.

Deschenes, O., 2014. Temperature, human health, and adaptation: A review of the empirical literature. Energy Economics 46, 606–619. URL: https://www.sciencedirect.com/science/article/pii/S0140988313002387, doi:https://doi.org/10.1016/j.eneco.2013.10.013.

Deschênes, O., Greenstone, M., 2007. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. American Economic Review 97, 354–385.

Donadelli, M., Jüppner, M., Paradiso, A., Schlag, C., 2021. Computing macro-effects and welfare costs of temperature volatility: A structural approach. Computational Economics 58, 347–394.

Donadelli, M., Jüppner, M., Riedel, M., Schlag, C., 2017. Temperature shocks and welfare costs. Journal of Economic Dynamics and Control 82, 331–355.

Faccia, D., Parker, M., Stracca, L., 2021. Feeling the heat: extreme temperatures and price stability. ECB Working Paper Series .

Ferrari, A., Nispi Landi, V., 2020. Whatever it takes to save the planet? Central banks and unconventional green policy. Working Paper Series 2500. European Central Bank. URL: https://ideas.repec.org/p/ecb/ecbwps/20202500.html.

Fisher, A.C., Hanemann, W.M., Roberts, M.J., Schlenker, W., 2012. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. American Economic Review 102, 3749–60.

Forni, M., Gambetti, L., 2016. Government spending shocks in open economy VARs. Journal of International Economics 99, 68–84.

Gallic, E., Vermandel, G., 2020. Weather shocks. European Economic Review 124, 103409.

Gertler, M., Karadi, P., 2015. Monetary policy surprises, credit costs, and economic activity. American Economic Journal: Macroeconomics 7, 44–76.

Graff Zivin, J., Neidell, M., 2014. Temperature and the allocation of time: Implications for climate change. Journal of Labor Economics 32, 1 – 26.

Griffit, W., Veitch, R., 1971. Hot and crowded: Influence of population density and temperature on interpersonal affective behavior. Journal of Personality and Social Psychology 17(1), 92–98.

Gurkaynak, R.S., Sack, B., Swanson, E., 2005. Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. International Journal of Central Banking 1.

Hansen, L.P., 2021. Central banking challenges posed by uncertain climate change and natural disasters. Journal of Monetary Economics .

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D.J., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larsen, K., Houser, T., 2017. Estimating economic damage from climate change in the united states. Science 356, 1362–1369. URL: https://www.science.org/doi/abs/10.1126/science.aal4369, doi:10.1126/science.aal4369, arXiv:https://www.science.org/doi/pdf/10.1126/science.aal4369.

Iacoviello, M., Navarro, G., 2019. Foreign effects of higher u.s. interest rates. Journal of International Money and Finance 95, 232–250.

Jin, Z., Li, F.W., Lin, Y., Zhang, Z., 2021. Do firms adapt to rising temperatures? evidence from establishment-level data. Available at SSRN URL: http://dx.doi.org/10.2139/ssrn.3573260.

Jorda, 2005. Estimation and Inference of Impulse Responses by Local Projections. American Economic Review 95, 161–182.

Kala, N., 2019. Learning, adaptation, and climate uncertainty: Evidence from indian agriculture. MIT Center for Energy and Environmental Policy Research Working 50.

Kelly, D., Kolstad, C., Mitchell, G.T., 2005. Adjustment costs from environmental change. Journal of Environmental Economics and Management 50, 468–495.

Kim, H.S., Matthes, C., Phan, T., 2021. Extreme Weather and the Macroeconomy. Working Paper 21-14. Federal Reserve Bank of Richmond. doi:10.21144/wp21-14.

Kotz, M., Wenz, L., Stechemesser, A., Kalkuhl, M., Levermann, A., 2021. Day-to-day temperature variability reduces economic growth. Nature Climate Change 11, 319–325.

Makkonen, A., Vallström, D., Uddin, G.S., Rahman, M.L., Haddad, M.F.C., 2021. The effect of temperature anomaly and macroeconomic fundamentals on agricultural commodity futures returns. Energy Economics 100, 105377.

McFarland, J., Z.Y.C.L.e.a., 2015. Impacts of rising air temperatures and emissions mitigation on electricity demand and supply in the United States: a multi-model comparison. Technical Report.

Miranda-Agrippino, S., Ricco, G., 2021. The transmission of monetary policy shocks. American Economic Journal: Macroeconomics 13, 74–107. URL: https://www.aeaweb.org/articles?id=10.1257/mac.20180124, doi:10.1257/mac.20180124.

Moore, F.C., 2017. Learning, adaptation, and weather in a changing climate. Climate Change Economics 08, 1750010. URL: https://doi.org/10.1142/S2010007817500105, doi:10.1142/S2010007817500105, arXiv:https://doi.org/10.1142/S2010007817500105.

Morss, R.E., Lazo, J.K., Demuth, J.L., 2010. Examining the use of weather forecasts in decision scenarios: results from a us survey with implications for uncertainty communication. Meteorological Applications 17, 149–162.

Mukherjee, K., Ouattara, B., 2021. Climate and monetary policy: do temperature shocks lead to inflationary pressures? Climatic change 167.

