Sentimental Business Cycles and the Protracted Great Recession

Christos Andreas Makridis*

December 27, 2017

In 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, exposure to different local shocks influences economic sentiments: a one percentage point rise in county employment and housing price growth is associated with a 0.24-0.57 and 0.21-0.40 standard deviation rise in perceptions about the current state of the economy. These estimates are not driven by unobserved local shocks to consumer demand. Second, economic sentiments influences consumption and hiring decisions: a standard deviation increase in sentiments is associated with a 0.06% rise in daily consumption of non-durables and a 0.17 standard deviation rise in hiring intensity. These estimates are robust to exploiting quasi-experimental variation in daily temperature as an instrument for sentiments and implementing a difference-in-difference estimator of around the 2016 Presidential election. A back-of-the-envelope calculation suggests that the decline in sentiment can account for 19.3% of the decline in consumption during the Great Recession.

Keywords: business cycle, consumption, uncertainty, sentiments.

JEL: E20, E21, E32

---

*Christos Makridis, Department of Economics and Department of Management Science & Engineering, Huang Engineering Center, 475 Via Ortega, Stanford, CA 94305, http://stanford.edu/~cmakridi/. Acknowledgements. Thank you to George-Marios Angeletos, Paul Beaudry, Nicholas Bloom, Joshua Gottleib, Jennifer La'O, Paige Ouimet, Andrew Oswald, Daniela Scur, and Basit Zafar. I also thank Kris Hodgins at Gallup and Ron Nakao at Stanford for taking the time to help me with the Gallup micro-data and to Gallup more broadly for the academic partnership.

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] Expectations play a major role in accounting for these sudden and sharp declines in asset prices (Adam et al., 2017), particularly housing (Kaplan et al., 2016), that prompted these consumption and employment declines. Was the formation of expectations among households linked to their exposure of different local shocks? If so, can the decline in economic sentiments 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; 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., 2016), and even the transmission of the Great Recession across countries (Bacchetta et al., 2012; Bacchetta and van Wincoop, 2016). The primary contribution of this paper is to provide microeconomic evidence behind the formation of these beliefs about the economy and the effect that these beliefs have on real economic activity.[3]

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.

[1] While there has been some dispute that the expansion of credit was not driven by sub-prime borrowers (Adelino et al., 2016; Ferreira and Gyourko, 2015), there is no debate that housing wealth declined across the board.

[2] Understanding the specific role of expectations is important in isolating the driver behind the surge in housing prices leading up to the financial crisis. While I do not take a stand on the underlying source of the run-up in housing prices, Adelino et al. (forthcoming) and Mian and Sufi (2016) for two competing accounts of the "expectations" and "subprime" view.

[3] 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.

While survey participants are different each day, the data offers three advantages that complement existing sources, like the University of Michigan Index of Consumer Sentiments. 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 not only sentiment indices about perceptions of both the current (one to four scale) and future (one to three scale) states of the economy, but also income bins, non-durable consumption expenditures, and hiring intensity (one to three scale). Both measures of sentiment have a high correlation with the volatility index and the Baker et al. (2016) index of policy uncertainty.[4] Third, respondents also report detailed geographic information, such as their zipcode and county, enabling me to distinguish between very local and aggregate shocks. Together with additional demographic and occupational information, I also examine how different types of individuals respond to the same local shock.

Two descriptive patterns are useful to characterize the empirical setting. First, not only did households become more pessimistic about the current and future state of the economy during the 2007-2009 recession, but also households grew increasingly uncertain about the state of the economy (see Figure 1). This is consistent with establishment evidence that first moment shocks are procyclical, whereas second moment shocks are countercyclical (Bloom et al., 2015). For example, 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. Second, sentiments are strongly correlated with consumption expenditures with a correlation of 0.42 (see Figure 2). This is consistent with related survey evidence from the University of Michigan Index of Consumer Sentiments (Carroll et al., 1994; Bram and Ludvigson, 1998).

[INSERT FIGURES 1 AND 2 HERE]

The second part of the paper estimates how individuals respond to local productivity shocks measured using county $\times$ quarter employment and housing price growth given that expectations over both the housing (Piazzesi and Schneider, 2009; Burnside et al., 2016; Glaeser and Nathanson, 2017; Bailey et al., forthcoming) and labor (Carroll and Dunn, 1997; Hendren, 2017) markets are important for understanding booms and busts. I find that exposure to a one percentage point

[4]The constructed measures of sentiments exhibit a striking resemblance with the measure of economic policy uncertainty developed in Baker et al. (2016). The correlation between mean (standard deviation) sentiments and their economic policy uncertainty index is 0.67 (0.43). I also validate my measure using a Google trends count index for the words "economic uncertainty", which produces similar correlations. Despite these similarities, there are two core advantages of the new measure: first, its micro-level variation (making it more plausibly exogenous) and second, its broader representation of perceptions of economic activity (rather than uncertainty).

(pp) higher employment and housing shock is associated with a 0.24-0.57 and 0.21-0.40 standard deviation rise in perceptions about the current state of the economy. However, these averages mask significant heterogeneity. For example, I find nearly four-times as large of an impact of employment growth on sentiment among construction workers in comparison with professional services workers. Given that the construction sector was especially impacted during the financial crisis, this result is consistent with recent evidence from Malmendier and Nagel (2016) and Kuchler and Zafar (2017) that personal experience matters for understanding the formation of expectations. Similarly, I find over a three-times as large of an impact of employment growth on sentiment among individuals who are liquidity constrained in comparison with those who are not.[5]

To address the concern that these sentiment gradients merely reflect reverse causality, I implement several additional exercises. First, the results are robust to controlling for county × quarter employment growth in the wholesale and retail trade sector, which is a proxy for consumer demand. If exposure to heterogeneous local shocks was simply a reflection of reverse causality arising from heightened consumer demand, then the gradient would be attenuated by the inclusion of this control. Second, the results are robust to estimating the sentiment gradient across the distribution of different jobs (e.g., as discussed in the case of construction and professional services workers). Again, if reverse causality were the primary culprit, then there should be little response among employees who work in the tradables sector and do not depend explicitly on local consumer demand. I also exploit monthly state variation in unanticipated mass layoffs and find that a 1% rise in mass layoffs is associated with a 0.041 and 0.016 decline in the standard deviation of perceptions about the current and future state of the economy, respectively. These results are consistent with related evidence on sentiments from Benhabib and Spiegel (2016) and Gillitzer and Prasad (2016).

The third part of the paper estimates a micro elasticity between sentiments and non-durables 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)). I find that a standard deviation rise in sentiments is associated with a 0.06% rise in daily consumption among households and a 0.17 standard deviation rise in hiring intensity among employers. A back-of-the-envelope calculation suggests that the decline in sentiment can account for 19.3% of the decline in consumption during the Great

[5]Although my measure of liquidity constraints is not perfect, individuals respond to the following question on a scale of one to five: "In the last seven days, you have worried about money." I take those reporting a four or five as liquidity constrained and those reporting one or two as not liquidity constrained.

Recession. 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. Motivated by recent work in environmental economics on the link between temperature and temperament (Baylis, 2015; Baylis et al., 2017), I also show that these estimates are robust to instrumenting for sentiments using maximum daily temperature: hotter days make individuals more irritable, changing the way they process information.

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. (2016) on the importance of expectations in explaining the decline in consumption during the financial crisis.

An earlier draft of this paper was also written around the same time as Benhabib and Spiegel (2016), Gillitzer and Prasad (2016), and Mian et al. (2017). Using the University of Michigan Index of Consumer Sentiments, 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 (2016) exploit changes in the government party in power to identify the effects of expectations on an intent to spend more in the future. Mian et al. (2017) also uses the University of Michigan survey, as well as the Gallup U.S. daily poll in the most recent 2017 draft of their paper, but focuses on how partisan shocks affect consumer spending. While they find that, for example, the 2016 Presidential election raised expectations among Republicans, they do not find any significant effect on consumer spending.

This paper provides a partial reconciliation of these contrasting views in the literature by providing an alternative source of variation and focusing on several dimensions of heterogeneity that matter. First, using quasi-experimental variation in climate (e.g., temperature), I show that abnormally hot days reduce individual perceptions about economic activity. This result is in line with an emerging body of evidence in environmental economics that temperature affects temperament (Baylis, 2015), especially since it affects the individual personally (versus indirectly through word of mouth or the news). Second, using heterogeneity in the self-reported intensity of political beliefs, I show that individuals identifying as very conservative increased their spending twice as much as nominal conservatives following the election.[6] This is consistent with evidence from the political science literature that intensity matters. Especially given the fact that Donald Trump was an outsider who disrupted many establishment Republicans, average Republican vote shares in Mian et al. (2017) might mask very heterogeneous beliefs.

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 estimates how exposure to different county employment and housing growth shocks affect perceptions about the current and future states of the economy. Section 4 estimates how sentiments affects real economic activity measured primarily through personal consumption expenditures and hiring intensity. Section 5 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.[7] Appendix Section 6.1 provides descriptive details about the data and compares daily consumption expenditures from Gallup with the BEA at a state $\times$ year aggregation.

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

[6] 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.

[7] Unfortunately, the data does not contain a panel component, so I cannot do the decomposition in Kimball et al. (2015).

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. The survey has changed in some dimensions since 2008 with the inclusion of detailed well-being related indices since 2014, but the main outcomes of the analysis are available throughout.

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.[8] While a limitation is that it only asks the respondent about their spending the day before, the Data Appendix compares the Gallup measure with the Bureau of Economic Analysis state-level personal consumption panel between 2008 and 2014. To provide a graphical characterization of the spatial variation, Figure 3 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 sentiments. Not surprisingly, Texas has the highest sentiments, followed by states like North Dakota, which experienced a shale gas boom.

