

# Research Statement

Joseph Vavra

July 2019

My research focuses on empirical questions in macroeconomics using various micro data sources, statistical methods and quantitative structural models to try to understand the implications of microeconomic heterogeneity for aggregate fluctuations and stabilization policy. Throughout my work, this mix of methodological approaches allows for identification which cannot be achieved using macro data alone while still allowing for counterfactual analysis which is crucial for evaluating alternative policies.

A central conclusion of my research is that predicting the aggregate response of the economy to shocks or changes in policy requires paying attention to heterogeneity in the economy and how it evolves across time. I have concentrated on these implications primarily in the context of pricing and housing markets. For example:

(1) Important implications for monetary transmission and the sources of business cycles arise from cyclical variation in price-setting behavior; (2) The sensitivity of the economy to aggregate house price driven booms and busts depends on time-variation in the distribution of household mortgage debt; (3) Macroeconomic effects arise from changes in the *local* distribution of house price growth and default risk and their interactions with mortgage market institutions; (4) Aggregate durable spending responds sluggishly to fiscal stimulus during recessions, since household durable spending is very lumpy and the distribution of desired durables varies over the cycle; and (5) Markup movements over the business cycle and over longer periods of time are affected by changes in household shopping behavior.

## I. Prices and Markups

My paper “**Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation**” (*QJE 2014*) argues that the real effects of monetary policy are substantially dampened during recessions due to increases in the volatility of price changes. The paper begins by documenting two new facts using restricted access BLS CPI micro data: 1) The cross-sectional standard deviation of price changes is strongly countercyclical: price changes become substantially more disperse during recessions. 2) The dispersion of price changes (conditional on price adjustment) commoves strongly with the frequency of price adjustment: when more prices are changing, there is more heterogeneity in the size of those changes.

The second half of the paper argues that these facts imply strong restrictions on price-setting models and that models consistent with these facts imply that aggregate prices are more flexible during recessions. Standard menu cost models include only *first moment* shocks (like changes in aggregate money) which shift all firms’ desired prices by the same amount. First moment shocks

can raise the frequency of adjustment by pushing firms out of an “inaction” region, but they must push all firms out of this inaction region in the same direction, which lowers price dispersion. This induces a robust counterfactual negative correlation between frequency and dispersion.

While first moment shocks cannot explain the data, I next show that introducing *second moment* shocks which increase firm-level volatility into a menu cost model allows it to match the data. These second moment shocks change the distribution of firms’ desired price changes: when volatility rises, more firms choose to adjust and by larger amounts so that there is a positive relationship between frequency and dispersion. I then show that the increased price churning that occurs with greater volatility leads to greater aggregate price flexibility and reduced real effects of monetary policy. The model implies that greater volatility in the 2001 and 2008 recessions reduced the real effects of nominal stimulus from their long-run average by more than half.<sup>1</sup>

There is now broad consensus that the dispersion of many economic variables is countercyclical. (See Bloom (2014) for a detailed review.) There is less consensus on *why* dispersion is countercyclical. In general, greater dispersion during recessions could arise from greater volatility of shocks or from agents responding more to shocks of constant size. Without data separately measuring exogenous shocks and endogenous responses, these alternative cannot be distinguished. In **Shocks vs. Responsiveness: What Drives Time-Varying Dispersion** (with Berger; *JPE forthcoming*), we show how a large and observable cost shock, the nominal exchange rate, can be used to differentiate these explanations in one environment.

In the first half of the paper, we use confidential BLS import price data to document that item-level price change dispersion is both countercyclical and is highly correlated with exchange rate passthrough. The first fact shows that the patterns found for retail prices in my QJE paper extend to import prices. The fact that price change dispersion is correlated with exchange rate passthrough in turn allows us to infer *why* dispersion changes. In particular, we estimate a workhorse open-economy price-setting model using indirect inference to show that these facts imply that greater responsiveness rather than larger shocks drives the rise in dispersion during recessions. Our results suggest that the theoretical literature trying to explain countercyclical dispersion has perhaps been too quick to embrace volatility/uncertainty shocks but we propose a promising alternative.<sup>2</sup>

---

<sup>1</sup> In **Dynamics of the U.S. Price Change Distribution** (with Berger; *EER 2018*) we show that while counterfactuals require a full structural model, just measuring aggregate price flexibility at a point in time can be done with more minimal assumptions. We apply this procedure to BLS CPI, PPI and IPP micro data and show that price flexibility is strongly countercyclical for each.