Nakamura, E., Steinsson, J., 2018. Identification in macroeconomics. Journal of Economic Perspectives 32, 59–86. URL: https://www.aeaweb.org/articles?id=10.1257/jep.32.3.59, doi:10.1257/jep.32.3.59.

National Weather Service, 1989. The december 1989 cold wave. URL: https://www.weather.gov/ilx/dec1989-cold.

Pankratz, N., Schiller, C., 2021. Climate Change and Adaptation in Global Supply-Chain Networks, in: Proceedings of Paris December 2019 Finance Meeting EUROFIDAI-ESSEC, European Corporate Governance Institute–Finance Working Paper.

Papoutsi, M., Piazzesi, M., Nber, May, 2021. How unconventional is green monetary policy?*.

Ramey, V., 2016. Chapter 2 - macroeconomic shocks and their propagation, Elsevier. volume 2 of *Handbook of Macroeconomics*, pp. 71 – 162.

Roth Tran, B., 2022. Sellin' in the Rain: Weather, Climate, and Retail Sales. Technical Report. Federal Reserve Bank of San Francisco Working Paper.

Schiermeier, Q., 2018. Droughts, heatwaves and floods: How to tell when climate change is to blame . Nature 560, 20–22.

Somanathan, E., Somanathan, R., Sudarshan, A., Tewari, M., 2021. The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing. Journal of Political Economy 129, 1797–1827.

Starr-McCluer, M., 2000. The effects of weather on retail sales. Technical Report.

Sterba, J.P., 1977. 1977, a year of weather extremes, causing concern among experts. URL: https://www.nytimes.com/1977/07/14/archives/1977-a-year-of-weather-extremes-causing-concern-among-experts-1977.html.

Appendix A. Diagnostics on the shocks

In order to check that our US-wide shock series has desirable properties to be used in macroeconomic analysis, we perform the Ljung-Box Q-test for serial correlation. We repeat the test 6 times including a different number of lags each time (1,4,8,12,16,20). All tests fail to reject the null hypothesis of no residual autocorrelation. The autocorrelation function is displayed in Figure Appendix A.1, showing that evidence for serial correlation is limited at best.

Figure Appendix A.1: **Autocorrelation function of the US-wide shock**

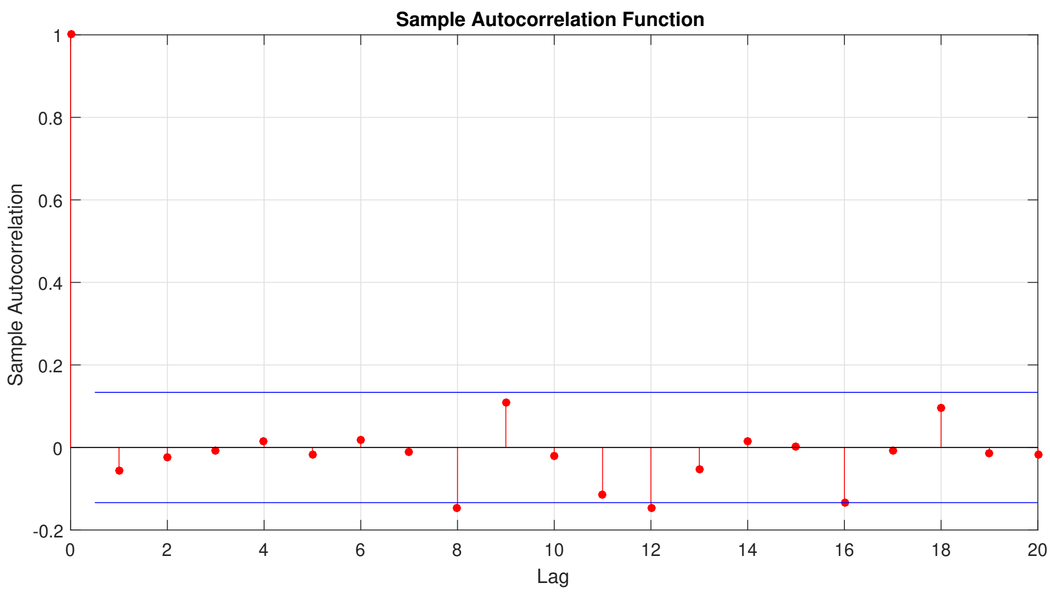

Appendix B. Exercise using weather forecasts

We conduct an illustrative exercise to investigate whether days that we label as surprisingly hot or cold have also been more difficult to predict. We employ data from the Model Output Statistics (MOS) database of the National Weather Service, collected from the Iowa State Mesonet archive and available since year 2000. For our forecasting exercise, we focus on Washington DC county (FIPS code 11001) as temperature forecasts are made through one weather station only (station ID: KDCA): in this way we avoid the need of making spatial interpolation of different forecasts. Forecasts are made six times per day and relate to temperatures for the same day and up to 3 days ahead, at 3-hour frequency. We retain the last forecasts made during the day (8 pm DC time) to construct predictions of one- and two-day ahead average daily temperatures. Using realized county temperatures, we construct daily ex-post forecast errors as

$$\mathrm{fe}_t^{\,j} = f_{t-j}^{\,j} - \mathrm{temp}_t$$

where f_{t-j}^{j} are forecasts made at time $t-j$ with horizon j, with $j = 1, 2$ days and $temp_t$ are the realized average daily temperatures used to construct our shock. To compare forecasts with surprises, we compute a dummy variable that equals one if the day is labeled as surprisingly hot or cold under our procedure, and 0 otherwise. We then estimate a probit regression model to test if high forecast errors (in absolute values) are correlated with the probability of a day being a surprise for agents. In formulas

$$\text{Prob}_t\ (\text{surprising day}_t = 1) = F\left[\alpha + \beta_j\ |\text{fe}_t{}^j|\right] \qquad j = 1, 2$$

