Sentimental Business Cycles and the Protracted Great Recession

Christos Andreas Makridis[*]

November 13, 2019

For Review

Abstract

Using newly licensed individual-level data from Gallup between 2008 and 2017, this paper provides microeconomic evidence that sentiments about economic activity played an important role in amplifying and propagating the Great Recession. First, after controlling for aggregate shocks, a 1pp rise in county employment and housing price growth is associated with a 0.30sd and 0.67sd rise in perceptions about the current state of the economy and a 0.12pp and 0.27pp rise in perceptions the economy is improving. Second, exploiting plausibly exogenous variation in the 2016 Presidential election, consumption of non-durable goods grew by 4.2% with a larger 10-12% increase among conservatives. The causal effect of sentiment on consumption is robust to three separate instrumental variables strategies: a state Bartik-like measure of gasoline price shocks, county fluctuations in daily temperature, and exposure to different housing price shocks through social networks. A back-of-the-envelope calculation suggests that the decline in sentiment can account for 34-68% of the decline in consumption during the Great Recession and an additional 14-43 months of delayed recovery.

Keywords: beliefs, business cycles, consumption, uncertainty, sentiments.

JEL: E21, E32, E71

1   Introduction

THE YEARS FOLLOWING THE GREAT RECESSION mark the slowest economic recovery in the United States post-WWII history [Taylor, 2014]. While there are many factors behind the sluggish recovery [Fernald et al., 2017], one prominent view is that the decline in housing wealth led to a large decline in consumption [Mian and Sufi, 2011, Mian et al., 2013] and subsequently employment [Mian and Sufi, 2014].[1] However, expectations also played a major role in accounting for these sudden and sharp declines in asset prices [Adam et al., 2017], particularly housing [Kaplan et al., 2019, Adelino et al., 2018a], that prompted these consumption and employment declines. If expectations among households were linked to their exposure to local shocks, then can the decline in economic sentiment ("beliefs") help explain the severity and length of the Great Recession?[2]

While the Keynesian [1936] insight that employment and production decisions are based on expected consumer demand is not new, there is little empirical evidence about the role that household sentiments play in potentially amplifying aggregate economic fluctuations. Macroeconomic theory models featuring incomplete information have generally focused on the formation of sentiment-driven equilibria based on waves of optimism or pessimism [Morris and Shin, 1998, 2002,

---

[1]While there is some debate about the origins of these housing wealth declines—that is, whether they were driven more by household expectations [Adelino et al., 2016, 2018a] versus an expansion of debt among sub-prime borrowers [Mian and Sufi, 2009, 2011]—this paper does not take a stance on their origins.

[2]Figure 7 in Appendix Section 7.1 uses the University of Michigan Survey of Consumers since 1978 to compare and contrast the 2008 financial crisis with other business cycles over the past four decades. Measurement of economic sentiment (beliefs) in my paper is based on a combination of survey responses to two separate questions: (i) "How would you rate economic conditions in this country today: as excellent, good, only fair, or poor?", and (ii) "Right now, do you think that economic conditions in this country, as a whole, are getting better or getting worse?" Measurement is discussed in more detail and summarized in Table 1.

Angeletos and La'O, 2013, Benhabib et al., 2015, Beaudry et al., 2011]. In particular, these models of self-fulfilling business cycles have, more generally, linked beliefs about low wealth with lower consumption [Farmer, 2012], volatility in asset prices and international credit [Bacchetta et al., 2012, Perri and Quadrini, 2016, Azariadis et al., 2016], consumption and housing prices [Kaplan et al., 2019], and even the transmission of the Great Recession across countries [Bacchetta et al., 2012, Bacchetta and van Wincoop, 2016]. The primary contributions of this paper are to quantify the effects of beliefs on real economic activity and examine the degree to which they can help account for the severe and protracted recovery from the Great Recession.

The first part of the paper introduces new, licensed micro-data from Gallup between 2008 and 2017 from their U.S. Daily Poll, containing survey responses from 1,000 people each day. The data offers three advantages, relative to traditional sources. First, given Gallup's infrastructure and specialization in survey methodology, they are able to launch large surveys with comparable questions over time. Second, the U.S. Daily contains sentiment indices about perceptions of the current (one to four scale) and future (one to three scale) state of the national economy, together with demographic characteristics, income, and non-durable consumption expenditures. Third, respondents also report detailed geographic information, such as their zipcode and county, enabling me to isolate the effects of local shocks, controlling for the aggregate state.

After comparing these data with the University of Michigan Survey of Consumer Sentiment, the volatility index, and the economic policy uncertainty (EPU) index from Baker et al. [2016], I show that local shocks affect individual beliefs about the national state of the economy: a 1pp rise in employment and housing price growth is associated with a 0.30 and 0.67 standard deviation rise in beliefs about the current state of the economy and a 0.12pp and 0.27pp rise in the probability that individuals report that the future state of the economy is improving. Although a number of

diagnostics are implemented to guarantee that these elasticities are not driven by omitted variables or reverse causality, the most restrictive specification compares renters and home owners in the same zipcode over time after controlling for all common state $\times$ time shocks and finds that these effects are concentrated among home owners. These results are inconsistent with fully rational expectation models, consistent instead with models where beliefs about housing [Piazzesi and Schneider, 2009, Burnside et al., 2016, Glaeser and Nathanson, 2017, Bailey et al., 2018b] / labor [Carroll and Dunn, 1997, Hendren, 2017] and personal experience play a driving role in belief formation [Malmendier and Nagel, 2016, Kuchler and Zafar, 2019, Bailey et al., 2018b].[3]

The second part of the paper estimates a micro elasticity between sentiments and non-durable consumption expenditures. The available macroeconomic evidence tends to rely on aggregate data and vector auto-regressions. These studies, however, have not yet produced a consensus since different methodologies generate different results (e.g., Barsky and Sims [2012] versus Beaudry and Portier [2006] and Beaudry et al. [2011]). Using the 2016 Presidential election as source of plausibly exogenous variation in expectations, I find that consumption grew by 4.2% in the cross-section and by 10-12% among conservatives, relative to their moderate and liberal counterparts. To guarantee the robustness of these results, I pursue three separate instrumental variables specifications: (i) a Bartik-like measure that exploits heterogeneity in state exposure to gasoline shocks, (ii) daily county temperature fluctuations, and (iii) housing price shocks experienced among socially connected counties. Although these strategies all produce similar estimates, my preferred specification using (iii) suggests that a standard deviation rise in perceptions about the state of the economy is associated with a 0.26% rise in daily non-durable consumption expenditures.

---

[3]While these results allow me to rule out a fully rational expectations model, I cannot distinguish between rational inattention [Sims, 2003] and near-rational models [Akerlof and Yellen, 1985].

The third part of the paper uses these elasticities to conduct a back-of-the-envelope calculation to gauge the aggregate importance of sentiments on consumption. Using the estimated elasticity, I find that the decline in economic optimism during the Great Recession accounts for 34-68% of the decline in consumption on non-durable goods. Moreover, I show that states that experienced a more severe sentiment shock also experience a slower recovery. Evaluating the reduced-form elasticities at the mean decline in sentiment during the crisis suggests that housing price growth and real GDP across states would have recovered 43 and 14 months sooner, respectively, in the absence of the decline in sentiment.[4] These estimates are consistent with descriptive evidence from Pistaferri [2016] that the sluggish consumption growth is best explained by low consumer confidence and high uncertainty, quantitative evidence from Milani [2017] that deviations in trend optimism can generate persistent movements in macroeconomic fundamentals, and from Burnside et al. [2016] that fluctuations in the cross-section of beliefs can generate booms and busts.

These results are also consistent macroeconomic models on unemployment risk and its effects on consumption [Carroll, 1992, Carroll and Dunn, 1997, Crossley and Low, 2014], aggregate demand and wealth [Ravn and Sterk, 2017, Challe et al., 2017, Beaudry et al., 2018, Heathcote and Perri, forthcoming], liquidity [Mertens and Ravn, 2014], and the amplification of fluctuations [Den Haan et al., 2017, Beaudry et al., 2018]. I specifically find that labor income risk is a bigger worry among higher skilled workers, whereas unemployment risk is a bigger worry among lower skilled workers. Both predict declines in daily consumption. The fact that exposure to local shocks affects the formation of expectations, and that these updated expectations influence real activity, is also consistent with recent quantitative work in Kaplan et al. [2019] on the importance of expectations

---

[4]Both of these back-of-the-envelope calculations are subject to the usual assumptions about taking elasticities from partial to general equilibrium, e.g., the role of allocation as emphasized by Beraja et al. [2019]. Future work should seek to incorporate realistic micro-elasticities of sentiment into aggregate models of cyclical dynamics. The inclusion of the aggregation exercise is simply for illustrative purposes.

in explaining the decline in consumption during the financial crisis.

This paper also relates directly with three studies that began concurrently with mine [Benhabib and Spiegel, 2016, Gillitzer and Prasad, 2018, Mian et al., 2017].[5] Using the University of Michigan Survey of Consumer Sentiment, Benhabib and Spiegel [2016] provide state-level evidence that changes in economic activity are correlated with changes in sentiment about national conditions. Using Australian micro-data, Gillitzer and Prasad [2018] exploit changes in the government party in power to identify the effects of expectations on an intent to spend more in the future. While my results are consistent with those from Benhabib and Spiegel [2016] and Gillitzer and Prasad [2018], Mian et al. [2017] focus more heavily on the role of partisan politics and household spending. Importantly, Mian et al. [2017] find that increases in the Republican vote share is associated with movements in economic sentiment, but not spending or vehicle purchases.[6] I reconcile these contrasting views by emphasizing the role of heterogeneity and providing alternative sources of plausibly exogenous variation. First, using heterogeneity in the self-reported intensity of political beliefs, I show that political affiliation is an important moderating variable for beliefs and expectations. For example, individuals identifying as very conservative raise consumption by 3-4pp more than their counterparts who identify as somewhat conservative. Second, using variation in exposure to different social networks, I show that connectivity-weighted housing price shocks affect individual perceptions about economy activity and, in turn, consumption. These results are

---

[5]Lagerborg et al. [2018] also provide evidence about the important role of sentiments using a vector auto-regressive approach, but using plausibly exogenous variation in public shootings to generate movements in sentiment.

[6]As I discuss later, there are several potential reasons for these differences. First, automobile purchases were nearly flat in a neighborhood around the 2016 Presidential election. Although vehicle purchases grew after the 2010 Cash for Clunkers program, they began to plateau. Second, since Mian et al. [2017] focus on zipcode Republican vote shares, their imputation bundles together Republicans and Democrats who respond heterogeneously to the 2016 Presidential election. For example, I show that conservatives raise consumption by 10-12% more than their left-of-center counterparts. Basit Zafar also alluded that this could be the case in an earlier write-up at the New York Federal Reserve: http://libertystreeteconomics.newyorkfed.org/2017/01/measuring-americans-expectations-following-the-2016-election.html.

robust to two separate instrumental variables strategies.

The structure of the paper is as follows. Section 2 introduces the data, compares the measure of sentiments with existing measures, and characterizes several descriptive features of the data. Section 3 quantifies how local housing and labor market shocks shape beliefs about the current and future state of the economy. Section 4 estimates how sentiments affects real economic activity measured primarily through personal consumption expenditures and hiring intensity. Section 5 conducts an aggregation exercise and examines how the fluctuations in beliefs may help account for the slow recovery from the Great Recession. Section 6 concludes.

2 Data and Measurement

2.1 Sources

Gallup Daily Polling Repeated Cross-section.—The primary source consists of newly licensed data with Gallup Inc. Gallup is the United States' premier polling service and conducts daily surveys of 1,000 U.S. adults on various political, economic, and well-being topics. Specifically, 200 Gallup interviewers conduct computer-assisted telephone interviewers with randomly sampled respondents (age 18 or over) from all 50 states and the District of Columbia. Detailed location data, such as the zip-code and metro area, is also available with corresponding sample weights.

Gallup's current polling relies on live, not automated, interviews with dual-frame sampling (including random-digit-dial [RDD]) landline and wireless phone sampling. Half of the respondents receive the "well-being track" version (with a 9% survey response) of the survey questions, whereas the other half receives the "politics and economy track" (with a 12% survey response). The two

surveys contain different topical questions, but both contain the same identifying demographic information. Gallup also conducts the survey in Spanish to record replies from those Spanish speakers who do not also speak English. The sampling methodology also uses a three-call design to reach respondents who do not pick up on the original attempt. Although there have been changes in the survey questions (e.g., a broader set of well-being indices are made available in 2014), the main results draw from consistently measured survey questions.

While the survey does not cover every county in the United States, it reaches 1514 counties (over a third of all counties) and 1089 counties with at least 300 respondents. Table 1 presents the Gallup questions used to recover information about sentiments, consumption, and well-being. A particularly unique feature of the data is its measurement of consumption on non-durable goods (see Figure 8 in Appendix Section 7.2.2 for a comparison with BEA data).[7]

[INSERT TABLE 1 HERE]

To provide a graphical characterization of the spatial variation, Figure 1 plots measures of economic sentiments across all U.S. states. Panels A and B plot the share of individuals in 2008 and 2015, respectively, reporting that the economy is getting worse. First, the share of individuals reporting that the economy is getting worse not surprisingly declines from nearly 90% of individuals in 2008 to 60-70% in 2015. Second, there is incredible spatial variation across the U.S. and it varies over time. Panels C and D plot the average z-score of perceptions about the current state of the economy. Unlike the former measure, higher levels of this index signal positive

---

[7]Since the measure is specifically about the individual's spending on the prior day, there is a potential selection problem: individuals who did not go shopping the day before will have zero consumption expenditures. Fortunately, the day of the week that the interview takes place appears to explain some of the censoring, meaning that these censored observations are not crucial for identification. For example, 13.5% of the consumption observations are missing for individuals interviewed on Sunday, 16.3% for Monday, 15.2% for Tuesday, 14.1% for Wednesday, 13.5% for Thursday, 14.2% for Friday, and 13.1% for Sunday. While one approach would be to use these fixed effects with a Heckman [1979] selection correction, they do not do a very good job. Censored values are, therefore, omitted, but day of the week fixed effects are included as controls to remove potential bias.

sentiments. Not surprisingly, Texas has the highest sentiments, followed by states like North Dakota, which experienced a shale gas boom.

[INSERT FIGURE 1 HERE]

To provide a graphical characterization of the time series variation, Figure 4 plots the distribution of sentiments across locations between 2008-2009 and 2014-2015. First, households become not only more pessimistic about the current and future state of the economy during the Great Recession, but also more uncertain about the state of the economy, consistent with the view that first moment shocks are procyclical, whereas second moment shocks are countercyclical [Bloom et al., 2015]. For example, Figure 4 For example, shows that the standard deviations of sentiments about the current and future state of the economy in 2008-2009 are 111% and 318% as large as their 2014-2015 counterparts, respectively. Over the time series, consumption expenditures on non-durable goods have a high 0.63 correlation with perceptions about the state of the economy (see Figure 5). While the correlation has no causal interpretation because of omitted macro shocks, the raw data from Gallup affirms the traditionally tight link between consumption and beliefs in the cross-section [Carroll et al., 1994, Bram and Ludvigson, 1998, Ludvigson, 2004].

[INSERT FIGURES 4 AND 5 HERE]

There, however, some limitations to subjective survey questions. The first is the "halo effect". Recipients answer different questions with the same mental state of mind, which produces a mood that can spill over from the answer in one question to another.[8] The second is the potential wedge

---

[8]Oswald [2008] examines the validity of subjective measures by leveraging information on individuals' self-reported measures of relative height to other individuals of the same gender (e.g., "how tall do you feel you are relative to your gender?"). Using the auxiliary height information, together with actual height, Oswald [2008] is able to measure the reporting function for subjective measures of well-being, finding that the subjective and objective measures have an approximately 0.80 correlation; see Oswald and Wu [2010] provide additional evidence in another setting. In this sense, while a concern may remain about subjective measures, it appears that they are capturing the underlying fundamentals.

between stated and revealed preferences. Recent evidence from Benjamin et al. [2012], for example, finds many individuals are not well-informed about their options and/or may respond very different to different sets of question cues. The third is that the ordering of questions in surveys—that is, whether questions about politics and the economy are asked first, for example—influences the elicited life satisfaction indices [Deaton, 2012]. In spite of these limitations, self-reported measures of well-being and sentiment still contain important information and validation studies have found that they tend to be sufficiently reliable [Krueger and Schkade, 2008]. Gallup prides itself on maintaining a professional and rigorous polling methodology, which helps obviate concerns about the underlying tone and sampling frame of the survey.

County Panel of Employment, Wages, and Housing.—Since the Gallup data contains significant geographic detail across both space and time, I subsequently match the micro-data with quarterly employment and wages data from the Quarterly Census of Employment and Wages (QCEW) and the Federal Housing Administration's (FHA) annual county housing price index. The administrative records from the QCEW cover 95% of jobs in the U.S. and are maintained in part by state agencies for tracking and distributing unemployment insurance. I use these data to compute county employment growth not only on average, but also in specific sub-sectors (e.g., trade) to proxy for consumer demand shocks. Similarly, I use the FHA housing price index to compute annual housing price growth. One of advantage of the index introduced by Bogin et al. [forthcoming] is that it is constructed using repeated sales transactions, purging variation in the quantity of housing and other time-invariant characteristics that are common with a location.[9]

---

[9]In an earlier draft, I also used the median housing price per square foot data from Zillow who use their proprietary technology ("Zestimates") to estimate housing prices using recent sale transactions within each micro-region. Within each micro-region, they feed the model various home attributes—including the time it is on the market, specific amenities in the home (e.g., number of bedrooms), and neighboring transaction prices—and predict the sales price. Although their Zestimates predictions correlate well with the FHA index and has been used in recent research (e.g., Bailey et al. [2018b]), the Zillow estimate is limited to a more restricted sample of counties and

2.2 Measuring Economic Sentiments

Economic sentiments are measured using perceptions about the state of the economy. The Gallup micro-data surveys individuals about both the current and future state. Individuals are asked to rank their perceptions of the current state of the economy based on one of four values, whereas they are asked to rank their perceptions of the future state of the economy based on one of three values (see Table 1 for the wording). While the conditional mean characterizes an important dimension of economic sentiments, the standard deviation could also be used to help characterize higher-order sentiments and uncertainty [Angeletos and La'O, 2009, Angeletos et al., 2014].

Before bench-marking these sentiment indices as measures of uncertainty, I begin by bench-marking the Gallup data with the University of Michigan Survey of Consumer Sentiment. Using data on perceptions about the current and future state of the economy aggregate at a monthly frequency between 2009 and 2017, Figures 9 and 10 in Appendix Section 7.2.3 show the close correspondence between the two datasets across different educational attainment and age brackets. While the correlation between the two datasets over the share of individuals reporting that the economy is worsening tends to have a correlation upwards of 0.50, the correlation over the share of individuals reporting about the current state of the economy tends to be upwards of 0.80. Moreover, Table 8 in Appendix Section 7.2.3 shows that the Gallup sentiment data does a better job at predicting real GDP growth than its University of Michigan counterpart.