[INSERT FIGURE 3 HERE]

[INSERT TABLE 1 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

[8] 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.

that can spill over from the answer in one question to another.[9] The second is the potential wedge 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, 2011). 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 quarterly median housing price per square foot data from Zillow over the same sample period. The administrative records from the QCEW covers 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, the median housing price per square foot data from Zillow uses 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. Their Zestimates predictions correlate well with the Federal Housing Administration's housing price index and has been used in recent research (e.g., Bailey et al. (forthcoming)).

[9]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.

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 can also be used to help characterize higher-order sentiments and uncertainty (Angeletos and La'O, 2009; Angeletos et al., 2014).

How do such measures of sentiments and uncertainty compare with existing measures? Figure 4 compares the standard deviation of the measure produced from the Gallup micro-data with a more standard measure of economic policy uncertainty from Baker et al. (2016). Whereas greater values of my index imply a more certain state of the economy, greater values of the Baker et al. (2016) index imply a more uncertain state of the economy. It is, therefore, remarkable that there is a -0.66 correlation between the two, suggesting that they are capturing similar dimensions of uncertainty in the U.S. economy.[10] Figure 4 also shows that dispersion in perceptions about the state of the economy has a -0.47 correlation with the volatility index. Given these assuring correlations at a national time series aggregation, it is again useful to emphasize that one of Gallup's main advantage is its micro-level and location-specific variation. I have also found a relatively strong correlation with investor sentiment from Baker and Wurgler (2006).

[INSERT FIGURE 4 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

[10]The first moment of perceptions about the economy is also correlated with their index (correlation of -0.57), but I report the dispersion to more accurately map into the measurement of uncertainty. My focus in this paper, however, is simply on sentiment.

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 5. 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 about the future state of the economy are much more bi-modal in 2008-2009. In other words, many individuals who thought the economy would get better quickly also thought the economy would be getting worse. These differences exist even within the same area.

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 5 HERE]

3 Sentiments and Exposure to Economic Shocks

3.1 Empirical Strategy

To understand the relationship between sentiments—using standardized perceptions of the current and future state of the economy—and real economic activity, I consider regressions 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 time (year and quarter). 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.

The main threat to this identification assumption is reverse causality. If sentiments affect consumer spending, then that feeds into aggregate demand conditions and, therefore, employment growth. Similarly, a positive rise in sentiments could lead individuals to become more optimistic about the housing market and, therefore, bid up local housing prices. Since there is no silver-bullet solution to these challenges, my strategy is to implement placebo exercises that help validate the baseline results. First, since employment in wholesale and retail trade is often used as a proxy for consumer demand, I control for it to examine whether exposure to county employment growth is simply masking a response to a demand shock. Second, to the extent these gradients represent pure demand-side factors that affect sentiments, then individuals in the tradables sector, such as manufacturing or professional services workers, should not respond much to local shocks.

3.2 Main Results

Table 2 documents the results associated with Equation 1. Columns 1-2 and 5-6 simply report the the gradients on employment and housing price growth in isolation. As the high gradients suggest, these are subject to significant upwards bias since both employment and housing prices were co-moving and affecting economic activity (Mian and Sufi, 2014). However, once they are both included together, column 3 suggests that an additional percentage point (pp) rise in employment and housing price growth is associated with a 0.57 and 0.40 standard deviation rise in perceptions about the current state of the economy. The fact that employment growth has a larger association with sentiments, relative to housing prices, reflects the fact that it is a potentially more salient shock—that is, local residents observe employment declines more than they observe housing price declines since the latter are not always physically visible.[11]

Turning towards perceptions about the future state of the economy, housing price growth has a larger and more consistent positive association with sentiments. For example, whereas a one pp rise in employment growth is not statistically associated with improved perceptions about the future state of the economy when both employment and housing growth are included (column 7), a one pp rise in housing price growth is associated with a 0.40 standard deviation rise. One reason that employment growth is insignificant in this specification arises from the fact that declines in housing prices led to an immediate decline in the expansion of credit and home equity, thereby accounting for most of the variation in expectations (Mian and Sufi, 2009; Mian et al., 2013).

Since there there are a number of potentially confounding state-level forces moderating how employment and housing price shocks, such as differences in state labor market institutions (Shoag and Veuger, 2016) and judicial foreclosure laws (Mian et al., 2015), columns 4 and 8 introduce state $\times$ year $\times$ quarter fixed effects to exploit within-county variation, controlling for all shocks common to individuals across states. Columns 4 and 8 suggests that a one pp rise in employment and housing price growth is associated with a 0.24 and 0.21 (0.22 and 0.26) standard deviation rise in perceptions about the current (future) states of the economy. The fact that housing price growth affect perceptions about the future more than comparable employment growth is consistent with the view that investors often look at housing prices as a signal for future economic activity, whereas employment growth (especially in wholesale and retail trade) is viewed as a metric for current economic activity.

[INSERT TABLE 2 HERE]

Macroeconomic models predict that liquidity constrained individuals are those who will respond most elastically fiscal stimulus payments (Kaplan and Violante, 2014). While an individual's monthly income is a useful heuristic for financial constraints, there are also many wealthy hand-to-mouth workers who earn and spend a lot, therefore holding little liquid wealth. Using information from one of Gallup's survey questions on whether the individual has worried about money in the

[11]Motivated by the fact that firms respond to both first and second moment shocks, I also used the QCEW's information on monthly employment growth to compute the standard deviation of employment growth within a quarter. I found that, even conditional on the first moments, increases in employment growth volatility are strongly negatively related with sentiment. I omit these formal tests since they are not central to the main message of the paper.

past seven days, I classify individuals as constrained and unconstrained and estimate the baseline specification again with income bins as an additional control.[12]

Table 3 documents these results. As expected, the elasticities are stronger for constrained individuals when the outcome variable is perception about the current state of the economy. For example, a one pp rise in employment and housing price growth is associated with a 0.52 and 0.18 standard deviation rise in sentiments among constrained individuals versus a 0.39 and 0.16 standard deviation rise among unconstrained individuals. However, the evidence is more mixed when the outcome is perception about the future state of the economy. In particular, while the employment growth elasticity is noisy and only marginally larger for constrained individuals, the housing price growth elasticity is larger and more precisely estimated for unconstrained individuals.[13] Appendix Section 6.2 also examines several additional dimensions of heterogeneity across the income distribution. I find that higher employment and housing price growth generally has a bigger effect on sentiments for higher income individuals, but a more attenuated effect on the probability for being liquidity constrained. These results are consistent with partial insurance against labor market risk for higher skilled households (Blundell et al., 2008).

[INSERT TABLE 3 HERE]

3.3 Exercises Mitigating Reverse Causality

Given that the central argument in this paper is that exposure to labor and housing market shocks affect individual expectations about future economic activity, thereby feeding into their consumption decisions and endogenously influencing aggregate growth, an important concern is that the results thus far merely reflect reverse causality. Since my variation in sentiments is at the individual-level, whereas the local shocks are at a county-level, controlling for individual income is a first-step in assuaging the concern that the labor and housing market fluctuations on the right-hand-side are not merely a function of individual sentiment changes. Nonetheless, I also implement several additional exercises that help further mitigate this concern.

My first exercise introduces employment growth in the wholesale and retail trade sector as an additional control and proxy for consumer demand. Since the subset of employment that should

[12] The survey question: "In the last seven days, you have worried about money." Between 2013-2016, the measure is on a one to five scale of intensity, whereas between 2009-2012 it is a binary variable. For 2013-2016, I classify those as constrained if they report a four or five and unconstrained if they report a one or two.

[13] One possibility is that my proxy for liquidity constraints is simply noisy.

be responding to sentiments is this sector, then controlling for it should attenuate and/or remove the estimated gradient on overall county employment growth if reverse causality is the primary culprit. However, in estimating Equation 1, I find that $\gamma = 0.54$ and $\phi = 0.39$—similar to the baseline results—both with p-values equal to 0.00. (The coefficient on employment growth in trade is 0.06.) Since the manufacturing and finance and real estate sectors were also hit hard during the crisis, I also control for employment growth in them and the results are not altered.

My second exercise exploits 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, than 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 6 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.

[INSERT FIGURE 6 HERE]

My third exercise exploits 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 and future state of the economy 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 0.41 (p-value = 0.00) and 0.16 (p-value = 0.00) standard deviation decline in current and future sentiments. These gradients are also larger among workers in services and production jobs, which are precisely the jobs that are more susceptible to mass layoffs. 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.

4 Quantifying the Real Effects of Sentiments

4.1 The Elasticity of Sentiments and Consumption & Hiring

To understand the relationship between sentiments and real economic activity, I focus on two measures of real economic activity—personal consumption expenditures on non-durable goods, which declined significantly during the financial crisis (Mian et al., 2013; Pistaferri, 2016), and hiring at the employee's company—through regressions of the form

$$y_{ict} = \zeta s_{ict} + \beta X_{it} + \omega w_{ict} + \gamma \Delta e + \phi \Delta h + + \eta_c + \lambda_t + \epsilon_{ict} \tag{2}$$

where y denotes the individual outcome variable, s denotes the measure of economic sentiments, 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, year, and quarter.[14] Standard errors are clustered again at the county-level.

The primary coefficient of interest, ζ, is identified off of comparisons among similar individuals within-county in their variation on sentiment fluctuations and how they differentially affect personal consumption expenditures and employer hiring. The identifying assumption is that unobserved shocks to consumption or hiring are uncorrelated with underlying sentiments, conditional on individual income and county employment and housing price growth. These controls are important for mitigating the obvious concern of reverse causality that would emerge from the effects of exposure to local shocks on sentiments. In this sense, identification of ζ comes from variation

[14]The Gallup survey on consumption is asked as follows: "Next, we'd like you to think about your spending yesterday, not counting the purchase of a home, motor vehicle, or your normal house bills. How much money did you spend or charge yesterday on all other types of purchases you may have made?" The survey question on hiring is asked as follows: "Now thinking more generally about the company or business you work in, including all of its employees. Based on what you know or have seen or would say, in general, your company or employer is: (i) hiring new people and expanding the size of its workforce, (ii) not changing the size of its workforce, (iii) letting people go and reducing the size of its workforce. I re-code the variable so higher levels of the index indicate expansion and lower levels indicate contraction. While one limitation is that consumption and sentiment data are only simultaneously available from 2008 to 2013 since the Gallup survey design changed in 2013—the information on consumption and information on sentiments were partitioned into two separate surveys—there is still variation.

in beliefs about the economy that are not driven by individual or local income shocks.