The estimates are carried out in the sample going from Jul 1, 2000 to December 31, 2019. Results show that the probability of a day being surprising is positively and significantly correlated with the size of forecasts errors, confirming that days with temperatures going beyond expectations are also more difficult to predict than those with milder temperatures.

Appendix C. Data sources

Table Appendix C.1: Data summary

Variable	Source	Frequency	Sample
Temperature shock			
gridded average daily temperatures	Northeast Regional Climate Center	daily	Jan 1,1970 - Dec 31,2019
population (by county)	Census Bureau	yearly	1969 - 2018
employment (by county)	Census Bureau	yearly	1969 - 2018
personal income (by county)	Census Bureau	yearly	1969 - 2018
land extension (by county)	Census Bureau	yearly	1969 - 2018
Response variables			
Real GDP (USGDP...D)	Refinitiv	quarterly	1975q1-2019q4
Real Personal Consumption Expenditures (USCNPER.D)	Refinitiv	quarterly	1975q1-2019q4
Real gross private domestic investment (USGDPRIND)	Refinitiv	quarterly	1975q1-2019q4
Consumer Price Index (USCONPRCE)	Refinitiv	quarterly	1975q1-2019q4
3-month rate (TB3MS)	Fred	quarterly	1975q1-2019q4
2-year Treasury yield (GS2)	Fred	quarterly	1975q1-2019q4
10-year Treasury yield (GS10)	Fred	quarterly	1975q1-2019q4
GDP nowcast (t)	Greenbook Forecasts	quarterly	1980q1-2015q4
GDP forecasts (t+1)	Greenbook Forecasts	quarterly	1980q1-2015q4
GDP forecasts (t+2)	Greenbook Forecasts	quarterly	1980q1-2015q4
CPI nowcast (t)	Greenbook Forecasts	quarterly	1980q1-2015q4
CPI forecasts (t+1)	Greenbook Forecasts	quarterly	1980q1-2015q4
CPI forecasts (t+2)	Greenbook Forecasts	quarterly	1980q1-2015q4
FOMC word count			
Transcripts of FOMC meetings	federalserve.gov	meeting calendar	1976 - 2015
Weather forecasts			
1-day ahead temperature forecasts (at 8 pm DC time)	Iowa State Mesonet archive	daily	Jul 1,2000 - Dec 31,2019
2-day ahead temperature forecasts (at 8 pm DC time)	Iowa State Mesonet archive	daily	Jul 1,2000 - Dec 31,2019

Appendix D. List of words for textual analysis

The three sets of words employed for textual analysis in FOMC meeting transcripts are detailed below:

- **Climate change**: climate change, climate crisis, climate emergency, climate breakdown, global warming, global heating, carbon emissions, greenhouse gas emissions
- **Temperatures**: extreme heat, heat wave, temperature, hot days, weather, cold days, drought, wildfires, heat stroke, sunstroke
- **Climate-related non-temperature events**: flood, landslide, tornado, hurricane, dustnado, snowfall, snow, rainfall, precipitation