How do these measures of sentiments and uncertainty compare with existing measures? Panel A of Figure 2 compares the mean perception of the current state of the economy from Gallup

---

implies that housing price growth affects beliefs about the economy roughly half as much as employment growth, whereas the baseline estimates suggest they affect beliefs about twice as much as employment growth.

with the economic policy uncertainty index from Baker et al. [2016]. Whereas greater values of Gallup imply a more optimistic state of the economy, greater values of the Baker et al. [2016] index imply a more uncertain state. There is a correlation of -0.59 between the two indices, suggesting that they are capturing similar dimensions of uncertainty in the U.S. economy. Panel B of Figure 2 also shows that optimism about the economy has a -0.64 correlation with the daily volatility index. In sum, measures from Gallup appear to reflect genuine movements in aggregate measures of economic activity and uncertainty, as well as containing important spatial variation over time.

[INSERT FIGURE 2 HERE]

2.3 Cross-Sectional Dispersion during Booms/Busts

While there is some literature characterizing the cross-sectional dispersion of life satisfaction indices (e.g., see Oswald and Wu [2011] and Glaeser et al. [2016]), there is little evidence on not only the dispersion of other indices, but also the dispersion during a boom versus a recession. Using information on the underlying state of the economy, perceptions of work place practices, and perception of city amenities—each of which are also available in the U.S. Daily Poll—I collapse across all individuals within the same metropolitan area using the national sample weights and plot the distribution of each variable across all metro areas. The state of the economy and future city prospect measures are in the form of an index (poor, only fair, good, excellent for state of the economy and getting worse, the same, and getting better for future city prospects), whereas the work-place practices and city satisfaction measures are binary indicators.

These distributions are each plotted in Figure 3. Beginning with the state of the economy, there is a stark dispersion both in current and future attitudes during and after the Great Recession.

While the distribution of sentiments is substantially shifted towards the right between 2014-2015, relative to 2008-2009, for perceptions about the current state of the economy, attitudes within metros about the future state of the economy are much more bi-modal in 2008-2009.

Turning towards the measures of work-place practices and city satisfaction, there is not a statistically significant difference in any of them pre and post recession. It is possible that these attitudes tend to be relatively sticky, or at least tied closely with the underlying job or location, rather than the business cycle. For example, Bloom et al. [2014] discuss that management practices tend to be sticky because of adjustment costs—values and norms in organizations take time to change (e.g., due to incumbents, organizational inertia, etc).

[INSERT FIGURE 3 HERE]

3   Belief Formation and Personal Experience

There is now an active literature on the role of personal experience in belief formation about, for example, inflation [Malmendier and Nagel, 2016] and housing prices [Kuchler and Zafar, 2019, Bailey et al., 2018b]. This section validates these results by examining how individual beliefs respond to county fluctuations in labor and housing market activity. While the estimates may not have a causal interpretation, this section shows that changes in local economic conditions have statistically and economically significant effects on individual beliefs about both the current and future state of the economy, conditional on demographic characteristics and location fixed effects.

3.1   Empirical Specification

My measures of economic sentiment consist of an index on a scale of one to four about the current state of the economy and an indicator for whether the future state of the economy is improving. I create a standard-normal z-score of the former and an indicator of the latter, estimating least squares and linear probability regressions, respectively, of the form:

$$s_{ict} = \gamma \Delta e_{ct} + \phi \Delta h_{ct} + \beta X_{it} + \eta_c + \lambda_t + \epsilon_{ict} \tag{1}$$

where s denotes the measure of sentiment, Δe and Δh denote the year-to-year growth rates of county employment and housing prices per square foot, X denotes a vector of individual covariates, and η and λ denote fixed effects on county and day of the year. Importantly, the day of the year fixed effects purge variation arising from aggregate shocks that are publicly visible to everyone (e.g., daily news). Standard errors are clustered at the county-level to allow for arbitrary degrees of autocorrelation in the errors over time in the same location [Bertrand et al., 2004].

The coefficients of interest in Equation 1 (γ and ϕ) are identified off of within-county fluctuations in labor and housing market conditions. The identifying assumption is that unobserved shocks to sentiment are uncorrelated with housing and labor market fluctuations. The inclusion of location and time fixed effects controls for the non-random sorting of individuals with different perceptions about the economy into locations with different economic growth rates. The inclusion of both Δe and Δh also helps control for potential confounders since both the housing and labor markets were in significant flux during the financial crisis. I also examine specifications containing state $\times$ year $\times$ quarter fixed effects to exploit within-county variation, controlling for all shocks

common to individuals across states and time since differences in state institutions [Shoag and Veuger, 2016] and judicial foreclosure laws [Mian et al., 2015] may have amplified and propagated sentiment fluctuations in ways that are correlated with real activity.

In addition to the usual concerns about omitted variables bias, the main threat to this identification assumption is reverse causality. For example, an unobserved shock to expectations might raise investor expectations about the housing market, which empirically played an important role in explaining housing price growth prior to the financial crisis [Adelino et al., 2018a]. Similarly, such a shock could also prompt companies to open up new plants and/or jobs within an area, thereby raising employment growth. While the goal in these exercises is not to recover a causal effect, but rather highlight the way that local shocks influence beliefs about the national state of the economy, I conduct several placebo exercises that highlight the mechanism of interest.

3.2   Main Results

Table 2 documents the results associated with Equation 1 with the z-score of the current state of the economy and an indicator that the future state of the economy will be better. As a starting point, columns 1-2 and 5-6 report the the gradients on employment and housing price growth in isolation, which may produce upwards bias for both of them given the co-movement between labor market and housing wealth shocks during the financial crisis [Mian and Sufi, 2014]. For example, these specifications suggest that a 1pp rise in employment (housing price) growth is associated with a 0.65sd (0.71sd) rise in perceptions about the current state of the economy and a 0.26pp (0.28pp) rise in the probability that individuals report that the economy is improving.

However, since labor and housing market shocks took place jointly over the crisis, the coef-

ficients on employment and housing price growth are likely upwards biased—that is, areas with housing price declines also exhibit larger employment declines, which drives down beliefs about the economy. Under the preferred specifications in columns 3 and 7, which contain all of the controls and fixed effects, a 1pp rise in employment and housing growth is associated with a 0.30sd and 0.67sd rise in perceptions of the current state of the economy and a 0.12pp and 0.27pp rise in the probability that individuals report that the economy is improving. Importantly, the magnitude of annual housing price growth is approximately twice as large as the magnitude of employment growth, consistent with the view that housing wealth declines drove down the demand for labor during the crisis, especially in the non-tradables sector [Mian and Sufi, 2014].

Since these coefficients may reflect unobserved shocks to local beliefs, I now control more parsimoniously for consumer demand shocks and exposure to the financial sector. Columns 4 and 8 estimate the baseline specification with the inclusion of employment growth in the wholesale & retail trade sectors and the finance & real estate sectors. However, their inclusion does not alter the coefficient estimates in any meaningful way. Moreover, since there there are a number of potentially confounding state-level forces moderating employment and housing price shocks, such as state labor market institutions [Shoag and Veuger, 2016] and judicial foreclosure laws [Mian et al., 2015], I also introduce state × year × quarter fixed effects, recovering gradients of 0.23 and 0.54 on employment and housing price growth, respectively (omitted from the table for brevity).

[INSERT TABLE 2 HERE]

3.3 Robustness Exercises

If these elasticities between beliefs and local employment and housing price shocks reflect a genuine causal relationship, then the effects of housing price shocks should be concentrated among home owners.[10] Appendix Section 7.3 presents difference-in-difference results that exploit a subset of the data containing information on home ownership. Although the direct effects of housing price growth are insignificant, they are concentrated among home owners: a 1pp rise in housing price growth is associated with a 0.28sd (p-value = 0.038) rise in perceptions about the current state of the economy and a 0.17pp (p-value = 0.00) rise in the probability the individual reports the economy is improving. Importantly, these elasticities are identified off of comparisons among individuals within the same zipcode after controlling for all state × time shocks and usual individual covariates. In addition to confirming the proposed mechanism about the role of personal experience, these results are consistent with Adelino et al. [2018b] who find that beliefs about housing price risk are less sensitive among renters.

Appendix Section 7.3 examines several additional diagnostic exercises that suggest that our coefficients do not simply reflect reverse causality or omitted variables, but rather the view that personal experience matters in forming beliefs about the future. First, Table 10 shows that the association of sector-specific labor market shocks are concentrated among individuals working within those specific sectors (e.g., manufacturing or finance & real estate), which is consistent with the view that individuals most exposed to the shocks should be the ones who respond most elastically. Moreover, Figure 11 shows that these gradients are statistically different from zero

---

[10]See Stroebel and Vavra [forthcoming] for a similar strategy that compares housing price shocks in areas with greater versus less home ownership.

in nearly every occupation, illustrating that these changes in beliefs are not driven exclusively by those working in the hardest hit sectors (e.g., finance and construction). Second, Table 11 shows that employment and housing price shocks also affect individual life satisfaction, which is a factor that moderates beliefs about the economy by altering the ways individuals acquire and process information, consistent with the importance of personal experience in the process of belief formation. Third, I follow Sullivan and Wachter [2009] by exploiting plausibly exogenous variation in unanticipated monthly state mass layoffs, finding that they are associated with similar declines in perceptions about the current and future state of the economy.

3.4 Heterogeneity

These results potentially mask important dimensions of heterogeneity. For example, macroeconomic models predict that wealthy hand-to-mouth, liquidity constrained individuals respond most elastically fiscal stimulus payments [Kaplan and Violante, 2014, Misra and Surico, 2014] and firms with low leverage respond most elastically to monetary policy [Ottonello and Winberry, 2018]. However, whether liquidity constrained individuals update their beliefs more than their counterparts in response to local shocks is an empirical question. On one hand, liquidity constrained individuals, much like their high leverage firm counterparts in Ottonello and Winberry [2018], might have less margin to update their beliefs and corresponding consumption. On the other hand, liquidity constrained individuals might have a greater return to updating their beliefs precisely because their marginal propensity to consume is higher.

To measure liquidity constraints, I draw on variables that gauge whether the respondent has worried about having enough money. Between 2009-2012, the question is measured through an

indicator in response to "You have more than enough money to do what you want to do"; between 2013-2017, the question is measured through a one-to-five scale in response to "You have enough money to do everything you want to do." Table 3 documents results associated with Equation 1 separately for types of liquidity constrained individuals. Appendix Section 7.3 also documents similar results when estimating Equation 1 separately by monthly income bracket.

Table 3 reveals some meaningful heterogeneity. A 1pp rise in employment growth is associated with a 0.47sd rise in beliefs about the current state of the economy among liquidity constrained individuals, but a 0.54sd rise among their counterparts. Similarly, a comparable 1pp rise in employment growth is associated with a 0.12pp rise in the probability that a liquidity constrained individual reports that the economy is improving, but a 0.16pp rise among their counterparts. While the gradients are not statistically different from one another, and recognizing the potential measurement error introduced through my classification of individuals as liquidity constrained, these results suggest that those who are not liquidity constrained might update their beliefs more in response to local shocks—these are precisely the individuals who face a tougher time smoothing consumption, so beliefs are more sensitive to economic conditions.

[INSERT TABLE 3 HERE]

4 Quantifying the Real Effects of Sentiments

4.1 The Identification Problem

To understand the relationship between sentiments and real economic activity, consider a naive model that relates belief shocks, denoted s, with personal consumption expenditures of non-

durable goods, denoted c, which declined significantly during the financial crisis [Mian et al., 2013, Pistaferri, 2016], through regressions of the form:

$$c_{ict} = \zeta s_{ict} + \beta X_{it} + \omega w_{ict} + \gamma \Delta e + \phi \Delta h + \eta_c + \lambda_t + \epsilon_{ict} \qquad (2)$$

where w denotes the individual's monthly income (bin), X denotes individual covariates, Δe and Δh denote county employment and housing price growth, and η and λ denote fixed effects on county and time. The inclusion of day of the year fixed effects purges all fluctuations in sentiment that are driven by aggregate shocks. Standard errors are clustered at the county-level.

Unfortunately, identifying ζ in Equation 2 is fraught with challenges. First, motivated by the earlier results that beliefs are influenced by local employment and housing price shocks, then any unobserved determinants of consumption are also likely to correlate with beliefs, thereby causing me to overestimate the effects of beliefs on consumption. While controlling for local employment and housing price growth, as well as individual income, helps, there are a wide array of potential omitted variables, especially unobserved individual-level heterogeneity that is not explained purely by income and demographic characteristics. Second, and potentially more threatening, is a more classic reflection problem. For example, Friedman [1992] remarked that consumer confidence indices "are mostly a reflection of what's going on rather than a cause."

Because there is no silver bullet approach for solving these two endogeneity problems, the remainder of the section takes a diversified approach: exploiting several sources of plausibly exogenous variation that imply similar estimates of the elasticity between beliefs and consumption. After showing that beliefs have a causal effect on consumption, I provide suggestive evidence that these belief shocks help explain the protracted recovery across locations: states that experienced

larger negative belief shocks also take longer to recover to their pre-crisis consumption trends.

4.2   The 2016 Presidential Election as an Event Study

Recent research from Gillitzer and Prasad [2018] and Mian et al. [2017] has exploited unanticipated political events as a source of plausibly exogenous variation in economic sentiment.[11] For example, Gillitzer and Prasad [2018] use Australian data to show that consumers report heightened spending intentions on household items and automobiles when the political party they support enters government, which generally lines up with postcode data that they have on automobile sales. Similarly, Mian et al. [2017] exploit the unanticipated victory of Donald Trump during the 2016 election as a source of plausibly exogenous variation in sentiment.[12]

Following the event study approach in Mian et al. [2017], I not only compare consumption before versus after the election, but also consumption among individuals with more versus less conservative worldviews before versus after the election. Since political views may influence the way individuals interpret publicly visible information, conservatives and liberals might respond differently to the same aggregate shock. To understand how individuals respond to changes in information, and the potential role that heterogeneity plays, I consider an event study that allows for heterogeneous treatment effect among individuals identifying as somewhat conservative and very conservative, relative to their liberal and moderate counterparts:

---

[11]Unlike Gillitzer and Prasad [2018], however, I cannot instrument for sentiment using voting intention because individuals are not simultaneously asked questions on consumption and sentiment after 2014. However, Equation 3 clarifies, variation in voting intention is explicitly incorporated into the specification, drawing on the election as a plausibly exogenous shock to expectations (i.e., a proxy for sentiment).

[12]Using elections as a source of quasi-experimental variation has precedent in the literature. For example, Kueng [2016] uses both the 1992 and 2000 elections to study the consumption response among higher-income households in response to news about expected changes in after-tax income. Admittedly, however, one problem with the shock is that it might represent more of a shock to political beliefs, rather than the economy. In the exercises that follow, I obtain fairly similar (and more favorable) results using separate identification strategies that focus on the shock to economic beliefs.

$$y_{it} = \alpha_1 Post_t + \alpha_2 SC_{it} + \alpha_3 VC_{it} + \xi_1(Post_t \times SC_{it}) + \xi_2(Post_t \times VC_{it}) + \beta X_{it} + \eta_c + \Phi_{st} + \epsilon_{it} \quad (3)$$

where y denotes logged daily consumption expenditures, $Post$ denotes after the November 8, 2016 election, SC denotes "somewhat conservative", VC denotes "very conservative", X denotes individual controls, η denotes county fixed effects, and Φ denotes state $\times$ year $\times$ month fixed effects. Equation 3 is restricted to the sample of moderates and conservatives; the coefficients of interest—ξ_1 and ξ_2—are interpreted relative to the consumption of a liberal or moderate.

Table 4 documents these results under two sample restrictions: a sample between July 2016 & May 2017 (columns 1-3) and a sample between September 2016 & March 2017 (columns 4-6).[13] Beginning with column 1, consumption grew 4.2% after the 2016 election, conditional on individual covariates (column 1). However, there is significant heterogeneity across individuals. In particular, allowing for heterogeneous treatment effects suggests that individuals identifying as somewhat conservative and very conservative report 7.2% and 10.5% higher consumption after the election, respectively (column 2). The estimated gradients are nearly identical even after controlling for both county and state $\times$ year $\times$ month fixed effects (columns 3), which exploits individual consumption after controlling for all shocks common to individuals within the same state over time. Moreover, these results are also robust to focusing on a more narrow three-

[13]I do not, however, use the election results of 2008 to conduct a comparable event study for two reasons. First, as discussed by Deaton [2012], there was an important difference in the ordering of the survey questions to accommodate additional inquiries about respondents' voting patterns for the election. By asking questions about politics early on, Deaton [2012] shows that individuals appear to have reported lower well-being scores, which may have also translated into lower beliefs about the economy. Second, since 2008 coincided with the financial crisis, consumption was low, relative to trend, because of the sheer magnitude of the aggregate shock. Since the variation in the event study is national, I cannot distinguish between the role of beliefs from the Great Recession more broadly.

month neighborhood around the 2016 election.[14] Finally, Appendix Section 7.4.1 illustrates that residualized non-durable consumption expenditures between conservative and liberals / moderates trend almost identically leading up to the election, but subsequently diverge.

The rise in consumption associated with economic optimism about the election is consistent with political science literature about expectations and economic behavior around elections. For example, Canes-Wrone and Park [2012] find that sectors that demand greater irreversible investments tend to experience a decline in the pre-election period when policy uncertainty about election outcomes are sufficiently high. Moreover, the fact that heterogeneity in consumption responses is so great is consistent with Gerber and Huber [2010] who provide similar evidence that partisanship influences economic assessments and sales growth using variation from a panel around the November 2006 election. However, McGrath [2017] finds that the result from Gerber and Huber [2010] was largely driven by a single state—Texas in 1996. One difference in our empirical approaches is my application of individual-level data on both consumption and political affiliation. Prior studies that have had to rely on geographic data may have experienced some attenuation bias arising from bundling together responses from heterogeneous individuals.

These estimates are also surprisingly similar, but lower in magnitude, to "macro elasticities" presented in Appendix Section 7.4.2. Using Bureau of Economic Analysis (BEA) consumption data, together with weighted averages of economic sentiment, Figure 15 plots the logged-level and growth rates of consumption and sentiment, displaying a gradient of 0.20 for the preferred growth rates specification when the outcome variable is in non-durables. Interestingly, even though housing price and consumption growth of non-durables are positively correlated (see Figure 16), the

---

[14]Appendix Section 7.4.1 also presents similar results using state-level changes in legislature and government as instruments for economic optimism, controlling for state and time fixed effects. While the results are much noisier since state elections may tend to be less polarizing and salient compared with federal elections, the implied elasticity is 0.354 (p-value = 0.097), qualitatively consistent with the aforementioned results.

gradient on housing price growth turns to -0.065 (p-value = 0.013) when controlling for economic sentiment, which produces a large gradient of 0.259 (p-value = 0.00). The fact that the gradient on economic sentiment rises after controlling for housing price growth suggests that expectations about housing price growth might have played a role in amplifying the decline in consumption.