Besides reverse causality, the primary threat to identification is that if daily consumption expenditures is explained by the sum of a transitory component of consumption and a permanent person-specific fixed effect, then ζ will be biased since that person-specific effect will also be correlated with perceptions about the economy. However, the direction of the bias is not obvious, especially since the coefficient is identified off of variation that already purges income fluctuations. I solve this concern by constructing a leave-one-out estimate of county $\times$ year $\times$ quarter sentiment as an alternative right-hand-side regressor. By construction, the measure does not contain individual heterogeneity and the location and time fixed effects remove heterogeneity due to non-random sorting. In the subsection that follows, I also explore robustness relating to reverse causality and other time-varying omitted variables in greater detail.

Table 4 documents these results. Using the individual-level measure of sentiments, I find that a standard deviation rise in perceptions about the current state of the economy is associated with a 0.02% and 0.13 standard deviation rise in daily consumption expenditures and hiring intensity, respectively. Once logged monthly income is added as a control, the gradient declines in half when consumption is the outcome variable, whereas it marginally rises for hiring intensity, reflecting the close connection between higher incomes and greater spending on non-durable goods. Indeed, a 1% rise in monthly income is associated with a 0.21% rise in daily consumption. One concern with the measure of hiring intensity is that it is a weak proxy since employees may not have information about the status of hiring at their company. To examine this concern, I restrict the sample to business owners ($N = 19,057$) and estimate the same specification, recovering a gradient of 0.152 (p-value = 0.00). Although it is slightly larger in magnitude, it is remarkably similar and, therefore, assuages the concern that hiring intensity among non-owners is simply noise.[15]

Turning towards the leave-one-out county average, I find larger gradients. For example, a standard deviation rise in sentiment is now associated with a 0.06-0.08% rise in daily consumption and a 0.16-0.17 standard deviation rise. Why does the gradient rise? If the individual sentiment

[15] I further examined the potential that this does not represent realized hiring, but rather stated aspirations for hiring, by drawing from the Bureau of Labor Statistic's monthly state employment and employment data. I consider regressions of logged state employment on the z-score of sentiments about the economy and find a coefficient of 0.198 (p-value = 0.467). When the outcome variable is the year-to-year growth in employment, I find a gradient of 0.082 (p-value = 0.016). Regressions are weighted by state population and controls include annual measures of: the share who are male, the share who are white, black, the share who are married, shares on the age distribution (between 0-17, 35-64, and 65+ normalized to 18-34 as the omitted group), shares on the education distribution (no high school, high school, some college, and graduate degrees normalized to the college degree as the omitted group), employment growth. Since hiring more closely resembles employment growth, but employment growth also pools separations, the similarity between these results is very assuring.

contains noise, then it will be more likely to produce bias, whereas the location-specific average produces a more reliable signal. Since the leave-one-out estimate only improves upon the individual measure for counties that have a sufficiently high number of observations in the data, I restrict the sample to counties with at least 50 county $\times$ year $\times$ quarter respondents.

[INSERT TABLE 4 HERE]

As a comparison to more reliable regional data, I now combine annual consumption expenditure data from the Bureau of Economic Analysis (BEA) between 2008 and 2016 and estimate regressions along the lines of Equation 2 at the state $\times$ year level. Sentiment is measured in two ways: (i) the average self-reported index of perceptions about the current state of the economy, and (ii) the share of individuals reporting that the economy is worsening. Since the BEA distinguishes among many types of consumption, I focus on four main measures: total per capita consumption expenditures, durables, non-durables, and services. All measures are deflated by the 2010 personal consumption expenditure index and logged. I also use the American Community Survey (ACS) from the Census Bureau to control semi-parametrically for the age distribution (four bins), gender, race (white and black), marital status, the education distribution (five bins), and the employment growth. These highly flexible controls help mitigate the concern that the sentiment gradient reflects changes in composition and/or reverse causality.

Table 5 documents these results along two specifications. The first simply presents the conditional correlations; the second estimates the specification in growth rates to remove time-invariant heterogeneity across states. Increases in sentiment are positively associated with consumption per capita in every instance, but they are largest for non-durable consumption expenditures. For example, a 1% rise in average state sentiment is associated with a 0.16% rise in per capita consumption on non-durables. When estimating the specification in growth rates, I find that increases in the growth in the share of individuals reporting that the economy is worsening is negatively associated with consumption per capita growth in ever case and (again) strongest for non-durables.[16] The results are also robust to including the share of individuals working in the construction sector, manufacturing, or the finance, insurance and real estate (FIRE) sectors.

[INSERT TABLE 5 HERE]

[16] While the results are also robust to simply adding state and year fixed effects, little variation remains and the R-squared is very high. Writing the specification in quasi-growth rate terms (outcome variable and main independent variable) achieves the same goal (removing time-invariant heterogeneity) while keeping more variation.

What do these elasticities imply about the effect of sentiment on aggregate consumption? Although a full general equilibrium model is required to make any definitive statements, I implement a simple back-of-the-envelope calculation. According to the Bureau of Economic Analysis, consumption of non-durables per capita fell by 5.2% from \$7,458 to \$7,069. This is comparable to the decline in overall consumption per capita tracked by the St. Louis Federal Reserve of \$10,151 to \$9,792 (3.5% decline). Since the U.S. Daily Poll begins in 2008, I have no pre-recession measure of sentiment and, therefore, draw from Benhabib and Spiegel (2016) who uses the University of Michigan's Index of Consumer Sentiment and reports that the share of individuals reporting that the economy is good fell from 50% in 2007 to 30% in 2008 (a decline of 66%). To make sure my elasticity is measured in the same units, I re-estimate Equation 2 using the share of individuals who report the economy is "fair" or "excellent" as the main right-hand-side variable, recovering a coefficient of 1.52%. Putting these numbers together, the decline in economic sentiment accounts for 19.3% ($= 1.52 \times 66/5.2$) of the decline in consumption expenditures on non-durables.

4.2 Instrumental Variables Robustness

Since a concern associated with the results thus far is that sentiments are simply responding to unobserved and time-varying location-specific shocks, I now exploit an alternative source of quasi-experimental variation in the form of daily temperature shocks.[17] The motivation for temperature as an instrument for sentiments comes from recent empirical work using billions of social media posts from twitter (Baylis, 2015; Baylis et al., 2017). Hotter days are likely to make individuals more irritable, which affect how individuals perceive and process information around them. I also leverage the fact that different types of workers will be heterogeneously exposed to temperature through their day-to-day activities, allowing me to interact daily temperature with occupation fixed effects. I estimate the following two-staged least squares specification

$$\begin{aligned} s_{ict} &= \rho T_{ct} + \textstyle\sum_k \delta^k (O^k_{ict} \times T_{ct}) + \beta X_{it} + \omega w_{ict} + \gamma \Delta e + \phi \Delta h + + \eta_c + \lambda_t + \epsilon_{ict} \\ y_{ict} &= \beta X_{it} + \zeta \widehat{s_{ict}} + \omega w_{ict} + \gamma \Delta e + \phi \Delta h + + \eta_c + \lambda_t + \epsilon_{ict} \end{aligned} \tag{3}$$

where X contains not only the usual controls, but also fixed effects on occupation, i.e. $\sum_k O^k$, and both precipitation and time allocated towards commuting to control for the fact that temper-

[17] I use daily county data between 2008 and 2013 assembled by Schlenker and Roberts (2009) from the PRISM database. I end the sample at 2013 since the consumption and sentiment micro-data from Gallup is not asked together with the same person in subsequent years after they revised the survey methodology.

ature might affect the timing of consumer purchases by delaying purchases (Agarwal et al., 2017). One concern about the exclusion restriction in Equation 3 is that temperature may have direct effects on productivity. To examine whether this concern is a major threat to identification, I draw from 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. I find no association between temperature and either weekly earnings or hours worked—the details of the sample construction and the full results are presented in Appendix Section 6.3.1.

Table 6 documents these results. Figure 7 begins by providing a non-parametric characterization of the first-stage: very cold and very hot days reduce economic sentiments. Although the first-stage F-statistic of four is below the conventional rule of thumb (Stock and Yogo, 2005) in the fully-specified 2SLS implementation, which is a feature of the limited within-county variation in temperature between 2008 and 2013, the point estimates are still statistically significant and close to the baseline. For example, a standard deviation rise in perceptions about the current state of the economy is associated with a 0.12% rise in daily consumption expenditures.

[INSERT TABLE 6 HERE]

[INSERT FIGURE 7 HERE]

4.3 Comparison with Recent Research and Implications for Macroeconomic Models

There has been a revitalization of interest in the relationship between consumption and sentiment in recent work by Benhabib and Spiegel (2016), Gillitzer and Prasad (2016), and Mian et al. (2017). However, my results differ slightly from Mian et al. (2017) who find that the Presidential election results increased economic expectations among Republicans, but did not lead to a concurrent rise in consumption expenditures. One important difference between these studies and mine is that their source of variation when instrumenting for sentiment is based on a political shock, i.e., the impact of an election on expectations among party affiliates versus their opponents. While Mian et al. (2017) found that political shocks influenced expectations about economic activity, they did not find that it translated into actual increases in automobile purchases at a zipcode-level.[18]

[18] Gillitzer and Prasad (2016), however, in the setting of Australia, found a delayed increase in consumption.