Appendix E. Robustness exercises

Figure Appendix E.2: **Breakdown of the shocks**

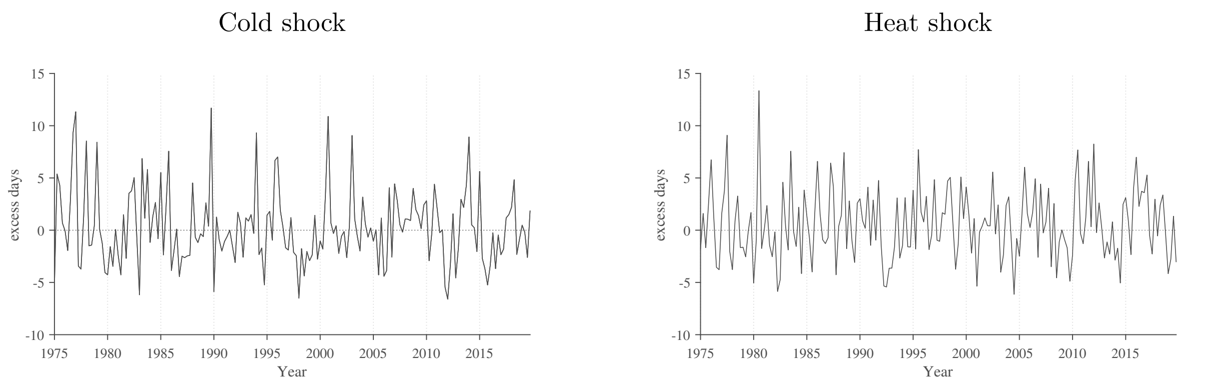

Figure Appendix E.3: **IRFs to heat shock**

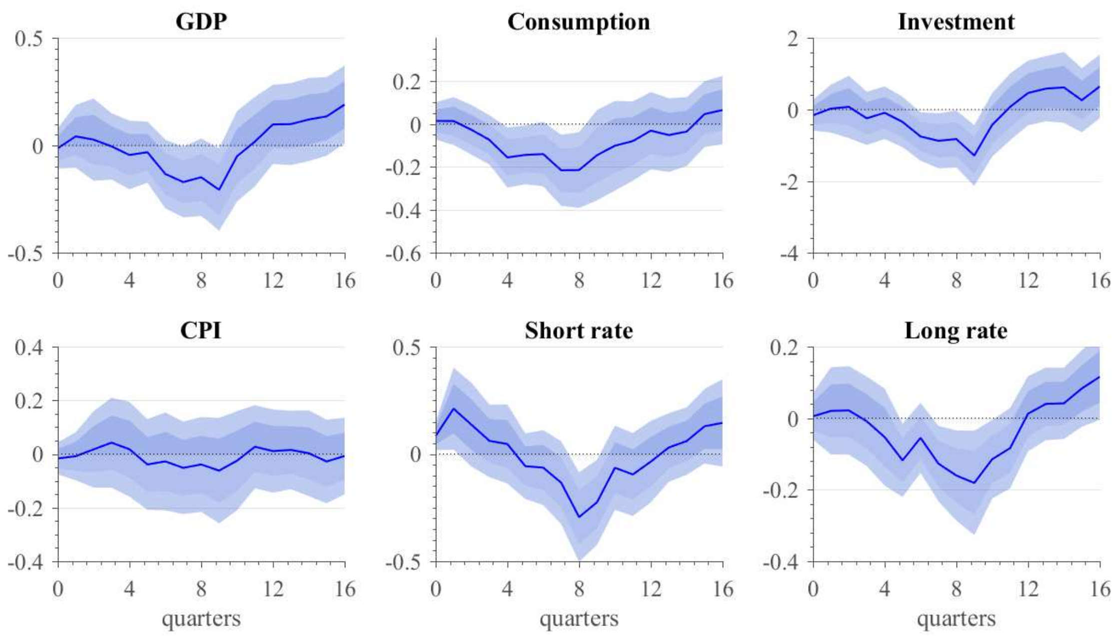

Notes: 68% and 90% confidence bands.

Figure Appendix E.4: **IRFs to cold shocks**

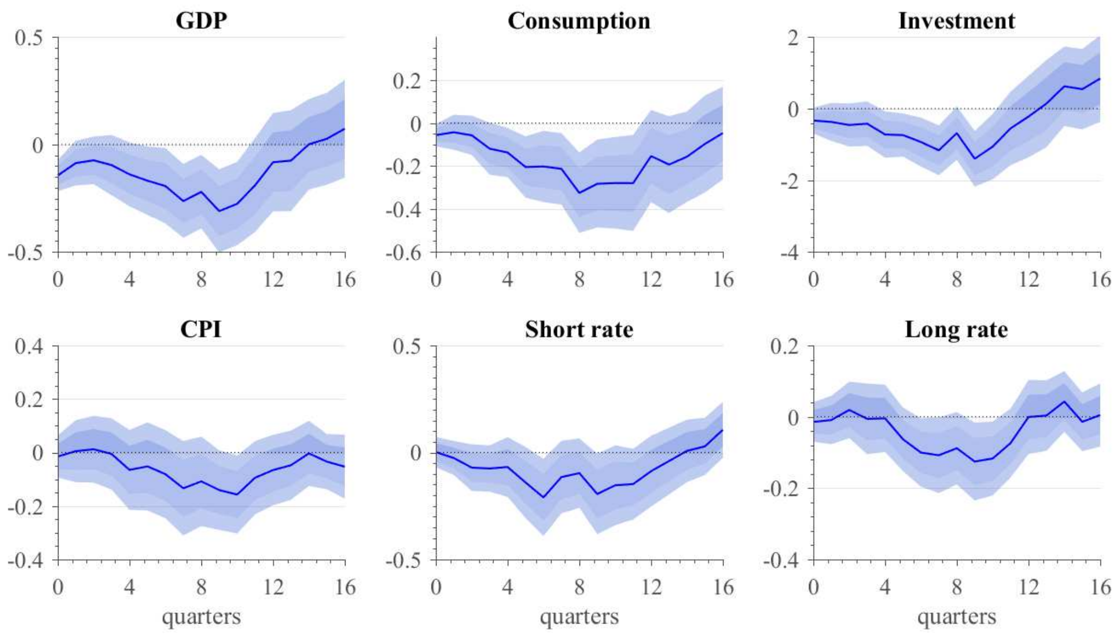

Notes: 68% and 90% confidence bands.