[INSERT TABLE 4 HERE]

Contrasting with the aforementioned results and those from Gillitzer and Prasad [2018] that an unanticipated election shock transmits to consumption, Mian et al. [2017] do not find an effect using zipcode automobile purchases as a proxy for consumption expenditures. There are two potential reasons for these contrasting estimates. First, since the baseline strategy in Mian et al. [2017] requires imputing zipcodes as either Republican or Democrat, measurement error can create bias, especially since it is a binary variable [Aigner, 1973]. The individual-level data from Gallup highlights that, while the average treatment effect on consumption after the election is roughly 4%, the treatment effect among conservatives is nearly 3x as high. In this sense, not finding an increase in automobile purchases following the election in Republican-dominated zipcodes might reflect measurement error in the assignment strategy.

Second, since vehicles are an infrequent durable purchase, they might not be as reliable of a metric for validating the role of sentiments over a short time window, especially after exhausting consumer demand following the spike in auto purchases through the 2010 Cash for Clunkers [Green et al., 2016]. For example, Figure 14 in Appendix Section 7.4.1 plots logged non-durables consumption expenditures (deflated with the 2009 personal consumption expenditure index) and logged vehicle sales (correlation = 0.017). While non-durables grows by 2.5% after the 2016 election, whereas vehicle sales does not grow after the election. Figure 13 in Appendix Section

7.4.1 also compares consumption among conservatives and their counterparts before and after the election, finding almost identical parallel trends leading up to it.

An important potential concern with Gallup's measure of self-reported daily non-durable consumption expenditures is that, in the presence of increasing political partisanship [Gentzkow, 2016], individuals might report higher levels of consumption simply because of optimism about their party holding the presidency even if real economic outcomes have not changed [Mian et al., 2017]. To understand whether these survey measures of economic activity in Gallup are reliable, I regress an indicator for whether an individual reports their employer is hiring and logged daily non-durables consumption expenditures on county employment growth, controlling for county and day of the year fixed effects and the usual set of individual covariates. Importantly, I estimate these regressions separately for conservatives, moderates, and liberals. Table 5 documents these results. Beginning with the pooled sample in columns 1 and 5, a 1pp rise in county employment growth is associated with a 0.337pp rise in the probability that the individual reports that their employer is hiring more workers and a 0.053% rise in daily non-durable consumption expenditures. Moreover, in unreported regressions, I also control for individual life satisfaction to address the potential for omitted variables bias in time-varying sentiment and/or person-specific heterogeneity; these regressions produced statistically indistinguishable results.

Turning towards the results estimated separately for different political affiliations, I find that a 1pp rise in county employment growth is associated with a 0.334pp rise in the probability that conservatives report that their employer is hiring, but a null association that their consumption rises. In contrast, the gradients for both outcome variables are statistically significant and positive among moderates and liberals. If increases in the perception of employer expansion and/or consumption simply reflected political partisanship, then there should be a null correlation in each

of these specifications. Moreover, the fact that conservatives respond least elastically to increases in local employment growth should bias the results from the event study down, making me less likely to find that their consumption rises following the 2016 Presidential election.[15]

[INSERT TABLE 5 HERE]

4.3 Robustness Using Alternative Strategies

One potential concern with these results is that they are confounded by other potential time-varying shocks to consumption. To examine the robustness, I draw on three alternative identification strategies. One limitation in each of two of these exercises, however, is that only the 2008 to 2013 sample from Gallup contains respondent replies for both the consumption and sentiment variables; following 2013, respondents answer either one set of the survey questions or the other, but not both. The reduction in sample size limits the amount of identifying variation, which will contribute to somewhat low first-stage correlations.

The first strategy leverages the fact that individuals reside in areas that are heterogeneously exposed to national gasoline price shocks. Exploiting cross-sectional variation in the share of output that comes from the energy sector, which is most exposed to gasoline price shocks, I construct a Bartik-like instrument of the form

$$Z_{st}^{BARTIK} = (GDP_{s,2006}^{ENERGY}/GDP_{s,2006})\Delta p_t \tag{4}$$

---

[15]As an additional diagnostic, I examine whether individuals who are liquidity constrained and/or report worrying more about money also report fewer consumption expenditures. Regressing logged non-durable consumption on my measure of liquidity constraints produces a gradient of -0.147 (p-value = 0.00), suggesting that these individuals consume 14.7% less than their counterparts. There is also little evidence of heterogeneous treatment effects when partitioning the sample by political worldview. For example, the gradient is -0.18 for conservatives, -0.102 for moderates, and -0.117 for liberals (all significant at the 1% level). In this sense, individuals reporting that they are worried about not having enough money also report lower daily consumption expenditures.

where $GDP_{s,2006}^{ENERGY}/GDP_{s,2006}$ denotes the energy share of output in a state (baseline year = 2006) obtained from the Bureau of Economic Analysis (BEA) and Δp_t denotes the growth rate of the national gasoline price obtained form the St. Louis Federal Reserve. The identifying assumption is that fluctuations in the national price of gasoline affect consumption only through their effects on beliefs about the economy, which is moderated by the fact that individuals exposed to greater gasoline price shocks become more pessimistic about the current and future state of the economy [Binder and Makridis, 2018]. Moreover, since the energy sector exhibits significant fixed costs, these differences across states are likely to emerge from largely historical and geographic factors that make some states better places to produce energy over others.

The second strategy leverages the fact that extreme temperatures affect individuals mood and, therefore, beliefs about the state of the economy [Makridis, 2018b]. Exploiting day-to-day fluctuations in local temperature, I instrument sentiment using average daily county temperature, its square, and interactions with occupational fixed effects to capture the fact that certain types of occupations are more exposed to outside temperatures (e.g., construction). The identifying assumption is that fluctuations in daily county temperature affect individual consumption only through their effects on beliefs about the economy, which is moderated by the fact that individuals exposed to more extreme temperatures become more pessimistic about the current and future state of the economy [Makridis, 2018a]. Importantly, precipitation is included in the vector of controls since rainfall may delay the timing of consumption expenditures [Agarwal et al., 2017].

The third strategy leverages the fact that individuals in different counties are exposed to different information based on heterogeneity in social networks across counties. For example, an individual living in San Jose, CA, will be exposed to different information shocks than an individual living in Phoenix, AZ, based on their social network. Using data from Facebook in April 2016,

Bailey et al. [2018a] introduce the social connectedness index (SCI), which captures the number of friends between county c and every other county c' in the United States. Motivated by the result that local housing price shocks affect beliefs about the economy from Section 3, as well as recent results from Bailey et al. [2018b] the diffusion of information on housing prices through social networks, I now construct a connectivity-weighted measure of housing price growth using the Federal Housing Administration's (FHA) housing price index. The identifying assumption is that housing price shocks to individuals in a connected county affect individual consumption only through their effects on individual beliefs about the economy. Importantly, a county's ties with other counties does not need to be random—only changes in the shocks to connected counties.[16]

Table 6 documents the results associated with each of these identification strategies with supplemental material about the first-stage correlations in Appendix Section 7.4.3. Starting with the Bartik instrument in column 1, a standard deviation rise in economic sentiment is associated with a .157% rise in daily non-durable consumption. The F-statistic is 25.5, which is well-above the recommendation from Stock and Yogo [2005]. One potential reason for the lower magnitude of the gradient arises from the negative impact that gasoline price shocks have on consumption through complementarities between the energy sector and, for example, manufacturing. Turning towards the temperature instrument in column 2, a standard deviation rise in economic sentiment is associated with a .223% rise in daily non-durable consumption. While the F-statistic is weak, the second-stage estimate on the consumption elasticity is still significant at a 1% level.

To further validate these results, column 3 presents the consumption elasticity when using SCI-

---

[16]The instrument shares many overlapping features with the logic in the well-known China syndrome shock from Autor et al. [2013] where counties differ in their exposure to import trade competition. Whereas they used Chinese exposure to import competition from foreign countries, which generates movements in China's exports and thus import competition among U.S. counties, my approach exploits exposure to counties that experienced larger versus smaller housing price fluctuations.

weighted housing price growth. Put differently, how does housing price growth in every county c' connected with c affect beliefs in county c after weighting appropriately based on the strength of each tie between c and c'? The identification strategy here follows from Bailey et al. [2018b] that the SCI is plausibly exogenous. Given that I am controlling for county fixed effects, I isolate variation in individuals' exposure to heterogeneous information based on the housing price shocks that connected counties are experiencing between 2008 and 2017. To maximize the variation in the sample, I produce a weighted average of beliefs at a county-level since starting in 2014 Gallup stops asking households questions about both consumption and economic optimism. After instrumenting for beliefs, I find that a standard deviation rise in beliefs about the current state of the economy is associated with a 0.26% rise in daily consumption.[17],[18]

[INSERT TABLE 6 HERE]

How does the consumption elasticity potentially vary based on whether a county experienced a relatively moderate decline or potential increase versus a large decline in housing prices during the financial crisis? Using the FHA housing price index, I compute growth for each county between 2007 and 2012, subsequently binning counties based on being above or below the median housing price growth of -8%. Estimating the consumption elasticities separately by group implies a gradient of 0.205 (p-value $= 0.231$, F-stat $= 22.9$) for counties above the median and 0.358 (p-value $= 0.063$,

---

[17]Controlling for a county's own housing price growth should not matter unless exposure to different housing price shocks through network effects affects housing prices directly. Although it is possible among the sample of real estate developers, they are more likely to make investment decisions based on their own experience. Nonetheless, adding county housing price growth as a control changes the gradient to 0.29, although it becomes less statistically precise (p-value $= 0.064$, F-stat $= 47.52$).

[18]Why are these estimates larger in magnitude than the earlier estimates from the event study? Since the event study leverage an unanticipated shock to expectations, it did measured not only beliefs about the economy, but also a broader set of potentially political, social, and cultural expectations associated with the election results. Depending on how these expectations are correlated with individual consumption, it could generate a higher or lower consumption elasticity. The estimates from Table 6, however, are based directly on plausibly exogenous variation in economic sentiment on individual consumption.

F-stat = 39) for counties below the median. While the estimates are not precise enough to reject the null that they are equal, the estimates suggest that decreases in housing prices may reduce consumption more than increases in housing prices raise consumption. In contrast, there is a marginally larger elasticity of 0.29 (p-value = 0.045), relative to 0.22 (p-value = 0.185), among counties that experienced above median employment growth of -4% between 2007Q1-2012Q1.

5 Understanding the Protracted Recovery and Discussion

What do these elasticities imply about the aggregate effects of sentiment on consumption? Although a full general equilibrium model is required to make definitive statements, I implement a simple back-of-the-envelope calculation. According to the St. Louis Federal Reserve, personal consumption expenditures on non-durable goods declined by 9.54% between the height of the boom and bottom of the bust from \$2.35 to \$2.12 billion (https://fred.stlouisfed.org/series/PCND). Using the Gallup World Poll, which provides a pre-2008 measure, I find that a similar measure of economic optimism declines by 27% from 67.39/100 to 48.74/100.[19] Putting these aggregate facts together with the consumption elasticity, the decline in economic sentiment accounts for between 34-68% ($= \varepsilon \times 0.27/.094$) of the decline in consumption during the Great Recession where ε is set to either 0.12 or 0.236 depending on the preferred elasticity or reference group.[20]

If local housing and labor market shocks affect individuals beliefs about the aggregate state

---

[19]The question is based off of respondent answers that are used to construct an "optimism index" that parallels the question in the U.S. Daily Poll. The index in 2006 is 74.60, in 2007 is 67.39, in 2009 is 48.75, in 2010 is 60.75, in 2011 is 59.77. Changes in the reference year will, not surprisingly, alter the back-of-the-envelope calculation. Therefore, I compare 2007 with 2009 as the baseline.

[20]What explains the potentially large aggregate effects of sentiment shocks? Although local shocks clearly matter for shaping beliefs, shocks to connected networks may also matter. Using the SCI introduced earlier, Table 13 Appendix Section 7.5 examines how labor and housing market shocks to connected counties affects individual beliefs even after controlling for a county's own shocks. I find that SCI-weighted housing price shocks are roughly half the magnitude of a county's own housing price shocks.

of the economy, and sentiments affect real consumption expenditures, then a natural question is whether fluctuations in economic sentiment can also help explain the protracted recovery from the Great Recession. To gauge the quantitative importance of sentiments on the persistent recovery, I measure the intensity of economic confidence between 2008 and 2009 across states and compute the number of months it takes for each state to return to its pre-crisis 2006 level of housing price growth and real GDP (starting in 2008Q1). Evaluating the decline in sentiment at the mean, I compare economic optimism in 2006 (index = 74.60) with its approximated level in 2010 (60.75) since 2009 is not available, which amounts to an 18.6% decline.

Figure 6 shows that a 10pp rise in economic confidence is associated with taking 23 additional months for housing price growth to recover to pre-crisis rates and 7.7 additional months for real GDP to recover to pre-crisis levels. In light of the mean decline of 18.6%, then, absent the decline in sentiment, housing prices would have recovered 43 months sooner ($= 2.3 \times 18.56$) and real GDP would have recovered 14.3 months sooner ($= 0.77 \times 18.56$). Figures 21 and 22 in Appendix Section 7.5 provides additional robustness using perceptions that the economy is improving. Although these are not causal pieces of evidence, they are consistent with the view that differences in economic sentiment help explain the delayed recovery from the Great Recession. Figure 23 in Appendix Section 7.5 also provides additional evidence that low economic sentiment in 2016 explains 63% of the voter turnout in the 2016 Presidential election.

[INSERT FIGURE 6 HERE]

5.1 Comparison with Recent Research and Implications for Macroeconomic Models

Ever since Carroll et al. [1994] and Bram and Ludvigson [1998], economists have recognized that consumer sentiment has predictive power of consumption even after controlling for a number of traditional macroeconomic factors. However, what was lacking was more credible microeconomic evidence. My results provide comprehensive and unified evidence for these macroeconomic models that feature sentiments and precautionary savings under uncertainty. For example, Benhabib et al. [2015] build a model featuring waves of optimism and pessimism under rational expectations, with no externalities, and with no non-convexities. Given an information friction—that firms cannot separately identify the component of demand stemming from consumer sentiments versus idiosyncratic preference shocks—sentiments that are unrelated to economic fundamentals can affect output and employment. Heathcote and Perri [forthcoming] also develop a model with sunspot-driven fluctuations where large declines in asset prices lead to a decline in confidence and subsequent decline in consumption, especially among low wealth households. Mertens and Ravn [2014] develop an alternative model that also emphasizes how a persistent decline in consumer confidence can generate a liquidity trap that depresses consumption and investment.

Turning towards broader work on precautionary savings, Carroll [1992] was among the first to show that an increase in uncertainty causes the level of consumption to fall as consumers build up their stock of assets. Challe et al. [2017] develop a general framework incorporate incomplete insurance and heterogeneous agents to understand how individuals undertake precautionary savings against unemployment risk. They find that the aggregate demand effect has largely contributed

to the amplification and propagation of the Great Recession. Ravn and Sterk [2017] also develop a heterogeneous agents model with incomplete markets and nominal rigidities. However, their focus is on the search process. An increase in job uncertainty can decrease aggregate demand, which feeds back into lower hiring and produces even greater uncertainty. Den Haan et al. [2017] show that the interaction between incomplete markets and nominal wages is important for capturing the effect of unemployment fears on precautionary sentiments, whereas they actually work in the opposite way when implemented in isolation. Beaudry et al. [2018] studies how past over-investment can lead to low productivity for a potentially prolonged period of time using precautionary behavior as a mechanism: given less demand in the durable sector, individuals fearing unemployment reduce demand for non-durables as well.

Perhaps most closely, however, my results provide specific estimates over a recent contribution from Kozlowski et al. [2017] that shocks to beliefs can have a persistent effect on macroeconomic aggregates, like real GDP. For example, since few people in 2006 expected a financial crisis around the corner, an unanticipated shocks to these beliefs can have a significant effect on individual beliefs about economic activity. Importantly, these beliefs may take time to update even after fundamentals have adjusted. Although Kozlowski et al. [2017] provide a theoretical model that illustrates these belief shocks can help explain the slow recovery from the Great Recession, this paper is the first to provide causal and specific microeconomic evidence.

Finally, my results also relate with an ongoing debate about how sentiments are related with stock prices. For example, more recent work from Lemmon and Portniaguina [2006] find that measures of consumer confidence from the University of Michigan's consumer confidence survey forecasts the size premium (i.e., the tendency of stocks for firms with a smaller market capitalization to outperform the stocks for firms with a larger market capitalization). Kaplanski and Levy

[2010] find that stock prices in the airline industry sharply "over-react" to plane crashes, consistent with behavioral models of salience. Baker and Wurgler [2006] find that sentiments affect stock prices for companies with which have ambiguous investor expectations.

To test how these measures of sentiments relate with fluctuations in stock prices, I now draw on a panel of monthly monthly stock returns from publicly traded companies and regress logged firm stock prices on a z-score of monthly perceptions about the current state of the economy. The implied coefficients producing coefficient on sentiment is 1.28 (p-value = 0.00) and 0.76 (p-value = 0.00) when fixed effects on firm, month, and year are included. Interestingly, a regression of the standard deviation of logged stock prices (across all publicly traded companies) on the standard deviation of perceptions about the state of the economy produces a gradient of -0.869 (p-value = 0.00), reflecting that greater sentiment volatility depresses stock prices. While these estimates are not causal, they are nonetheless consistent with the types of exercises implemented by Carroll et al. [1994] and Bram and Ludvigson [1998].

6 Conclusion

Housing wealth and employment declined significantly during the Great Recession, remaining low for more years than any prior recession in the U.S. post-war era. Using new micro-data from Gallup's U.S. Daily Poll, this paper provides microeconomic evidence in support of recent macroeconomic models that have emphasized the role of sentiments as a source behind sluggish consumer demand despite the improvement in the unemployment rate and stock market.

Consistent with an emerging evidence that individuals form beliefs based on their personal experience, I begin by providing additional empirical support behind an emerging consensus that

personal experience matters in forming beliefs about economic activity. To understand how fluctuations in economic optimism affects consumption expenditures on non-durable goods, I exploit several sources of plausibly exogenous variation. For example, using the 2016 Presidential election as an event study, I show that consumption rises by 4.2% even after controlling for stated life satisfaction. However, these estimates are concentrated among conservatives whose consumption grew by 10-12% after the election. Turning towards three separate instrumental variables specifications, I find that a standard deviation rise in beliefs is associated with a 0.15-0.26% rise in consumption. A back-of-the-envelope calculation with these elasticities suggests the decline in beliefs during the financial crisis accounts for 34-68% of the decline in aggregate consumption of non-durable goods. Moreover, areas with lower sentiment take significantly longer to recover from the financial crisis measured in terms of housing price growth, real output, and consumption.