One potential reason for the null association with consumer spending is that there are heterogeneous treatment effects. In particular, since Mian et al. (2017) identify zipcodes based on the share of Republican voters, these areas might pool "establishment Republicans" with those in the "silent majority"—political intensity might be a moderating variable for expectations. To understand the potential role that this heterogeneity plays, I examine how individual consumption expenditures on non-durables behaves for three groups of workers—those who self-report as moderates, somewhat conservative, and very conservative—implemented through the following difference-in-difference estimator

$$y_{it} = \alpha_1 Post_t + \alpha_2 SC_{it} + \alpha_3 VC_{it} + \xi_1 (Post_t \times SC_{it}) + \xi_2 (Post_t \times VC_{it}) + \beta X_{it} + \eta_c + \Phi_{st} + \epsilon_{it} \quad (4)$$

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 4 is restricted to the sample of moderates and conservatives; the coefficients of interest—ξ_1 and ξ_2—are interpreted relative to the consumption expenditures of a moderate.

Table 7 documents these results. First, although there is an upward trend in consumption expenditures overall, it becomes statistically insignificant after the inclusion of fixed effects on year and month. Second, both somewhat and very conservative individuals do not have statistically different spending patterns than their moderate counterparts. Third, and most importantly, there is a strong increase in consumption expenditures among conservatives, especially very conservative individuals: after the election, very conservative individuals spend 12% more on consumption items than their moderate counterparts. The results are robust to restricting the sample to four months before and after the election. Whereas one of the challenges that Mian et al. (2017) face in working with zipcode-level data is the geographic aggregation, this data enables me to exploit high frequency variation among individuals with heterogeneous political preferences.

[INSERT TABLE 7 HERE]

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.

Finally, my results also related 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).

5 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.

The first main set of results highlight how exposure to local labor and housing market shocks affect individual sentiments about economic activity. In particular, a one percentage point rise in quarterly county employment and housing price growth is associated with a 0.57 and 0.40 standard deviation rise in perceptions about the current state of the economy. These results are robust to introducing state $\times$ year $\times$ quarter fixed effects and exploiting only within-county variation. I find similar results for sentiments about the future state of the economy. I subsequently show that these estimates are not driven by reverse causality through three supplementary exercises: (i) the estimates are robust to including a control for consumer demand proxied via employment growth in the wholesale and retail trade sectors, (ii) the estimates are robust even among workers in jobs that are concentrated in the tradables sector and not dependent directly on the local economy, and (iii) sentiments also respond to mass layoffs at a month $\times$ state frequency, which represent plausibly exogenous variation in labor market conditions.

The second main set of results highlight how these sentiments influence real economic activity measured through daily consumption expenditures and perceived hiring intensity. In particular, a standard deviation rise in perceptions about the current state of the economy is associated with a 0.06% rise in consumption expenditures and a 0.17 standard deviation rise in hiring intensity—even after controlling for individual income fluctuations and county employment and housing conditions. These estimates suggest that consumption and employment respond to sentiments even after purging variation in realized economic activity. A back-of-the-envelope calculation suggests that the decline in sentiment can account for 19.3% of the decline in consumption during the Great Recession. To deal with the potential for reverse causality, I also exploit changes in daily temperature as an instrument for sentiment. I provide new evidence that temperature fluctuations shape perceptions about the current state of the economy. I also reconcile these results with conflicting evidence that have used political shocks to identify the impact of sentiments on consumption by showing how pooling together politically heterogeneous individuals attenuates estimates.

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 **?**. Integrating microeconomic data with these macroeconomic heterogeneous agent models will be an essential step forward in understanding these broader phenomena.

References

Adam, K., A. Marcet, and J. Beutel (2017): "Stock price booms and expected capital gains," *American Economic Review*, 107, 2352–2408.

Adelino, M., A. Schoar, and F. Severino (2016): "Loan originations and defaults in the mortgage crisis: The role of the middle class," *Review of Financial Studies*, 29, 1635–1670.

——— (forthcoming): "Dynamics of housing debt in the recent boom and bust," *NBER Macro Annual 2017*.

Agarwal, S., J. B. Jensen, and F. Monte (2017): "The geography of consumption," *NBER working papers*.

Angeletos, G.-M., F. Collard, and H. Dellas (2014): "Quantifying confidence," *Working paper*.

Angeletos, G.-M. and J. La'O (2009): "Incomplete information, higher-order beliefs and price inertia," *Journal of Monetary Economics*, 56.

——— (2013): "Sentiments," *Econometrica*, 81, 739–779.

Azariadis, C., L. Kaas, and Y. Wen (2016): "Self-fulfilling credit cycles," *Review of Economic Studies*, 83, 1364–1405.

Bacchetta, P., C. Tille, and E. van Wincoop (2012): "Self-fulfilling risk panics," *American Economic Review*, 102, 3674–3700.

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

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (forthcoming): "The economic effects of social networks: Evidence from the housing market," *Journal of Political Economy*.

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

Baker, S. R. (forthcoming): "Debt and the consumption response to household income shocks," *Journal of Political Economy*.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring economic policy uncertainty," *Quarterly Journal of Economics*, forthcoming.

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

Baylis, P. (2015): "Temperature and temperament: Evidence from a billion tweets," *Berkeley Job Market Paper*.

Baylis, P., N. Obradovich, Y. Kryvasheyeu, H. Chen, L. Coviello, E. Moro, M. Cebrian, and J. H. Fowler (2017): "Weather impacts expressed sentiment," *Working paper*.

Beaudry, P., D. Galizia, and F. Portier (2018): "Reconciling Hayek's and Keynes' views of recessions," *Review of Economic Studies*, 85, 119–156.

Beaudry, P., D. Nam, and J. Wang (2011): "Do mood swings drive business cycles and is it rational?" *Working paper*.

Beaudry, P. and F. Portier (2006): "Stock prices, news and economic fluctuations," *American Economic Review*, 96, 1293–1307.

Benhabib, J. and M. M. Spiegel (2016): "Sentiments and economic activity: Evidence from U.S. states," *Working paper*.

Benhabib, J., P. Wang, and Y. Wen (2015): "Sentiments and aggregate demand fluctuations," *Econometrica*, 83, 549–585.

Benjamin, D. J., O. Heffetz, M. S. Kimball, and A. Rees-Jones (2012): "What do you think would make you happier? What do you think you would choose?" *American Economic Review*, 102, 2083–2110.

——— (2014a): "Can marginal rates of substitution be inferred from happiness data? Evidence from residency choices," *American Economic Review*, 104, 3498–3528.

Benjamin, D. J., O. Heffetz, M. S. Kimball, and N. Szembrot (2014b): “Beyond happiness and satisfaction: Toward well-being indices based on stated preference,” *American Economic Review*, 104, 2698–2735.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): “How much should we trust differences-in-differences estimates?” *Quarterly Journal of Economics*, 119, 249–275.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta, and S. Terry (2015): “Really uncertain business cycles,” *Econometric, RR.*

Bloom, N., R. Lemos, R. Sadun, D. Scur, and J. Van Reenen (2014): “The new empirical economics of management,” *Journal of European Economic Association*, 12.

Blundell, R., L. Pistaferri, and I. Preston (2008): “Consumption inequality and partial insurance,” *American Economic Review*, 98, 1887–1921.

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

Burnside, C., M. Eichenbaum, and S. Rebelo (2016): “Understanding booms and busts in housing markets,” *Journal of Political Economy*, 124, 1088–1147.

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

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

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

Challe, E., J. Matheron, X. Ragot, and J. F. Rubio-Ramirez (2017): “Precautionary saving and aggregate demand,” *Quantitative Economics*, 8, 435–478.

Crossley, T. and H. Low (2014): “Job loss, credit constraints and consumption growth,” *Review of Economics and Statistics*, 9, 876–884.

Deaton, A. (2011): “The financial crisis and the well-being of Americans,” *Oxford Economic Papers*, 64, 1–26.

Den Haan, W. J., P. Rendahl, and M. Riegler (2017): “Unemployment (fears) and deflationary spirals,” *Journal of the European Economic Association.*

Deryugina, T. and S. M. Hsiang (2017): “The marginal product of climate,” *NBER working paper.*

Deschenes, O. and M. Greenstone (2011): “Climate change, mortality and adaptation: Evidence from annual fluctuations in weather in the US,” *American Economic Journal: Applied Economics*, 3, 152–185.

di Tella, R., R. MacCulloch, and O. andrew (2003): “The macroeconomics of happiness,” *Review of Economics and Statistics*, 85, 809–827.

di Tella, R., R. J. MacCulloch, and O. andrew J. (2001): “Preferences over inflation and unemployment: Evidence from surveys of happiness,” *American Economic Review*, 91, 335–341.

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

Fernald, J. G., R. E. Hall, J. H. Stock, and M. W. Watson (2017): “The disappointing recovery of output after 2009,” *NBER working paper.*

Ferreira, F. and J. Gyourko (2015): "A new look at the U.S. foreclosure crisis: Panel data evidence of prime and subprime borrowers from 1997 to 2012," *NBER working paper*.

Gillitzer, C. and N. Prasad (2016): "The effect of consumer sentiment on consumption," *Working paper*.

Glaeser, E. L., J. D. Gottleib, and O. Ziv (2016): "Unhappy cities," *Journal of Labor Economics*, 34, S129–S182.

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

Guiso, L., P. Sapienza, and L. Zingales (2015): "Time varying risk aversion," *Working paper*.

Heathcote, J. and F. Perri (forthcoming): "Wealth and volatility," *Review of Economic Studies*.

Heathcote, J., K. Storesletten, and G. L. Violante (2014): "Consumption and labor supply with partial insurance: An analytical framework," *American Economic Review*, 104, 1–52.

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

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

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

Kaplan, G., K. Mitman, and G. Violante (2016): "Consumption and house prices in the Great Recession: Model meets evidence," *Working paper*.

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

Kaplanski, G. and H. Levy (2010): "Sentiment and stock prices: The case of aviation disasters," *Journal of Financial Economics*, 95, 174–201.

Keynes, J. M. (1936): "The general theory of employment, interest and money," *Palgrave Macmillan*.

