Debt Constraints and Employment*

Patrick Kehoe† Virgiliu Midrigan‡ Elena Pastorino§

November 2017

Abstract

During the Great Recession, the regions of the United States that experienced the largest declines in household debt also had the largest drops in consumption, employment, and wages. Employment declines were larger in the nontradable sector. Motivated by these findings, we develop a search and matching model with credit frictions. In the model, tighter debt constraints raise the cost of investing in new job vacancies and thus reduce worker job-finding rates and employment. The key new feature of our model, on-the-job human capital accumulation, is critical to generating sizable drops in employment. On-the-job human capital accumulation increases the duration of the flows of benefits from posting vacancies and, in our quantitative model, amplifies the employment drop from a credit tightening tenfold relative to the standard Diamond-Mortensen-Pissarides model. We show that our model reproduces well the salient cross-regional features of the U.S. economy during the Great Recession.

Keywords: Search and Matching, Employment, Debt Constraints, Human Capital, Great Recession, Regional Business Cycles.

JEL classifications: E21, E24, E32, J21, J64.

---

*We thank Eugenia Gonzalez Aguado, Julio Andres Blanco, Sonia Gilbukh, and Sergio Salgado Ibanez for excellent research assistance, and Joan Gieseke for invaluable editorial assistance. We are especially grateful to Erik Hurst for guiding us in the replication of the wage measures in Beraja, Hurst, and Ospina (2016). We are indebted to Moshe Buchinsky, Pawel Krolikowski, and Ayşegül Şahin for sharing their data and codes with us. Kehoe thanks the National Science Foundation for financial support. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

†Stanford University, UCL, and Federal Reserve Bank of Minneapolis, pkehoe@stanford.edu.

‡New York University, virgiliu.midrigan@nyu.edu.

§Hoover Institution, Stanford University, and Federal Reserve Bank of Minneapolis, epastori@stanford.edu.

During the Great Recession, the regions of the United States that experienced the largest declines in household debt also had the largest drops in consumption, employment, and wages. A popular view of this cross-sectional evidence is that large disruptions in the credit market played a critical role in generating the differential cross-regional declines in output and employment. This view is motivated by the regional patterns documented in recent work by several authors. Mian and Sufi (2011, 2014) show that U.S. regions that experienced the largest declines in household debt and housing prices also saw the greatest drops in consumption and employment, especially in the nontradable goods sector. Beraja, Hurst, and Ospina (2016) show that discount factor shocks can account for the vast bulk of the cross-regional variation in employment in the United States during the Great Recession. Moreover, these authors document that wages were moderately flexible; that is, the cross-regional decline in wages was almost as large as the decline in employment.

We develop a version of the Diamond-Mortensen-Pissarides (*DMP*, henceforth) model with risk-averse agents, borrowing constraints, and human capital accumulation to investigate how the interplay between credit and labor market frictions can account for these cross-sectional patterns. Our exclusive focus is on showing how shocks to credit can account for the cross-regional patterns observed during the Great Recession; we do not attempt to account for the time series patterns of aggregates. In large part, this focus is motivated by the findings of Beraja, Hurst, and Ospina (2016), who show that shocks to discount factors account for little of the aggregate employment decline, despite their accounting for most of the cross-sectional variation in employment.

Our analysis builds on the idea that hiring workers is an investment activity: the costs of posting job vacancies are paid up-front, whereas the benefits, as measured by the flows of surplus from the match between a firm and worker, accrue gradually over time. Like for any investment activity, a credit tightening generates a fall in such investment and, hence, a drop in employment in the aggregate. Although this force is present in any search model, we show that the drop in employment following a credit tightening is very small in the textbook version of the DMP model without human capital accumulation. In such a model, a large fraction of the present value of the benefits from forming a match accrues early in the match. Indeed, according to a standard measure of the timing of such flows—the *Macaulay duration*—these flows have a very short duration of two to three months. The resulting small drop in employment in the DMP model following a credit tightening is then reminiscent of standard results in corporate finance, according to which a tightening of credit has little impact on investments with low-duration cash flows. (See, for example, Eisfeldt and Rampini (2007) and the references therein.)

The flows of benefits from forming a match, in contrast, have a much longer duration in the presence of human capital accumulation on the job. In this case, a match not only produces current output but also augments a worker's human capital, with persistent effects on a worker's future output flows. For a sense of the magnitude of these effects, a quantified version of our model, consistent with the evidence on the dynamics of wages with tenure and experience and across employment spells, generates surplus flows with a duration of about 10 years. This significantly longer duration amplifies the drop in employment from a credit contraction by a factor of 10 relative to that implied by the DMP model.

To illustrate the workings of our new mechanism, we first consider a one-good model. To build intuition, consider a firm's incentives to post vacancies after a credit tightening that leads to a temporary fall in consumption. When consumers have preferences for consumption-smoothing, the shadow price of goods increases after a credit tightening and then mean-reverts. This temporary increase in the shadow price of goods has two countervailing effects. First, it increases the cost of posting vacancies by raising the shadow value of the goods used in this investment. Second, it increases the surplus from a match by raising the shadow value of the surplus flows produced by a match. Since the cost of posting new vacancies is incurred immediately when goods are especially valuable, whereas the benefits accrue gradually in the future, the cost of posting vacancies rises by more than the benefits. As a result, firms post fewer vacancies and employment falls. The resulting drop in vacancies is larger the longer is the duration of the surplus flows from a match.