<sup>2</sup> Importantly, increases in volatility and responsiveness both have the same implications for aggregate price flexibility, so this does not change implications in my QJE paper, but it is important for understanding the source of business cycle fluctuations and what policies can mitigate them.

Having shown that responsiveness and thus firms' desired markups vary cyclically, the next question is why? In **House Prices, Local Demand and Retail Prices** (with Stroebe, *JPE* 2019) we provide new empirical evidence on the response of prices and markups to demand shocks. In particular, we link micro retail scanner price data and household purchase data to zip-code-level house prices to identify the response of price-setting and shopping behavior to house-price-induced local demand shocks. We first argue for a large causal response of retail prices to house prices. We do so using two complementary identification strategies: the first, more standard approach, uses housing supply elasticity as an instrument for house price movements while the second, more novel approach, uses variation in local homeownership rates. We next argue using numerous pieces of evidence that this causal retail price response is largely driven by markups rather than marginal costs. In the final part of the paper we provide empirical evidence that homeowners become less price-sensitive when house prices rise, which is why raising markups is profitable.

These empirical results have important implications for interpreting existing aggregate evidence on markup cyclicality. In simple New Keynesian models, all movements in markups are driven by sticky prices interacting with shocks to nominal marginal cost. Our evidence shows that markups can also move for a second reason as countercyclical household shopping intensity pushes firms to choose relatively higher markups in booms. It is important to account for this channel when interpreting aggregate time-series evidence on the cyclicality of markups. If firms' desired markups are procyclical while the sticky-price induced component of markups is countercyclical, then the total markup measured in the data will depend on the strength of these two forces. Furthermore, the strength of these two forces will depend on the overall flexibility of prices, which I argue in my earlier research varies across time. This can help explain conflicting empirical evidence on markup cyclicality in the literature, and it provides another example of the conclusion that accurate interpretation of macroeconomic fluctuations requires a careful accounting of time-varying microeconomic patterns.<sup>3</sup>

Recent decades have given rise to substantial changes in the technology of how goods are designed, produced, marketed and delivered. The implications of these forces for macroeconomic aggregates and welfare is typically explored using producer-, retailer-, or broader industry-level data. In my newest working paper, **The Rise of Niche Consumption** (with Neiman), we provide a *household* level perspective on the rapidly changing retail environment. We show that over the last 15 years, the typical household has increasingly concentrated its spending on a few preferred products. However, this is not driven by "superstar" products capturing growing market shares. Instead, households are increasingly segmenting in which products they purchase and aggregate product concentration is declining. In the second half of the paper, we develop a new heterogeneous house model to understand and interpret this rise of "niche consumption". We use the model to conclude that increasing product variety is a key force driving the rise of niche

---

<sup>3</sup> For example, Gali, Gertler and Lopez-Salido (2007) find that markups fall in response to expansionary monetary policy. However, using an identical methodology, Nekarda and Ramey (2013) show that this result changes when using more recent data.

consumption and that this raises welfare as households can better select products suited to their particular tastes. In addition, the model delivers a new tractable form of variable markups because producers of popular products care more about maximizing profits from existing customers while producers of less popular niche products care more about expanding their set of customers. Surprisingly, however, our model can match the observed trends in household and aggregate concentration without any resulting change in aggregate market power.

## II. Macro Implications of Housing Markets and Durables

Mian and Sufi (2013), my paper with Stroebel and many others show large consumption responses to house price changes in the data. However, this is at odds with a prominent theoretical view which, using the logic of the permanent income hypothesis (PIH), argues that consumption responses should be small. In **House Prices and Consumer Spending** (with Berger, Guerrieri, and Lorenzoni; *Restud 2018*) we show that a more realistic workhorse incomplete markets model calibrated to rich cross-sectional micro facts actually implies large consumption responses, in line with the data. To explain this result, we show that consumption responses to house price shocks in this model can be approximated by a simple rule-of-thumb: the marginal propensity to consume out of temporary income times the value of housing. This rule-of-thumb formula allows us to map the theoretical analysis of house price effects to the much better understood analysis of MPCs, and it provides simple intuition for why our model implies large house price effects and PIH models do not. In our model, consumption responses depend on a number of factors such as the level and distribution of debt, the size and history of house price shocks, house price expectations, and the level of credit supply. Each of these effects is naturally explained with our simple formula.