Figure Appendix E.5: **IRFs to heat shocks by season (GDP, CPI and short rate)**

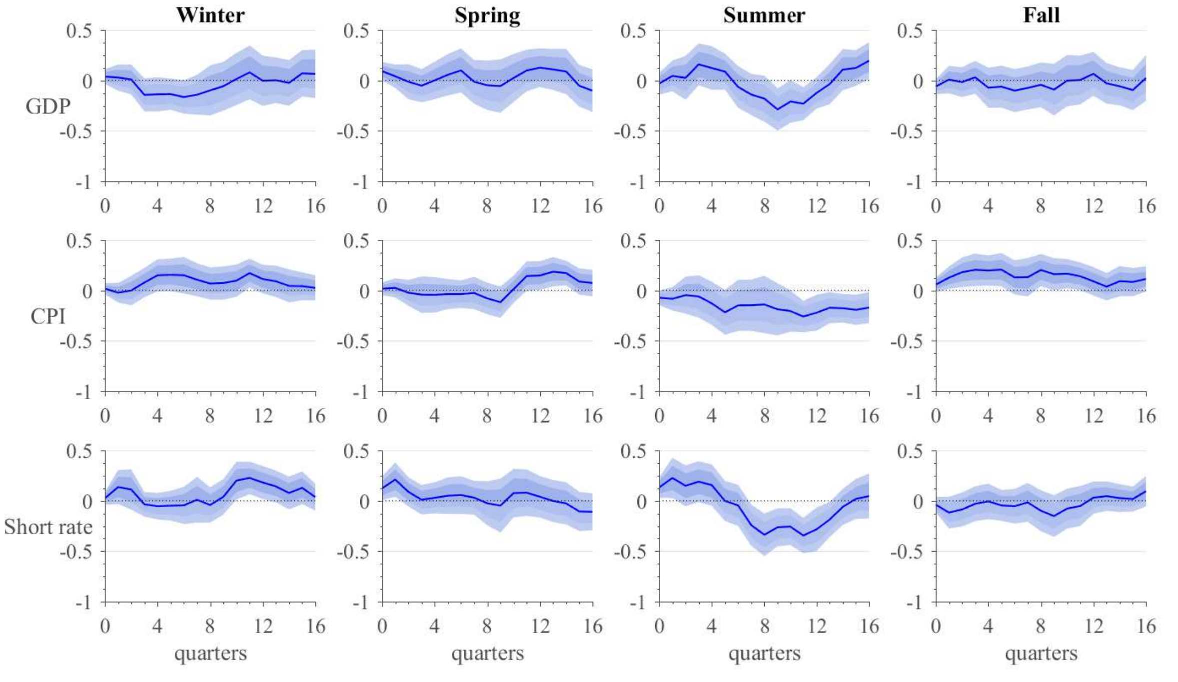

Notes: 68% and 90% confidence bands.

Figure Appendix E.6: **IRFs to cold shocks by season (GDP, CPI and short rate)**

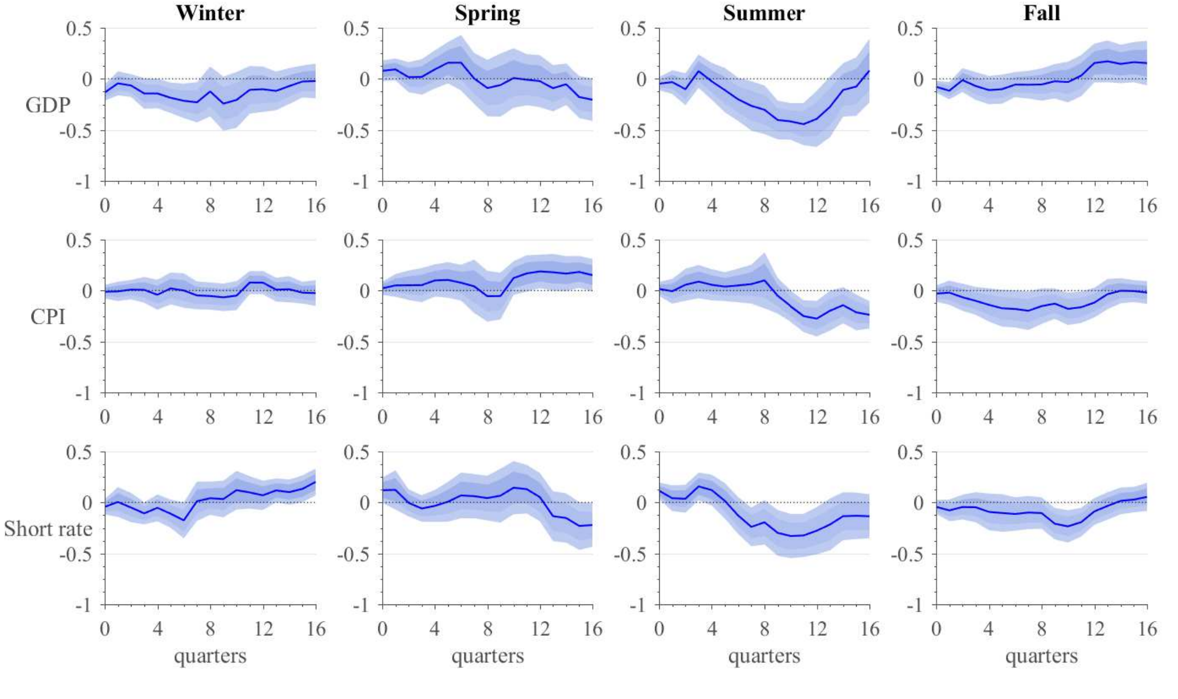

Notes: 68% and 90% confidence bands.