To understand why durations are short in the DMP model and longer in our model, note that the surplus flows, defined as the net benefits to a worker and a firm from forming a match, can be expressed as the difference between the average streams of output produced by a consumer who begins a new match and those produced by an otherwise identical consumer who is currently unmatched. These average streams incorporate the transition rates between employment and nonemployment as well as the Nash bargaining rule.

Without human capital accumulation, the surplus flows end when either the initially matched consumer separates or the initially unmatched consumer becomes employed. When job-finding rates are high, as they are in the data, durations are short because the initially unmatched consumer quickly finds a job. Hence, a temporary increase in the shadow price of goods increases the present value of benefits by about as much as the costs, leading to only a small drop in vacancy creation.

In sharp contrast, when we introduce human capital accumulation and allow some of the acquired capital to be transferable across matches, the surplus flows from a given match extend beyond a particular employment relationship. Since the human capital acquired in a match increases the output that the initially employed consumer produces in all subsequent matches, a substantial fraction of the present value of the surplus flows accrues far into the future. Hence, a temporary increase in the shadow price of goods increases the present value of the benefits of posting vacancies by much less than their costs, leading to a large drop in vacancy creation.

We model human capital as partially transferable across matches by assuming that consumers accumulate two types of human capital: *general human capital* that is fully transferable across matches and *firm-specific human capital* that fully depreciates when a match dissolves. We show that general rather than firm-specific human capital accumulation is mostly responsible for the amplification of the employment response to a credit tightening in our model relative to the DMP model. The key intuition behind this result is that with general human capital, surplus flows last beyond the current employment relationship, whereas with firm-specific human capital, these flows end when the current employment relationship terminates.

To shed light on this critical feature of our model, we consider simplified versions of our model with constant accumulation of either general or firm-specific human capital, which admit closed-form solutions for surplus flows and durations. In the standard DMP model, the surplus flows from a match t periods after it is formed follow a first-order difference equation with solution $s_{t+1} = c\delta^t$, in which the root of the difference equation, referred to as the *DMP root*, is $\delta = 1 - \sigma - \gamma\lambda_w$, where σ is the match destruction rate, γ is the worker's bargaining weight, and λ_w is the worker's job-finding rate. For standard parameterizations, the DMP root is substantially smaller than one, implying that the surplus flows decay quickly—at a rate of about 25% per month—and, thus, have a short duration.

Adding either form of human capital accumulation implies that surplus flows follow a second-order difference equation with solution $s_{t+1} = c_s\delta_s^t + c_\ell\delta_\ell^t$, where δ_s and δ_ℓ are, respectively, the *small* and *large* roots of the equation, and c_s and c_l are the corresponding weights on these roots. In the firm-specific human capital model, the small root is equal to the DMP root, whereas the large one is equal to $(1-\sigma)(1+g_h)$, where g_h is the growth rate of firm-specific human capital. The large root is less than one for reasonable parameterizations of the rate of human capital accumulation, implying that firm-specific human capital also generates short durations. In contrast, in the general human capital model, although the small root is approximately equal to the DMP root, the large root is always greater than one and increasing in the growth rate of general human capital. This large root is the source of the much longer duration of surplus flows in the general human capital model compared to either the firm-specific human capital model or the DMP model.

We quantify the parameters governing human capital accumulation based on two sources of data: cross-sectional data from Elsby and Shapiro (2012) on how wages increase with experience and longitudinal data from Buchinsky et al. (2010) on how wages grow over an employment spell. Our model not only matches explicitly targeted moments from these data but is also broadly consistent with evidence from the Panel Study of Income Dynamics (PSID) on wage declines upon separation, how these wage declines vary with tenure in the previous job, the distribution of durations of nonemployment spells, and the evidence on wage losses from displaced worker regressions in the spirit of Jacobson, LaLonde, and Sullivan (1993).

We show that our results on employment declines in response to a credit tightening are robust to a range of estimates of wage growth in the literature, including estimates of how wages increase with experience from Rubinstein and Weiss (2006) and estimates of how wages grow over an employment spell from Altonji and Shakotko (1987) and Topel (1991). More generally, we find that the employment response is determined almost entirely by the amount of life-cycle wage growth that workers experience and is highly nonlinear in this growth. In particular, as long as life-cycle wage growth exceeds a threshold of about 1% per year, further increases in the amount of life-cycle wage growth have little effect on how employment responds to a credit tightening. This 1% threshold is consistent with essentially all of the estimates in the literature, including those for workers with different levels of education. (See the survey by Rubinstein and Weiss (2006) and Buchinsky et al. (2010).)

This nonlinearity stems from two of our model's implications. First, the duration of the flows of benefits from a match is a concave function of the rate of human capital accumulation, so that above a certain accumulation rate, the increase in duration becomes small. Second, a credit tightening leads to a transitory drop in consumption, so that flows received after some future date do not experience a large increase in their valuation, regardless of how distant in time they are. Because of this nonlinearity, our results are robust not only to estimates of wage growth from the studies cited earlier but also to any estimate of life-cycle wage growth above our 1% threshold.