The evidence here only touches the surface and towards several fruitful areas of further research. First, how are happiness and sentiments related? Motivated by a large literature on the cyclicality of happiness [di Tella et al., 2001, 2003], there is recent interest in using well-being data to infer information about marginal rates of substitution [Benjamin et al., 2014a,b]. One possibility is that economic sentiments are in part influenced about the individual's underlying hedonic state. Second, even after controlling for income and individual covariates, there is a great deal of residual variation in sentiment. Is all of this residual variation explained by worry over the business cycle, or are there other potentially important cyclical determinants, like time-varying risk aversion [Guiso et al., 2015]? Third, what are the mechanisms through which sentiments affect real economic activity? While an obvious channel that was tested here is the decline in consumer spending, heterogeneity in beliefs about the economy may also affect stock market participation and the over-accumulation of capital as in Perri and Quadrini [2016] and Beaudry et al. [2018]. Integrating

microeconomic data with these macroeconomic heterogeneous agent models will be an essential

step forward in understanding these broader phenomena.

References

Klaus Adam, Albert Marcet, and Johannes Beutel. Stock price booms and expected capital gains. *American Economic Review*, 107(8):2352–2408, 2017.

Manuel Adelino, Antoinette Schoar, and Felipe Severino. Loan originations and defaults in the mortgage crisis: The role of the middle class. *Review of Financial Studies*, 29(7):1635–1670, 2016.

Manuel Adelino, Antoinette Schoar, and Felipe Severino. Dynamics of housing debt in the recent boom and bust. *NBER Macro Annual 2017*, 32, 2018a.

Manuel Adelino, Antoinette Schoar, and Felipe Severino. Perception of house price risk and homeownership. *Working paper*, 2018b.

Sumit Agarwal, J. Bradford Jensen, and Ferdinando Monte. The geography of consumption. *NBER working papers*, 2017.

D. Aigner. Regression with a binary independent variable subject to errors of observations. *Journal of Econometrics*, 1:49–60, 1973.

George A. Akerlof and Janet L. Yellen. A near-rational model of the business cycle with wage and price inertia. *Quarterly Journal of Economics*, 100:823–838, 1985.

George-Marios Angeletos and Jennifer La'O. Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics*, 56(S1), 2009.

George-Marios Angeletos and Jennifer La'O. Sentiments. *Econometrica*, 81(2):739–779, 2013.

George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Quantifying confidence. *Working paper*, 2014.

David H. Autor, David Dorn, and Gordon H. Hanson. The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.

Costas Azariadis, Leo Kaas, and Yi Wen. Self-fulfilling credit cycles. *Review of Economic Studies*, 83: 1364–1405, 2016.

Philippe Bacchetta and Eric van Wincoop. The Great Recession: A self-fulfilling global panic. *American Economic Journal: Macroeconomics*, 8(4):177–198, 2016.

Philippe Bacchetta, Cedric Tille, and Eric van Wincoop. Self-fulfilling risk panics. *American Economic Review*, 102(7):3674–3700, 2012.

Mike Bailey, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, 32(3):259–280, 2018a.

Mike Bailey, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, 126(6):2224–2276, 2018b.

Malcolm Baker and Jeffrey Wurgler. Investor sentiment and the cross-section of stock returns. *Journal of Finance*, 61(4):1645–1680, 2006.

Scott R. Baker. Debt and the consumption response to household income shocks. *Journal of Political Economy*, forthcoming.

Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

Robert B. Barsky and Eric R. Sims. Information, animal spirits and the meaning of innovations in consumer confidence. *American Economic Review*, 102(4):1343–1377, 2012.

Patrick Baylis. Temperature and temperament: Evidence from Twitter. *Working paper*, 2018.

Patrick Baylis, Nick Obradovich, Yury Kryvasheyeu, Haohui Chen, Lorenzo Coviello, Esteban Moro, Manuel Cebrian, and James H. Fowler. Weather impacts expressed sentiment. *PLoS One*, 13(4), 2018.

Paul Beaudry and Franck Portier. Stock prices, news and economic fluctuations. *American Economic Review*, 96(4):1293–1307, 2006.

Paul Beaudry, Deokwoo Nam, and Jian Wang. Do mood swings drive business cycles and is it rational? *Working paper*, 2011.

Paul Beaudry, Dana Galizia, and Franck Portier. Reconciling Hayek's and Keynes' views of recessions. *Review of Economic Studies*, 85(1):119–156, 2018.

Jess Benhabib and Mark M. Spiegel. Sentiments and economic activity: Evidence from U.S. states. *Working paper*, 2016.

Jess Benhabib, Pengfei Wang, and Yi Wen. Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2):549–585, 2015.

Daniel J. Benjamin, Ori Heffetz, Miles S. Kimball, and Alex Rees-Jones. What do you think would make you happier? What do you think you would choose? *American Economic Review*, 102(5):2083–2110, 2012.

Daniel J. Benjamin, Ori Heffetz, Miles S. Kimball, and Alex Rees-Jones. Can marginal rates of substitution be inferred from happiness data? Evidence from residency choices. *American Economic Review*, 104(11):3498–3528, 2014a.

Daniel J. Benjamin, Ori Heffetz, Miles S. Kimball, and Nichole Szembrot. Beyond happiness and satisfaction: Toward well-being indices based on stated preference. *American Economic Review*, 104(9): 2698–2735, 2014b.

Martin Beraja, Erik Hurst, and Juan Ospina. The aggregate implications of regional business cycles. *Econometrica*, forthcoming, 2019.

Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275, 2004.

Carola Conces Binder and Christos A. Makridis. Gas prices, sentiment, and consumption. *Working paper*, 2018.

Nicholas Bloom, Renata Lemos, Raffaella Sadun, Daniela Scur, and John Van Reenen. The new empirical economics of management. *Journal of European Economic Association*, 12(4), 2014.

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta, and Stephen Terry. Really uncertain business cycles. *Econometrica, R&R*, 2015.

Alexander N. Bogin, William M. Doerner, and William D. Larson. Local house price dynamics: New indices and stylized facts. *Real Estate Economics*, forthcoming.

Jason Bram and Sydney C. Ludvigson. Does consumer confidence forecast household expenditure? A sentiment index horse race. *Federal Reserve Bank of New York: Economic Policy Review*, 4(2):59–78, 1998.

Craig Burnside, Martin Eichenbaum, and Sergio Rebelo. Understanding booms and busts in housing markets. *Journal of Political Economy*, 124(4):1088–1147, 2016.

Brandice Canes-Wrone and Jee-Kwang Park. Electorical business cycles in OECD countries. *American Political Science Review*, 106(1):103–122, 2012.

Christopher Carroll. The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity*, 23(2):61–156, 1992.

Christopher Carroll and Wendy Dunn. Unemployment expectations, jumping (S,s) Triggers and household balance sheets. *NBER Macroeconomics Annual 1997*, 12:165–230, 1997.

Christopher D. Carroll, Jeffrey C. Fuhrer, and David W. Wilcox. Does consumer sentiment forecast household spending? If so, why? *American Economic Review*, 84(5):1397–1408, 1994.

Edouard Challe, Julien Matheron, Xavier Ragot, and Juan F. Rubio-Ramirez. Precautionary saving and aggregate demand. *Quantitative Economics*, 8(2):435–478, 2017.

Thomas Crossley and Hamish Low. Job loss, credit constraints and consumption growth. *Review of Economics and Statistics*, 9(5):876–884, 2014.

Angus Deaton. The financial crisis and the well-being of Americans. *Oxford Economic Papers*, 64(1): 1–26, 2012.

Wouter J. Den Haan, Pontus Rendahl, and Markus Riegler. Unemployment (fears) and deflationary spirals. *Journal of the European Economic Association*, 2017.

Tatyana Deryugina and Solomon M. Hsiang. The marginal product of climate. *NBER working paper*, 2017.

Olivier Deschenes and Michael Greenstone. Climate change, mortality and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, 3:152–185, 2011.

Rafael di Tella, Robert J. MacCulloch, and Andrew J. Oswald. Preferences over inflation and unemployment: Evidence from surveys of happiness. *American Economic Review*, 91(1):335–341, 2001.

Rafael di Tella, Robert MacCulloch, and Andrew Oswald. The macroeconomics of happiness. *Review of Economics and Statistics*, 85:809–827, 2003.

Roger E. Farmer. The stock market crash of 2008 caused the Great Recession: Theory and evidence. *Journal of Economic Dynamics and Control*, 36:693–707, 2012.

John G. Fernald, Robert E. Hall, James H. Stock, and Mark W. Watson. The disappointing recovery of output after 2009. *NBER working paper*, 2017.

Milton Friedman. Interview. *The Commonwealth, a magazine of the Commonwealth Club of California*, 1992.

Matthew Gentzkow. Polarization in 2016. *Toulouse Network of Information Technology, white paper*, 2016.

Alan S. Gerber and Gregory A. Huber. Partisanship, political control, and economic assessments. *American Journal of Political Science*, 54(1):153–173, 2010.

Christian Gillitzer and Nalini Prasad. The effect of consumer sentiment on consumption: Cross-sectional evidence from elections. *American Economic Journal: Macroeconomics*, 10(4):234–269, 2018.

Edward L. Glaeser and Charles G. Nathanson. An extrapolative model of house price dynamics. *Journal of Financial Economics*, 126(1):147–170, 2017.

Edward L. Glaeser, Joshua D. Gottlieb, and Oren Ziv. Unhappy cities. *Journal of Labor Economics*, 34 (2):S129–S182, 2016.

Daniel Green, Brian T. Melzer, Jonathan A. Parker, and Arcenis Rojas. Accelerator or broke? Cash for clunkers, household liquidity and aggregate demand. *Working paper*, 2016.

Luigi Guiso, Paola Sapienza, and Luigi Zingales. Time varying risk aversion. *Working paper*, 2015.

Jonathan Heathcote and Fabrizio Perri. Wealth and volatility. *Review of Economic Studies*, forthcoming.

James Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.

Nathaniel Hendren. Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7):1778–1823, 2017.

Daniel Kahneman and Angus Deaton. High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, 107(38):16489–16493, 2010.

Greg Kaplan and Giovanni L. Violante. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239, 2014.

Greg Kaplan, Kurt Mitman, and Gianluca Violante. Consumption and house prices in the Great Recession: Model meets evidence. *Journal of Political Economy, R&R*, 2019.

Guy Kaplanski and Haim Levy. Sentiment and stock prices: The case of aviation disasters. *Journal of Financial Economics*, 95:174–201, 2010.

John M. Keynes. The general theory of employment, interest and money. *Palgrave Macmillan*, 1936.

Lutz Kilian. The economic effects of energy price shocks. *Journal of Economic Literature*, 46(4):871–909, 2008.

Julian Kozlowski, Laura Veldkamp, and Venky Venkateswaran. The tail that wags the economy: Beliefs and persistence stagnation. *Working paper*, 2017.

Alan B. Krueger and David A. Schkade. The reliability of subjective well-being measures. *Journal of Public Economics*, 92:1833–1845, 2008.

Theresa Kuchler and Basit Zafar. Personal experiences and expectations about aggregate outcomes. *Journal of Finance*, 2019.

Lorenz Kueng. Tax news: The response of household spending to changes in expected taxes. *Working paper*, 2016.

Andresa Lagerborg, Evi Pappa, and Morten Ravn. Sentimental business cycles. *Working paper*, 2018.

Michael Lemmon and Evgenia Portniaguina. Consumer confidence and asset prices: Some empirical evidence. *Review of Financial Studies*, 19(4):1499–1529, 2006.

Sydney C. Ludvigson. Consumer confidence and consumer spending. *Journal of Economic Perspectives*, 18(2):29–50, 2004.

Christos Makridis. Hotter Days, Better Days? The Effect of Temperature on Economic Sentiment. *Working paper*, 2018a.

Christos Makridis. Can You Feel the Heat? Extreme Temperatures, Stock Returns, and Economic Sentiment. *Working paper*, 2018b.

Christos Makridis and Tyler Ransom. Beating the Heat: Temperature and Spatial Reallocation over the Short and Long-run. *Working paper*, 2017.

Ulrike Malmendier and Stefan Nagel. Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87, 2016.

Mary C. McGrath. Economic behavior and the partisan perceptual screen. *Working paper*, 2017.

Karel R. S. M. Mertens and Morten O. Ravn. Fiscal policy in an expectations-driven liquidity trap. *Review of Economic Studies*, 81(4):1637–1667, 2014.

Atif Mian and Amir Sufi. The consequences of mortgage credit expansion: Evidence from the U.S. mortgage default crisis. *Quarterly Journal of Economics*, 124(4):1449–1496, 2009.

Atif Mian and Amir Sufi. House prices, home equity-based borrowing and the US household leverage crisis. *American Economic Review*, 101(5):2132–56, 2011.

Atif Mian and Amir Sufi. What explains the 2007-2009 drop in employment? *Econometrica*, 82(6): 2197–2223, 2014.

Atif Mian, Kamalesh Rao, and Amir Sufi. Household balance sheets, consumption and the economic slump. *Quarterly Journal of Economics*, 128(4):1687–1726, 2013.

Atif Mian, Amir Sufi, and Francesco Trebbi. Foreclosures, house prices and the real economy. *Journal of Finance*, 70(6):2587–2634, 2015.

Atif R. Mian, Amir Sufi, and Nasim Khoshkhou. Partisan bias, economic expectations, and household spending. *Working paper*, 2017.

Fabio Milani. Sentiment and the U.S. business cycle. *Journal of Economic Dynamics & Control*, 82: 289–311, 2017.

Kanishka Misra and Paolo Surico. Consumption, income changes, and heterogeneity: Evidence from fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4):84–106, 2014.

Stephen Morris and Hyun-Song Shin. The social value of public information. *American Economic Review*, 92:1521–1534, 2002.

Stephen Morris and Hyung Song Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–597, 1998.

Andrew J. Oswald. On the curvature of the reporting function from objective reality to subjective feelings. *Economics Letters*, 100(3):369–372, 2008.

Andrew J. Oswald and Stephen Wu. Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science*, 327(5965):576–579, 2010.

Andrew J. Oswald and Stephen Wu. Well-being across America. *Review of Economics and Statistics*, 93 (4):1118–1134, 2011.

Pablo Ottonello and Thomas Winberry. Financial heterogeneity and the investment channel of monetary policy. *Econometrica, R&R*, 2018.

Fabrizio Perri and Vincenzo Quadrini. International recessions. *Working paper*, 2016.

Monika Piazzesi and Martin Schneider. Momentum traders in the housing market: Survey evidence and a search model. *American Economic Review*, 99(2):406–411, 2009.

Luigi Pistaferri. Why has consumption remained moderate after the Great Recession. *Working paper*, 2016.

Morten O. Ravn and Vincent Sterk. Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125–141, 2017.

Daniel Shoag and Stan Veuger. Uncertainty and the geography of the Great Recession. *Journal of Monetary Economics*, 84:84–93, 2016.

Christopher A. Sims. Implications of rationale inattention. *Journal of Monetary Economics*, 50(3): 665–690, 2003.

J. H. Stock and M. Yogo. Testing for weak instruments in linear IV regression. *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 2005.

Johannes Stroebel and Joseph Vavra. House prices, local demand and retail prices. *Journal of Political Economy*, forthcoming.

Daniel Sullivan and Till Wachter. Job displacement and mortality: An analysis using administrative data. *Quarterly Journal of Economics*, 124(3):1265–1306, 2009.

John B. Taylor. The role of policy in the Great Recession and the weak recovery. *American Economic Review*, 104(5):61–66, 2014.

7 Tables and Figures

Variable	Survey Question	Rating
Life Satisfaction	Please imagine a ladder with steps numbered from zero at the bottom to ten at the top. The top of the ladder represents the best possible life for you and the bottom of the ladder represents the worst possible life for you. On which step of the ladder would you say you personally feel you stand at this time?	0-10 scale
Perception of Current Economic Activity	How would you rate economic conditions in this country today: as excellent, good, only fair, or poor?	1-4 scale
Perception of Future Economic Activity	Right now, do you think that economic conditions in this country, as a whole, are getting better or getting worse?	1-3 scale
Hiring	Now thinking more generally about the company or business you work for, including all of its employees. Based on what you know or have seen, would you say that, in general, your company or employer is (a) hiring new people and expanding the size of its workforce, (b) not changing the size of its workforce, or (c) letting people go and reducing the size of its workforce.	1-3 scale
Non-durables consumption expenditures	Next, we'd like you to think about your spending yesterday, not counting the purchase of a home, motor vehicle, or your normal household bills. How much money did you spend or charge yesterday on all other types of purchases you may have made.	Continuous

Table 1: Main Gallup Survey Questions

Notes.–Sources: Gallup. The table reports the survey questions and associated rating index used by Gallup when speaking with respondents.
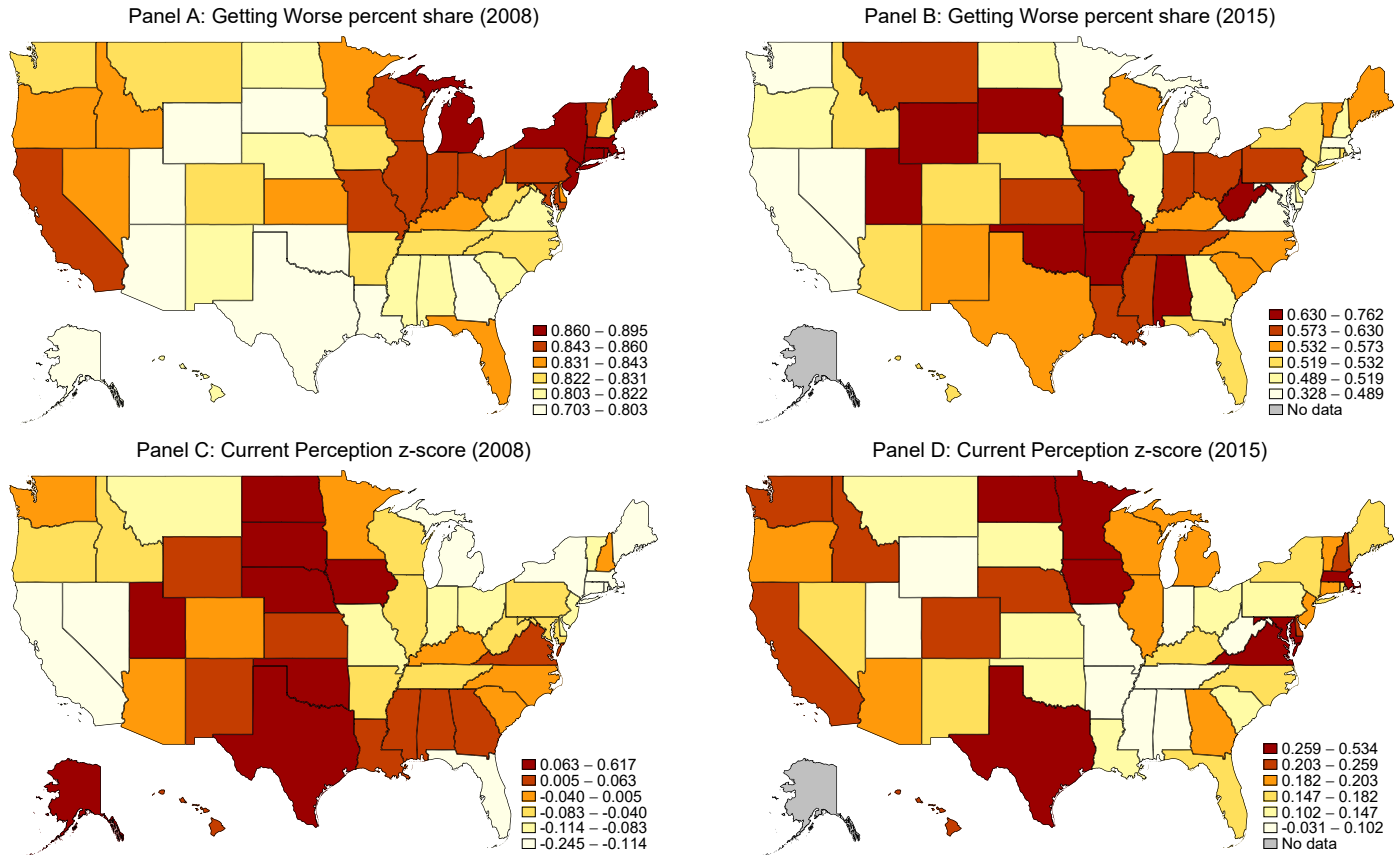
Figure 1: Spatial Variation in Economic Sentiments

Notes.–Sources: Gallup. The figure plots the spatial variation across states for two sets of questions. The first (used to produce Panels A and B) asks participants: "Right now, do you think that economic conditions in this country, as a whole, are getting better or getting worse?" I subsequently compute the share of individuals in a state who report getting worse. The second (Panels C and D) asks participants: How would you rate economic conditions in this country today: as excellent, good, only fair, or poor?" I subsequently compute the z-score of the one to four index (coded so that higher is better). Sample weights are used to produce the state averages in the different years.