Figure Appendix E.7: **IRFs to shock computed with different thresholds**

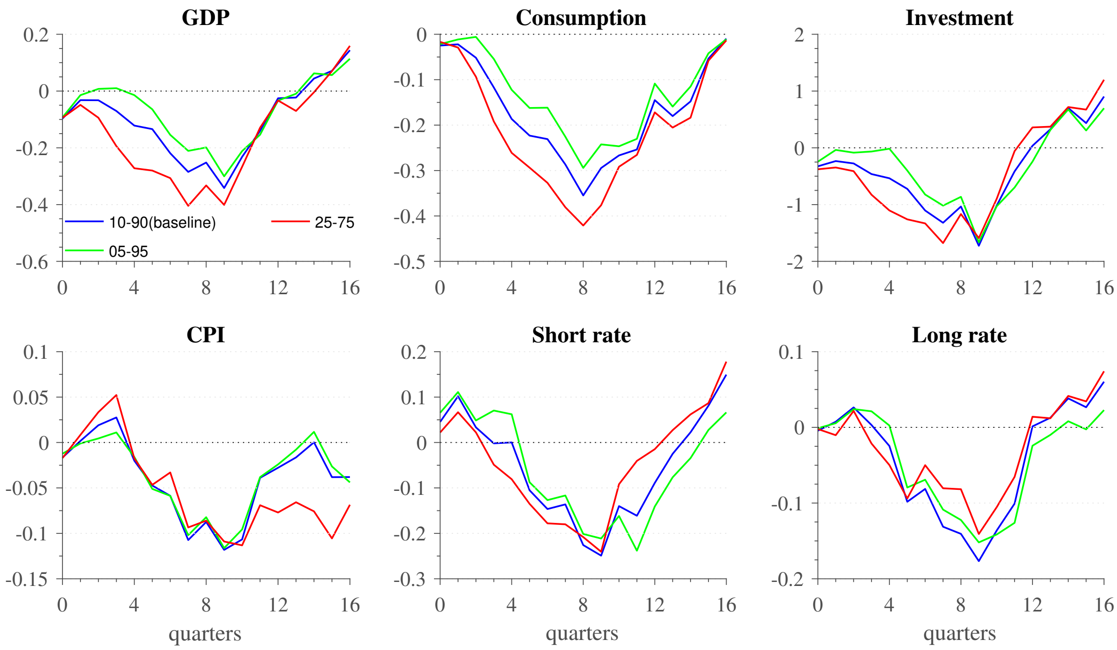

Notes: 68% and 90% confidence bands.

Figure Appendix E.8: **IRFs to shock computed under different weighting schemes**

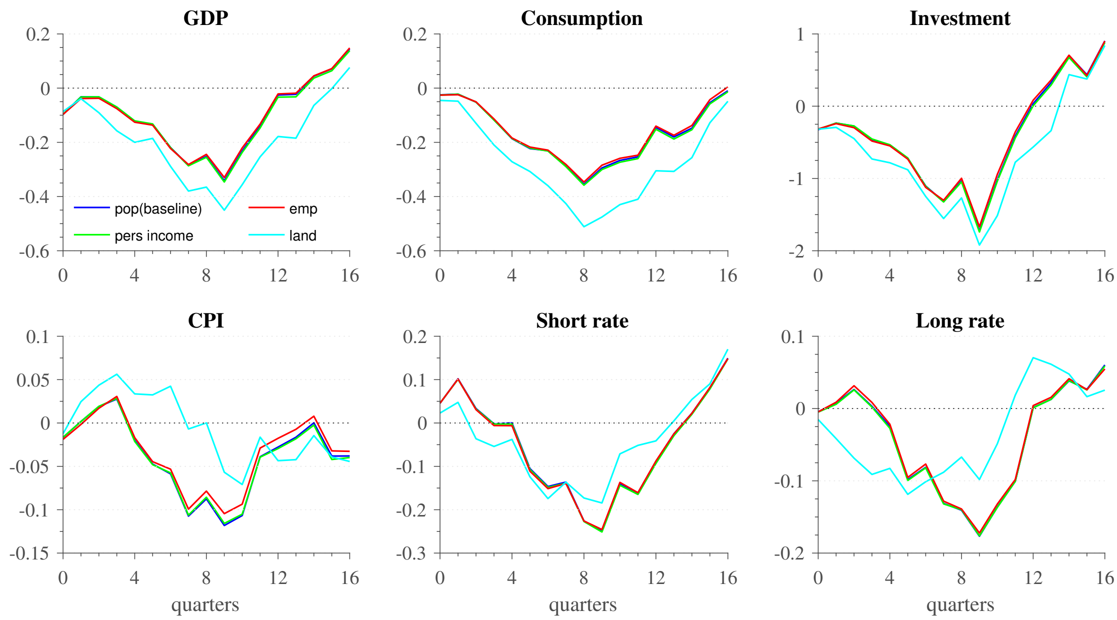

Notes: 68% and 90% confidence bands.