We study a simple model of credit frictions but show that for a large class of models, if credit shocks produce the same paths for consumption and, hence, the shadow prices of goods, then these models produce the same paths for labor market variables. Specifically, we first show that our model has implications for employment and wages that are identical to those of an economy with housing in which debt is collateralized by the value of a house. This result allows us to rationalize the drop in employment as driven by a tightening of collateral constraints arising from a fall in the price of housing.

We then show that our model has implications for employment and wages that are also identical to those of an economy with illiquid assets in which a tightening of debt constraints reduces the consumption of even wealthy households. This result allows us to interpret our model as applying to (wealthy) net savers rather than only to net borrowers. More generally, these equivalence results formalize the view of Beraja, Hurst, and Ospina (2016) that discount factor shocks are an appropriate reduced-form representation of a tightening of household borrowing limits in the sense that this representation is consistent with a range of primitive models. Overall, our equivalence results show that the robust link across models is the one between consumption and labor market outcomes rather than the one between either house prices or levels of net assets and labor market outcomes. Motivated by these results, we focus our quantitative work on this robust link.

To confront the regional patterns discussed earlier that motivate our work, we extend our economy to include a large number of islands, each of which produces a nontradable good that is consumed only on the island and a tradable good that is consumed everywhere in the world. Each worker is endowed with one of two types of skills that are used with different intensities in the tradable and nontradable goods sectors. Labor is immobile across islands but can switch between sectors. Importantly, the differential intensity of the use of skills across sectors generates a cost of reallocating workers between sectors.

In this economy, an island-specific credit tightening has two effects. The first, the *investment* effect, is similar to that in the one-good model: the cost of posting vacancies increases by more than the benefits, leading to a reduction in the number of vacancies and, hence, to a drop in overall employment on that island. The second effect, the *relative demand* effect, is due to the reduction in the demand for the nontradable goods produced on the island. This drop in demand for nontradable goods, in turn, leads to a drop in demand for workers by that sector, which leads those workers to reallocate to the tradable goods sector. The smaller is the cost of reallocating workers, the larger is the reallocation and, thus, the larger is the drop in nontradable employment and the smaller is the drop in tradable employment.

We find that our extended model reproduces well the regional patterns of the U.S. economy during the Great Recession. In particular, in the data, a credit tightening that leads to a 10% fall in consumption across U.S. states between 2007 and 2009 is associated with a fall in nontradable employment of 5.5% and a negligible increase in tradable employment of 0.3% across states. Our model has similar predictions: the same fall in consumption is associated with a fall in nontradable employment of 5.7% and a negligible increase in tradable employment of 0.3% across states. Furthermore, our model accounts for most of the resulting change in overall employment: a 10% drop in consumption is associated with a 3.8% drop in employment in the data and a 3.3% drop in the model.

Critically, our model is also consistent with the Beraja, Hurst, and Ospina (2016) observation that in the cross section of U.S. states, wages are moderately flexible: a 10% drop in employment is associated with a fall in wages of 7.8% in both the data and the model. Thus, our model predicts sizable employment changes even though wages are as flexible as they are in the data.

Other Related Literature. In incorporating human capital accumulation into a search model, we build on the work of Ljungqvist and Sargent (1998, 2008), who extend McCall's (1970) model to include stochastic human capital accumulation on the job and depreciation off the job. Since they retain key features of the McCall model, such as linear preferences and an exogenous distribution of wages, the forces behind our results are not present in their framework.

Our work is complementary to that of Hall (2017), who studies the effects of changes in the discount rate in a search model. Hall's model features no human capital accumulation, and as such, implies short durations of the flows of benefits from matches between firms and workers. In contrast to our model, which assumes that wages are determined through Nash bargaining, Hall (2017) assumes the bargaining protocol in Hall and Milgrom (2008), which implies that wages do not fall much in response to shocks. In our model, wages fall only moderately in response to a credit tightening even with the standard Nash bargaining protocol. In this sense, our mechanism is complementary to that in Hall's work.

Our work is also closely related to that of Krusell, Mukoyama, and Şahin (2010) on the interaction between labor market frictions and asset market incompleteness. Their work focuses on an economy's response to aggregate productivity shocks but abstracts from human capital accumulation, whereas we focus on an economy's response to regional credit shocks in a model in which human capital accumulation plays a critical role.

Our work is related to a burgeoning literature that links a worsening of financial frictions on the consumer side to economic downturns. In particular, Guerrieri and Lorenzoni (2015), Eggertsson and Krugman (2012), and Midrigan and Philippon (2016) study macroeconomic responses to a household-side credit crunch. All three of these papers find that a credit crunch has only a minor impact on employment unless wages are sticky. Our analysis complements this work by exploring a mechanism that does not impose sticky wages but rather generates an employment decline within a search model of the labor market in which wages are renegotiated every period through Nash bargaining.

Our model is also related to the work of Itskhoki and Helpman (2015) on sectoral reallocation in an open economy model with search frictions in the labor market, of Pinheiro and Visschers (2015) on endogenous compensating differentials and unemployment persistence in a labor market model with search frictions, as well as the work on house prices, credit, and business cycles of Ohanian (2010), Head and Lloyd-Ellis (2012), and the work comprehensively surveyed by Davis and Van Nieuwerburgh (2014).

