

The Channels of Financial Distress
During the Great Recession:
Some Evidence on the Aggregate Effects

Mark Gertler and Simon Gilchrist
New York University and NBER

October 2019

Abstract

This paper presents evidence on the aggregate employment effects of the household balance sheet channel versus banking distress during the Great Recession. We make use of a quarterly panel of state level data, which allows us to exploit both cross-sectional and time series methods to identify both shocks and causal relationships in the data. In doing so we are able to identify the contribution of house price versus financial shocks to the aggregate as well as regional employment variation. We find that both shocks were important during Great Recession. House prices shocks help account for both the regional and aggregate variation of employment. The financial shock helps account for the unusual depth and persistent of the Great Recession. On the methodological side, with our panel data approach, we are able to separately identify the partial equilibrium (within state) versus general equilibrium (aggregate) effects of shocks.

1 Introduction

While the global financial crisis of 2007-09 was largely unanticipated, we have now had a decade to analyze what happened. Over this time, the profession has reached a broad understanding of the forces that were at work. At the risk of some oversimplification, the literature has emphasized two main channels via which financial distress was transmitted to the real economy.¹ The first involved the impact of the house price bust on household balance sheets and spending. The second involved the transmission of banking distress to the real economy via the disruption of credit supply. To be sure, nonfinancial firms also felt the brunt of financial distress. However, the buildup in financial vulnerability that instigated the crisis arose mainly in the household and banking sectors.

Identification of these financial distress channels is a challenge due to simultaneity: While financial conditions may affect real activity, real activity in turn can influence financial conditions. Accordingly, beginning with Mian and Sufi (2013), a vast literature has creatively exploited cross-sectional variation to identify causality. As a result there is a wealth of evidence based on variation across regions that declining house prices weaken consumer spending and in turn employment via the impact on household balance sheets. Similarly, there is a host of evidence based on variation across banks' exposure to financial distress that weakening of banks' financial health reduced the economic activity of their respective borrowers. Overall, this literature has provided evidence that both mechanisms were operative during the Great Recession. However, because the empirical analyses typically involve cross-sectional analysis, this literature is largely silent on the aggregate impact.

The goal of this paper is to present evidence on the aggregate employment effects of each channel. To do so we make use of a quarterly panel of state level data. The advantage of the panel is that we can exploit both cross-sectional and time series methods to identify causal relationships in the data. As in the literature we can use the cross region variation to help identify the household balance sheet channel. In turn we can use time series methods to identify orthogonal shocks to house prices (both aggregate and local) and to aggregate financial conditions. We can then combine our econometric framework with the identified shocks to perform a historical decomposition of the data to assess the relative importance of housing shocks and the house-

¹See Gertler and Gilchrist (2018) and Kashyap et. al. (2019) for overviews.

hold balance sheet channel versus shocks to financial conditions. In doing so we can analyze the sources of variation in both aggregate time series and the cross-section. In addition, given our panel data approach, we are able to separately identify the partial equilibrium (within state) versus general equilibrium (aggregate) effects of shocks.

Overall, we find that both channels were important but that banking distress was key in turning the recession into the Great Recession. Absent variation in banking distress, the recession would only been half as deep at the trough and the return to trend much faster. As we show, banking distress helps account for the enhanced and protracted decline in employment. Consistent with the literature, shocks to housing prices help account for the cross state variation during the Great Recession. We find further that roughly half the impact of house price shocks works through the household balance sheet channel, while the other half likely reflects factors such as the impact on construction. We also find that aggregate house price shocks have stronger effects on employment than local shocks, which is consistent with the theory: Local shocks affect only nontraded employment while aggregate shocks affect traded as well as non-traded.

Our paper is related to several strands of the literature that provide an empirical analysis of the transmission of financial distress during the Great Recession. Directly relevant are the cross-sectional studies of household balance sheet channel (e.g., Mian and Sufi, 2014, Kaplan et. al.2017) and the banking distress channel (e.g.. Chodorow-Reich, 2014, Huber, 2018). Also, relevant are time series studies that provide a historical decomposition of the main forces at work in the aggregate data. Christiano, Eichenbaum and Trabandt (2015) use a structural model to identify the key shocks. Gertler and Gilchrist (2018) and Bernanke (2018) use time series methods and both argue that financial shocks played an important role. In this paper we differ by taking a unified approach toward explaining the cross-sectional and regional data.

On the methodological side, our paper is also related to a recent literature that examines how to draw inferences from cross-sectional data about implications for aggregate data. A fairly common approach is to employ a macroeconomic model disciplined by the cross-sectional to infer general equilibrium effects (e.g. Steinsson and Nakamura, 2014). Of course the end product of the analysis should always be a model: Otherwise any kind of counterfactual analysis is not possible. However, it is also desirable to have some evidence that does not rely on a tightly structured model. To this

extent our paper relies on less structured time series methods to draw inferences about aggregate effects, as well as disentangle parameters that reflect local versus general equilibrium effects. In this regard our paper is related to Sarto (2018), who employs a factor analysis to distinguish between aggregate versus local fiscal shocks. We address a different issue and we also use a different identification strategy.

2 Descriptive and Cross-sectional Evidence

In order to provide context for the empirical work that follows, we first provide a brief synopsis of the descriptive evidence in the aggregate data for the household balance sheet and banking distress channels. We also review some of the cross-sectional evidence that provides formal support for each of these mechanisms.

Partly because of how the data was unfolding in real time during the crisis, much of the early focus was on the household balance sheet channel. The origins of crisis involved an extraordinary housing boom, featuring a dramatic runup in house prices and mortgage debt. Among the factors triggering the boom were: a secular decline in mortgage rates due to a combination of declining long term interest rates and innovation in mortgage finance; relaxation of lending standards; and widespread optimism about housing prices.

The descriptive evidence suggests that a household balance sheet channel was a key conduit of financial distress during the Great Recession. The top panel in Figure 1 shows the behavior of household debt to income (the solid line) versus debt to assets (the dotted line) over the period 2004 to 2012. The bottom panel shows house prices (the solid line) and consumer durable consumption (the dotted line). In each panel the gray area is the recession and the solid vertical line is the date of the Lehman bankruptcy. Preceding the crisis there is a roughly twenty percent run-up of household debt to income (due mainly to an increase in mortgage debt). Debt to assets remains stable until 2007, reflecting that house prices increase along with debt. However, as house prices decline starting in early 2007, household debt to assets sharply increase. The weakening of household balance sheets, in turn, leads to a sharp drop in consumer durable spending.

With aggregate data, however, it is difficult to identify causality. Aggregate housing prices could be responding to the decline in real activity,

as opposed to influencing it. In a series of highly influential papers, Atif Mian and Amir Sufi use cross-sectional data to identify the household balance sheet channel. They first show that regions which experienced the largest run up in house prices and mortgage debt in the years prior to the crisis suffered the largest drops in house prices and real activity once the crisis hit. For the crisis period, they estimate cross-sectional regressions that relate some measure of real activity, for example consumption or non-tradable employment, to the decline in household net worth. The latter is measured by the rate of decline in house prices weighted by household leverage at the beginning of the crisis. They identify exogenous variation in household net worth by using an instrumental variable based on the local land supply elasticity.

In particular, let $e_i^{NT} \equiv$ nontradable employment in region i ; $p_i^H \equiv$ housing prices; and $p_i^H H_i / N_i^H \equiv$ household leverage in 2006 (the year prior to the crisis), specifically the ratio housing values to housing equity. All variable are in logs and Δ is the first difference operator. Then Mian and Sufi (2014) estimate the following cross-sectional regression over 2007-09

$$\Delta e_i^{NT} = \alpha + \eta \left(\frac{p_i^H H_i}{N_i^H} \cdot \Delta p_i^H \right) + \epsilon_i \quad (1)$$

The independent variable, the percent change in housing prices interacted with initial household leverage, is the regional change in housing net worth. Because this change is regional and not aggregate, it affects only the demand for nontraded goods within the region (since the demand for traded goods with depend on the aggregate change in housing net worth). Hence, the dependent variable is nontraded employment, given that it is this employment measure that responds to the demand for nontradables. Finally, as noted earlier, the authors use the a measure of land supply elasticity (based on Saiz) as an instrumental variable.