Figure Appendix E.9: **Shock computed with 10-year reference distribution**

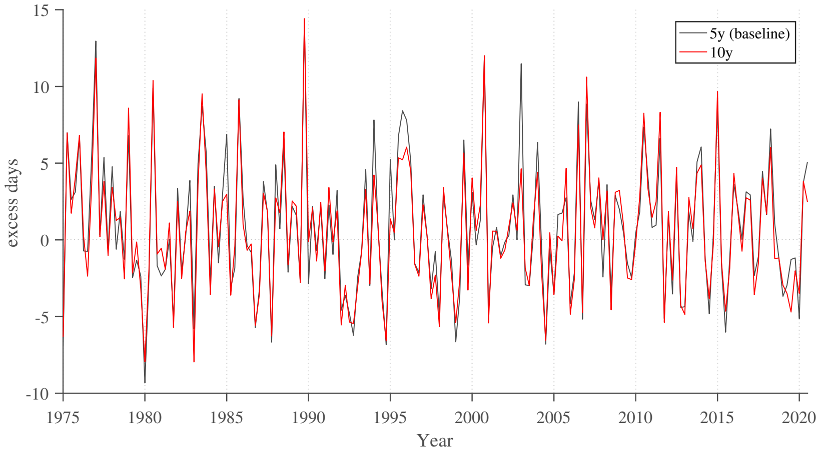

Figure Appendix E.10: **Two-lag controls**

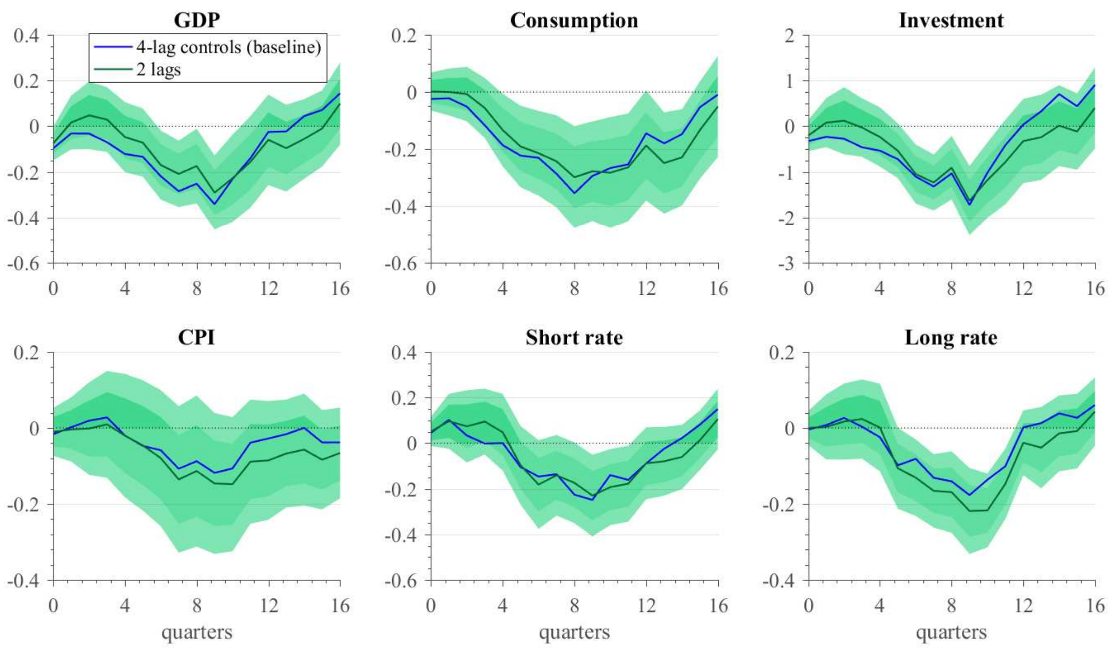

Notes: 68% and 90% confidence bands.