Finally, our work is related to the large literature on financial intermediation, dating back at least to Bernanke and Gertler (1989), Kiyotaki and Moore (1997), and Bernanke, Gertler, and Gilchrist (1999). More recent work includes Mendoza (2010), Gertler and Karadi (2011), Gertler and Kiyotaki (2010), and Gilchrist and Zakrajšek (2012). This literature focuses on how credit frictions amplify the response of physical capital investment to shocks. Our work, instead, focuses on how credit shocks amplify employment responses in a model with labor market frictions and human capital accumulation. Moreover, this literature studies the overall effects of aggregate shocks, whereas we focus on the effects of regional shocks.

1 A One-Good Economy

We consider a small open economy, one-good DMP model. The economy consists of a continuum of firms and consumers. Each consumer survives from one period to the next with probability ϕ. In each period, a measure $1-\phi$ of new consumers is born, so that there is a constant measure one of consumers. Individual consumers accumulate general and firm-specific human capital and are subject to idiosyncratic shocks. Firms post vacancies in markets indexed by a consumer's general human capital. Consumers are organized in families that own firms and insure against idiosyncratic risks. Each family is subject to time-varying debt constraints.

1.1 Technologies

Consumers are indexed by two state variables that summarize their ability to produce output. The variable z_t, referred to as *general human capital*, captures returns to experience and stays with the consumer even after a job spell ends. The variable h_t, referred to as *firm-specific human capital*, captures returns to tenure and is lost every time a job spell ends. A consumer with state variables (z_t, h_t) produces $z_t h_t$ when employed and $b(z_t)$ when not employed. When the consumer is employed, general human capital evolves according to

$$\log z_{t+1} = (1-\rho)\log \bar{z}_e + \rho \log z_t + \sigma_z \varepsilon_{t+1}, \tag{1}$$

where ε_{t+1} is a standard Normal random variable, whereas when the consumer is not employed, it evolves according to

$$\log z_{t+1} = (1-\rho)\log \bar{z}_u + \rho \log z_t + \sigma_z \varepsilon_{t+1}. \tag{2}$$

We assume that $\bar{z}_u < \bar{z}_e$. Newborn consumers start as nonemployed with general human capital z, where $\log z$ is drawn from $N(\log \bar{z}_u, \sigma_z^2/(1-\rho^2))$. This specification of human capital is in the spirit of that in Ljungqvist and Sargent (1998). We denote the Markov processes in (1) and (2) as $F_e(z_{t+1}|z_t)$ and $F_u(z_{t+1}|z_t)$ in what follows. The consumer's firm-specific human capital starts at $h_t = 1$ whenever a job spell begins and then evolves on the job according to

$$\log h_{t+1} = (1-\rho)\log \bar{h} + \rho \log h_t, \tag{3}$$

with $\bar{h} > 1$. The assumption that $\bar{z}_u < \bar{z}_e$ implies that when a consumer is employed, on average, the variable z_t drifts up toward a high level of productivity $\bar{z}_e$ from the low average level of productivity $\bar{z}_u$ of newborn consumers. Similarly, when the consumer is not employed, on average, the variable z_t depreciates and hence drifts down toward a low level of productivity, $\bar{z}_u$, which we normalize to 1. The assumption that $\bar{h} > 1$ implies that when the consumer is employed, firm-specific human capital increases from $h = 1$ toward $\bar{h}$ over time.

The parameter ρ governs the rate at which general and firm-specific human capital converge toward their means. The higher is ρ, the slower both types of capital accumulate during employment and the slower general human capital depreciates during nonemployment. For simplicity, here we assume that these rates are the same for all three laws of motion mentioned earlier, but in Section 5, we explore the implications of alternative rates of decay for the nonemployed. Allowing for idiosyncratic shocks ε_{t+1} to general human capital allows the model to reproduce the dispersion in wage growth rates observed in the data. For simplicity only, we assume that the process of firm-specific human capital accumulation is deterministic.

We represent the insurance arrangements in the economy by imagining that each consumer belongs to one of a large number of identical families, each of which has a continuum of household members. Each family as a whole receives a deterministic amount of income in each period generated by its working and nonworking members. Risk sharing within a family implies that at date t, each household member consumes the same amount of goods, c_t, regardless of the idiosyncratic shocks that such a member experiences. (This type of risk-sharing arrangement is familiar from the work of Merz (1995) and Andolfatto (1996).) Each family is subject to debt constraints.

Given this setup, we can separate the problem of a family into two parts. The first part determines the common consumption level of every family member. The second part determines the vacancies created and the matches continued by each firm owned by the family, as well as the employment and nonemployment status of each consumer in the family.

1.2 A Family's Problem

We purposely consider a simple model of a family's consumption-savings choice in order to focus attention on the interaction between credit and labor market frictions. In this model, the family trades a single risk-free security and faces debt constraints. We later show that this *economy with debt constraints* has implications for consumption, employment, and wages that are equivalent to those of richer models in which either debt is collateralized by housing and debt constraints tighten as house prices fall, or families are debt constrained even though they have savings (in an illiquid asset) and are net savers.