The authors main finding is that parameter η , which measures the sensitivity of nontraded employment growth to the percent change in housing net worth, is both statistically and economically important. The results accordingly provide persuasive evidence that an operative household balance sheet channel can help account for the regional variation in nontraded employment during the crisis. But there are two main caveats to this finding. Because the regression is purely cross-sectional, it is not possible to identify the aggregate effects of the household balance sheet channel. Second, at the same time household balance sheets were weakening, the financial sector was succumbing to stress. The declining in housing prices following the peak in

late 2006/early 2007 led to a rise in defaults by the summer of 2007. In turned, the defaults induced losses on mortgage related securities, weakening the balance sheets of the financial institutions that held them. The stress in financial markets built steadily over the recession, culminating in the Lehman Brothers collapse. Hence, the aggregate decline in consumer durable spending portrayed in figure 1 for example, could reflect the influence of banking distress as well as a household balance sheet channel. In assigning aggregate effects, it is important to disentangle these forces.

Indeed the descriptive evidence is also consistent with the existence of a banking distress channel in the aggregate. The mirror image of the sharp increase in household indebtedness portrayed in Figure 1 was a sharp increase in the leverage of the banking system, particularly the shadow banking system that operated outside the direct regulatory control of the Federal Reserve. The left panel in Figure 2 illustrates the behavior of the liabilities of broker-dealers, i.e. investment banks which were the main actors in the shadow banking sector. The sharp growth in liabilities financed the sharp increase in mortgage related securities, a product of the housing boom described earlier. Importantly, while the assets held by these institutions were mainly long term, the liabilities were mostly short term. With the benefit of hindsight, this maturity mismatch made them vulnerable to panics. The downturn in house prices portrayed in figure 1 not only weakened household balances sheets: It also induced losses in mortgage related securities held by both shadow and commercial banks. The highly levered and lightly regulated shadow banking sector was particularly vulnerable. The losses on mortgage related securities led to panic in markets for wholesale short term funding, culminating with the failure of Lehman Brothers and the subsequent collapse investment banking. The collapse in broker dealer liabilities portrayed in the figure was the product of these events.

The banking distress channel emphasizes that the weakening of bank balance sheets over the crisis induced a contraction in intermediation, raising the cost of credit and thus weakening real activity. As with the household balance sheet channel, the aggregate data provides some suggestive support. The right panel in figure 2 plots GDP growth against a measure of financial distress, the financial excess bond premium (EBP) developed by Simon Gilchrist and Egon Zakresjek. The financial EBP measures the spread between the rate of return on corporate bonds for financial companies and the rate on similar maturity government bonds, but with the default premium removed. The latter adjustment implies that increases in the EBP reflects

elevation in the cost of credit for financial companies as opposed to simply a signal of increasing default. As the figure shows, the beginning of the recession features relatively modest declines in output growth and increases in the EBP. By the summer of 2008 the recession appeared like it would be similar the relatively mild downturns of 1990-91 and 2001-02. However, as figure 2 makes clear, closely correlated with the Lehman collapse is a sharp increase in the EBP along with a sharp contraction in GDP growth. This broad connection between disruption of banking, rising credit costs for financial institutions and declining real activity is highly suggestive of a banking distress channel.

In our empirical work that follows, we use the financial excess bond premium as an indicator of financial conditions. Accordingly, it is useful to characterize broadly how it is determined and how it influences the cost of credit for non-financial borrowers. The dotted line in the top panel of figure 3 is the financial EBP and the solid line is the market leverage (assets to equity) of the fifteen primary dealers (major commercial and investment banks). The sample is 2006 to 2012, the darkened area is the recession and the solid vertical line is the Lehman collapse. As bank balance sheets weaken (i.e. market leverage increases), the financial EBP increases. In this important respect, the data is consistent with the theory: The weakening of bank balance sheets makes banks less credit worthy, which increases the cost of raising funds, captured by the increase in the financial EBP. Note that both market leverage and the financial EBP peak around the time of the Lehman collapse, which also corresponds to the sharp contraction in output portrayed in Figure 1.

We next present descriptive evidence that the increase in the cost of funds faced by banks affected the terms of credit face by non-financial borrowers, The bottom panel in Figure 3 plots the financial EBP (the dotted line) against the change in banking lending standards from the senior loan officer survey. As the panel shows, the tightening of lending standards over the Great Recession is closely connected to the increase in the financial EBP, suggesting that increases in the costs of credit for banks are indeed passed on to nonfinancial borrowers. Other measures of credit costs for non-financial borrowers show a similar connection to the movement of the financial EBP. For example, the financial EBP also co-moves strongly with the spreads on mortgages, consumer loans, commercial paper and corporate debt. In the latter case, the non-financial EBP followed closely the increase in the financial EBP over the crisis. To the extent the increase in the financial EBP induced

a general increase in credit costs, this channel can help account for the sharp and prolonged output contraction over this period.

Just as with the household balance sheet channel, there is evidence based on cross-sectional data that is consistent with the banking distress channel. The identification challenge is the simultaneous relation between a bank's financial health and that of its nonfinancial borrowers. A weakening of a bank's balance sheet may raise the cost of credit for its nonfinancial borrowers for the reasons we just described. However, defaults by its' nonfinancial borrowers will weaken the bank's financial health, making causality difficult to disentangle. The approach in the literature is to isolate variation in bank financial health that is orthogonal to the economic prospects of the borrowers connected to the bank. Two prominent examples are Chodorow-Reich (2015) and Huber (2018). In the former case, a bank's exposure to Lehman Brothers prior to the collapse provides the source of exogenous variation in bank's financial health. In the latter it is the exposure of German banks to losses from U.S. mortgage related securities. In each case, the cross-sectional variation in banks financial health is orthogonal to the banks' borrowers financial positions. In turn, each finds that as banks' financial health weakens there is a significant contraction in both borrowing and employment by their respective non-financial borrowers, as the banking distress channel would suggest. However, as with the housing related channel, because the evidence comes from the cross-section, there is no direct measure of the aggregate importance.

3 Econometric Framework

Our econometric framework is a panel data version of Mian and Sufi's (2014) cross-sectional analysis described earlier. We differ by exploiting both cross-sectional and time series methods. We also allow for shocks to financial conditions as well as housing conditions.

3.1 Data and Descriptive Evidence

The data is quarterly over the sample 1990:Q1 to 2012:Q3. There are three main state level series: employment, house prices, and the mortgage to income ratio (our measure of regional household leverage). Employment is from the Bureau of Labor Statistics Employment and Unemployment Re-

ports. House prices are a Purchase Only Index from the Federal Housing Authority. Mortgage data are from the New York Federal Reserve Bank/Equifax Consumer Credit Panel (note, this excludes HELOCs). Personal income data are from the Bureau of Economic Analysis Regional Accounts. In addition, to the state level data, we use the Gilchrist/Zakrasjek financial excess bond premium to capture financial conditions in banking.

To foreshadow the main results that will come from the formal econometric analysis, it is useful to begin with a descriptive analysis of the raw data. We divide the states into three regions based on the percentage decline in house prices from the peak within the state to the trough. The first category includes the four “sand” states (Arizona, California, Florida and Nevada), which each experienced a house price drop of more than forty percent. These states account for roughly twenty percent of overall employment. The “middle” category includes states which experienced a drop between twenty and forty percent. These states account for roughly thirty percent of total employment. As we will see, they behave similarly to the national average. Finally, our “low” category includes states experiencing a drop of less than twenty percent. These states account for roughly half of total employment. Within each region we construct population weighted averages of house prices, the mortgage to income ratio and employment.

The top left panel in Figure 4 shows the behavior of house prices in each of the three regions over the period 2004 to 2014. The dotted line shows house prices for the sand states, the dashed line is for the middle group, the line with crosses is for the low group, and the solid line gives the national average. As one would expect, the sand states display the largest boom and bust in house prices. House price dynamics for the middle group are close to the national average. The average drop in house prices for the low group is quite mild.

The right panel illustrates the behavior of the mortgage to income ratio. Consistent with Mian and Sufi, the mortgage to income ratio grows the most rapidly in the sand states (and starts at the highest level of the three regions.) The middle states experience a milder though nontrivial rise, one that is close to the national average. Finally, the low states experience a very mild increase.

The lower left panel shows the behavior of employment. The cross-region behavior of employment is consistent with Mian and Sufi. The regions which suffered the largest boom and bust in house prices and housing leverage also experience the largest boom and bust in employment. In particular,

the employment contraction is largest in the sand states, followed by the middle states, which again behave close to the national average. The low states experience the mildest decline: Nonetheless there is still a significant employment contraction in this group of states.