Kimball, M., R. Nunn, and D. Silverman (2015): "Accounting for adaptation in the economics of happiness," *Working paper*.

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

Kuchler, T. and B. Zafar (2017): "Personal experiences and expectations about aggregate outcomes," *Journal of Finance, R&R*.

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

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage risk and employment risk over the life cycle," *American Economic Review*, 100, 1432–1467.

Malmendier, U. and S. Nagel (2016): "Learning from inflation experiences," *Quarterly Journal of Economics*, 131, 53–87.

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

Mian, A., K. Rao, and A. Sufi (2013): "Household balance sheets, consumption and the economic slump," *Quarterly Journal of Economics*, 128, 1687–1726.

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

——— (2011): "House prices, home equity-based borrowing and the US household leverage crisis," *American Economic Review*, 101, 2132–56.

——— (2014): "What explains the 2007-2009 drop in employment?" *Econometrica*, 82, 2197–2223.

——— (2016): "Who bears the cost of recessions? The role of house prices and household debt," *Handbook of Macroeconomics*.

Mian, A., A. Sufi, and F. Trebbi (2015): "Foreclosures, house prices and the real economy," *Journal of Finance*, 70, 2587–2634.

Mian, A. R., A. Sufi, and N. Khoshkhou (2017): "Partisan bias, economic expectations, and household spending," *Working paper*.

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

Morris, S. and H.-S. Shin (2002): "The social value of public information," *American Economic Review*, 92, 1521–1534.

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

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

——— (2011): "Well-being across America," *Review of Economics and Statistics*, 93, 1118–1134.

Perri, F. and V. Quadrini (2016): "International recessions," *Working paper*.

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

Pistaferri, L. (2016): "Why has consumption remained moderate after the Great Recession," *Working paper*.

Ravn, M. O. and V. Sterk (2017): "Job uncertainty and deep recessions," *Journal of Monetary Economics*, 90, 125–141.

Schlenker, W. and M. J. Roberts (2009): "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change," *Proceedings of the National Academy of Sciences*, 106, 15594–15598.

Shoag, D. and S. Veuger (2016): "Uncertainty and the geography of the Great Recession," *Journal of Monetary Economics*, 84, 84–93.

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

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

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

6 Tables and Figures

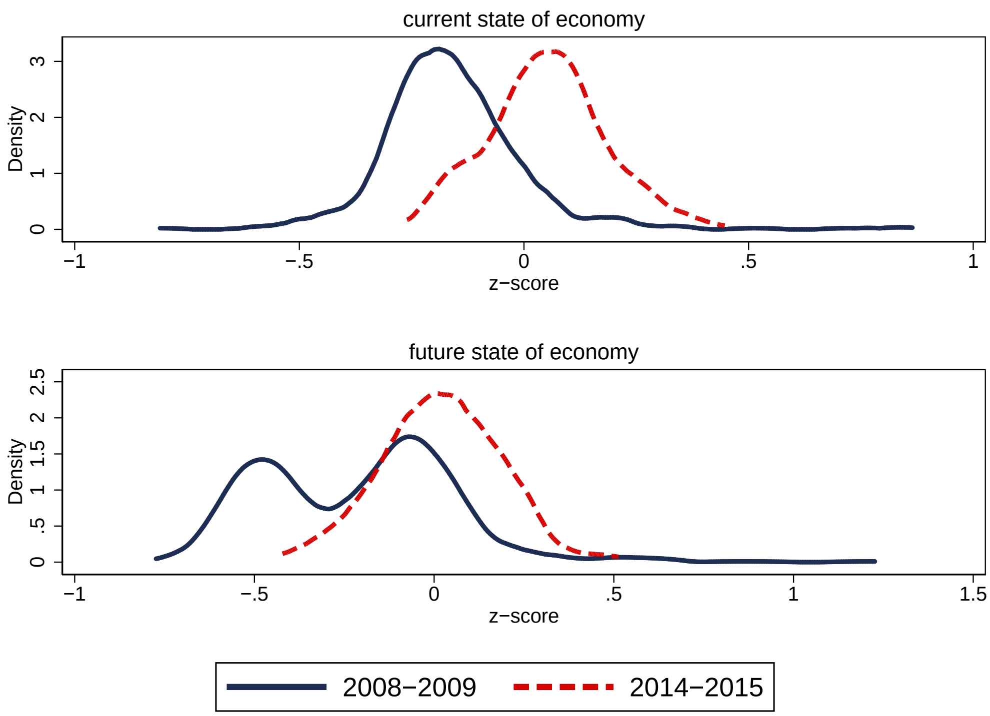

Figure 1: Distribution of Sentiments about the State of Economy

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.

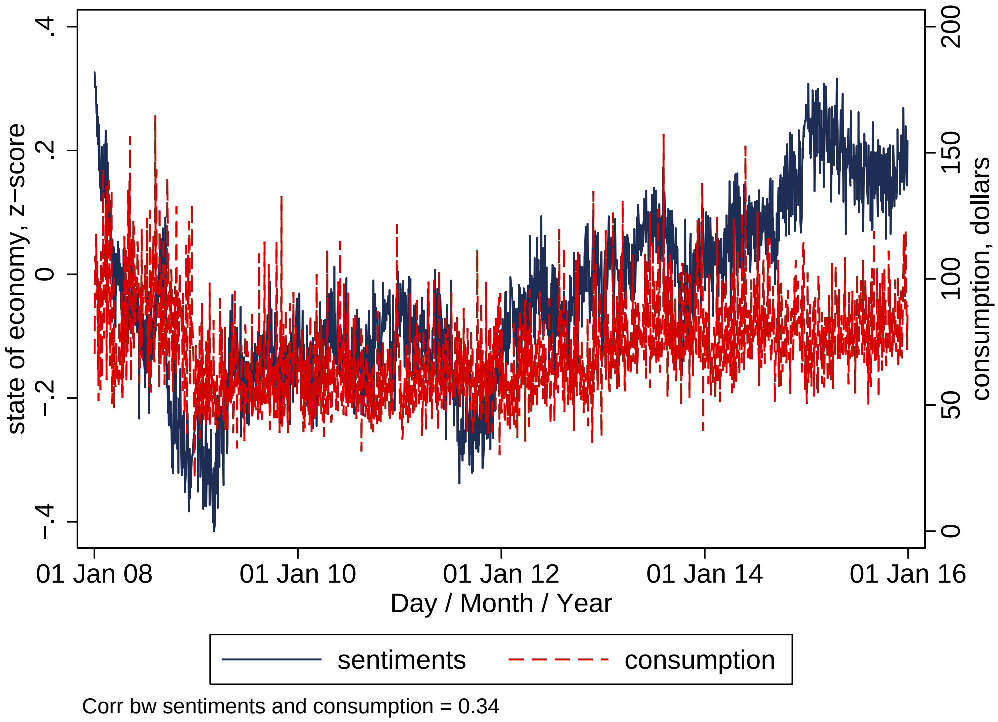

Figure 2: Sentiments and Real Consumption Expenditures

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.

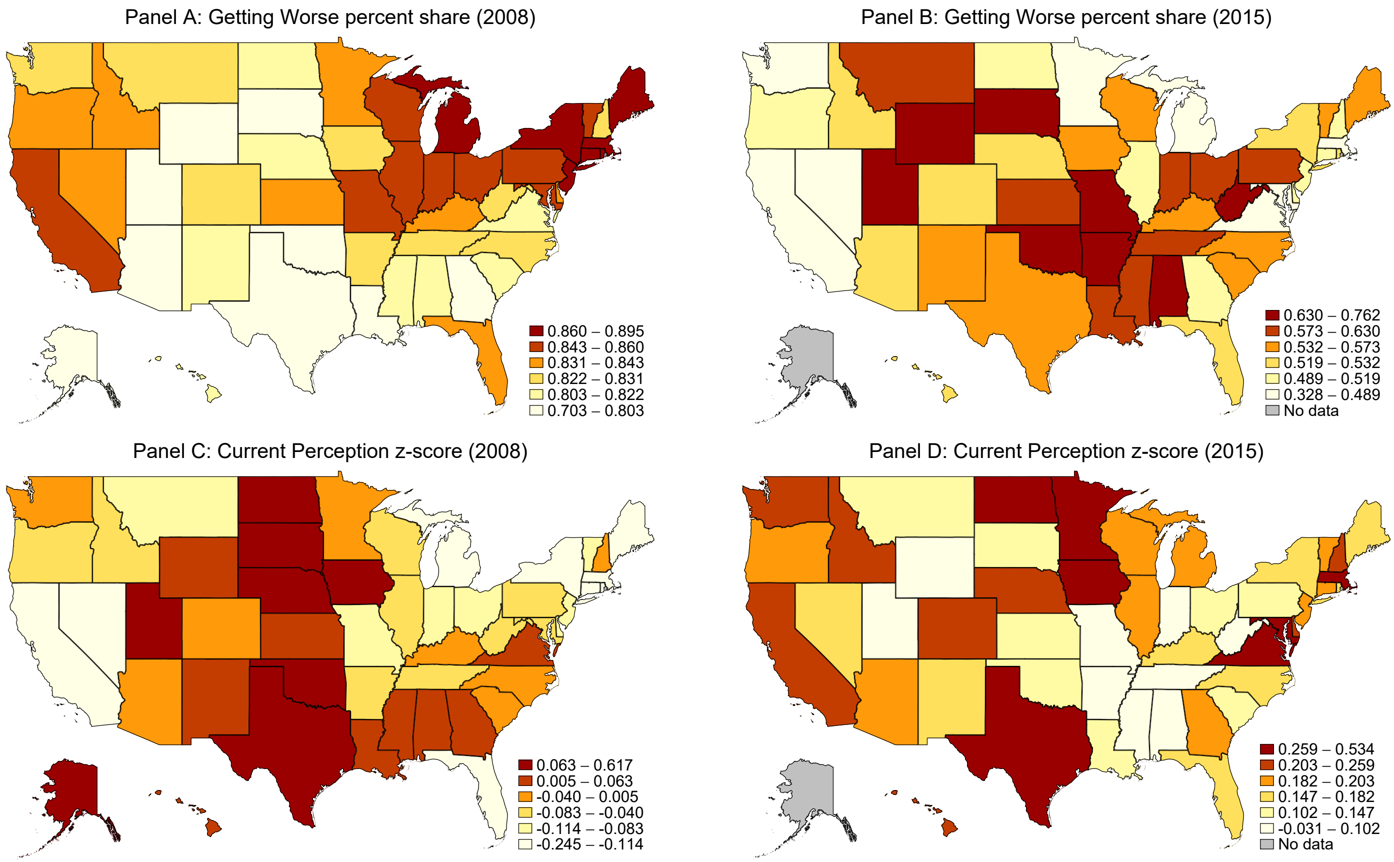

Figure 3: Spatial Variation in Economic Sentiments

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.