The consumption allocation problem of a family is given by

$$\max_{c_t, a_{t+1}} \sum_{t=0}^{\infty} \beta^t u\left(c_t\right), \tag{4}$$

subject to the budget constraint

$$c_t + q a_{t+1} = y_t + d_t + a_t \tag{5}$$

and the debt constraint

$$a_{t+1} \geqslant -\chi_t. \tag{6}$$

Here β is the discount factor of the family, c_t is consumption, a_{t+1} are savings, y_t represents the total income from the wages of the employed members of the family and home production of its nonemployed members, and d_t are the profits from the firms the family owns. The family saves or borrows at a constant world bond price $q > \beta$ subject to an exogenous deterministic sequence $\{\chi_t\}$ of positive borrowing limits. Because the bond price and debt limits are deterministic, and there is a continuum of family members who face only idiosyncratic risk, the family's problem is deterministic.

The family values one unit of goods at date t at the *shadow price* of $Q_t = \beta^t u'(c_t)/u'(c_0)$ units of date 0 goods. The Euler equation for consumption is

$$q Q_t = Q_{t+1} + \theta_t,$$

where θ_t is the multiplier on the debt constraint. A tightening of debt constraints—a fall in χ_t—raises the value of date t consumption goods by forcing the family to repay its debt and temporarily reduce consumption.

We next describe the second part of the family problem, which consists of firms' choices of vacancy creation and match destruction, as well as consumers' choices between employment and nonemployment.

1.3 An Individual Firm's Problem

We posit and then later characterize equilibrium wages as the outcome of a (generalized) Nash bargaining problem that yields a wage $w = \omega_t(z, h)$. For a given wage w, the present value of profits earned by a firm matched with a consumer with human capital levels z and h, expressed in date 0 consumption units, is given by

$$\tilde{J}_t\left(w, z, h\right) = Q_t(zh - w) + (1-\sigma)\,\phi \int \max\left[J_{t+1}\left(z', h'\right), 0\right] \mathrm{d}F_e\left(z'|z\right). \tag{7}$$

The flow profits are simply the difference between the amount zh the firm produces and the wage w it pays the consumer. Since the firm is owned by the family, it values date t

profits using the family's shadow price Q_t. Note that the maximum operator on the right side of (7) reflects the firm's option to destroy an unprofitable match. Given the function $w = \omega_t(z, h)$ from the Nash bargaining problem discussed later, the firm's value is defined as $J_t(z, h) = \tilde{J}_t(\omega_t(z, h), z, h)$.

1.4 An Individual Consumer's Values

The consumer's value in any period depends on whether the consumer is employed. The present value of an employed consumer's earnings, expressed in date 0 consumption units, is

$$\begin{aligned} \tilde{W}_t(w, z, h) = Q_t w + \phi(1-\sigma) \int \max\left[W_{t+1}(z', h'), U_{t+1}(z')\right] \mathrm{d}F_e(z'|z) \\ + \phi\sigma \int U_{t+1}(z') \mathrm{d}F_e(z'|z), \end{aligned} \tag{8}$$

where $U_{t+1}(z')$ denotes the present value at $t+1$ of a nonemployed consumer's earnings with general human capital z', general human capital evolves according to the law of motion $F_e(z'|z)$ in (1), and firm-specific human capital evolves according to the law of motion in (3). The continuation value reflects the consumer's survival probability ϕ, the exogenous match separation probability σ, and the possibility of endogenous match separation. Given the wage function $w = \omega_t(z, h)$ from the Nash bargaining problem below, the consumer's value of working is defined as $W_t(z, h) = \tilde{W}_t(\omega_t(z, h), z, h)$.

The present value of a nonemployed consumer's earnings, expressed in date 0 consumption units, is

$$\begin{aligned} U_t(z) = Q_t b(z) + \phi\lambda_{wt}(z) \int \max\left[W_{t+1}(z', 1), U_{t+1}(z')\right] \mathrm{d}F_u(z'|z) \\ + \phi[1-\lambda_{wt}(z)] \int U_{t+1}(z') \mathrm{d}F_u(z'|z). \end{aligned} \tag{9}$$

Here $\lambda_{wt}(z)$, described in full later on, is the probability that a consumer with general human capital z is matched with a firm at date t, in which case the consumer's state at $t+1$ consists of general human capital z' and firm-specific human capital $h' = 1$. The continuation value reflects the consumer's survival probability ϕ, the consumer's matching rate $\lambda_{wt}(z)$, and the endogenous match acceptance decisions.[1]

1.5 Matching, Nash Bargaining, and Vacancy Creation

We now consider the matching technology, the determination of wages through Nash bargaining, the vacancy creation problem of firms, and the resulting steady-state measures of

[1]The only time a match is not accepted is when a worker with general human capital z in period t draws a sufficiently low shock so that the resulting z' at $t+1$ leads to a negative surplus. In our quantitative analysis, nearly all matches are indeed accepted.

employed and nonemployed consumers.