In addition to the clear cross-sectional pattern seen in earlier literature, however, there are also some important common temporal movements in employment across the three regions. This common movement in employment is most clear around the time of the Lehman collapse: Employment in each of the three regions collapses at a roughly similar rate from the second quarter of 2008 to the business cycle trough. Based on this evidence and the earlier descriptive evidence, it seems quite plausible that financial conditions played an important role in the common cross-region employment collapse.

The common temporal movements in employment across regions become clearer when we remove construction employment and instead plot the behavior of non-construction employment. Much of the differential movement in employment prior the start of the recession reflects the boom and bust in residential investment, which initially affected mainly construction employment. Indeed, as the lower right panel shows, leading into the recession, the growth of non-construction employment was similar across the three regions. The same cross-sectional patterns in the employment decline emerge, but the differences in the employment contraction across regions are a bit smaller than for total employment. It can also be seen that the common decline in non-construction employment occurs a bit prior to Lehman, roughly in the second quarter of 2008. What all this suggests, is that to understand the overall employment contraction it is important to identify the common forces as well as the cross-sectional ones.

3.2 Identification of Financial and House Price Shocks

We now turn to our formal econometric framework. Our methodological approach involves two stages. First, using time series methods we identify orthogonal shocks to housing prices and financial conditions. We then develop a relation for state level employment growth that depends on these shocks. Following Mian and Sufi, we allow for a household balance sheet channel: That is, we allow for the impact of housing price shocks during the Great Recession period to depend on household leverage at the onset. Because the household balance sheet channel introduces a nonlinearity, we have to step outside the standard linear vector autoregression framework, as we discuss.

We begin with the identification of the financial and housing price shocks. We essentially use timing restrictions to identify these shocks. It is convenient to begin with the shock to financial conditions. Let s_t denote the financial excess bond premium (or “spread,”), p_t the log aggregate house price, e_t log aggregate employment and ε_t^s the structural shock to the spread. Then we assume the following equation gives the relation for s_t .

$$s_t = \eta_{s0}\Delta e_t + \gamma_{s0}\Delta p_t + \sum_{i=1}^4(\alpha_{si}s_{t-i} + \eta_{si}\Delta e_{t-i} + \gamma_{si}\Delta p_{t-i}) + \varepsilon_t^s \quad (2)$$

where, as before, Δ is the difference operator. We assume that the spread depends on the current values of employment growth and house prices as well as four lags of these variables and four own lags. Embedded in this formulation is the following key timing assumption: Current employment growth and current house prices growth can have an immediate impact on the spread. But movements in the spread can only affect house prices and employment with a one period lag. As a result, the residual ε_t^s provides our measure of the shock to financial conditions. It is the innovation in the spread that is orthogonal to the other right hand side variables, including employment and house prices.

An example of the shock to financial conditions is the unpredictable component of the rise in the financial excess bond premium at the time of the Lehman Brothers failure, as portrayed in Figure 3. To be clear, we are not arguing that shocks to financial conditions come completely out of thin air. At a complex nonlinear level, house prices influence this shock. That is, because at a complex nonlinear level, the decline in house prices eventually led to defaults on mortgages, which in turn induced losses on mortgage related securities, weakening bank balance sheets.. As described in the previous section, the weakening of bank balance sheets ultimately raises credit costs and so on. What we are assuming, accordingly, is that the way mortgage defaults driven by declining house prices affected the real economy was via tightening credit conditions, as measured by the innovation in the financial excess bond premium, ε_t^s .

We next turn to identifying the house price shocks, both for the aggregate and local level. Let p_{jt} be the log house price for state j , e_{jt} log state j employment, ε_j^p a fixed effect for house price growth in state j and ε_{jt}^p the shock to state j house price growth. Then the relation for a state’s house price growth is given by

$$\Delta p_{jt} = \eta_{p0} \Delta e_{jt} + \sum_{i=1}^4 (\alpha_{pi} s_{t-i} + \alpha_{pi} \Delta e_{jt-i} + \gamma_{pi} \Delta p_{jt-i}) + \varepsilon_j^p + \varepsilon_{jt}^p \quad (3)$$

Here we assume that Δp_{jt} depends on current and lagged state employment growth, lags of financial conditions and lags of itself. Here we are assuming that employment can have an immediate impact on house prices, but that house prices can affect employment only with a one period lag. We are also maintaining our early assumption that financial conditions affect house prices only with a one period lag. Under our identifying assumptions, ε_{jt}^p , is an orthogonal shock to house prices in state j .

In sum, we are identifying the structural shocks through the following three timing restrictions: (i) s_t depends on $\Delta e_t = \omega_j \Delta e_{jt}$ and $\Delta p_t = \omega_j \Delta p_{jt}$, where ω_j is state j 's population share; (ii) Δp_{jt} depends on Δe_{jt} but NOT s_t ; and (iii) Δe_{jt} depends neither on Δp_{jt} Nor s_t . We note that by “ordering” the spread after house prices, we are giving the edge to house prices in terms of overall explanatory power. At the same time, as we show, our main results are robust to variations on our timing assumptions.

Next, we disentangle housing price shocks at the state level into an aggregate and idiosyncratic component. Let ε_t^p be the aggregate shock to house prices and ψ_{jt}^p the idiosyncratic component of the state house price shock. Then we can express the state j level housing price shock as the following sum of aggregate and idiosyncratic components:

$$\varepsilon_{jt}^p = \theta_j \varepsilon_t^p + \psi_{jt}^p \quad (4)$$

with

$$\omega_j \theta_j = 1$$

Since the θ_j may differ across states, we are allowing for differential sensitivities of ε_{jt}^p to ε_t^p , due for example to due to differential land supply elasticities, etc. Note that under a certain normalization, ε_t^p corresponds to the common factor in ε_{jt}^p while θ_j is the state j factor loading

3.3 State Employment Growth Conditional on Housing Price and Financial Shocks

Let $M_j \equiv$ mortgage debt/income ratio in state j in 2006 (the year before the crisis) and let $I(crisis) \equiv$ crisis dummy (2007:1- 2009:4). Next, let

$e_{jt+h} - e_{jt}$ be employment growth in state j over h quarters and let $e_{t+h} - e_t$ be aggregate employment growth over h quarters. Then we assume that we can express state employment growth over horizon h , conditional on the state house price shock, ε_{jt}^p , and the aggregate house price and financial shocks, ε_t^p and ε_t^s respectively, as follows.

$$E_t\{e_{jt+h} - e_{jt} \mid \varepsilon_{jt}^p, \varepsilon_t^p, \varepsilon_t^s\} = \{\beta_{ph} + \beta_{mh}[I(crisis)]M_j\}\varepsilon_{jt}^p + \beta_{sh}\varepsilon_t^s \quad (5) \\ + \lambda_{hj}E_t\{e_{t+h} - e_t \mid \varepsilon_t^p, \varepsilon_t^s, \}$$

The terms on the top right of equation (5) reflect the local (partial equilibrium) effects on state j employment growth of shocks to state house prices and to financial conditions. We allow for both a linear effect and a nonlinear balance effect of the state house price shock. The linear effect represents, for example, the effect of state house prices on state housing construction. The nonlinear effect represents the impact of house prices on spending during the crisis via balance sheets conditions, as emphasized by Mian and Sufi. The last term on the right reflects the effect of financial conditions. Each of the shocks affects spending by households and firms within the state. As emphasized by Mian and Sufi, holding constant aggregate conditions, the effects on spending affect non-traded employment. Tradable employment demand, by contrast, depends on economy wide aggregate spending.

The bottom term on the right captures the effect on state employment growth of aggregate conditions, as measured by aggregate employment growth conditional on aggregate house price and financial shocks. Here we capture the spillover effect of aggregate shocks on state level employment that operates via the impact on tradable goods. An aggregate shock affects spending economy wide and will thus impact employment in each state j , depending on its tradable goods share. Note the term λ_{hj} reflects the aggregate spillover on state j employment growth. For the purpose of exposition we will assume in our baseline framework that the spillover effect is the same across states, i.e., $\lambda_{hj} = \lambda_h$. We will subsequently show that the results are robust to allowing the spillover effects to vary across states according to their respective tradable goods share (proxied for by the manufacturing share.)

Our goal is to obtain estimates of the “local effect” parameters, $\beta_{ph}, \beta_{mh}, \beta_{sh}$ and the “spillover effect parameter” λ_h over each horizon. Doing so will permit us to measure the contribution of each of the shocks to employment variation at both the aggregate and state level. To accomplish this task

we first need to find an expression for aggregate employment growth conditional on the aggregate housing and financial shocks, $E_t\{e_{t+h} - e_t \mid \varepsilon_t^s, \varepsilon_t^p\}$.