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?	1-10 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.

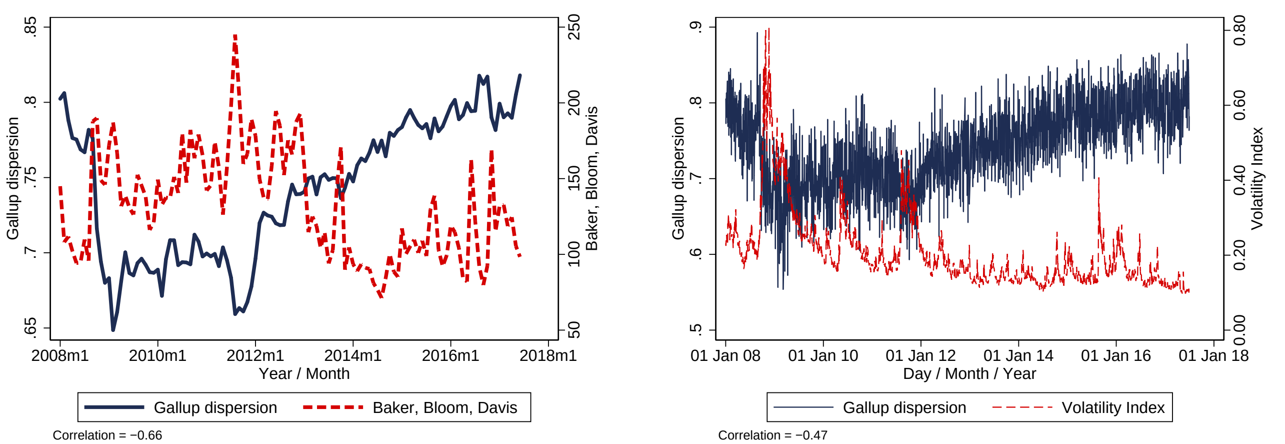

Figure 4: Measuring Uncertainty, Comparison Between Gallup and Baker et al. (2016)

Notes.–Sources: Gallup, 2008-2017. The figure plots the standard deviation of 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 correlation between the dispersion in Gallup sentiment and economic policy uncertainty (VIX) is -0.67 (-0.47).

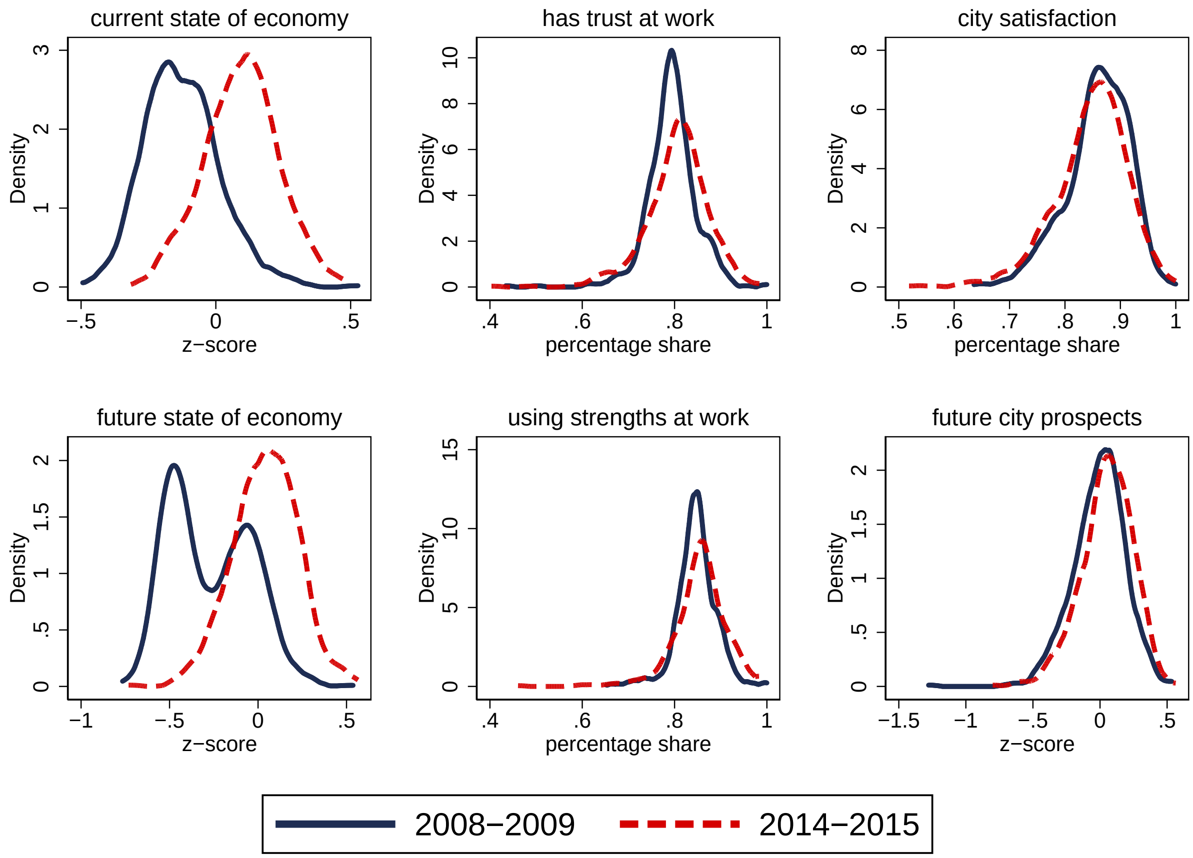

Figure 5: Dispersion in Sentiments Across Metro Areas, 2008-2009 and 2014-2015

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.

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

Dep. var. =	current state of the economy (z-score)				future state of the economy (z-score)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ ln(employment)	.92***		.57***	.24***	.28***		-.07	.22***
	[.10]		[.09]	[.07]	[.08]		[.08]	[.07]
Δ ln(housing price)		.45***	.40***	.21***		.39***	.40***	.26***
		[.04]	[.04]	[.03]		[.04]	[.04]	[.04]
R-squared	.06	.06	.06	.07	.09	.09	.09	.10
Sample Size	1243142	1243142	1243142	1243135	1222041	1222041	1222041	1222034
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year/Qtr FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State x Time FE	No	No	No	Yes	No	No	No	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The table reports the coefficients associated with regressions of standardized (z-score) individual sentiments on the year-to-year employment growth rate, median housing price per square foot growth rate, individual controls, and fixed effects on county, year, and quarter. Sentiments are measured through a one to four index of perceptions about the current state of the economy and a one to three index of the future state of the economy. Controls include: day of the week fixed effects, 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)		future state of the economy (z-score)	
	liquidity constrained	liquidity unconstrained	liquidity constrained	liquidity unconstrained
Δ ln(employment)	.37***	.38**	.18	.10
	[.14]	[.15]	[.18]	[.15]
Δ ln(housing price)	.16***	.13***	.13	.25***
	[.05]	[.04]	[.08]	[.06]
R-squared	.06	.11	.06	.06
Sample Size	188431	286212	186397	283064
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 sentiments on the year-to-year employment growth rate, median housing price per square foot growth rate, individual controls, and fixed effects on county, year, and quarter. Sentiments are measured through a one to four index of perceptions about the current state of the economy and a one to three index of the future state of the economy. Controls include: day of the week fixed effects, an indicator for whether the individual is employed, a quadratic in age, male, education fixed effects, race (black/white). The log of average monthly income (based on ten income bins) available until 2015. Liquidity constrained individuals are those who report having trouble making payments, which is measured as a binary variable from 2009-2012 and on a one to five scale from 2013 onward; for those latter years, I treat liquidity constrained individuals who report four or five on the five-point scale and unconstrained individuals who report one or two. Standard errors are clustered at the county-level and sample weights are used.

Panel A: Employment Growth Gradient

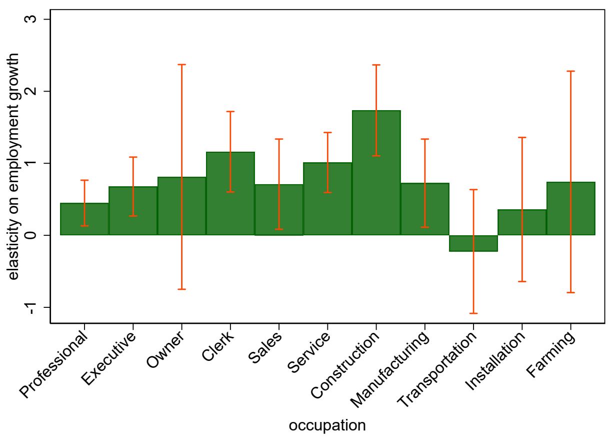

Panel A: Housing Price Growth Gradient

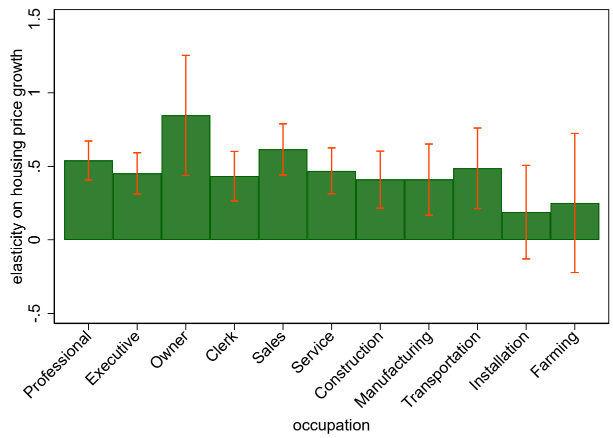

Figure 6: Heterogeneity in Employment and Housing Gradients, by Major Occupation

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, year, and quarter across each occupation. Controls include: day of the week fixed effects, 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.