Figure 2: Measuring Uncertainty, Comparison Between Gallup and Baker et al. [2016]

Notes.–Sources: Gallup, 2008-2017. The figure plots the mean perceptions of the state of the economy with the measure of policy uncertainty from Baker et al. [2016] at a monthly frequency and with the volatility index (provided through the St. Louis Federal Reserve) at a daily frequency. The figure shows that there is a strong negative correlation between the measure of current perception of the economy and both economic policy and aggregate uncertainty.
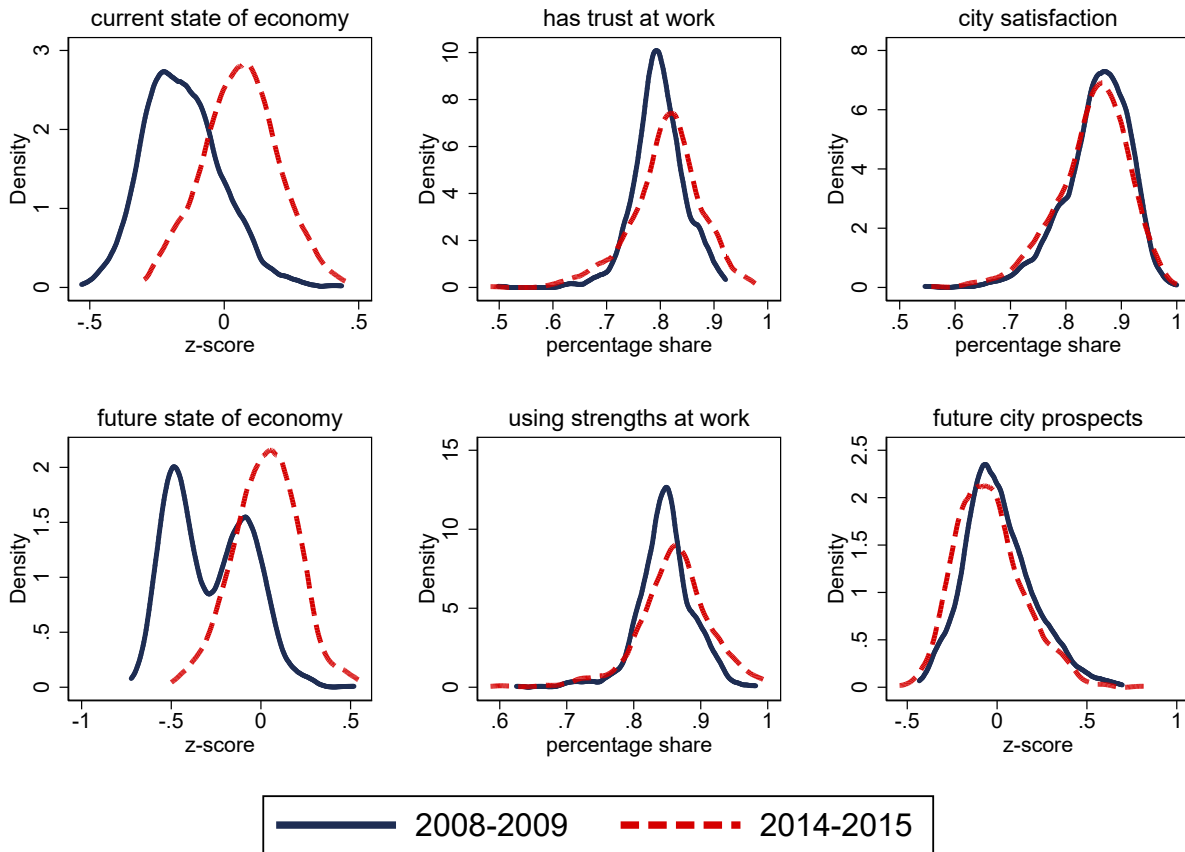
Figure 3: Dispersion in Sentiments Across Metro Areas, 2008-2009 and 2014-2015

Notes.–Sources: Gallup. The figure plots (i) the dispersion of standardized *z*-scores for the state of the economy, current and future in the first column, (ii) the dispersion of the fraction of people reporting (in a metro area) that they perceive trust at work and are able to leverage their strengths at work in the second column, and (iii) the dispersion of the fraction of people reporting that they are satisfied with their city and the standardized *z*-score for perceptions of future city prospects (getting worse, staying the same, getting better). in the third column. The index for the state of the economy ranges between 1 and 4: poor, only fair, good, and excellent. The work place practices measures are indicators, so their collapsed measures represent percent shares. City satisfaction is also an indicator, but future city prospects is an index (getting worse, the same, getting better). Each plot collapses across individuals within a metro area and secondly plots the kernel density across all metro areas. The sample is restricted to those metropolitan areas with at least 200 survey respondents in the data.

Figure 4: Distribution of Sentiments about the State of Economy

Notes.–Sources: Gallup. The figure begins by computing the *z*-score of the current and future state of the economy across years. The current state of the economy is an index with four values (poor, fair, good, excellent) and the future state of the economy is an index with three values (getting worse, staying the same, getting better). The variables are made continuous by averaging across all individuals within the same metro area. The sample is restricted to metro areas with over 250 observations. and collapsing to a metro-level by year with the survey sample weights. The figure subsequently plots the distribution of these values across metro areas.

Figure 5: Sentiments and Real Consumption Expenditures

Notes.–Sources: Gallup. The figure plots the daily *z*-score of the current state of the economy with daily real consumption expenditures on non-durable goods averaged across 1,000 individuals at a daily frequency. Nominal consumption is deflated using the 2009 real personal consumption expenditure index.

	Bartik Gasoline Shock	Temperature Shock	SCI Housing Shock
ln(non-durable consumption)	0.157***	0.223***	0.262***
	(0.0649)	(0.094)	(0.097)
F-statistic	25.5	3.54	87.34
R-squared	0.0256	0.024	0.089
N	622,830	561,900	850,532

Table 6: Alternative Identification Strategies to the Consumption-Sentiment Elasticity

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Federal Housing Administration, 2008-2017. The table reports the coefficients associated with regressions of logged daily consumption on non-durable consumption expenditures (deflated with the 2009 personal consumption expenditure index) on a *z*-score of perceptions about current economic activity, conditional on controls and both county and day of the year fixed effects, where perceptions about the economy are instrumented under three different specifications. The Bartik gasoline shock interacts energy output as a share of overall state output in 2006 interacted with weekly growth in the gasoline price, which captures the effect of gasoline price fluctuations on sentiment for individuals who are more versus less exposed to these shocks. The temperature shock uses a quadratic in daily average county temperature interacted with occupation fixed effects, which captures the effect of temperatures (and their non-linearities) on beliefs for individuals who are more versus less exposed to outside activities. The SCI housing shock uses the social connectedness index (SCI) from Bailey et al. [2018a] to produce a weighted housing price shock (year-to-year growth) for each county based on their connectivity with every other county in the United States, which captures the fact that information diffuses heterogeneously to individuals who are more versus less connected to areas with different housing price fluctuations. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level and sample weights are used.

Table 2: Baseline Results from Exposure to Local Labor and Housing Market Shocks

Dep. var. =	current state of the economy (z-score)				1[economy is improving]			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ ln(employment)	.65***		.30***	.30***	.26***		.12***	.13***
	[.06]		[.05]	[.06]	[.03]		[.03]	[.03]
Δ ln(housing price)		.71***	.67***	.67***		.28***	.27***	.27***
		[.06]	[.06]	[.06]		[.03]	[.03]	[.03]
Δ ln(trade employment)				-.01				.00
				[.03]				[.01]
Δ ln(finance employment)				-.01				-.00
				[.02]				[.01]
R-squared	.09	.09	.09	.09	.09	.09	.09	.09
Sample Size	1676462	1677240	1676462	1663434	1676462	1677240	1676462	1663434
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Federal Housing Administration, 2008-2017. The table reports the coefficients associated with regressions of standardized (*z*-score) perceptions about the current state of the economy (a one to four index with higher values implying a better state) and an indicator denoting that the economy is improving on the year-to-year county × quarter employment growth rate and annual county housing price growth from Bogin et al. [forthcoming], conditional on individual controls, and fixed effects on county and day of the year. Columns 4 and 8 also include quarterly county employment growth in the trade sectors (NAICS 42, 44-45) and finance & real estate sectors (NAICS 52, 53). Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level and sample weights are used.

Table 3: Heterogeneity in the Exposure to Local Labor and Housing Market Shocks

Dep. var. =	current state of the economy (z-score)		1[state of the economy improving]	
	liquidity constrained	liquidity unconstrained	liquidity constrained	liquidity unconstrained
Δ ln(employment)	.47***	.54***	.12*	.16**
	[.14]	[.14]	[.07]	[.07]
Δ ln(housing price)	.29***	.28***	.15***	.19***
	[.06]	[.06]	[.03]	[.04]
R-squared	.08	.08	.07	.07
Sample Size	251529	348954	251529	348954
Controls	Yes	Yes	Yes	Yes
Year/Qtr FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2009-2015. The table reports the coefficients associated with regressions of standardized (*z*-score) individual beliefs of the current state of the economy (one to four index with higher values being better) and an indicator for whether the economy is improving on the year-to-year employment growth rate, median housing price per square foot growth rate, individual controls, and fixed effects on county and day of the year fixed effects. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Liquidity constrained individuals are those who report not having enough money to do the things they want to do, which is measured as a binary variable from 2009-2012 and on a one to five scale from 2013 onward. Standard errors are clustered at the county-level and sample weights are used.

Table 4: Exploiting the 2016 Presidential Election to Estimate the Consumption Response

Dep. var. =	ln(daily consumption expenditures)					
	(1)	(2)	(3)	(4)	(5)	(6)
1[post2016]	.042***	.014	.049	.042***	.016	.051
	[.008]	[.011]	[.034]	[.011]	[.014]	[.034]
1[somewhat conserv.]		.010	.013		-.019	-.012
		[.015]	[.015]		[.020]	[.021]
1[very conserv.]		.021	.030		.025	.034
		[.025]	[.025]		[.034]	[.035]
1[somewhat conserv.] × 1[post2016]		.072***	.072***		.084***	.080***
		[.019]	[.019]		[.024]	[.025]
1[very conserv.] × 1[post2016]		.105***	.113***		.088**	.099**
		[.033]	[.033]		[.042]	[.043]
R-squared	.04	.04	.08	.04	.04	.10
Sample Size	115476	109377	109201	72292	68515	68211
Controls	Yes	Yes	Yes	Yes	Yes	Yes
County FE	No	No	Yes	No	No	Yes
State x Year x Month FE	No	No	Yes	No	No	Yes
Sample	Jul-May	Jul-May	Jul-May	Sep-Mar	Sep-Mar	Sep-Mar

Notes.–Sources: Gallup, 2016-2017. The table reports the coefficients associated with regressions of logged daily consumption (spending on non-durables yesterday deflated using the 2010 personal consumption expenditure index) on an indicator for post-election periods (after November 8, 2016), an indicator for being somewhat conservative, an indicator for being very conservative (normalized to liberals and moderates), their interactions, individual controls, and county, and × year × month fixed effects. The sample is restricted to a neighborhood around the 2016 election, which is defined from July 2016 to May 2017 in columns 1-3 and from September 2016 to March 2017 in columns 4-6.Controls include: a quadratic in age, education fixed effects, race fixed effects, employment status, and gender. Standard errors are heteroskedasticity-robust and sample weights are used.

Table 5: Evaluating the Reliability of Reported Employer Hiring and Consumption

Dep. var. =	1[employer's firm is expanding]				ln(non-durables consumption)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ ln(employment)	.337***	.334***	.423***	.400***	.053**	-.017	.133***	.163**
	[.028]	[.041]	[.051]	[.077]	[.025]	[.045]	[.051]	[.067]
R-squared	.09	.08	.12	.15	.04	.05	.07	.08
Sample Size	1509328	336991	301076	203749	1136151	303750	268070	188232
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample	All	Conservative	Moderate	Liberal	All	Conservative	Moderate	Liberal

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2017. The table reports the coefficients associated with regressions of an indicator for whether the individual reports that their employer is expanding and logged non-durables consumption expenditures on the year-to-year county × quarter employment growth rate, conditional on individual controls, and fixed effects on county and day of the year. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level and sample weights are used.
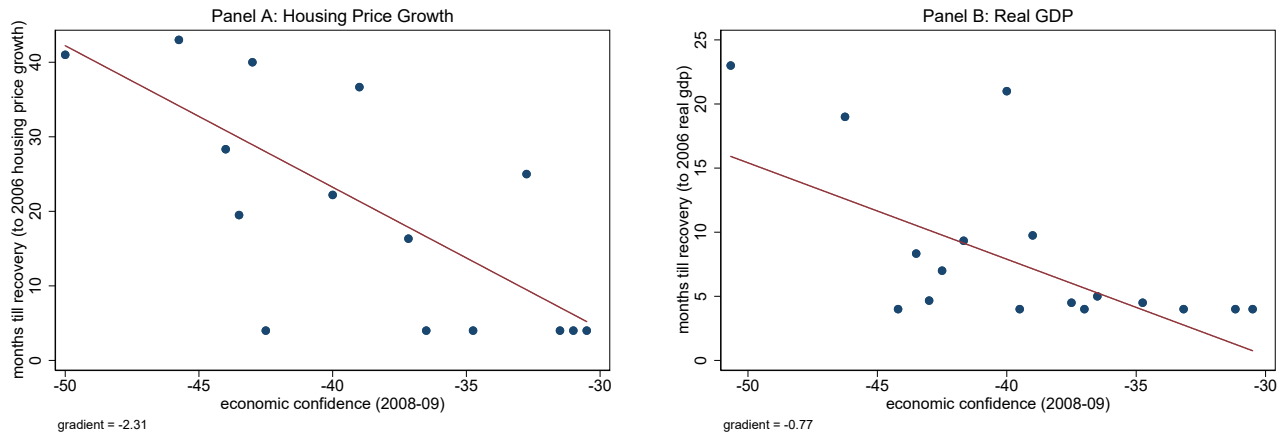
Figure 6: Belief Shocks and the Protracted Recovery (Housing and Real GDP)
Notes.–Sources: Gallup, Bureau of Economic Analysis, and Zillow, 2008-2017. The figure plots with `binscatter` average economic confidence between 2008-09 across states with the number of months it takes for housing prices and real GDP to recover to their pre-crisis rates and levels. The latter is computed by measuring the number of months it takes for a state to return to its 2006Q1 housing price growth and real GDP rate or level where counting begins in 2008Q1.