Let $\overline{M\theta} =_j \omega_j M_j \theta_j$, a weighted average of the product of state j mortgage to income ratio M_j and the state j sensitivity of house prices to the aggregate house price shock θ_j . Then using equation (5) to take a weighted average of state level employment growth yields an expression for conditional aggregate employment growth

$$E_t\{e_{t+h} - e_t \mid \varepsilon_t^p, \varepsilon_t^s\} = \{\{\beta_{ph} + \beta_{mh}[I(crisis)]\overline{M\theta}\}\varepsilon_t^p + \beta_{sh}\varepsilon_t^s\} + \lambda_h E_t\{e_{t+h} - e_t \mid \varepsilon_t^p, \varepsilon_t^s\} \quad (6)$$

The top term on the right is the average local effect of the aggregate shocks across states, while as before the bottom term is the spillover effect of aggregate conditions (via tradable goods demand)². We can rearrange to obtain the following expression for conditional aggregate employment growth

$$E_t\{e_{t+h} - e_t \mid \varepsilon_t^s, \varepsilon_t^p\} = \frac{1}{1 - \lambda_h} \{\{\beta_{ph} + \beta_{mh}[I(crisis)]\overline{M\theta}\}\varepsilon_t^p + \beta_{sh}\varepsilon_t^s\} \quad (7)$$

The multiplier $\frac{1}{1 - \lambda_h}$ translates the local effects of the aggregate shocks, captured by the term in brackets, into the aggregate general equilibrium effects. From equation (7) alone, though, it is not possible to disentangle the local and spillover effects of the shocks. We turn to this issue next.

It is first useful to decompose the state house price shock ε_{jt}^p into an aggregate component ε_t^p and a local component $\tilde{\varepsilon}_{jt}^p$. Given the process for ε_{jt}^p defined by equation (4), we can write

$$\varepsilon_{jt}^p = \varepsilon_t^p + \tilde{\varepsilon}_{jt}^p \quad (8)$$

with

$$\tilde{\varepsilon}_{jt}^p = (\theta_j - 1)\varepsilon_t^p + \psi_{jt}^p \quad (9)$$

According to equation (9), the local variation in the state j house price is the sum of the differential response to the aggregate shock $(\theta_j - 1)\varepsilon_t^p$ and the pure idiosyncratic shock ψ_{jt}^p .

From equations (7), (8) and (9), we can express state employment growth conditional on the local component of the state house price shock and

²Note that we are temporarily imposing the assumption of our baseline that the spillover effect λ_h is the same across states.

the aggregate house price and financial shocks as follows:

$$\begin{aligned}
E_t\{e_{jt+h} - e_{jt} \mid \widehat{\varepsilon}_{jt}^p, \varepsilon_t^p, \varepsilon_t^s\} &= \{\beta_{ph} + \beta_{mh}[I(crisis)]M_j\}\widehat{\varepsilon}_{jt}^p & (10) \\
&+ \frac{1}{1 - \lambda_h}\{\{\beta_{ph} + \beta_{mh}[I(crisis)]\overline{M\theta}\}\varepsilon_t^p + \beta_{sh}\varepsilon_t^s\}
\end{aligned}$$

Equation (10) relates conditional state employment growth to the sum of the pure local effect (the top term on the right side) and the aggregate general equilibrium effect (the bottom term). Note that the cross-sectional studies (e.g. Mian/Sufi etc.) just consider the local variation. (Either they consider a pure cross section or wash out the aggregate effects with time dummies). Accordingly, these studies cannot identify the general equilibrium effect summarized by the multiplier $1/(1 - \lambda_h)$. However, because we are able to disentangle local and aggregate variation, we can identify both the partial equilibrium parameters (the β s) and the general equilibrium parameter λ_h . In addition we can allow for a financial shock as well as housing price shocks.

3.4 Regression Strategy and Variance Decomposition

Let φ_j be a state fixed effect and φ_{jht} a residual error term. Then it follows from equation (10) that we can express state employment growth over h quarters as the following projection equation

$$\begin{aligned}
E_t\{e_{jt+h} - e_{jt} \mid \widehat{\varepsilon}_{jt}^p, \varepsilon_t^p, \varepsilon_t^s\} &= \beta_{ph}\widehat{\varepsilon}_{jt}^p + \beta_{mh}[I(crisis)]M_j\widehat{\varepsilon}_{jt}^p \\
&+ \frac{\beta_{ph}}{1 - \lambda_h}\varepsilon_t^p + \frac{\beta_{mh}}{1 - \lambda_h}[I(crisis)]\overline{M\theta}\varepsilon_t^p \\
&+ \beta_{sh}\varepsilon_t^s + \varphi_j + \varphi
\end{aligned}$$

Because the right hand side variables are orthogonal shocks, we can estimate the equation for each horizon h using ordinary least squares (following Jorda (2005)). As mentioned earlier we can also separately identify both the general equilibrium parameter as well as the partial equilibrium effect. Inspecting equation (11) reveals that over each horizon h , there are four parameters we wish to identify: the three local elasticities, $\beta_{ph}, \beta_{mh}, \beta_{sh}$; and the spillover effect λ_h . We can exploit the fact that we have both local variation in house price shocks, $\widehat{\varepsilon}_{jt+\bar{h}-h}^p$, and aggregate variation $\varepsilon_{t+\bar{h}-h}^p$ to disentangle the parameters that measure the local versus aggregate general equilibrium effects. Put differently, there are parameter restrictions on the coefficients

that we can impose that will permit separately identifying the local versus general equilibrium effects.

By estimating equation (11) over horizons that vary from $h = 1$ to $h = \bar{h}$, we can construct impulse responses for both state employment growth and aggregate employment growth over the horizon \bar{h} that capture the dynamic effects of each of the shocks. In addition we can also identify the contribution of housing price and financial shocks to employment variation over any horizon \bar{h} , both for aggregate and state level variation.

Rather than run \bar{h} separate regressions for each of the three shocks, we collapse everything into a single step as follows. Let β_{jph}^* be the composite effect of local house price variation on employment growth over horizon h and let β_{ph}^* be the composite effect of aggregate house price variation.

$$\begin{aligned}\beta_{jph}^* &= \beta_{ph} + \beta_{mh}[I(crisis)]M_j \\ \beta_{ph}^* &= \beta_{ph} + \beta_{mh}[I(crisis)]\overline{M\theta}\end{aligned}\quad (11)$$

Then we can express employment growth over horizon \bar{h} as a function of each of the house price and financial shocks from t to $t + \bar{h} - 1$, as follows

$$\begin{aligned}e_{jt+\bar{h}} - e_t &= \sum_{h=1}^{\bar{h}} \left\{ \beta_{jph}^* \tilde{\varepsilon}_{jt+\bar{h}-h}^p + \frac{\beta_{ph}^*}{1 - \lambda_h} \varepsilon_{t+\bar{h}-h}^p + \frac{\beta_{sh}}{1 - \lambda_h} \varepsilon_{t+\bar{h}-h}^s \right\} \\ &+ \sum_{i=1}^4 \delta_i \Delta e_{jt-i} + \varepsilon_j^e + \varphi_{jt,t+h}^e\end{aligned}\quad (12)$$

Note that we also allow for four lags of employment growth (the term $\sum_{i=1}^4 \delta_i \Delta e_{jt-i}$) to control for the effects of shocks prior to t on $e_{jt+\bar{h}} - e_t$ ³. The term on the upper right gives the contribution to $e_{jt+\bar{h}} - e_t$ of each of the shocks. The first element gives the contribution of local house price variation. The last two elements give the contribution of the aggregate house price and financial shocks.

Given the parameter estimates from equation (12), we can construct a historical decomposition of aggregate employment growth as follows: First, take a weighted average of equation (12) to obtain the following relation for aggregate employment growth over the horizon \bar{h} as a function of the

³By doing so, we control for low frequency effects on \bar{h} ahead employment growth. Further, by estimating the effects of the shock in a single step, we control for these low frequency effects in a consistent manner across horizons

aggregate shocks and lagged employment growth:

$$\begin{aligned}
e_{jt+\bar{h}} - e_t &= \sum_{h=1}^{\bar{h}} \left\{ \frac{\beta_{ph}^*}{1 - \lambda_h} \varepsilon_{t+\bar{h}-h}^p + \frac{\beta_{sh}}{1 - \lambda_h} \varepsilon_{t+\bar{h}-h}^s \right\} \\
&+ \sum_{i=1}^4 \delta_i \Delta e_{t-i} + \varepsilon^e + \varphi_{t,t+h}^e
\end{aligned} \tag{13}$$

Since the local variation in house prices washes out in the aggregate, employment growth can be expressed as a function of the aggregate house price and financial shocks.