Table 4: The Micro-elasticity of Sentiments and Real Economic Behavior

Dep. var. =	ln(daily consumption)				hiring intensity, z-score			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
sentiment (person)	.02***	.01***			.13***	.14***		
	[.00]	[.00]			[.00]	[.00]		
sentiment (county)			.08***	.06**			.16***	.17***
			[.02]	[.03]			[.01]	[.01]
ln(monthly income)		.21***		.21***		.03***		.04***
		[.00]		[.00]		[.00]		[.00]
R-squared	.05	.07	.04	.06	.08	.08	.06	.06
Sample Size	494789	414761	373113	311588	658169	534507	491342	398467
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year/Qtr FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Income Control	No	Yes	No	Yes	No	Yes	No	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The table reports the coefficients associated with regressions of logged consumption (spending on non-durables yesterday) and a z-score of hiring intensity (=3 if company is expanding, =2 if staying the same, and =1 if contracting) on a z-score of perceptions about the current the economy (one to four index) at both the individual and leave-one-out county × year × quarter level, controls, county × year × quarter employment and housing price growth, and county and time fixed effects. Controls include: 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.

Table 5: Robustness of the Consumption and Sentiment Elasticity using State Data, 2008-2016

Dep. var. =	ln(total)		ln(durable)		ln(nondurable)		ln(service)	
	level	growth	level	growth	level	growth	level	growth
sentiment index	.08**		.13***		.16**		.08**	
	[.03]		[.04]		[.08]		[.04]	
share worsening		-.02**		.01**		-.04**		-.01**
		[.01]		[.01]		[.02]		[.01]
R-squared	.91	.25	.72	.18	.53	.37	.92	.35
Sample Size	348	298	348	298	348	298	348	298
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.–Sources: Gallup, Bureau of Economic Analysis, American Community Survey, 2008-2016. The table reports the coefficients associated with regressions of logged consumption expenditures per capita (deflated using the 2010 personal consumption expenditure index) (and in the even columns, the growth rate of consumption per capita) on a z-score of perceptions about the current the economy (one to four index) (and in the even columns, the growth rate in the share of individuals reporting the economy is worsening), conditional on controls. Controls include: the share who are male, the share who are white, black, the share who are married, shares on the age distribution (between 0-17, 35-64, and 65+ normalized to 18-34 as the omitted group), shares on the education distribution (no high school, high school, some college, and graduate degrees normalized to the college degree as the omitted group), employment growth. Standard errors are clustered at the state-level and observations are weighted by state population.

Table 6: Instrumental Variables Robustness on the Micro-elasticity of Sentiments and Consumption

Dep. var. =	ln(daily consumption)	
	(1)	(2)
sentiments	.09*	.12**
	[.06]	[.05]
ln(commute time)		.01***
		[.00]
R-squared	-.02	-.06
Sample Size	322635	251603
F-statistic	3.41	4.21
Controls	Yes	Yes
County FE	Yes	Yes
Time FE FE	Yes	Yes

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The table reports the coefficients associated with regressions of logged consumption (spending on non-durables yesterday) on a z-score of perceptions about the current the economy (one to four index), individual controls, county × year × quarter employment and housing price growth, and county and time fixed effects. Sentiments is instrumented using logged average daily temperature and its interactions with occupation fixed effects to capture how temperature shocks might affect workers exposed to outside temperatures differently. 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.

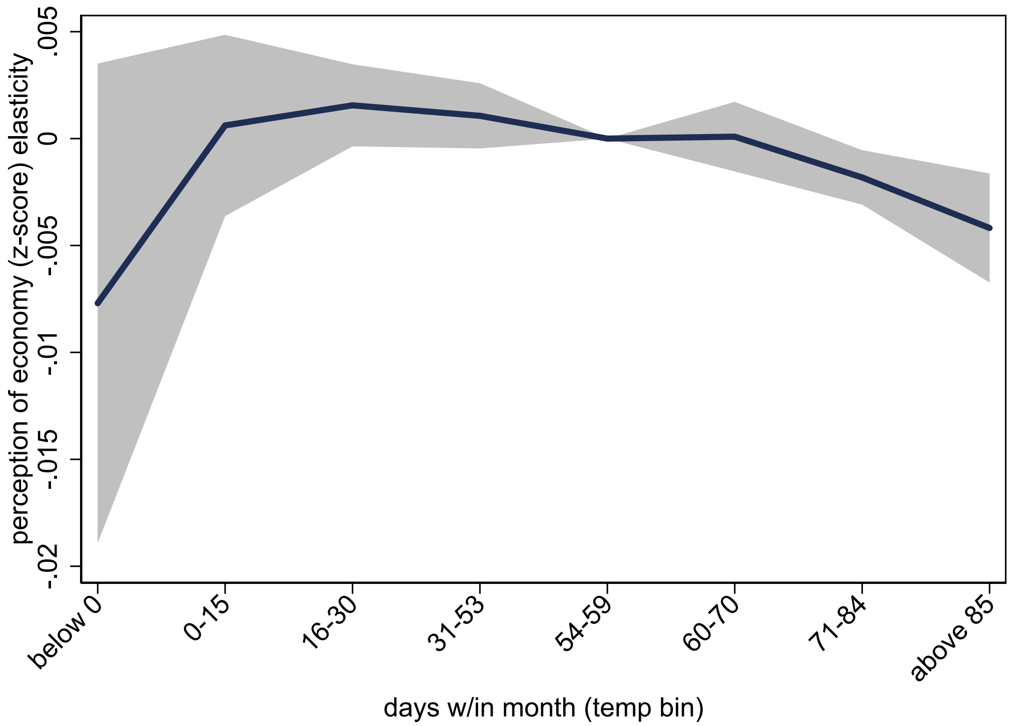

Figure 7: First-stage Relationship between Temperature and Economic Sentiments
Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The table reports the coefficients associated with regressions of the z-score of the sum of perception about the current and future state 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 $\times$ year $\times$ 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.

Online Appendix (Not For Print)

6.1 Supplement to Data and Measurement

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 8 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

Table 7: A Difference-in-difference Application to the 2016 Presidential Election

Dep. var. =	ln(daily consumption expenditures)			
	all	all	narrow	narrow
1[post]	.06***	.04	.05***	.05
	[.01]	[.04]	[.02]	[.04]
1[somewhat conserv.]	-.02	-.01	-.01	-.00
	[.01]	[.01]	[.02]	[.02]
1[very conserv.]	-.02	.00	-.02	.00
	[.02]	[.02]	[.03]	[.03]
1[somewhat conserv.] × 1[post]	.06***	.06***	.05**	.05**
	[.02]	[.02]	[.02]	[.02]
1[very conserv.] × 1[post]	.11***	.12***	.10***	.11***
	[.03]	[.03]	[.04]	[.04]
R-squared	.03	.07	.03	.09
Sample Size	142247	142145	66267	65987
Controls	Yes	Yes	Yes	Yes
County FE	No	Yes	No	Yes
State x Year x Month FE	No	Yes	No	Yes

Notes.–Sources: Gallup, 2016-2017. The table reports the coefficients associated with regressions of logged consumption (spending on non-durables yesterday) 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 being moderate—liberals are omitted from the sample), their interactions, individual controls, and county and time (year and month) fixed effects. Income is not measured in 2016 and, therefore, omitted as a control. 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.

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.

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

Table 8: 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]	[.00]	[.01]	[.01]	[.01]	[.01]
monthly income, 6500-8500	.90***	1.29***	.61***	.89***	.17***	.23***	.14***	.18***	.40***	.46***
	[.01]	[.01]	[.01]	[.01]	[.01]	[.01]	[.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]	[.01]
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		
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.

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.

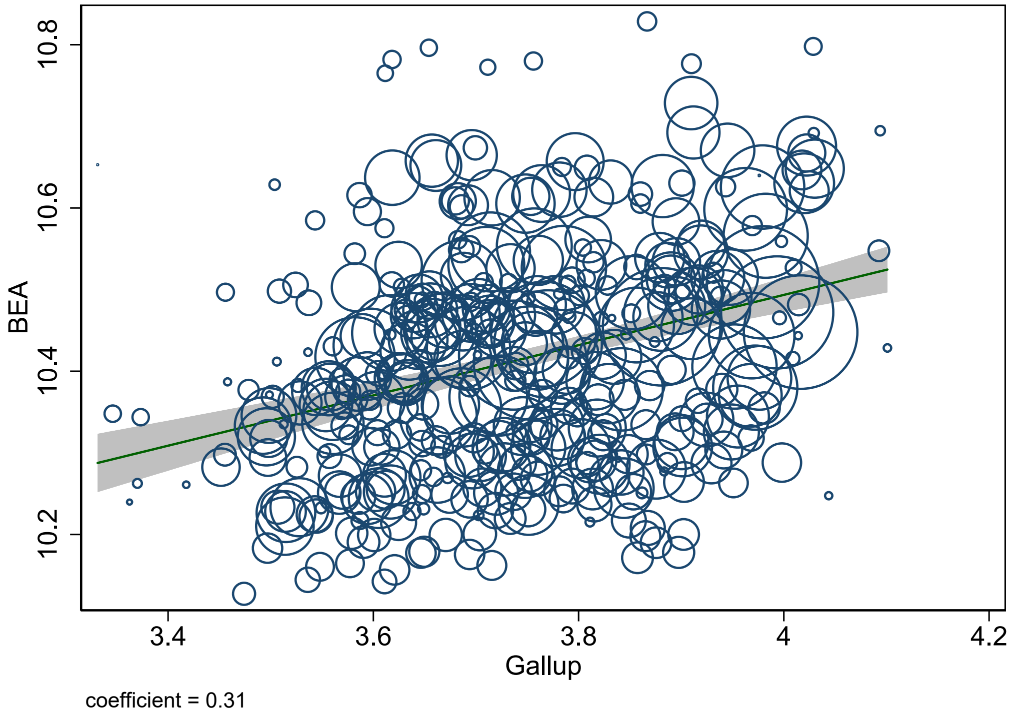

Figure 8: Comparison of Consumption from the BEA and Gallup

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.