Matching and Nash Bargaining. Firms can direct their search for consumers market by market, by posting vacancies for nonemployed consumers of a given level of general human capital z. Let $u_t(z)$ be the measure of nonemployed consumers with human capital z and $v_t(z)$ the corresponding measure of vacancies posted by firms for consumers in market z. The measure of matches in this market is generated by the matching function $m_t(z) = u_t(z)\, v_t(z) / [u_t(z)^\eta + v_t(z)^\eta]^{\frac{1}{\eta}}$, as in den Haan, Ramey, and Watson (2000) and Hagedorn and Manovskii (2008). We use this matching function to ensure that job-finding rates are between 0 and 1. Specifically, the probability that a nonemployed consumer of type z matches with a firm in market z is

$$\lambda_{wt}(z) = m_t(z) / u_t(z) = \theta_t(z) / [1 + \theta_t(z)^\eta]^{\frac{1}{\eta}},$$

where $\theta_t(z) = v_t(z) / u_t(z)$ is the vacancy-to-nonemployment ratio for consumers of type z, or *market tightness*, and the parameter η governs the sensitivity of $\lambda_{wt}(z)$ to $\theta_t(z)$. The probability that a firm posting a vacancy in market z matches with a consumer in this market is

$$\lambda_{ft}(z) = m_t(z) / v_t(z) = 1/[1 + \theta_t(z)^\eta]^{\frac{1}{\eta}}.$$

The Nash bargaining problem, which determines the wage $w = \omega_t(z, h)$ in any given match, is

$$\max_w \big[\tilde{W}_t(w, z, h) - U_t(z)\big]^\gamma \tilde{J}_t(w, z, h)^{1-\gamma},$$

where γ is a consumer's bargaining weight. Defining the surplus of a match between a firm and a consumer with human capital (z, h) as $S_t(z, h) = W_t(z, h) - U_t(z) + J_t(z, h)$, Nash bargaining implies that firms and consumers split this surplus according to

$$W_t(z, h) - U_t(z) = \gamma S_t(z, h) \text{ and } J_t(z, h) = (1 - \gamma) S_t(z, h).$$

Vacancy Creation. Consider the firm's choice of vacancy creation. The cost of posting a vacancy in any market z is κ units of goods. The free-entry condition in market z is then given by

$$Q_t\kappa \geq \phi\lambda_{ft}(z) \int \max\left[J_{t+1}(z', 1), 0\right] \mathrm{d}F_u(z'|z), \tag{10}$$

with equality if vacancies are created in *active market* z in that $v_t(z) > 0$. Since the surplus from a match increases with z, $F_u(z'|z)$ shifts to the right with z, and the firm's value is proportional to the surplus, there is a cutoff level of general human capital, z_t^*, such that firms post vacancies in all markets with $z \geq z_t^*$ and none in any market with $z < z_t^*$. This result arises because in markets with $z < z_t^*$, the value of expected profits conditional on matching is not sufficient to cover the fixed cost of posting a vacancy, even if a vacancy leads to a match with probability 1. The cutoff z_t^* then satisfies

$$Q_t \kappa = \phi \int \max \left[J_{t+1} \left(z', 1 \right), 0 \right] \mathrm{d}F_u \left(z' | z_t^* \right). \tag{11}$$

1.6 The Workings of the Model

Here we discuss how our model works. We first describe the model's steady-state properties and then the economy's response to a debt tightening.

Steady-State Properties. Panels A and B of Figure 1 display the steady-state measures of the employed $e(z)$ and nonemployed $u(z)$ as a function of general human capital, whereas panels C and D of this figure display the firm and consumer matching probabilities in market z. (We generate these figures by using the parameter values described later.)

As discussed, there is a cutoff level of z, z^*, such that in markets $z < z^*$, firms post no vacancies and consumers have a zero matching probability. To the right of z^*, the consumer job-finding probability increases with z because firms matched with consumers with higher levels of z earn higher profits and thus have greater incentives to post vacancies aimed at attracting such consumers. These incentives ensure that market tightness $v(z)/u(z)$ increases with z and the firm matching probability decreases with z, so that the expected value of posting a vacancy is the same in all active markets and equal to the cost of posting a vacancy.

A Tightening of Debt Constraints. Consider next how a tightening of debt constraints affects firms' incentives to post vacancies and thus employment in equilibrium. As we discuss later, such a tightening leads to a temporary decrease in the family's consumption as the family repays its debt to reduce its debt position. Hence, the debt tightening leads to a temporary increase in the family's marginal utility of consumption and so the shadow price of goods Q_t. Because the drop in consumption is transitory, the shadow prices Q_t, Q_{t+1}, Q_{t+2},..., initially increase above their steady-state levels and then revert back to these levels as consumption mean-reverts to its steady-state level.

To understand how this temporary increase in the shadow price Q_t affects firms' incentives to post vacancies, consider the free-entry condition. Since Nash bargaining implies that a firm's value is a constant fraction of match surplus, we can write the free-entry condition for active markets as

$$Q_t \kappa = \phi \lambda_{ft}(z)(1-\gamma) \int \max \left[S_{t+1} \left(z', 1 \right), 0 \right] dF_u \left(z' | z \right). \tag{12}$$

Here the cost of posting vacancies, $Q_t\kappa$, on the left side is equal to the benefit, namely the product of the firm's matching probability $\lambda_{ft}(z)$, a decreasing function of market tightness $\theta_t(z) = v_t(z)/u_t(z)$, and a term that just depends on to the expected surplus from a match.

The temporary increase in Q_t has two effects on the free-entry condition. First, it raises the benefits of posting vacancies by increasing the expected surplus from a match. The surplus increases because a match produces a greater flow of output than does nonemployment, and the family values this net flow more when its consumption is lower. Second, a higher Q_t directly raises the cost of posting vacancies, $Q_t\kappa$. Importantly, the second effect dominates the first, so that the cost increases by more than the benefit.