These results also lead to yet another situation in which time-varying microeconomic distributions imply time-variation in macroeconomic reactions to shocks: in our model, the response of aggregate consumption to house prices varies across time with the state of the economy and the resulting endogenous shifts in the distribution of household MPCs and housing values. The fact that the endogenous distribution of housing and debt matters for the strength of house price effects is in turn important for interpreting empirical evidence on consumption responses to house prices from different time periods and when contemplating potential policy interventions into housing markets. Both shocks and policy interventions may have effects on consumption that differ dramatically with the distribution of household state-variables at the time the policy is enacted.

In **Consumption Dynamics During Recessions** (with Berger; *Econometrica 2015*), we argue that the lumpy and infrequent nature of household level durable spending leads to a decline in the number of households buying/selling durables during recessions. This in turn leads aggregate durable spending to respond sluggishly to fiscal stimulus during recessions and implies that standard VAR approaches overstate the effectiveness of durable stimulus during recessions.

We start by using indirect inference to formally estimate a heterogeneous agent model with fixed costs of durable adjustment to match PSID micro data. To do, we develop a novel general methodological procedure for empirically inferring desired durable holdings, which cannot be observed directly in the data. Overall this model can explain more than 75% of actual household durable adjustment, and it implies a key role for changes in the cross-sectional distribution of desired durable holdings in recessions and booms. We then aggregate this model in general equilibrium and show that it generates a very procyclical response of aggregate durable spending to aggregate shocks and show this is driven by the cyclical variation in the household distribution of desired durable holdings, in line with the data. Finally, we provide additional reduced form evidence for this channel using cross-region auto spending responses to regional business cycles.<sup>4</sup>

A rapidly growing literature uses regional variation to inform macro models. While much of the empirical literature studying regional heterogeneity uses these shocks as a source of identification to inform aggregate phenomena, I have several recent papers arguing that regional shocks have important aggregate implications in their own right. Specifically, my research shows that accounting for the interaction of locally segmented housing markets with institutional features of US mortgage markets leads to important aggregate consequences of even purely regional shocks.

**In Regional Heterogeneity and the Refinancing Channel of Monetary Policy** (with Beraja, Fuster and Hurst; *QJE* 2019) we argue that the time-varying regional distribution of housing equity influences monetary policy through its effect on refinancing. Using detailed new loan-level data, we show that interest rate declines immediately after QE mostly stimulated refinancing activity (both cash-out and rate refi) in the regions with the smallest declines in house prices (which also had the smallest increases in unemployment). In contrast, we show that in the 2001 recession, when regional house price growth and resulting housing equity was mostly uncorrelated with unemployment, refinancing is strongest in high unemployment regions. These housing equity effects arise naturally from institutional features of the refinancing process: lenders generally require a minimum level of equity in order to allow borrowers to refinance, even if they are not extracting equity. In addition, the level of equity potentially extracted during refinancing (cashout) clearly depends on the existing level of equity in the house prior to refinancing.

We also explore the aggregate implications of this cross-region evidence. Our main empirical results measure only relative differences across regions, so any common aggregate effects are necessarily differenced out. In order to explore aggregate implications of our regional evidence for monetary policy, we thus need to develop a heterogeneous household macro model of refinancing. Our model implies that interest rate cuts in 2008 indeed had the smallest effects on depressed regions. More importantly, the regional distribution of housing equity in 2008 substantially dampened the aggregate effects of monetary policy. Since the distribution of equity varies across time and changes both the aggregate effects of monetary policy as well as the trade-

---

<sup>4</sup> In **Measuring How Fiscal Shocks Affect Durable Spending in Recessions and Expansions** (with Berger; *AER P&P* 2014) we also use STVARs to provide aggregate time-series evidence that durable spending responds more to identified government spending shocks in expansions than in recessions.

off between aggregate and regional inequality, we conclude that it is important for policy makers to track this variation.

**In Mortgage Prepayment and Path-Dependent Effects of Monetary Policy** (with Berger, Milbradt, and Tourre; *R&R at AER 2019*) we explore another aspect of the interaction between refinancing of fixed rate mortgages and monetary transmission, which focuses on rate incentives rather than effects of equity. Specifically, we argue that fixed rate pre-payable mortgages means that how much room the Fed has to stimulate the economy by cutting rates depends not just on how far current rates are from zero but also on their previous path. That is “policy space” cannot be determined by looking only at current rates. This is because if rates were low in the recent past, many households will have already locked in rates below the current rate and will not respond to cuts today. Monetary policy “ammunition” is used up by past rate cuts.