Online Appendix (Not For Print)

7.1 Was the Great Recession Different?

There is increasing evidence that the Great Recession was different from many other business cycles in the post-WWII era for the United States. Although it was not as severe as the Great Depression in the 1930s, it was effectively a financial crisis that led to one of the slowest recoveries for the U.S. [Taylor, 2014]. In particular, policymakers were especially concerned during the height of the crisis that individuals and institutional investors were losing confidence in the entire financial system, triggering a great deal of uncertainty.[21]

Drawing on the University of Michigan Survey of Consumers between 1978 and 2018, I investigate differences in sentiment over different business cycles over the past four decades. Figure 7 documents sentiment for four different measures. Panel A, for example, shows that individual

---

[21]https://money.cnn.com/2014/08/27/news/economy/ben-bernanke-great-depression/index.html

attitudes about business conditions relating to favorable, relative to unfavorable, news hit an all-time low during the Great Recession—even more than in the early 1980s. Panel B shows that individuals also feared at least as large of an increase in unemployment risk during the Great Recession, relative to prior recessions. Panels C and D show fairly similar trends for sentiment about the government response to the crisis and business conditions over the upcoming year. Notably, however, the Great Recession exhibits two large dips in pessimism about the government response to the crisis—once at the onset at the start of 2008 and again at the start of 2012. Although there are a lot of similarities between the Great Recession and other business cycles the U.S. has experienced, the University of Michigan survey data suggests that the recent recession was unique in at least a few key ways.
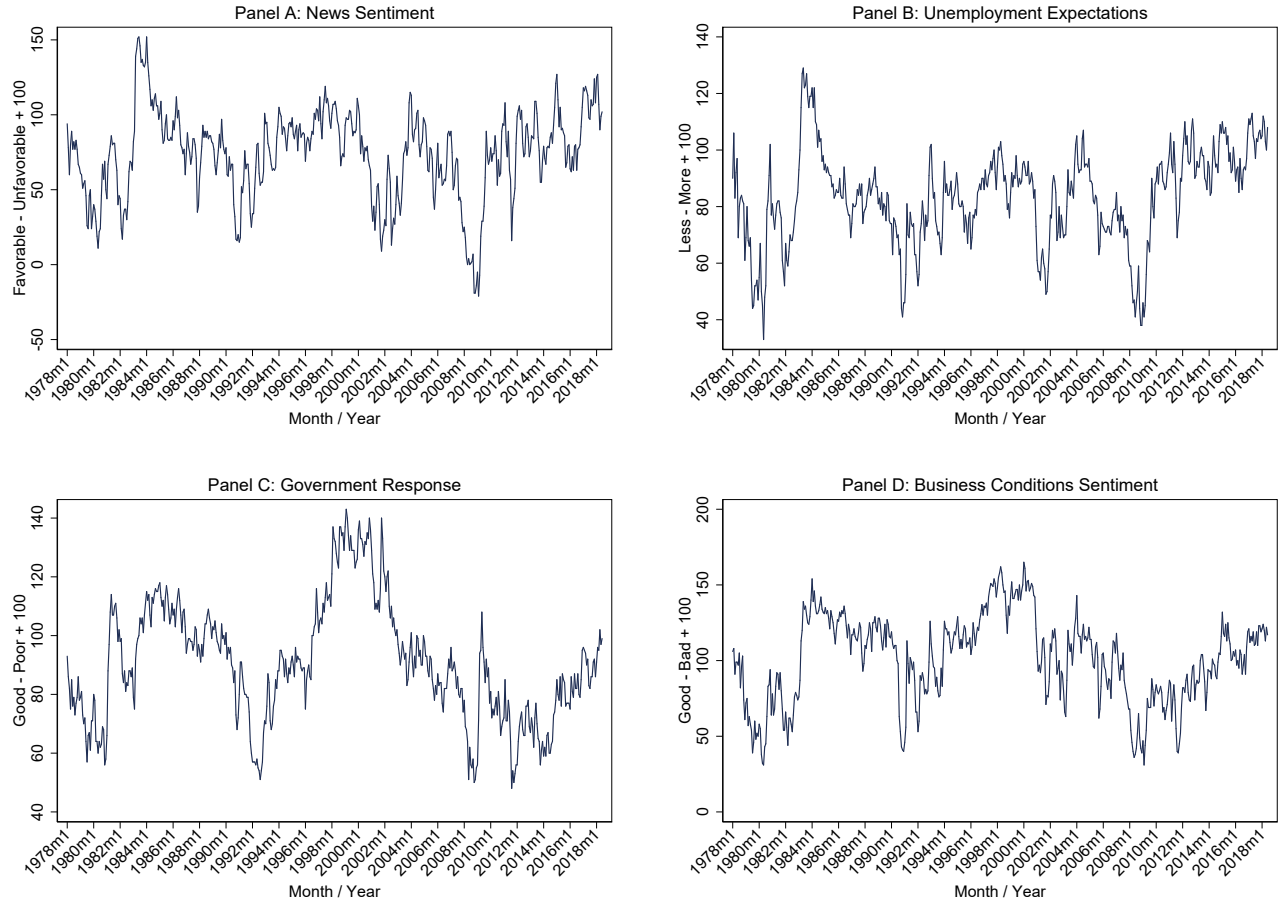
Figure 7: Historical Evaluation of Economic Sentiment, 1978-2018

*Notes.–*Sources: University of Michigan Survey of Consumers, 1978-2018. The figure plots four different measures of survey of consumer sentiment since 1978 to compare and contrast beliefs over different business cycles. The indices in Panels A-D are based off of responses as follows: (i) During the last few months, have you heard of any favorable or unfavorable changes in business conditions? (y-axis = favorable - unfavorable + 100), (ii) How about people out of work during the coming 12 months do you think that there will be more unemployment than now, about the same, or less? (y-axis = less + more + 100), (iii) As the economic policy of the government mean steps taken to fight inflation or unemployment would you say the government is doing a good job, only fair, or a poor job? (y-axis = good - poor + 100), (iv) Now turning to business conditions in the country as a whole do you think that during the next 12 months we'll have good times financially, or bad times, or what? (y-axis = good - bad + 100).

7.2 Supplement to Data and Measurement

7.2.1 Summary Statistics

Although the Gallup micro-data is constructed to be nationally representative, I begin by providing

a baseline characterization of the data, focusing on how different demographic groups have different

self-reported sentiments and well-being. I specifically regress z-scores of current and future life satisfaction, current and future economic sentiments, and an indicator of having a good standard of living on monthly income bin dummies, college attainment, a body mass index, an indicator for being male, an indicator for being white, and age. Table 7 documents these. I also estimate these specifications separately for employed and non-employed individuals.

I find that that individuals with higher monthly income have higher well-being and economic sentiments. For example, compared to the baseline of less than $1,500/month in income, those who are employed (non-employed) with over $8,500 have 111% (142%) higher current life satisfaction, 76% (93%) higher expected future life satisfaction, 23% (24%) higher economic sentiments about the current state of the economy, 20% (19%) higher sentiments about the future state of the economy, and 41% (55%) higher standards of living. The fact that well-being is increasing so rapidly across the income distribution suggests that there might not be a hump-shaped profile as some have suggested [Kahneman and Deaton, 2010]. I also find that college degree workers have roughly 20% higher life satisfaction and 5% higher economic sentiments. Employed males have 17% lower life satisfaction and unemployed males have 35% lower life satisfaction, but they are more optimistic about the future and have higher economic sentiments. Surprisingly, whites also have lower life satisfaction and economic sentiments. Age is negatively associated with life satisfaction and economic sentiments. These coefficient estimates are broadly comparable to those from Oswald and Wu [2011] who use the Behavioral Risk Factor Surveillance System (BRFSS) data.

52

Table 7: Descriptive Statistics of Well-being and Economic Sentiment Indices on Demographic Characteristics

Dep. var. =	current life satisfaction		future life satisfaction		current economic sentiment		future economic sentiment		standard of living	
	emp	unemp	emp	unemp	emp	unemp	emp	unemp	emp	unemp
monthly income, 1500-2500	.14***	.44***	.15***	.27***	.02***	.04***	.03***	.06***	.06***	.20***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.00]	[.01]	[.01]	[.01]	[.01]
monthly income, 2500-3500	.35***	.71***	.28***	.46***	.04***	.09***	.05***	.07***	.18***	.36***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.00]	[.01]	[.01]	[.01]	[.01]
monthly income, 3500-5500	.50***	.92***	.35***	.61***	.07***	.13***	.06***	.11***	.25***	.40***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]
monthly income, 5500-6500	.70***	1.13***	.48***	.78***	.12***	.18***	.10***	.14***	.36***	.47***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]
monthly income, 6500-8500	.90***	1.29***	.61***	.89***	.17***	.23***	.14***	.18***	.40***	.46***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.00]	[.01]	[.01]	[.01]	[.01]
monthly income, 8500+	1.11***	1.42***	.76***	.93***	.23***	.24***	.20***	.19***	.41***	.55***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]
college attainment	.21***	.15***	.15***	.29***	.05***	.04***	.07***	.10***	.37***	.23***
	[.01]	[.01]	[.01]	[.01]	[.00]	[.00]	[.00]	[.00]	[.01]	[.01]
body mass index	-.02***	-.03***	-.01***	-.01***	-.00***	-.00***	-.00***	-.00***	-.02***	-.02***
	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]
male	-.17***	-.35***	-.31***	-.34***	.03***	-.01***	.05***	.06***	.04***	-.02***
	[.01]	[.01]	[.01]	[.01]	[.00]	[.00]	[.00]	[.00]	[.01]	[.00]
white	-.13***	-.28***	-.34***	-.43***	-.09***	-.17***	-.24***	-.25***	.12***	.07***
	[.01]	[.01]	[.01]	[.01]	[.00]	[.00]	[.00]	[.00]	[.01]	[.01]
age	-.00***	.01***	-.03***	-.04***	-.00***	-.00***	-.01***	-.00***	-.00***	.01***
	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]	[.00]
R-squared	.06	.08	.08	.13	.02	.02	.03	.03		.03
Sample Size	634465	730967	624070	696765	402092	476590	398162	470200	650985	747394

Notes.–Sources: Gallup, 2008-2015. The table reports the coefficients associated with regressions of *z*-scores on current life satisfaction (one to ten index), expected future life satisfaction (one to ten index), perceptions about the current state of the economy (one to four index), perceptions about the current state of the economy (one to three index), and an indicator for a good standard of living all on demographic characteristics. Standard errors are heteroskedasticity-robust robust and sample weights are used.

7.2.2 Comparison of Gallup and BEA Consumption

One of the important features of the Gallup micro-data is that it contains information about daily consumption expenditures. However, one concern is that it contains significant measurement error since it does not represent a cumulative amount over an entire, for example, month or quarter. One of the potential concerns discussed in the main text is the validity of Gallup's measure of non-durables consumption expenditures. To provide an appropriate point of comparison between it and national account data, I consider regressions of the form

$$c_{st}^{BEA} = \beta_0 + \beta_1 c_{st}^{GALLUP} + \epsilon_{st}$$

where c_{st}^{BEA} denotes logged per capita consumption expenditures at a state-by-year level (2008-2016) in the Bureau of Economic Analysis (BEA) national accounts series and c_{st}^{GALLUP} denotes average logged consumption expenditures averaged across all individuals with non-zero consumption in the Gallup data. Figure 8 compares the Gallup state $\times$ year logged average daily consumption with the regional BEA consumption expenditure data and finds a strong positive correlation.

There are at least two reasons that the fit is not perfect. First, there is sampling variability in the Gallup data. Since different numbers of people are surveyed each year in different states, composition effects will tend to weaken the correlation with the true consumption series. Second, while consumption measured in Gallup "generally" involves non-durables goods, it is not as explicitly defined as the BEA national accounts data, making it difficult to map the two together perfectly. In this sense, while the two series do not predict each other perfectly, they do seem to be sufficiently correlated to take the Gallup series as plausible.

Figure 8: Comparison of Consumption from the BEA and Gallup

Notes.–Sources: Gallup and Bureau of Economic Analysis, 2008-2016. The figure plots logged daily consumption expenditures from the Gallup micro-data averaged by state × year against logged total consumption expenditures the BEA regional data. Both are deflated by the 2010 personal consumption expenditure index.

7.2.3 Comparison of Gallup and Survey of Consumer Sentiment

How does the Gallup micro-data compare with the more commonly cited Survey of Consumer Sentiment (SCS) from the University of Michigan used in, for example, Benhabib and Spiegel [2016]? Using data from the SCS, I focus on comparisons between it and the Gallup data over two variables: perceptions about the current and future state of economic activity. Specifically, I compute the share of individuals reporting that the economy is bad and the share of individuals reporting that the economy is worsening under different partitions of the population. Although the questions in the two datasets that lead to these survey responses differ in some ways, they are

relatively close and comparable; see the notes in Figures 9 and 10 for more detail.

Beginning with Figure 9, I compare the share of individuals reporting that the current and future states of the economy are bad and worsening, respectively. There is a particularly high correlation for the former correlations, but a lower correlation for the latter. For example, when pooling across the entire population, there is a 0.80 (Panel A) correlation between the two datasets over the the share of individuals reporting that the economy is poor. When looking at college educated workers, the correlation is slightly lower at 0.73 (Panel C), whereas the correlation is 0.76 (Panel E) among those with some college education. One reason for the higher correlation among less educated workers arises from the fact that the University of Michigan Survey might have more sampling variability at higher levels of education, especially disbursed across space. We find somewhat weaker correlations between the two datasets over the share of individuals reporting that the economy is worsening with a correlation of 0.59 (Panel B) when pooling across all individuals and correlations of 0.58 (Panel D) and 0.59 (Panel F) among individuals with a college degree or some college education, respectively. Figure 10 conducts the same analysis across the age distribution and finds similar correlations.
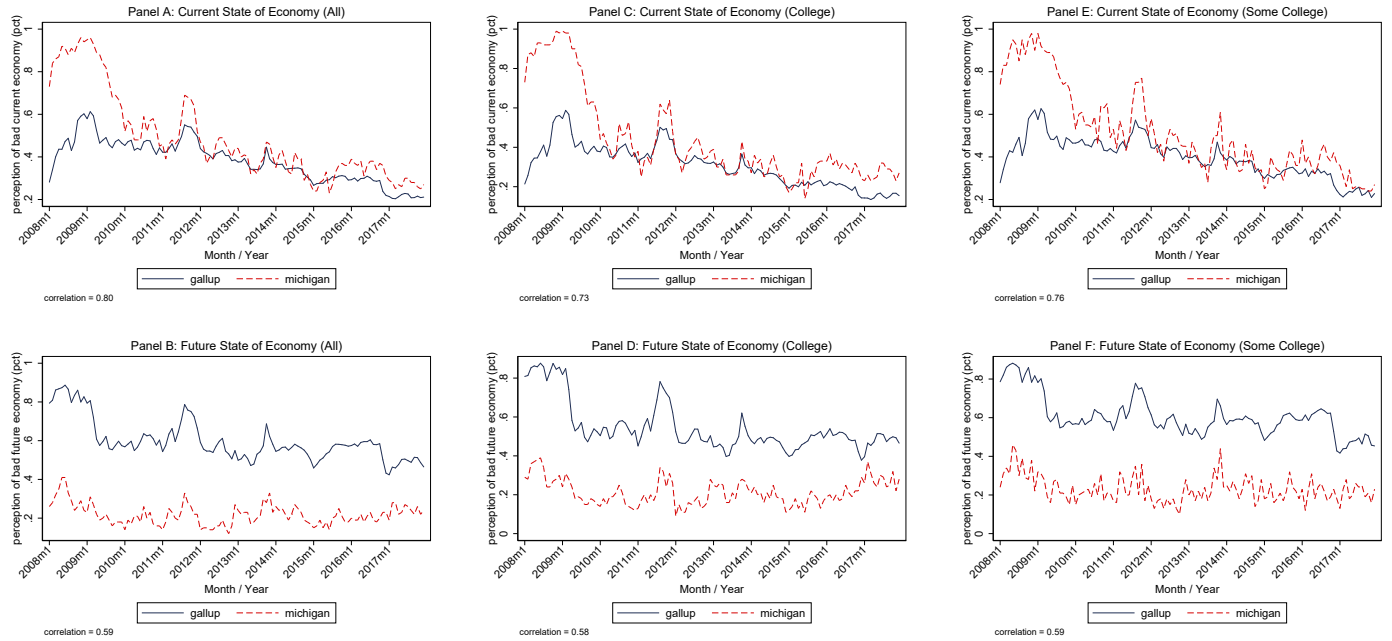
Figure 9: Comparison of Gallup and Survey of Consumer Sentiment (All + Education Brackets)

Notes.—Sources: Gallup U.S. Daily Poll and Michigan Survey of Consumers. The series of figures plot the share of individuals reporting that the economy is in a negative state and the share of individuals reporting that the economy is worsening in the Gallup and University of Michigan surveys. The closest analogue in the Survey of Consumer Sentiment is: "Would you say that at the present time business conditions are better or worse than they were a year ago?" The reply "worse" is used to compute the share of individuals. The corresponding Gallup question is: "How would you rate the economic conditions in this country today: as excellent, good, only fair, or poor?" For the question about the future, the closest analogue in the Survey of Consumer Sentiments is: "And how about a year from now, do you expect that in the country as a whole business conditions will be better, or worse than they are at present, or just the same?" The corresponding Gallup question is: "Right now, do you think that economic conditions in this country, as a whole, are getting better or getting worse?" Panels A and B plot the two time series when pooling across all individuals in the micro-data; panels C and D plot the two time series for college educated workers; panels E and F plot the time series for individuals with some college education. The plots cumulatively highlight the strong positive correlations between the two datasets.
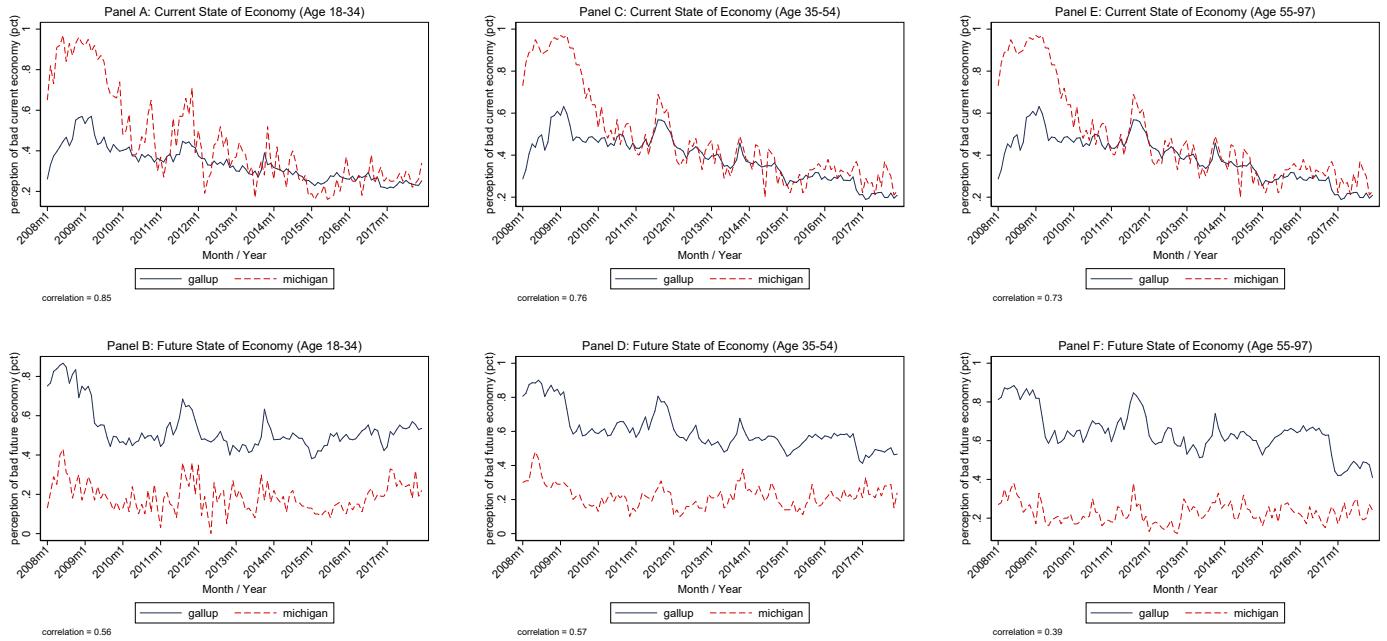
Figure 10: Comparison of Gallup and Survey of Consumer Sentiment (Age Brackets)

Notes.—Sources: Gallup U.S. Daily Poll and Michigan Survey of Consumers. The series of figures plot the share of individuals reporting that the economy is in a negative state and the share of individuals reporting that the economy is worsening in the Gallup and University of Michigan surveys. The closest analogue in the Survey of Consumer Sentiment is: "Would you say that at the present time business conditions are better or worse than they were a year ago?" The reply "worse" is used to compute the share of individuals. The corresponding Gallup question is: "How would you rate the economic conditions in this country today: as excellent, good, only fair, or poor?" For the question about the future, the closest analogue in the Survey of Consumer Sentiments is: "And how about a year from now, do you expect that in the country as a whole business conditions will be better, or worse than they are at present, or just the same?" The corresponding Gallup question is: "Right now, do you think that economic conditions in this country, as a whole, are getting better or getting worse?" Panels A and B plot the two time series for individuals between ages 18-34; panels C and D plot the two time series for individuals between ages 35-54; panels E and F plot the time series for individuals between ages 55-97. The plots cumulatively highlight the strong positive correlations between the two datasets.