Next, define $\Gamma_{t,t+\bar{h}}$ as the unpredictable aggregate employment growth over horizon \bar{h} (implied by equation (13)):

$$\Gamma_{t,t+\bar{h}} = (e_{t+\bar{h}} - e_t) - \left(\sum_{i=1}^4 \delta_i \Delta e_{t-i} + \varepsilon^e \right) \tag{14}$$

The unpredictable component is the difference between the realized employment growth and the component that is predicted by lagged employment growth and the constant trend growth ε^e .

Combining equations (13) and (14) then relates $\Gamma_{t,t+\bar{h}}$ to the sequences of aggregate house price and financial shocks from t to $t + \bar{h} - 1$, as follows

$$\Gamma_{t,t+\bar{h}} = \sum_{h=1}^{\bar{h}} \left\{ \frac{\beta_{ph}^*}{1 - \lambda_h} \varepsilon_{t+\bar{h}-h}^p + \frac{\beta_{sh}}{1 - \lambda_h} \varepsilon_{t+\bar{h}-h}^s \right\} + \varphi_{t,t+\bar{h}}^e \tag{15}$$

Given the parameter estimates and the identified shock series, we can then construct a historical decomposition of unanticipated aggregate employment growth. To do a decomposition of regional employment growth, we repeat the same steps, except we start with a weighted average of state employment growth within a region as opposed to the country as a whole. Note that in the case of regional employment variation, the state level local shocks do not average out to zero.

4 Results

We estimate equation (13) for each state using ordinary least squares. We weight each state's observation by its beginning of period employment. As

discussed earlier, we use the local versus aggregate variation in house prices along with the coefficient restrictions in equation (13) to separately identify the local versus general equilibrium effects. We pick the horizon for employment growth \bar{h} to be twelve quarters, a period long enough to capture the Great Recession period but short enough to isolate business cycle dynamics. Hence, our variance decomposition will focus on explain the sources of variation in the time series of unanticipated three year employment growth rates.

We begin with the estimates of the sensitivity of employment growth to each of the three shocks (local house price, aggregate house price and aggregate financial shock.) For a given shock, the parameter estimates for each horizon h trace out an impulse response of employment growth. We report the results in Figure 5.

The top left panel in Figure 5 gives the impulse response of employment to a one hundred basis point increase in the financial excess bond premium. Aggregate employment begins to contract two quarters after the shock. It reaches a trough eight quarters with a cumulative decline of roughly three percentage points. The error bands give 90 percent confidence bands. The next two panels report the effects of one percentage point declines in aggregate and local house prices, respectively. (In the latter case, we measure a weighted average of the local effect across states). In each case employment declines, reaching a trough roughly four to six quarters out. Consistent with the theory, the aggregate house price shock produces a larger employment decline, given it incorporates the general equilibrium spillover effect

Table 1 reports the estimates of the effects of: the composite (linear and nonlinear) local house price shocks (the top row); the composite aggregate house price shocks (middle row); and the spillover effect λ_h (bottom row). Note from equation (12) the ratio of the composite aggregate effect to the composite local effect equals the general equilibrium multiplier $1/(1 - \lambda_h)$. Over the twelve quarter horizon, the spillover effect averages roughly 0.70, with a peak of 0.8 after six quarters and a trough of 0.29 after twelve. The implied general equilibrium multiplier is an average of three and a third, with a range between five and one and a half. At first blush, a multiplier effect of 3 or more appears large, given that tradables are on average thirty percent of output within a state. However, it is important to keep in mind that a large share of tradables are durable goods as compared to nontraded goods which consist mostly of nondurable goods and services. Given that durables (and hence tradables) are more highly cyclical than nontradables, a large general

equilibrium multiplier effect is plausible.

We next decompose the effect of aggregate house prices on aggregate employment between the “linear” effect that arises independently of household financial distress (e.g., the impact on construction) versus that due to the household balance sheet channel being operative during the Great Recession. The left panel of Figure 6 reports the overall response of employment to the house price shock. The next two panels then report the decomposition between the employment drop due to the linear effect and the drop due to the household balance sheet channel (in the panel labeled “nonlinear.” Interestingly, each channel is of roughly equal importance in the overall employment contraction. Given the balance sheet channel is only operative during the Great Recession, the implication is that during the crisis, the impact of the housing price decline on employment was roughly double its amount during normal times.

We now analyze how aggregate house price and financial shocks account for employment variation over the Great Recession. The dark line in Figure 7 plots the unexpected three year growth rate in employment ($\Gamma_{t,t+12}$ in equation (14)) over the period 1998:1 through 2016:4. The red line is the combined contribution of the aggregate house price and financial shocks, as estimated from equation (15). Notice that the red line tracks the blue line fairly closely between 2007:1 and 2011:4, a period that encompasses the Great Recession. While there is some co-movement, the red line does not track the blue line as closely in the periods before 2007:1 and after 2011:4. The implication is that the house price and financial shocks account for nearly all the aggregate employment decline during the crisis period. Before and after the crisis the contribution is smaller.

In Figure 8 we analyze the relative contributions of the housing price versus financial shocks to the employment contraction over the Great Recession period. The solid line is the unanticipated three year growth rate from 2007:1 to 2011:1. The dashed line is the combined contribution of housing and financial shocks. The green line is the contribution of housing shocks alone while the red line is the contribution of the financial shock. The vertical dashed line denotes the Lehman Brothers bankruptcy.

What the figure suggests is that both shocks contributed to the employment contraction. Consistent with conventional wisdom, the house price shock accounts for nearly all of the employment contraction through the fourth quarter of 2008, a drop of roughly six percentage points. After this point the financial shock takes over, contributing significantly to both the

depth and persistence of the downturn. First, it adds another five percentage points to the employment contraction, which peaks in 2010Q:1. Second, absent the financial shock, employment would have recovered by four percentage points by 2011:1, just two percentage points below trend. However, the financial shock keeps employment growth down six percentage. In this way it also contributes to the persistence of the downturn.

Figure 9 disaggregates employment into construction versus non-construction. In each case, the impact of house prices relative to that of the financial shock is similar to the effect on aggregate employment: House prices explain most of the contraction up through 2008:4 but then the financial shock takes over. However, house prices are relatively more important in accounting for the construction employment drop as compared to the baseline in Figure 7. Conversely, the financial shock becomes important for the contraction in non-construction employment, accounting for roughly six percentage points in the drop at the trough in 2010:1 versus five percentage points. It also becomes more important for the persistence of the decline in non-construction employment, as compared to the housing price shock.

We next turn to accounting for regional variation. As in Figure 4, we divide the states into three regions according to the magnitude of the house price boom and bust. The left panel in Figure 10 characterizes the high boom/bust states, the middle panel the medium/boom, and the right panel the low boom/bust. In each panel, the solid line is the unexpected three year growth of employment from 1998:1 to 2011:4 and the red line is the combined contribution of house prices shocks (aggregate and local) and the financial shock. Overall, the two types of shocks account well for the regional behavior of employment over the crisis period, 2007:1 to 2011:4. As with aggregate employment in Figure 8, house price and financial shocks do not tightly track employment outside the crisis period.

Figure 11 shows the contribution of each shock to regional variation. Given our identification, the financial shock (given by the red line), has the same impact on employment within each region. Accordingly, the housing shock must account for all the regional variation in the employment contraction. Overall, the housing shock does well in account for the employment drop in the “middle” boom/bust and “low” boom/bust states (given by the middle and right panels). House prices also account well for the employment drop in the “high” boom/bust states through 2008:4. But they do not account fully for the sustained drop in employment afterward, suggesting that other forces may be at work (e.g. regional heterogeneity in the impact of the

financial shock).

Figure 12 show that house prices do slightly better in accounting for the cumulative drop in retail employment in the high boom/bust states, relative to the drop in total employment. House prices also account well for the retail employment drop in the middle and low regions. On the other hand, as Figure 13 shows, house prices do not do well in account for the relative drop in manufacturing employment, Since regional house price variation contains a significant local component, the relative differences between the response of retail and manufacturing employment are consistent with theory. In particular, as Mian and Sufi emphasize, local house price variation should affect mainly nontradable employment (retail) and not tradable (e.g. manufacturing.)