The intuition is simple. The cost is paid at t, when consumption is the lowest and goods are most valuable, but the benefits accrue in future periods when consumption has partially recovered, and so goods are less valuable than they are at date t. Because the cost of posting vacancies increases by more than the expected surplus from the match, the firm's matching probability, $\lambda_{ft}(z)$, must increase after a debt tightening to ensure that the free-entry condition holds. Intuitively, firms post fewer vacancies because the cost of investing in new vacancies increases by more than the returns to such investments. This is a familiar effect from a large class of models in which a worsening of financial frictions leads to lower investment. To see this intuition more formally, rewrite the free-entry condition in an active market z at t as

$$Q_t\kappa = \phi\lambda_{ft}(z_t)\times(1-\gamma)\times\left[Q_{t+1}\mathbb{E}_t s_{t+1}(z_{t+1})+Q_{t+2}\mathbb{E}_t s_{t+2}(z_{t+2})+Q_{t+3}\mathbb{E}_t s_{t+3}(z_{t+3})+\ldots\right]. \quad (13)$$

The term $Q_t\kappa$ on the left side is the cost of posting a vacancy, whereas the terms on the right side are the benefits of posting a vacancy. The benefits are defined by the product of the probability of consumer survival, ϕ, the probability of filling a vacancy, $\lambda_{ft}(z_t)$, the firm's bargaining weight, $(1-\gamma)$, and the surplus from a match, S_{t+1}, given by the term in brackets. Here $\mathbb{E}_t s_{t+k}(z_{t+k})$ denotes the expected *flow surplus* produced in period $t+k$. The expectation operator takes into account all of the uncertainty concerning a match, including variations in flow surplus due to shocks to a consumer's general human capital during employment as well as the possibility that a match dissolves because of death or other reasons. (See the Appendix for details on how we compute these components.)

Consider how a mean-reverting shock to the shadow price of goods affects the cost and benefits of posting a vacancy. Specifically, let $\log Q_t$ increase by e on impact and then mean-revert at rate ϱ so that $\mathrm{d}\log Q_{t+k} = \varrho^k e$. Clearly, the cost of posting a vacancy, given by $Q_t\kappa$, increases by e log points.[2] As for the benefits, the surplus S_{t+1} from the match increases by

$$\mathrm{d}\log S_{t+1} = \left[\varrho \frac{Q_{t+1}\mathbb{E}_t s_{t+1}(z_{t+1})}{S_{t+1}} + \varrho^2 \frac{Q_{t+2}\mathbb{E}_t s_{t+2}(z_{t+2})}{S_{t+1}} + \varrho^3 \frac{Q_{t+3}\mathbb{E}_t s_{t+3}(z_{t+3})}{S_{t+1}} + \dots\right] e. \quad (14)$$

We can rewrite this increase as

$$\mathrm{d}\log S_{t+1} = \sum_{k=0}^{\infty} \varrho^{k+1}\omega_{k+1} e, \quad (15)$$

where the weight $\omega_k = Q_{t+k}\mathbb{E}_t s_{t+k}(z_{t+k})/S_{t+1}$ is the share of the surplus received in the k-th period of the match. Note that in (14) we hold these weights ω_k fixed at their steady-state values $\omega_k = s_k / \sum_{j=1}^{\infty} \beta^{j-k} s_j$ to keep the algebra simple. Intuitively, the change in surplus is a weighted average of the amount by which Q_{t+k} changes in response to the credit tightening, that is, $\varrho^k = \mathrm{dlog}\, Q_{t+k}/\mathrm{d}e$ in period $t+k$, where the weight ω_k is the share of surplus accruing in that particular period. Define the surplus to be more *front-loaded* the larger $\sum_{k=0}^{\infty} \varrho^{k+1}\omega_{k+1}$ is: since the weights ω_k sum to one, a more front-loaded surplus is characterized by a greater share of the total surplus accruing early in a match. Equation (15) implies that the more front-loaded the surplus from a match is, the more the present value of the surplus increases after a given mean-reverting increase in the shadow price of goods.

To infer the implications of these changes in the cost and benefits of posting vacancies for the worker-finding rate, we totally differentiate the free-entry condition (13) and substitute $\mathrm{dlog}\, Q_t = e$ and the expression for $\mathrm{dlog}\, S_{t+1}$ in (15) to obtain

$$\mathrm{d}\log \lambda_{ft}(z) = \mathrm{d}\log Q_t - \mathrm{d}\log S_{t+1} = \left[1 - \sum_{k=0}^{\infty} \varrho^{k+1}\omega_{k+1}\right] e. \quad (16)$$

The expression in brackets in (16) provides an alternative version of the concept of *Macaulay duration* commonly used in finance to determine how the present value of a stream of payments changes in response to permanent changes in one-period discount rates.