6.2 Supplement to Sentiments and Exposure to Economic Shocks

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. Figure 9 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. I find that there is not a monotone pattern for the employment gradient, but housing price growth has a larger positive effect on higher income individuals. This result might emerge from an income effect—housing price appreciation on a larger and more expensive home has a bigger effect. It could also be capturing differences between renters and owners, however.

Figure 10 subsequently considers logit regressions for whether the individual is liquidity constrained on employment and housing price growth, conditional on the individual controls. I find that higher employment and housing price growth have a much weaker and attenuated association with liquidity constraints among higher income individuals, which likely reflects the fact that these individuals are less likely to face liquidity problems. Overall, these results provide microeconomic evidence behind the mechanism in models about the self-fulfilling effects of expectations (e.g., Heathcote and Perri (forthcoming)) and partial insurance against cyclical risk (e.g., Blundell et al. (2008), Low et al. (2010), and Heathcote et al. (2014)).

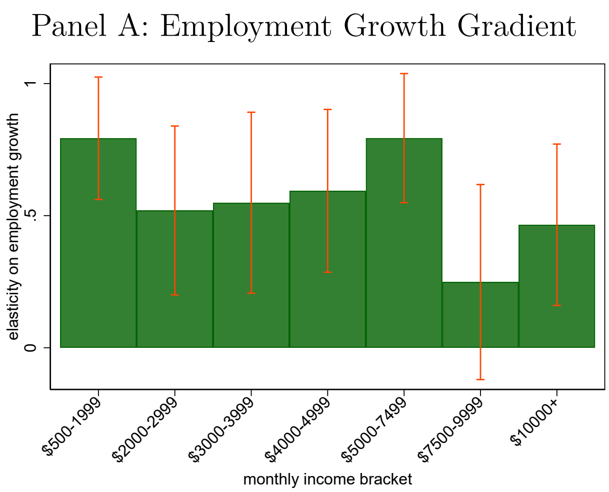

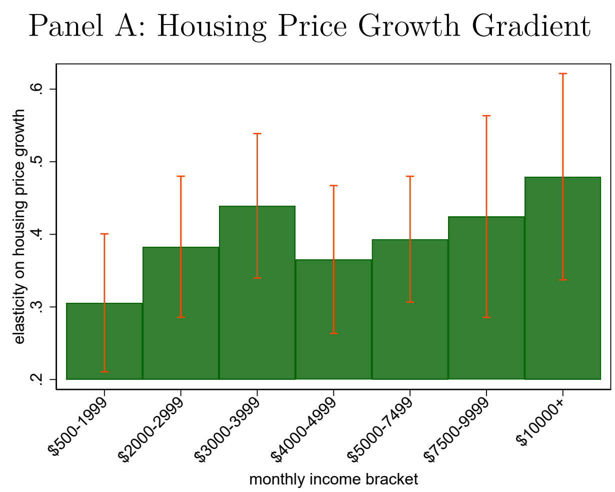

Figure 9: Sentiment Heterogeneity in Employment and Housing Gradients, by Income Bracket

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, year, and quarter across monthly income bins. Controls include: day of the week fixed effects, 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.

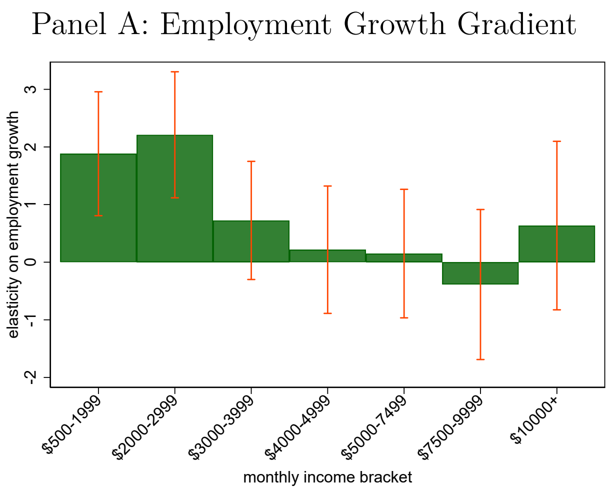

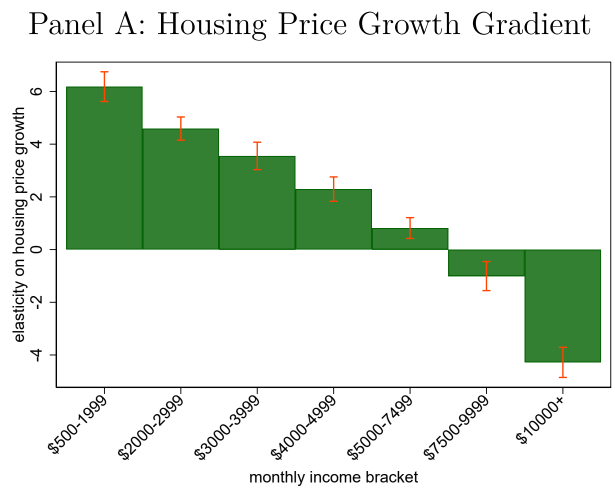

Figure 10: Liquidity Constraint Heterogeneity in Employment and Housing Gradients, by Income Bracket

Notes.–Sources: Gallup, Quarterly Census of Employment and Wages, Zillow, 2008-2016. The figure plots the coefficients associated with logit regressions of an indicator for whether the individual is liquidity constrained on the year-to-year employment growth rate, median housing price per square foot growth rate, and individual controls across monthly income bins. Controls include: day of the week fixed effects, 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.

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

6.3.1 Instrumental Variables Robustness

The motivation for using daily temperature as an instrument for sentiments comes from recent empirical work by Baylis (2015) and Baylis et al. (2017) who document strong correlations using social media data on positive and negative emotions. However, one potential concern is that temperature shocks also affect productivity in a direct way, thereby violating the exclusion restriction. To address this concern, I draw from 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. The sample is 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 9 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 9: 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.

However, average temperature in a month might mask significant heterogeneity in daily temperature over the course of a month. As an alternative test, I now partition temperature into several bins given by the following function

$$f(T_{jt}) = \sum_k \psi_k T_{jt}^k \tag{5}$$

T^k denotes the number of days in location j and period t that fall between the range $T^k \in [\underline{T}^k, \overline{T}^k]$. While the intervals on T^k in Equation 5 can be be arbitrarily flexible, we follow the literature in using 15 degrees Fahrenheit intervals, which is in the neighborhood of the norm in the literature (Deschenes and Greenstone, 2011; Deryugina and Hsiang, 2017). Figure 11 plots the corresponding gradients across the distribution of temperature with logged weekly earnings as the outcome variable. The relationship between the two is effectively flat and statistically indistinguishable from zero.

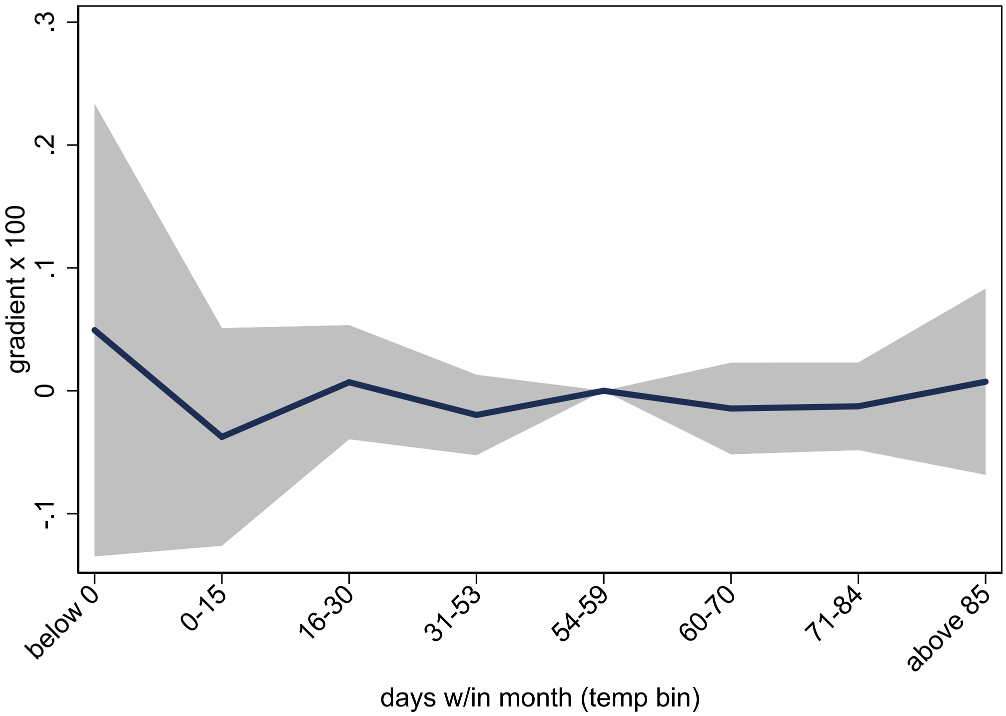

Figure 11: Semiparametric Response of Weekly Earnings to Temperature, 1994-2014

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.