To understand further our historical decomposition results in Figures 8,9, 11, 12 and 13, Figure 14 reports the time series of the aggregate house price shock and the financial shock over the crisis period. As the top panel shows, just following the peak of house prices in late 2006, there is a sustained sequence of negative house price shocks during the lead up to the Lehman collapse (denoted by the vertical line.) Put differently, the fall in house prices after the peak was faster than expected. Following the Lehman collapse, house prices recovered relative to expectation, leading to a series of positive house price shocks. The figure makes clear why the house price shocks were so important during the first phase of the crisis, Conversely, financial shocks pick up in mid 2007, around the time of the first wave of mortgage defaults. These shocks then increase further around the time of the Lehman collapsed. The lag in the financial shock relative to the house price shock helps account for why the former dominates in the post-Lehman phase of the crisis.

Finally, our results highlight the relative contribution of house price versus financial shocks to overall employment variation. Given the reduced form nature of the exercise, they do not provide explicit evidence about propagation mechanisms. For example, could the financial shock, in addition to generating credit market distress, also be triggering the household balance sheet channel by reducing house prices? Conversely, could the drop in house prices be triggering weakness in bank balance sheets and hence financial distress, leading to a rise in the financial excess bond premium. We address these issues in Figure 15.

The top panel of Figure 15 provides a historical decomposition of house prices while the bottom panel does so for the financial excess bond premium. The top panel shows that the financial shock has vitrually no effect on the

variation in house prices: House price variation is generated mainly by the house price shock. Thus, it is very unlikely that the household balance sheet channel is a factor in the transmission of the financial shock. Conversely, the bottom panels shows that house price shocks have a non-trivial impact on the financial excess bond premium, particularly around the time of the Lehman bankruptcy. The house price shock accounts for roughly half the increase in the financial excess bond premium during this period. Thus, financial distress could potentially be playing an important role in transmitting the impact of the house price shock on aggregate employment. For this reason, our estimates of the overall impact of financial distress from the variance decompositions are likely to be conservative estimates of the relative importance of this channel.

5 Concluding Remarks

We present evidence on the channels of financial distress to the real economy during the Great Recession. To do so we use a panel of state level data. We employ both cross-sectional and time series methods to identify shocks and causality. Our methods allow us to identify both local and aggregate general equilibrium effects of shocks.

Our main finding is that both the household balance sheet and banking distress channels played an important role in the contraction. House price shocks and shocks to the financial excess bond premium each account for roughly half the cumulative employment decline, with the former important in the early phase (pre-Lehman default) and the latter dominant post-Lehman. Roughly half the effect of the house price shock is transmitted via the household balance sheet channel while the other half reflects factors also operative outside the Great Recession, such as the impact of house prices on construction. Finally, the house price shock helps account for the regional variation in employment, consistent with earlier literature.

Our estimates of the effects of financial distress, further, are likely to be conservative for two main reasons. First, as we have discussed, house price shocks had a non-trivial impact on the excess bond premium at the height of the crisis, while the financial shock had virtually no impact on house price dynamics. Second, we do not allow for regional heterogeneity of the financial shock, implying that we are not allowing for the possibility that financial factors could also have helped explain regional variation.

Accordingly, we see two immediate directions for future work. The first is a structural model to help identify propagation mechanisms. The second is finding a creative way to allow for regional heterogeneity in the impact of the financial shock.

References

- [1] Adrian, Tobias and Hyun Shin, 2010, “The Changing Nature of Financial Intermediation and Financial Crisis of 2007-2009,” *Annual Review of Economics*, 603-618.
- [2] Adrian, Tobias, Paolo Colla, and Hyun Shin, 2012, “Which Financial Frictions? Parsing the Evidence from the Financial Crisis of 2007 to 2009,” In Daron Acemoglu, Jonathan Parker and Michael Woodford (eds), *NBER Macroeconomics Annual*, 27, 159-214. Chicago: University of Chicago Press.
- [3] Arellano, Cristina, Yan Bai and Patrick Kehoe, 2016, “Financial Frictions and Fluctuations in Volatility,” mimeo.
- [4] Becker, Bo and Victoria Ivashina, 2014, “Cyclicality of credit supply: Firm level evidence,” *Journal of Monetary Economics* 62: 76-93.
- [5] Benmelech, Efraim, Ralf Meisenzahl and Rodney Ramacharan, 2017, “The Real Effects of Liquidity During the Financial Crisis: Evidence from Automobiles,” *Quarterly Journal of Economics*, 132(1): 317-365.
- [6] Berger, David, Veronica Guerrieri, Guido Lorenzoni and Joseph Vavra, 2017, “House Prices and Consumer Spending,” *Review of Economic Studies*, forthcoming.
- [7] Bernanke, Ben, 1983, “Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression,” *American Economic Review* 73(1): 257-276.
- [8] Bernanke, Ben, 2010, “Causes of the Recent Financial and Economic Crisis”. Statement before the Financial Crisis Inquiry Commission, Washington, September 2.
- [9] Bernanke, Ben and Mark Gertler, 1989, “Agency Costs, Net Worth and Business Fluctuations,” *American Economic Review* 79(1): 14-31.
- [10] Bernanke, Ben and Mark Gertler, 1995, “Inside the Black Box: The Credit Channel of Monetary Policy Transmission,” *Journal of Economic Perspectives* 9(4): 27-48.

- [11] Bernanke, Ben, Mark Gertler and Simon Gilchrist, 1999, “The Financial Accelerator in a Quantitative Business Cycle Framework,” Handbook of Macroeconomics, Ben Friedman and Michael Woodford editors.
- [12] Brunnermeier, Markus and Yuliy Sannikov, 2014, “A Macroeconomic Model with a Financial Sector,” *American Economic Review* 104(2): 379-421.
- [13] Brunnermeier, M. K., T. Eisenbach and Y. Sannikov, 2013, “Macroeconomics with Financial Frictions: A Survey.” In D. Acemoglu, M. Arelano and E. Dekel (Eds.) *Advances in Economics and Econometrics: Tenth World Congress of the Econometric Society*. Vol. 2: 3-94. Cambridge: Cambridge University Press.
- [14] Chen, Brian, Samuel Hanson and Jeremy Stein, “The Decline of Big Bank Lending to Small Businesses: Impacts on Local Credit and Labor Markets,” mimeo.
- [15] Chodorow-Reich, Gabriel, 2014, “The Employment Effects of Credit Market Disruptions,” *Quarterly Journal of Economics* 129 (1):1-59.
- [16] Chodorow-Reich, Gabriel and Antonio Falato, 2017, “The Loan Covenant Channel: How Bank Health Transmits to the Economy,” mimeo.
- [17] Christiano, Lawrence, Roberto Motto and Massimo Rostagno, 2014, “Risk Shocks,” *American Economic Review* 104(1): 27-65.
- [18] Covitz, Daniel, Nellie Liang and Gustavo Suarez, 2013, “Evolution of a Financial Crisis: Collapse of the Asset-Backed Commercial Paper Market,” *Journal of Finance*, 68(3): 815-848.
- [19] Diamond, Douglas, and Philip H. Dybvig, 1983, “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy* 91(3): 401-419.
- [20] Eichengreen, Barry, 1992, *The Gold Standard and the Great Depression, 1919-1939* (New York: Oxford University Press.)
- [21] Eggertsson, Gauti and Paul Krugman, 2012, “Debt, Deleveraging, and Liquidity Trap,” *Quarterly Journal of Economics* 127(3): 1469-1513.