Finally, I now turn towards an alternative exercise that examines how the Gallup and University of Michigan survey data predict real GDP growth. Unfortunately, the comparison is only defined over the 2008 and 2017 sampling frame given the constraints from Gallup. I correlate two measures of sentiment with three lags and three forward variables with real GDP growth (normalized to 2009 prices) for both Gallup and University of Michigan. These results are displayed in Table 8. Starting with the first measure from Gallup ("aggregate"), which combines perceptions about the current and future state of the economy on a one to seven index, and from University of Michigan ("Index of Consumer Sentiment"), which combines all of the survey information relating

to current economy sentiment, the Gallup data does a better job of predicting real GDP growth in terms of the lags, but a slightly worse job in terms of the forward variables. Turning towards the second measure from Gallup ("future"), which measures the share of individuals who think the economy is improving, and from the University of Michigan ("Index of Consumer Expectations"), which combines all the survey information relating to future economic expectations, a similar phenomenon is visible: Gallup does a better job of predicting real GDP growth in terms of lags, but slightly worse in terms of the forward variables. Although the data is not definitive and is subject to a somewhat small sample of 33 quarters, it highlights the important predictive role that Gallup sentiment data can play for understanding movements in real economic activity.

	$t-3$	$t-2$	$t-1$	$t=0$	$t+1$	$t+2$	$t+3$
Gallup							
Aggregate	0.59	0.68	0.66	0.57	0.44	0.31	0.27
Future	0.77	0.76	0.68	0.52	0.31	0.10	0.05
UMichigan							
ICS	0.42	0.54	0.57	0.56	0.49	0.41	0.38
ICE	0.43	0.51	0.51	0.46	0.38	0.28	0.27

Table 8: Comparing Gallup and University of Michigan Indices for Predicting Real GDP Growth
Notes.—Sources: Gallup U.S. Daily Poll and Michigan Survey of Consumers. The table reports the correlations between lags and leads of Gallup's economic sentiment (sum of perceptions of the current and future state of the economy, creating a one to seven index) and future economic sentiment (share of individuals reporting the economy is improving) with the University of Michigan surveys on consumer sentiment and consumer expectations. The figure shows that Gallup does a much better job predicting real GDP growth when focusing on the lags, whereas University of Michigan does a better job when focusing on the leads.

7.3   Supplement to Belief Formation and Local Shocks

The main text argues that local shocks to employment and housing prices affect beliefs about the aggregate state of the economy. While the preferred specifications control for composition effects, one potential concern is reverse causality: individual beliefs drive an expansion of employment or bidding up of housing prices. I begin by presenting results from a difference-in-difference estimator that compares renters and homeowners in the same zipcode over time subject to different housing

price shocks, conditional on state $\times$ time fixed effects. Importantly, identification here comes off of comparisons among observationally equivalent individuals subject to different housing price shocks in the same zipcode, controlling for all the potential labor market shocks common within the same state over time. In this sense, the specification exploits the fact that individuals live in more segmented housing markets, whereas the labor markets are at a broader aggregation.

Table 9 documents these results for three outcome variables. Column 1 shows that 1pp rise in housing price growth is associated with a 0.29sd rise in perceptions about the current state of the economy among home owners, but a null association for renters. Although home owners generally have higher economic sentiment, which may reflect lower risk aversion associated with owning a home, the elasticity between housing shocks and beliefs is concentrated among precisely the set of individuals that would be expected to be most sensitive to the shocks—home owners should not be affected since they do not experience a gain or loss as a result of housing price changes. Column 2 shows that a 1pp rise in housing price growth is associated with a 0.18pp rise in the probability an individual reports the economy is improving, consistent with the prior result. Column 3 flips the outcome variable to current life satisfaction, but finds a statistically insignificant estimate, suggesting that the primary way housing price fluctuations affect individual risk perceptions is through their beliefs about the economy.

To examine this concern more carefully, I now turn towards to proxies for labor market shocks: employment growth in the finance and real estate (FRE) and trade sectors. In particular, I ask whether beliefs among individuals working in these sectors respond more elastically to these sector-specific shocks. I estimate Equation 1 without housing price growth and with an indicator for the type of worker interacted with the sector-specific employment shock on the right hand side restricted to the sample of employed individuals. Table 10 documents these results. I find

Table 9: Housing Price Shocks Among Renters Versus Owners

Dep. var. =	current state of the economy	1[economy is improving]	current life satisfaction
	(1)	(2)	(3)
1[home owner]	.03***	.01	.20***
	[.01]	[.00]	[.01]
Δ ln(housing price)	-.03	.01	.11
	[.15]	[.06]	[.14]
× 1[home owner]	.29**	.18***	.04
	[.14]	[.05]	[.12]
R-squared	.19	.21	.21
Sample Size	159167	157881	176278
Controls	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Zipcode FE	Yes	Yes	Yes
State x Year FE	Yes	Yes	Yes

Notes.–Sources: Gallup, Federal Housing Administration, 2009, 2010, 2012, 2013, 2014. The table reports the coefficients associated with regressions of standardized (*z*-score) perceptions about the current state of the economy (a one to four index with higher values implying a better state) and an indicator for whether the individual reports the economy is improving on year-to-year housing price growth, an indicator for whether the individual owns their home, and their interaction, conditional on individual controls, and fixed effects on zipcode and state × year and day of the year fixed effects. Controls include: a quadratic in age, male, education fixed effects, race (black/white). The sample is restricted to individuals who are employed. Standard errors are clustered at the zipcode-level and sample weights are used.

that individuals in the finance or real estate sectors have 4% higher perceptions about the state of the economy, relative to their peers. Interestingly, however, a 1pp rise in employment growth in the FRE sector is associated with a 0.14sd decline in the perception of the current state of the economy overall, but a 0.37sd rise among those working in the FRE sector, implying that these latter workers exhibit a net 0.22sd ($= 0.37 - 0.14$) rise in sentiment. Turning towards business owners and workers in the retail / wholesale trade sectors, I find that they have 7% lower sentiment, relative to their peers. I also find that a 1pp rise in employment growth in the wholesale and retail trade sector is associated with a 0.18sd rise overall and an additional 0.14sd rise among workers in the sector, relative to their peers, implying a net 0.32sd ($= 0.18 + 0.14$) rise in perceptions about the current state of the economy. However, the interaction effect is not statistically significant at conventional levels. Nonetheless, these results are consistent with the mechanism of interest—that personal experience is an important ingredient in the formation of beliefs.

Table 10: Examining the Effects of Sector-specific Labor Market Shocks

Dep. var. =	current state of the economy (z-score)	
	(1)	(2)
1[FRE]	.04***	
	[.00]	
Δ ln(FRE employment)	-.14***	
	[.05]	
$\times$ 1[FRE]	.37***	
	[.08]	
1[owner or trade]		-.07***
		[.01]
Δ ln(trade employment)		.18**
		[.09]
$\times$ 1[owner or trade]		.14
		[.17]
R-squared	.09	.09
Sample Size	481202	481592
Controls	Yes	Yes
Time FE	Yes	Yes
County FE	Yes	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages,2008-2017. The table reports the coefficients associated with regressions of standardized (*z*-score) perceptions about the current state of the economy (a one to four index with higher values implying a better state) on the year-to-year county × quarter employment growth rate in the finance and real estate (FRE) sector (NAICS 52, 53) and trade sector (NAICS 42, 44-45), conditional on individual controls, and fixed effects on county and day of the year. Controls include: a quadratic in age, male, education fixed effects, race (black/white). The sample is restricted to individuals who are employed. Standard errors are clustered at the county-level and sample weights are used.

Motivated by a similar argument using occupational variation, I also exploit the fact that individuals in jobs that are concentrated in the tradables sector should not be impacted by local shocks. In particular, if these estimates are driven purely by reverse causality, then an individual in, for example, a professional services job should not change their expectations about economic activity since their output does not directly affect local activity. While I do not have a measure of the industry the individual works in, I do have a broad occupational classification that is informative—as in the case of a professional service versus construction job. Figure 11 plots the estimated employment and housing price growth gradients separately by occupation. Despite the presence of heterogeneity, jobs that have no direct connection to local economic prospects also respond significantly to local shocks.

Panel A: Employment Growth Gradient   Panel A: Housing Price Growth Gradient

Figure 11: Heterogeneity in Employment and Housing Gradients, by Major Occupation
Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Federal Housing Administration, 2008-2016. The figure plots the coefficients associated with regressions of standardized (*z*-score) individual perceptions about the current state of the economy (one to four index) on the year-to-year employment growth rate, median housing price per square foot growth rate, individual controls, and fixed effects on county and day of the year across each occupation. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level.

To further examine that the mechanism of interest relates to the role of personal experience, I now consider regressions of the form in Equation 1 with a *z*-score of both current and expected future (in five years) life satisfaction (zero to ten scale). Table 11 documents these results. Under

the preferred specifications when both employment and housing price are put together as controls, a 1pp rise in employment and housing price growth is associated with a 0.17sd and 0.08sd rise in current life satisfaction. Interestingly, the gradient on employment growth is now larger than the gradient for housing price growth. However, when looking at changes in beliefs about expected future life satisfaction, the gradients are not statistically significant. In this sense, these housing and labor market shocks may affect future well being primarily through a channel of economic optimism, rather than some non-pecuniary well-being factors.

Table 11: Response of Life Satisfaction to Local Labor and Housing Market Shocks

Dep. var. =	current life satisfaction (z-score)			future life satisfaction (z-score)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ ln(employment)	.37***		.27**	.21**		.21**
	[.10]		[.11]	[.10]		[.10]
Δ ln(housing price)		.12***	.09**		.02	.00
		[.04]	[.04]		[.03]	[.03]
R-squared	.05	.05	.05	.13	.13	.13
Sample Size	779435	779441	779435	755728	755733	755728
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Federal Housing Administration, 2008-2017. The table reports the coefficients associated with regressions of standardized (*z*-score) current life satisfaction (a zero to ten index with higher values implying a better state) and expected life satisfaction in five years (measured similarly) on the year-to-year county × quarter employment growth rate and annual county housing price growth from Bogin et al. [forthcoming], conditional on individual controls, and fixed effects on county and day of the year. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level and sample weights are used.

As a final exercise, I exploit monthly state variation in mass layoffs, which are frequently unanticipated to residents. The Bureau of Labor Statistic's Mass Layoff Statistics (MLS) program collects reports on mass layoff actions that result in workers being separated from their jobs (https://www.bls.gov/mls/). These numbers are collected from establishments that have at least 50 initial claims for unemployment insurance filed during a five-week period. Although the data stopped being collected mid-2013, there is significant variation: the mean is 87.15 layoff events, the

median is 39, and the standard deviation is 121. These types of mass layoffs have been exploited as a source of "surprise" variation in recent work (e.g., Sullivan and Wachter [2009] and Baker [forthcoming]).

Using these data, I consider analogous regressions of the z-score on perceptions about the current state of the economy and an indicator for improving economic conditions on individual covariates, logged monthly layoff events, and state, year and month fixed effects. I find that a 10% rise in monthly layoff events is associated with a 0.15sd (p-value $= 0.047$) rise in perceptions about the current state of the economy and a 0.054pp (p-value $= 0.039$) rise in perceptions that the economy is improving. These results are merely a heuristic to highlight the fact that plausibly unanticipated shocks to regional employment outcomes affect sentiments—if anything, they are likely an underestimate since the specification treats all mass layoffs as homogeneous in intensity.

Turning towards heterogeneity, the main text shows that there are some differences among those who are liquidity constrained, relative to their peers. I now examine heterogeneity across the income distribution. Using individuals' reported monthly income bins, I partition individuals into those who earn between $500-1,999, $2,000-2,999, $3,000-3,999, $4,000-4,999, $5,000-7,499, $7500-9,999, and $10,000+ in labor income per month. While liquidity constraints and income are negatively correlated, they are capturing different economic behaviors.

Figure 12 plots the coefficients associated with regressions of perceptions about the current state of the economy on employment and housing growth across the income distribution, conditional on county and time fixed effects. Interestingly, very high income individuals do not appear to update their beliefs much in response to local employment shocks (their estimates are not statistically different from zero), which could reflect the fact that many high income earners were insulated to labor market shocks. However, the housing price elasticity is increasing and monotone

in income with high earners responding the most to an increase in housing price growth, reflecting the fact that their housing wealth may have fallen a lot during the financial crisis.

Panel A: Employment Growth Gradient   Panel A: Housing Price Growth Gradient

Figure 12: Sentiment Heterogeneity in Employment and Housing Gradients, by Income Bracket
Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The figure plots the coefficients associated with regressions of standardized (z-score) individual perceptions about the current state of the economy (one to four index) on the year-to-year employment growth rate, median housing price per square foot growth rate, individual controls, and fixed effects on county and day of the year across monthly income bins. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level.

7.4   Supplement to Quantifying the Effects of Sentiment on Real Economic Activity

7.4.1   Supplement to the 2016 Event Study

The main text shows that consumption increased, following the unanticipated victory of Donald Trump during the 2016 election. Moreover, conservatives increased their consumption by 10-12%, relative to their liberal and moderate counterparts, following the victory. These results are robust to using a more narrow three-month window around the election, rather than the baseline five-month window. An important assumption for interpreting these estimates as causal is that conservative and non-conservative respondents had similar pre-trends in daily consumption prior

to the election.

After residualizing consumption using the standard demographic covariates, Figure 13 plots a smoothed time series of consumption between the two groups across daily observations between January 2016 and December 2017. Importantly, while there is a near parallel trend between the two groups prior to the election, their series begin to diverge afterwards. One possibility behind the uptick before the election results is that conservatives in the "silent majority" already expected a victory after Donald Trump forcefully cleared the playing field following the primaries.

Figure 13: Examining Evidence of Parallel Trends for the Event Study

Notes.–Sources: Gallup U.S. Daily, 2016-2017. Using a lowess non-parametric smoother, the figure plots daily non-durables consumption expenditures between conservatives and their counterparts before and after the 2016 election where consumption is residualized using gender, a quadratic in age, education fixed effects, and race fixed effects.

However, there is some ambiguity about how to interpret these results in light of, for example, evidence from Mian et al. [2017] that automobile purchases did not increase. To understand whether automobile purchases are a suitable proxy for consumption, Figure 14 draws on monthly

data on non-durable consumption expenditures and vehicle purchases from the St. Louis Federal Reserve. Non-durable consumption is deflated using the 2009 personal consumption expenditure index. While there was a generally high degree of correlation between the two series over 2010-2015, the correlation is only 0.17 from 2015-2017. Restricting the sample to 2016-2018, a regression of logged non-durable consumption on an indicator for post-November 2016 produces a gradient of 0.0255 (p-value $= 0.00$), whereas a similar regression for vehicle purchases produces a gradient of -0.0086 (p-value $= 0.485$). Focusing on automobile purchases over a narrow window might bias downwards estimates of the causal effect of sentiment on real activity since autos are a type of infrequently purchased consumer durable.

correlation 2015-2018 = 0.17

Figure 14: Comparison of Non-durable Consumption and Vehicle Sales

Notes.−Sources: St. Louis Federal Reserve, 2000-2018. The figure plots monthly logged non-durables consumption expenditures (deflated using the 2009 personal consumption expenditure index) and total vehicle sales in millions (seasonally adjusted). The figure shows that there is a 2.5% increase in non-durables consumption after the 2016 election (when restricting to a 6 month window around the election), whereas there is no increase in vehicle sales.

As a final exercise for assessing the validity of the event study results, I turn towards state-

level data on legislative control (http://www.ncsl.org/research/about-state-legislatures/partisan-composition.aspx). In particular, I measure the number of Republicans and Democrats in both the House and Senate, as well as identify the years when the state governor or legislature switches party control. The results here are quite noisy, but they provide qualitative support for the main results. For example, using the logged number of Republicans and Democrats in the House and Senate at a state × year frequency between 2009-2017 as instruments for current perceptions of the economy implies that a standard deviation rise in economic sentiment is associated with a 0.28% rise in consumption (p-value = 0.205). While the estimate is not statistically significant at conventional levels, it nonetheless highlights that a similar type of more disaggregated variation implies similar results.

7.4.2   Comparison to Mian, Sufi, and Khoshkhou (2018)

How do sentiments influence consumption expenditures? Using consumption data from the Bureau of Economic Analysis (BEA), Figure 15 implements two sets of results at a state × year aggregation between 2008 and 2016. First, Panels A and B display a strong positive association between higher economic sentiment and vehicle purchases and non-durable consumption: a unit increase in perceptions about the state of the economy in a state are associated with 0.64% and 0.12% higher purchases of vehicles and non-durables consumption, respectively. Second, Panels C and D display these results in growth rates, highlighting that the correlations in Panels A and B are not driven by time-invariant differences in the sample or beliefs of households in different states. The gradients are not as strong, potentially because of measurement error in the sentiment data due to imperfect samples, but they still highlight the positive correlations. Although not reported, there

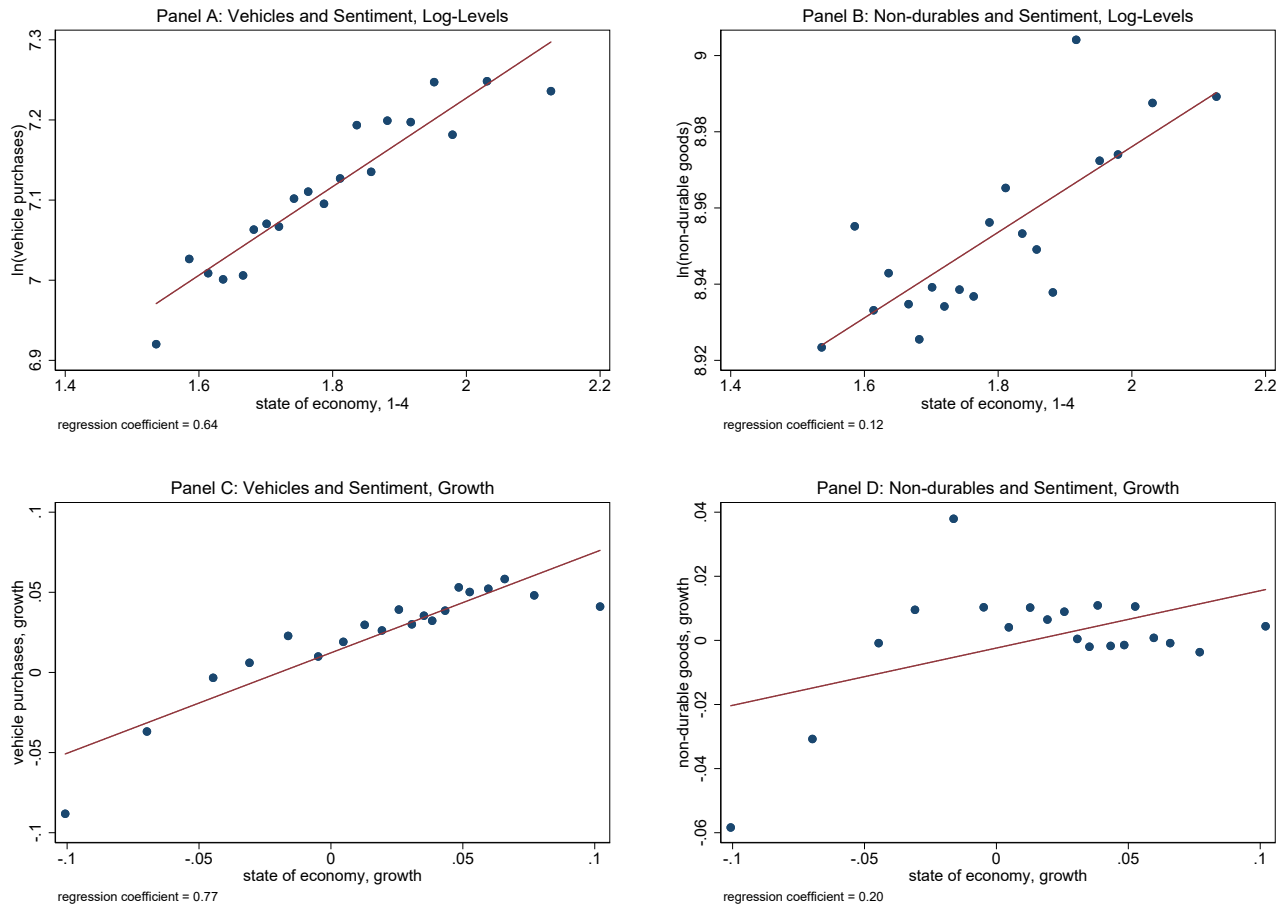is also a similar correlation between overall consumption and sentiment.