Our paper begins by using loan-level data to refine the well-established fact that interest rate incentives are an important factor for predicting refinancing. While it is well known that loans with interest rates higher than the current market rate (i.e. positive “gaps”) are more likely to refinance, the first contribution of the paper is to show that after controlling for a variety of non-rate related factors (such as equity or credit score), the probability a loan refinances is approximately determined just by whether its gap is positive or negative and *not* by how big the gap is. This surprising fact means that predicting the influence of the gap distribution on refinancing rates does not require information on the entire gap distribution and instead only requires a simple summary statistic: the fraction of loans with positive gaps ( $\text{frac} > 0$ ).

In the second half of the paper, we argue that this can naturally be explained using a model with “inattentive” refinancing. We embed this model of refinancing into an otherwise standard incomplete markets setup and show that it implies important path-dependent effects of interest rates and that recent interest rate paths will generate substantial headwinds for future monetary stimulus. All of these effects are naturally explained by the dynamics of the rate gap distribution and resulting ( $\text{frac} > 0$ ).

**In Regional Redistribution Through the U.S. Mortgage Market**, (with Hurst, Keys and Seru; *AER 2016*) we use data on mortgage loans, which represent the bulk of household borrowing, to document two new facts. First, risk-adjusted rates on most loans are not equalized across locations with the U.S.: there is large regional variation in predictable default risk, but there is no regional variation in mortgage contract rates for loans securitized by GSEs. Since GSEs securitize most loans, this constant *contract* rate in the face of spatial variation in default risk implies that most borrowers face *risk-adjusted* rates which vary with their locations. Second, this lack of risk-based pricing does not occur because this risk cannot be observed: we show that otherwise similar non-GSE loans securitized in the private market do increase rates when local default risk rises.

Since most mortgage rates do not respond to regional shocks that increase default risk, individuals in those regions face lower borrowing costs than they would if this risk was priced into rates, which in turn offsets some of the negative regional shock that increased default risk in the first

place. Thus, the constant interest rate “policy” followed by the GSEs results in state-contingent regional transfers. In the second half of the paper we use a structural model to quantify the size and welfare implications of these implicit transfers. Overall the results of this model imply that the magnitude of redistribution induced by the GSEs through the mortgage market is economically meaningful when comparing to other policies like unemployment insurance or tax rebates, which have received vastly more attention. Thus, looking at the local distribution of shocks and risk and how it interacts with mortgage markets has important implications for aggregate redistribution policies.

## References

### References to My Own Work:

#### Published Papers

Berger and Vavra (forthcoming) “Shocks vs. Responsiveness: What Drives Time-Varying Dispersion”, *Journal of Political Economy* (scheduled Vol. 127, Issue 4, 2019).

Stroebel and Vavra (2019) “House Prices, Local Demand and Retail Prices”, *Journal of Political Economy*, Vol. 127, Issue 3.

Beraja, Fuster, Hurst and Vavra (2019) “Regional Heterogeneity and the Refinancing Channel of Monetary Policy”, *Quarterly Journal of Economics*, Vol. 134, Issue 1.

Berger, Guerrieri, Lorenzoni and Vavra (2018) “House Prices and Consumer Spending”, *Review of Economic Studies*, Vol. 85, Issue 3.

Berger and Vavra (2018) “Dynamics of the U.S. Price Distribution”, *European Economic Review*, Vol 103, Issue 4.

Hurst, Keys, Seru and Vavra (2016) “Regional Redistribution through the U.S. Mortgage Market”, *American Economic Review*, Vol 106. No 10.

Berger and Vavra (2015) “Consumption Dynamics During Recessions,” *Econometrica*, Vol. 83, No. 1.

Berger and Vavra (2014) “Measuring How Fiscal Shocks Affect Durable Spending in Recessions and Expansions,”, *AER Papers and Proceedings*, Vol. 104. No 5.

Vavra (2014) “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *Quarterly Journal of Economics*, Vol. 129. No. 1.

## **Working Papers**

Neiman and Vavra (2019) “The Rise of Niche Consumption”, *Mimeo*

Berger, Milbradt and Tourre (2018) “Mortgage Prepayment and Path-Dependent Effects of Monetary Policy”. *NBER Working Paper 25157*. Revision requested at *American Economic Review*.

“Time-Varying Phillips Curves”, 2013. *NBER Working Paper 19790*.

## **References to Others’ Work:**

Bloom (2014): “Fluctuations in Uncertainty,” *Journal of Economic Perspectives*.

Gali, Gertler and Lopez-Salido (2007): “Markups, Gaps, and the Welfare Costs of Business Fluctuations”, *Review of Economics and Statistics*

Mian, Rao and Sufi (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*.

Nekarda and Ramey (2013) “The Cyclical Behavior of the Price-Cost Markup”, *NBER Working Paper 19099*.