Figure Appendix E.11: **Alternative, one-sided shock**

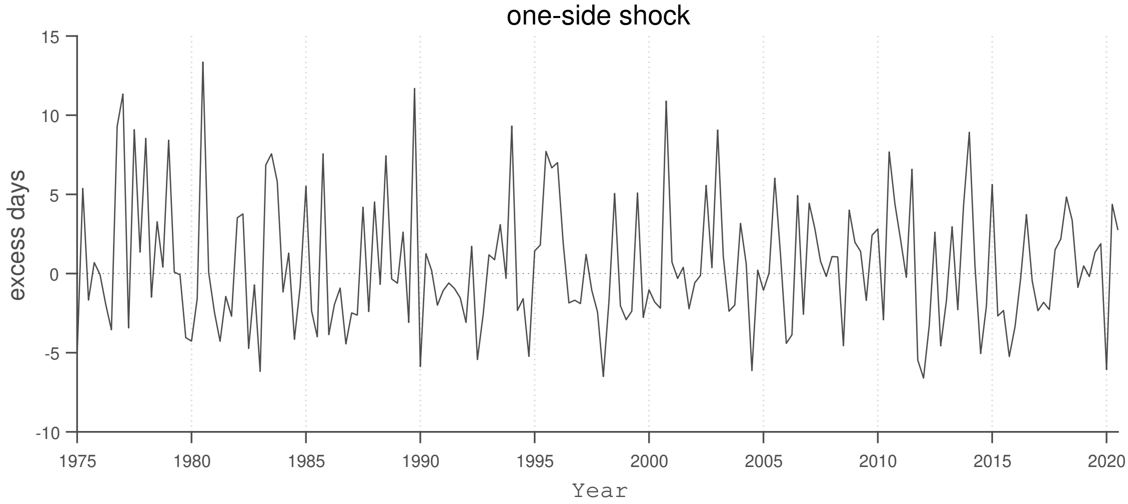

Figure Appendix E.12: **IRFs to one-sided shock**

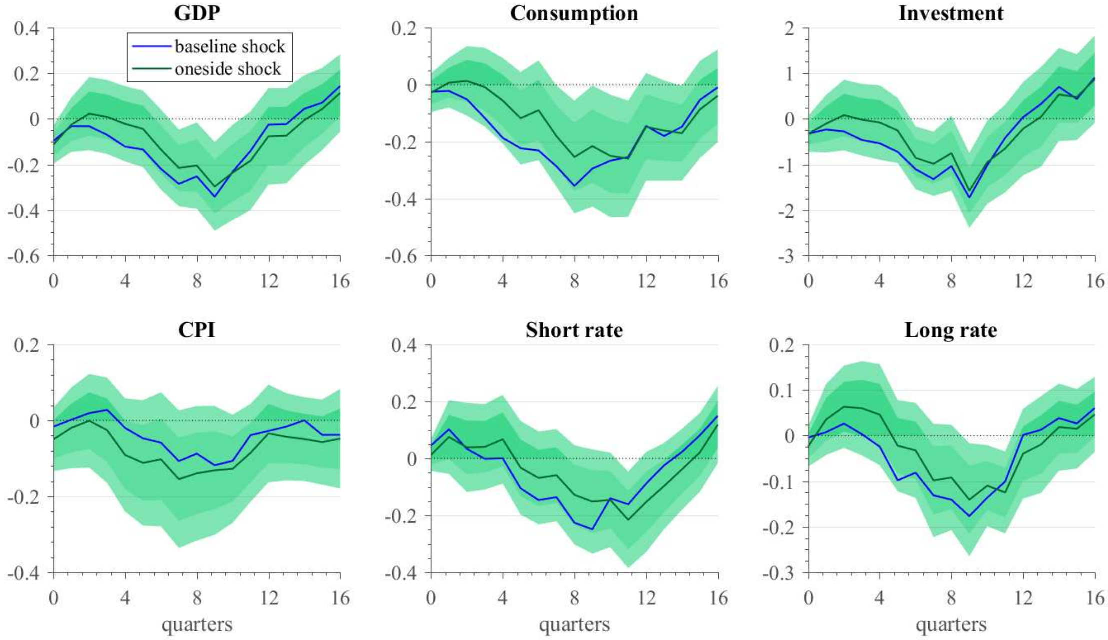

Notes: 68% and 90% confidence bands.

Figure Appendix E.13: **Positive vs negative shocks**

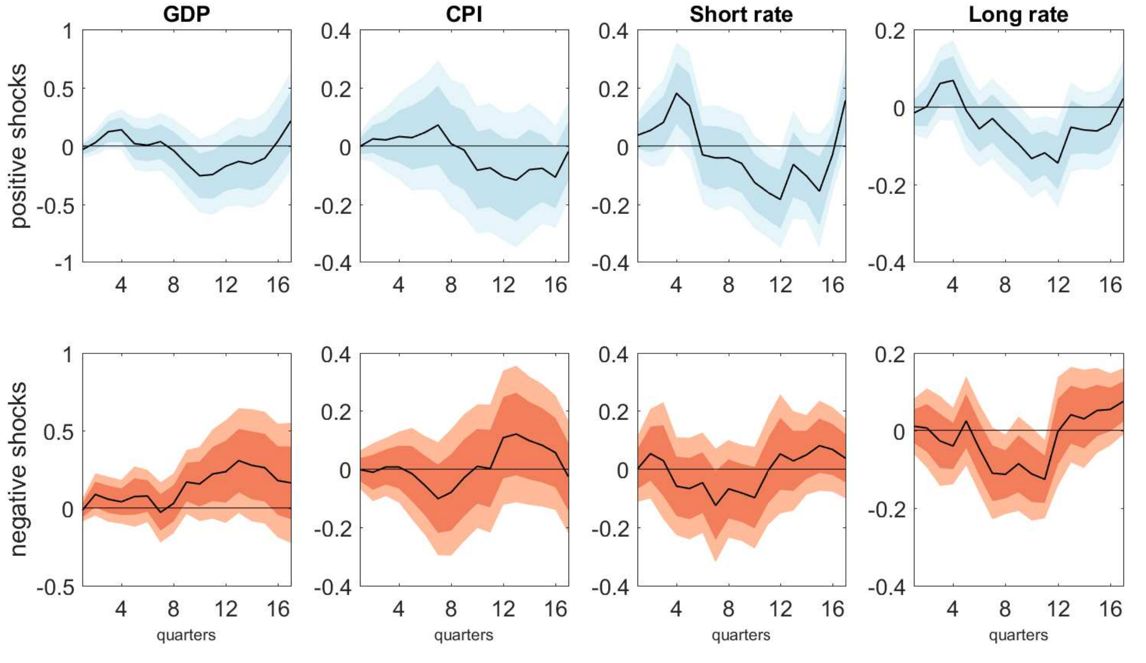

Notes: 68% and 90% confidence bands.