Figure 15: Correlations between Levels and Growth Rates of Consumption & Sentiment

Notes.–Sources: Gallup, Bureau of Economic Analysis, 2008-2016. Panels A and B plot the logged consumption expenditures in vehicle purchases / parts and non-durable goods in real dollars (using the 2009 personal consumption expenditure index) against the mean sentiment (on a one to four scale) about perceptions on the current state of the economy. Panels C and D plot the year-to-year growth rates of the two variables. Regression coefficients in table notes are reported using state population as a weight; the plots are represented through `binscatter`.

Turning from these consumption and sentiment regressions, Figure 16 now plots the correlations

between consumption and housing price growth. For example, a 1pp rise in housing price growth

is associated with a 0.36pp and 0.16pp decline in vehicle purchases and non-durable consumption

expenditures, respectively. These results are consistent with Mian et al. [2017] who show that

housing price growth is positively correlated with consumption, but these results differ in that

sentiment growth is also correlated with consumption growth across a variety of consumption

categories. Why do is there a correlation with vehicle purchases here, whereas there was not in the event study? Vehicle purchases were bid up as a result of the Cash for Clunkers program, but by the later years they began to plateau and, in particular, flat during the neighborhood of the election event study.


Figure 16: Correlations between Consumption and Housing Price Growth

Notes.–Sources: Gallup, Federal Housing Administration, 2008-2016. Panels A and B plot the growth of consumption expenditures in vehicle purchases / parts and non-durable goods in real dollars (using the 2009 personal consumption expenditure index) against the growth of the state housing price index. Regression coefficients in table notes are reported using state population as a weight; the plots are represented through `binscatter`.

7.4.3 Supplement to Alternative Instrumental Variables Strategy

The main text presents estimates of the consumption elasticity under three separate identification strategies. The first strategy exploits a result from Binder and Makridis [2018] that high-frequency state gasoline price fluctuations affect beliefs about the state of the economy. The intuition here is that gasoline price shocks are salient since individuals regularly have to fill their cars up with gasoline or pass gasoline stations where prices are clearly visible. While it is not possible to fully rule out a disposable income channel—that is, higher gasoline prices reduce consumption on other

non-durable goods by reducing disposable income, which in turn leads to substitution away from other goods—it is unlikely the primary culprit because of inelastic demand for gasoline [Kilian, 2008].

Given the connection between state gasoline price shocks and beliefs, why not use the same approach in the main text? Unfortunately, the Energy Information Administration (EIA) stopped collecting high-frequency gasoline prices for every state, so defaulting to the state price data would require a significant reduction in sample size. As an alternative strategy, I create a Bartik-like instrument that exploits the fact that individuals in states with greater exposure to the energy sector will be more heavily affected than their counterparts in other states. To measure exposure, I use the Bureau of Economic Analysis (BEA) state data on output, computing the share of output from the energy and transportation sectors. In this sense, individuals residing in states where the sector operates a bigger share will be more sensitive to national gasoline price shocks.

Figure 17 plots the first-stage relationship between the two, displaying a strong correlation. Although the first-stage F-statistic is well above the conventional rule of thumb, the more challenging assumption required for estimating a causal effect is that the Bartik-like shock affects individual consumption only through its effects on beliefs. An important ingredient for the assumption to hold, for example, is that the measure of state exposure to the energy sector is not correlated with unobserved shocks to consumption between 2008 and 2017. One way of validating whether that is the case is by examining the correlation between the 2006 state exposure shares and growth in, for example, housing prices between 2007 and 2012. Figure 18 plots the relationship and illustrates that there is only an imprecise positive association.
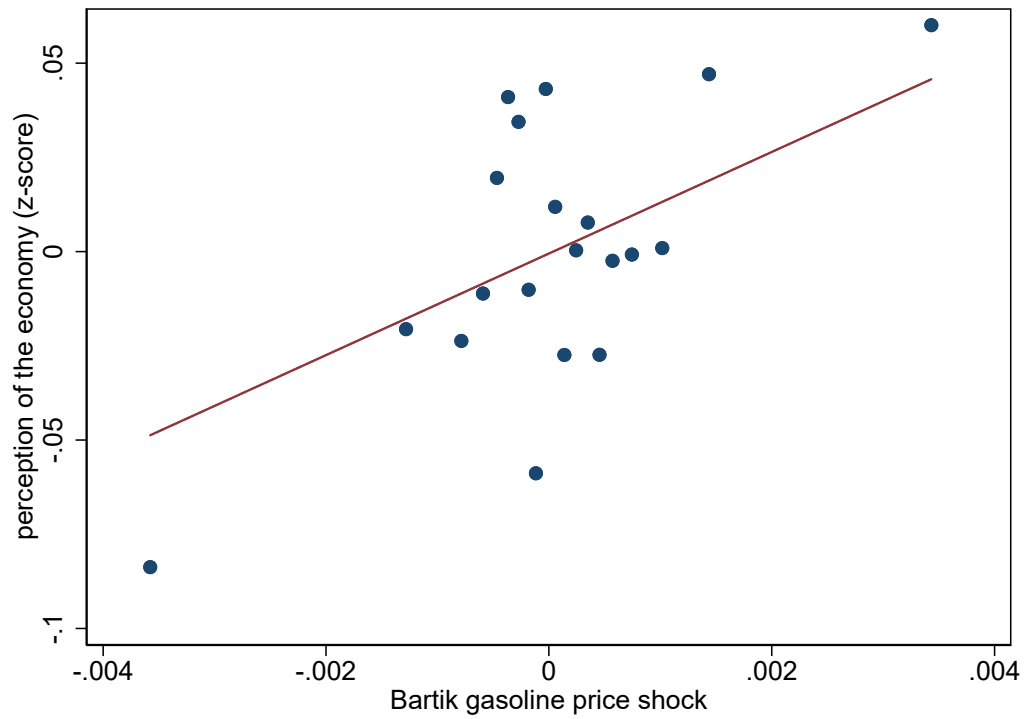
Figure 17: First-stage Relationship between Energy Bartik and Economic Sentiment

Notes.–Sources: Bureau of Economic Analysis and Gallup, 2008,2017. The figure plots the first-stage correlation between perceptions of the state of the economy and the energy Bartik instrument, which is obtained by computing the state output share of energy using 2006 data and interacting it with the national weekly gasoline price. The figure shows that there is a strong first stage correlation between the two.

Figure 18: Relationship between Housing Price Shocks and Energy Exposure Across States
Notes.–Sources: Bureau of Economic Analysis and Federal Housing Administration. The figure plots housing price growth between 2007Q1-2012Q1 and the energy share of output across states (defined as output from the energy and transportation sectors as a share of the total).

The second strategy exploits the fact that shocks to the external environment can shape the information that individuals pay attention to and process. For example, companion work in Makridis [2018a] shows that extreme temperatures at the bottom and top of the temperature distribution affect beliefs about the current and future state of the economy. Figure 19 replicates the results here. Although the main text does not allow for separate treatment effects on sentiment across the temperature distribution, I use a quadratic in daily temperature and interactions with different occupation fixed effects to capture the potential non-linear relationship and how it varies for individuals who are heterogeneously exposed to temperature (e.g., working inside versus outside). These results are consistent with Baylis [2018] and Baylis et al. [2018] who document strong correlations using social media data on positive and negative emotions.

Figure 19: First-stage Relationship between Temperature and Economic Sentiment

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The figure reports the coefficients associated with regressions of the *z*-score of the sum of perception about the current and future state of the economy (a one to seven index) on counts of the number of days in a month that fall between an upper and lower temperature bound (below 0, 0-15, 16-30, 31-53, 60-70, 71-84, and 85+ normalized to 54-59 as the omitted group), individual controls, county × year × quarter employment and housing price growth, and county and time fixed effects. Controls include: occupation fixed effects a quadratic in age, education fixed effects, race fixed effects, employment status, and gender. Standard errors are clustered at the county-level and sample weights are used.

However, an important assumption for the exclusion restriction to hold is that temperature does not have direct effects on productivity. If it did, then consumption could decline in response to changes in local temperature-induced productivity shocks, rather than economic sentiment. To illustrate that there is no evidence of these temperature-induced productivity effects, I replicate results in companion work [Makridis and Ransom, 2017]. Drawing on the monthly Current Population Survey (CPS) between 1994 and 2015 accessed through the Integrated Public Use Microdata (IPUMS) data portal at the University of Minnesota restricted to full-time workers between ages 20 and 65 with over \$5,000 in annual labor income and over \$2 hourly wages (both deflated using the 2010 real personal consumption expenditure index), I match monthly average temperatures at

a metropolitan level of aggregation and examine the association between logged weekly earnings on logged average temperature, conditional on individual covariates and metro and year / month fixed effects.

Table 12 documents these results under various specifications. Columns 1 and 4 show that the unconditional correlation between temperature and income (hours worked) is negative (positive). However, the unconditional correlation could be confounded by a wide array of omitted variables—most notably the fact that different types of individuals sort into areas that might have different temperatures and different labor markets. Once basic individual covariates are introduced, such as age and education, the correlation with income vanishes (column 2). Furthermore, once location and time fixed effects are introduced, the correlation also vanishes with hours worked.

Table 12: Examining the Relationship between Monthly Temperature and Income/Hours

Dep. var. =	ln(weekly earnings)			ln(hours worked)		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(monthly temperature)	-.041***	-.003	-.001	.015***	.017***	.002
	[.010]	[.008]	[.004]	[.003]	[.003]	[.001]
R-squared	.00	.27	.29	.00	.07	.08
Sample Size	2090105	2090105	2090105	1964763	1964763	1964763
Controls	No	Yes	Yes	No	Yes	Yes
Year / Month FE	No	No	Yes	No	No	Yes
Metro FE	No	No	Yes	No	No	Yes

Notes.–Sources: NOAA, Current Population Survey (CPS) monthly from 1994-2014. The figure reports the coefficients associated with least squares and fixed effects regressions of logged weekly earnings (deflated using the 2010 personal consumption expenditure index) and separately of logged weekly hours worked on logged monthly average temperature at the metropolitan level. Controls include: a quadratic in both educational attainment and age, number of children, race, gender, marital status, and race. Standard errors are clustered at the metropolitan level (2013 OMB codes) and observations are weighted by the survey sample weights.

To further guarantee that the average temperature is not masking heterogeneity over the course of a month, I allow for separate effects of temperature on income across the temperature distribution as in prior environmental economics literature [Deschenes and Greenstone, 2011, Deryugina and Hsiang, 2017]. Figure 20 plots the corresponding gradients across the distribution of temperature with logged weekly earnings as the outcome variable. The relationship between temperature

and weekly earnings is effectively flat and statistically indistinguishable from zero.

Figure 20: Semiparametric Response of Weekly Earnings to Temperature, 1994-2014

Notes.–Sources: NOAA, Current Population Survey (CPS) monthly from 1994-2014. The figure reports the coefficients associated with least squares and fixed effects regressions of logged weekly earnings (deflated using the 2010 personal consumption expenditure index) and separately of logged weekly hours worked on bins for the number of days in a month that fall within the corresponding bin (e.g., number of days in a month below zero degrees Fahrenheit), which is unique at the metropolitan level. Controls include: a quadratic in both educational attainment and age, number of children, race, gender, marital status, and race. Standard errors are clustered at the metropolitan level (2013 OMB codes) and observations are weighted by the survey sample weights.

7.5   Supplement to Aggregation Results

The main text shows that the fluctuations in sentiment can explain a large share of the aggregate

decline in consumption. To understand the mechanism behind these effects, it is useful to explore

the role that social networks play in amplifying shocks. Using the social connectedness index (SCI)

from Bailey et al. [2018a], I estimate the following regressions:

$$s_{ict} = \gamma^{OWN}\Delta e_{ct}^{OWN} + \phi^{OWN}\Delta h_{ct}^{OWN} + \gamma^{CON}\Delta e_{ct}^{CON} + \phi^{CON}\Delta h_{ct}^{CON} + \beta X_{it} + \eta_c + \lambda_t + \epsilon_{ict}$$

where OWN and CON denote a county's own employment and housing price shock versus their connectivity-weighted shock. The usual controls and fixed effects are included. If the diffusion of information through social networks is not important in shaping beliefs, then γ^{CON} and ϕ^{CON} should be small and potentially also statistically indistinguishable from zero. Even though the variation in friendship ties (i.e., connected counties) is not plausibly exogenous [Bailey et al., 2018b], the identifying assumption here is that changes in the exposure to different shocks among connected counties is quasi-random.

Table 13 documents these results. Columns 1 and 3 simply report the baseline results with the two different outcome variables, not controlling for SCI-weighted shocks as a reference point. Once the SCI-weighted shocks are included in columns 2 and 4, we see that they are quantitatively important, especially so in the context of employment shocks. For example, I find that a 1pp rise in a county's own employment growth is associated with a 0.191sd and 0.094pp rise in beliefs about the current state of the economy and the probability of reporting that the economy is improving, whereas a comparable 1pp rise in SCI-weighted employment growth is associated with a large 0.387sd and 0.134pp rise. The fact that the SCI-weighted employment shock generates an elasticity that is greater than the direct effect is interesting and open to future work. However, the SCI-weighted housing price shock generates more plausible estimates of 0.21sd and 0.088pp, which are roughly half of the corresponding direct effect of a county's own housing price shock on beliefs. In either case, the elasticities are precisely estimated and suggest that beliefs to connected

counties can indirectly influence an individual's beliefs.

Table 13: Social Network Amplification of Housing and Labor Market Shocks

Dep. var. =	current state of the economy (z-score)		1[economy is improving]	
	(1)	(2)	(3)	(4)
Δ ln(employment)	.167**	.191***	.069**	.094***
	[.065]	[.061]	[.029]	[.026]
Δ ln(housing price)	.611***	.438***	.296***	.177***
	[.069]	[.059]	[.032]	[.026]
Δ SCI ln(employment)		.387***		.134***
		[.124]		[.042]
Δ SCI ln(housing price)		.209***		.088***
		[.037]		[.019]
R-squared	.06	.07	.11	.11
Sample Size	1138002	781306	1138002	781306
Controls	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Federal Housing Administration, 2008-2017. The table reports the coefficients associated with regressions of standardized (z-score) perceptions about the current state of the economy (a one to four index with higher values implying a better state) and an indicator denoting that the economy is improving on the year-to-year county × quarter employment growth rate, annual county housing price growth from Bogin et al. [forthcoming], and their social connectedness index (SCI) weighted averages, conditional on individual controls, and fixed effects on county and day of the year. The SCI-weighted employment and housing price growth averages are constructed by matching for every county c the corresponding employment and housing price growth at the connected county c' and taking the weighted average (by county population) over the product of the shock and strength of the connectivity tie based on the SCI from Bailey et al. [2018a]. Controls include: an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). Standard errors are clustered at the county-level and sample weights are used.

The main text also provides descriptive evidence that states with lower sentiment during the start of the financial crisis took longer to recover to their pre-crisis housing price growth and real GDP rates and levels. One of the reasons these two series are used is because the quarterly frequency provides additional variation for measuring the delay of the recovery, whereas years would restrict me to a maximum of nine years up until 2017. However, one limitation is that the bulk of the paper is about consumption, rather than housing prices or real GDP.

Although both GDP and housing price growth are highly correlated with consumption of non-durable goods, they nonetheless capture different features of real economic activity. Despite these potential differences, Figure 21 shows that a similar result about the negative association

between economic confidence and the delayed recovery holds up for consumption of non-durable goods despite lacking as much variation, although the results are not as statistically significant or large. Moreover, Figure 22 replicates these plots for housing price growth and real GDP by plotting against the share of individuals who believe that the economy is improving. The negative association between economic optimism and recovery times remains.


gradient = -0.06

Figure 21: Belief Shocks and the Protracted Recovery of Consumption

Notes.−Sources: Gallup, Bureau of Economic Analysis, and Zillow, 2008-2017. The figure plots with `binscatter` average economic confidence between 2008-09 across states with the number of months it takes for consumption of non-durable goods per capita to recover to their pre-crisis levels. It is computed by measuring the number of months it takes for a state to return to its 2006Q1 housing price growth and real GDP rate or level where counting begins in 2008Q1.

Figure 22: Belief about the Future Shocks and the Protracted Recovery

Notes.–Sources: Gallup, Bureau of Economic Analysis, and Zillow, 2008-2017. The figure plots with `binscatter` the share of individuals reporting that the national state of the economy is improving between 2008-09 across states with the number of months it takes for housing prices and real GDP to recover to their pre-crisis rates and levels. The latter is computed by measuring the number of months it takes for a state to return to its 2006Q1 housing price growth and real GDP rate or level where counting begins in 2008Q1.

As a final exercise, I examine how economic sentiment in 2016 explains variation in voting patterns in the 2016 Presidential election. Figure 23 plots the share of individuals who voted for Donald Trump with economic confidence in 2016 across states. There is a remarkably strong negative relationship—so much so that economic sentiment explains 63% of the variation in voting patterns. In this sense, economic sentiment has a wide array of real economic implications, ranging from impacts on aggregate consumption to housing prices to election results.

Figure 23: Economic Confidence and the 2016 Presidential Election

Notes.–Sources: Gallup, 2016. The figure plots confidence about the state of the economy in 2016 and the 2016 Republican vote share for Donald Trump using `binscatter` across all states.