- [22] Fisher, Irving, 1933, “The Debt-Deflation Theory of the Great Depression,” *Econometrica* , 1(4): 337-357.
- [23] Gertler, Mark, and Karadi, Peter., 2011, “A Model of Unconventional Monetary Policy,” *Journal of Monetary Economics* 58(1): 17-34.
- [24] Gertler, Mark and Nobuhiro Kiyotaki, 2011, “Financial Intermediation and Credit Policy in Business Cycle Analysis”. In Friedman, B., and Woodford, M. (Eds.), *Handbook of Monetary Economics*. Elsevier, Amsterdam, Netherlands.
- [25] Gertler, Mark and Nobuhiro Kiyotaki, 2015, “Banking, Liquidity and Bank Runs in an Infinite Horizon Economy,” *American Economic Review*, forthcoming.
- [26] Gilchrist, Simon and Egon Zakrajsek, 2012, “Credit Spread and Business Cycle Fluctuations,” *American Economic Review* 102(4): 1692-1720.
- [27] Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajsek, 2017, “Inflation Dynamics During the Financial Crisis,” *American Economic Review* 107(3): 785-823.
- [28] Gilchrist Simon, Michael Siemer and Egon Zakrajsek, 2017, “The Real Effects of Credit Booms and Busts: A County Level Analysis,” mimeo.
- [29] Giroud, Xavier and Holger Mueller, 2017, “Firm Leverage, Consumer Demand, and Unemployment During the Great Recession,” *Quarterly Journal of Economics*, forthcoming.
- [30] Gorton, Gary, 2010. *Slapped by the Invisible Hand: The Panic of 2007*. Oxford University Press.
- [31] Guerrieri, Veronica and Guido Lorenzoni, 2017, “Credit Crises, Precautionary Saving and the Liquidity Trap,” *Quarterly Journal of Economics* 132(3): 1427-1467.
- [32] He, Zhiguo and Arvind Krishnamurthy, 2013. “Intermediary Asset Pricing,” *American Economic Review* 103(2): 732-770.
- [33] He, Zhiguo and Arvind Krishnamurthy, 2014, “A Macroeconomic Framework for Quantifying Systemic Risk,” Working Paper, University of Chicago and Stanford University.

- [34] Jermann, Urban and Vincenzo Quadrini, 2012, “Macroeconomic Effects of Financial Shocks,” *American Economic Review* 102(1): 238-271.
- [35] Jorda, Oscar, 2005, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review* 95(1): 161-182.
- [36] Kacperczyk, Marcin and Philipp Schnabl. 2010, “When Safe Proved Risky: Commercial Paper during the Financial Crisis of 2007-2009,” *Journal of Economic Perspectives*, 24(1), 29-50.
- [37] Kaminsky, Graciela and Carmen Reinhart, 1999, “The Twin Crises: The Causes of Banking and Balance of Payments Problems,” *American Economic Review* 89(3), 473-500.
- [38] Kaplan, Greg, Kurt Mitman and Gianluca Violante, 2016a, “Nondurable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data,” mimeo.
- [39] Kaplan, Greg, Kurt Mitman and Gianluca Violante, 2016b, “The Housing Boom and Bust: Model Meets Evidence,” mimeo.
- [40] Kashyap, Anil, Jeremy Stein and David Wilcox, 1993, “Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance,” *American Economic Review* 83(1): 78-98.
- [41] Kiyotaki, Nobuhiro, and John Moore, 1997, “Credit Cycles,” *Journal of Political Economy* 105(2): 211-248.
- [42] Krishnamurthy, Arvind, Stefan Nagel and Dimitri Orlov. 2014, “Sizing Up Repo,” *Journal of Finance* 69(6): 2381-2417.
- [43] Mian, Atif and Amir Sufi, 2014, “What Explains the 2007-2009 Drop in Employment,” *Econometrica* 82, 2197-2223.
- [44] Mian, Atif, Kamalesh Rao and Amir Sufi, 2013, “Household Balance Sheets, Consumption and the Economic Slump,” *Quarterly Journal of Economics*.
- [45] Mian, Atif, Amir Sufi and Emil Verner, 2017, “Household Debt and Business Cycles Worldwide,” *Quarterly Journal of Economics* 132, 1755-1817.

- [46] Midrigan, Virgiliu and Thomas Phillipon, 2016, “Household Balance Sheets and The Recession,” mimeo.
- [47] Muir, Tyler and Arvind Krishnamurthy, 2017, “How Credit Cycles Across a Financial Crisis,” mimeo.
- [48] Reinhart, Carmen and Kenneth Rogoff, 2009, This Time is Different, Eight Centuries of Financial Folly, Princeton.

Table 1: Local vs Aggregate House Price Effects and Spillovers

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|-----------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|
| Local | 0.04 (0.03) | 0.07 (0.03) | 0.10 (0.03) | 0.11 (0.03) | 0.11 (0.03) | 0.09 (0.03) | 0.09 (0.03) | 0.09 (0.03) | 0.11 (0.03) | 0.14 (0.03) | 0.18 (0.03) | 0.21 (0.03) |
| Aggregate | 0.14 (0.07) | 0.27 (0.08) | 0.38 (0.08) | 0.46 (0.08) | 0.49 (0.08) | 0.49 (0.08) | 0.46 (0.07) | 0.43 (0.07) | 0.40 (0.07) | 0.36 (0.07) | 0.33 (0.07) | 0.30 (0.06) |
| Spillover | 0.73 | 0.74 | 0.74 | 0.76 | 0.78 | 0.80 | 0.81 | 0.78 | 0.71 | 0.61 | 0.46 | 0.29 |

Figure 1: Household Balance Sheets, House Prices and Spending

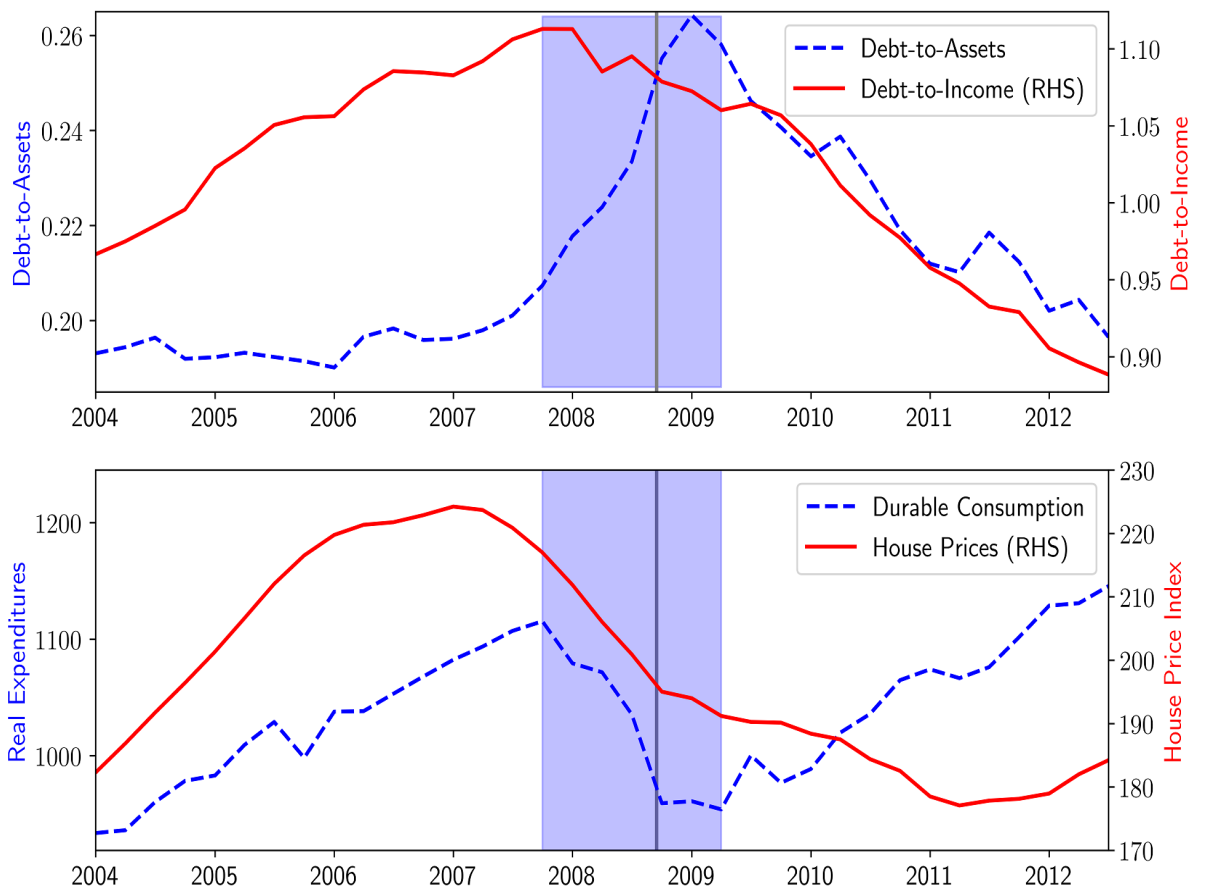
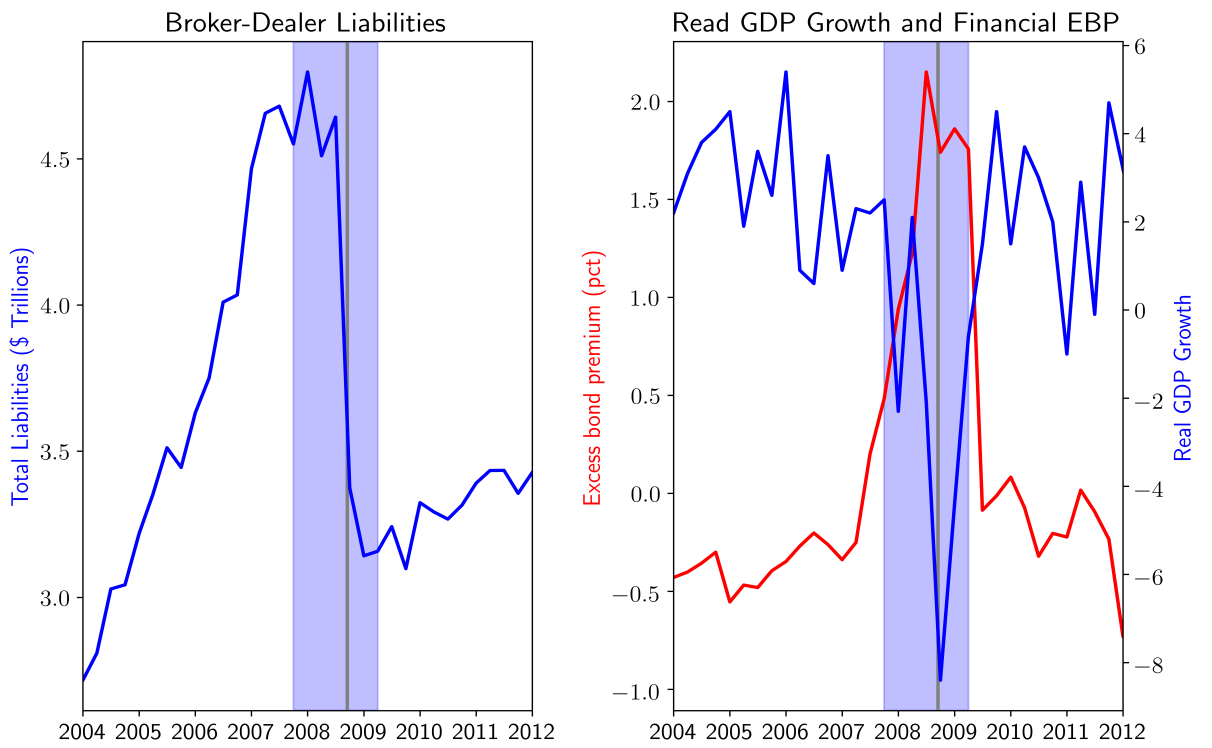


Figure 2: Banking Distress and Real Activity



Excess Bond Premium = rate of return on corporate bonds minus return on similar maturity government bonds, with default premium removed.

Figure 3: Bank Market Leverage, Financial EBP and Lending Standards




Figure 4: Regional and Temporal Variation


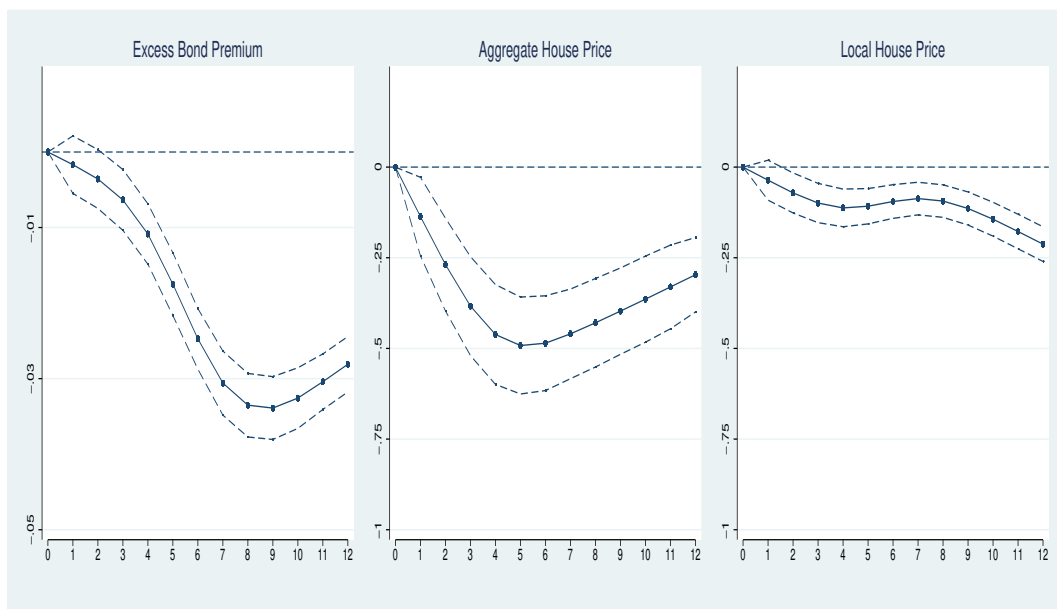


Figure 5: Employment Response to Financial vs. Housing Price Shocks



Left panel: shock is 100bp in Financial EBP; Middle and Right: 100bp decrease in house price

Figure 6: Employment Response to Aggregate House Price Shock:
Linear vs. Nonlinear Effects





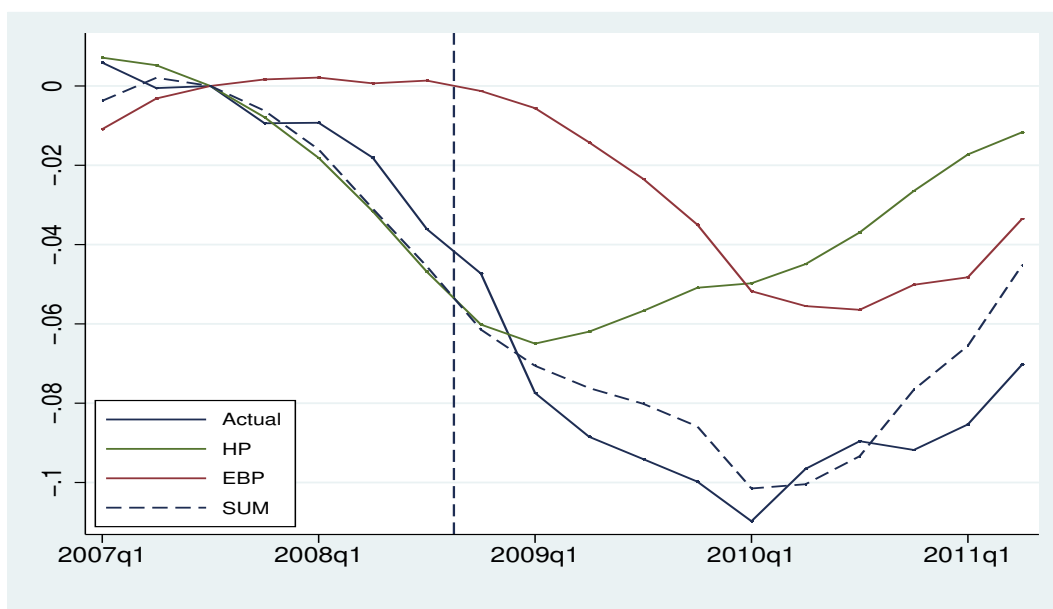
Figure 7: Employment: Actual vs. Fitted



Actual: Unexpected 3 year growth


Fitted: Component explained by financial and house price shocks

Figure 8: Contribution of Housing Price vs. Financial Shocks to Employment




Actual: Unexpected 3 year employment growth

Figure 9: Construction vs. Ex-Construction Employment:
House Price vs EBP effect




Actual: Unexpected 3 year employment growth

Figure 10: Regional Employment: Actual vs. Fitted




Actual: Unexpected 3 year growth. Fitted: Component explained by financial and house price shocks

Figure 11: House Price vs. Financial Shocks and Regional Employment




Actual: Unexpected 3 year employment growth

Figure 12: House Price vs. Financial Shocks and Regional Retail Employment



Actual: Unexpected 3 year employment growth

Figure 13: House Price vs Financial Shocks and Regional Manufacturing Employment



Actual: Unexpected 3 year employment growth

Figure 14: Aggregate Housing Price and Financial Shocks




Figure 15: What Drives Housing Prices and the EBP?