To see this connection, rewrite (16) in terms of one-period shadow discount rates instead of shadow goods prices. To do so, note that in our exercise, the shadow discount rate increases on impact by $\Delta r_t = (1-\varrho)e$ at t and then mean-reverts at rate ϱ so that the expected future short-term rates are $\Delta r_{t+k} = \varrho^k \Delta r_t$. Substituting $\Delta r_t/(1-\rho)$ for e in (16) yields that the change in the firm's job-finding rate following such a change to discount rates is

$$\mathrm{d}\log \lambda_{ft}(z) = \mathrm{d}\log Q_t - \mathrm{d}\log S_{t+1} = \left[\frac{1 - \sum_{k=0}^{\infty} \varrho^{k+1}\omega_{k+1}}{1-\varrho}\right] \Delta r_t. \quad (17)$$

[2] Because we have no aggregate uncertainty, throughout we use that in period t, the long-term discount rate from period t to any period $t+k$ is the product of the one-period discount rates from period t to period $t+k$, that is, the pure expectation hypothesis holds. Also, the durations we compute below, which reflect changes in the surplus from persistent shocks to consumption, can equally well be thought of as reflecting the change in the present value of a stream of payments in t from a change in the term structure of long-term discount rates in period t.

The term in brackets on the far right side of (17) is an alternative Macaulay duration for nonparallel shifts in the term structure of discount rates, which we refer to as *alternative Macaulay duration* for brevity. This measure of duration extends the standard notion of Macaulay duration, $\sum_{k=0}^{\infty}(k+1)\omega_{k+1}$, used to calculate the effects of parallel shifts in the term structure of discount rates.[3] Our alternative measure modifies this notion to allow for a particular type of nonparallel shift in the term structure, namely, that implied by our temporary credit tightening, which leads short-term discount rates to increase more than long-term ones.[4]

The higher this duration, the smaller the increase in match surplus after a credit tightening and, hence, the larger the increase in the firm's worker-finding rate and, since $\lambda_{wt}(z)^{\eta} = 1 - \lambda_{ft}(z)^{\eta}$, the larger the decline in the worker's job-finding rate. In this case, a credit tightening leads to a large fall in employment. Later, we develop this intuition further in the context of two simplified versions of the economy.

2 Equivalence Results

In our economy the labor market outcomes, including employment, nonemployment, vacancies, and wages, are uniquely determined by the sequence of shadow prices $\{Q_t\}$. Because of this feature, many alternative setups for the family problem yield equivalent outcomes for consumption and the associated shadow prices of goods, and hence for all labor market outcomes, although they may have different implications for other variables. To illustrate this point, we show the equivalence between our economy and an *economy with housing* and an *economy with illiquid assets*.

In our economy with housing, debt is collateralized by the value of a house. Our equivalence result allows us to rationalize the drop in consumption and resulting labor market outcomes in our baseline economy as actually driven by a tightening of collateral constraints arising from a fall in the price of housing, as many have argued occurred during the Great Recession.

In our economy with illiquid assets, consumers are net savers with the rest of the world rather than net borrowers, as in our baseline model. Nevertheless, since savings are illiquid, a tightening of debt constraints reduces consumption because of liquidity constraints, as in the work of Kaplan and Violante (2014) on wealthy but liquidity-constrained households. Our equivalence result allows us to rationalize the drop in consumption and resulting labor

[3]Obviously, $\sum_{k=1}^{\infty} k\omega_k = \sum_{k=0}^{\infty}(k+1)\,\omega_{k+1}$ and $\frac{1-\sum_{k=1}^{\infty}\varrho^k\omega_k}{1-\varrho} = \frac{1-\sum_{k=0}^{\infty}\varrho^{k+1}\omega_{k+1}}{1-\varrho}$.

[4]Note that the expression in (17) arises from a *fixed* change in interest rates of Δr that decays at rate ϱ, rather than a given change in e. Holding fixed that change in Δr as we vary ϱ, the expression in brackets in (17) converges to the Macaulay duration, $\sum_{k=0}^{\infty}(k+1)\omega_{k+1}$, as ϱ converges to 1.

market outcomes in our baseline economy as arising in a model in which consumers are net savers.

These results make clear that the robust nexus across models is the one between consumption and labor market outcomes, not the one between either house prices or levels of net asset positions and labor market outcomes. Motivated by these results, in our quantitative exercise we choose sequences of state-level shocks to reproduce the state-level consumption paths observed in the data and then study the resulting implications for labor market outcomes.

2.1 An Economy with Housing

Consider an economy in which families own houses and their borrowing is subject to collateral constraints based on the value of their houses. The preferences of the family are

$$\max_{c_t, h_{t+1}} \sum_{t=0}^{\infty} \beta^t [u(c_t) + \psi_t v(h_t)], \tag{18}$$

where c_t is consumption and h_t is the amount of housing consumed at date t. The family faces a budget constraint,

$$c_t + q a_{t+1} + p_t h_{t+1} = y_t + d_t + a_t + p_t h_t. \tag{19}$$

Here a family owns a house of size h_t with value $p_t h_t$, chooses its next period housing level h_{t+1}, and faces a collateral constraint that limits the amount it can borrow to a fraction $\bar{\chi}$ of the value of the family's house,

$$a_{t+1} \geqslant -\bar{\chi} p_t h_{t+1}.$$

The housing supply is fixed, and each unit of housing delivers one unit of housing services each period. Note that the parameter ψ_t in the utility function governs the relative preference for housing. This parameter varies over time and is the source of changes in house prices and thus, through the collateral constraint, in the amount the family can borrow.

Let $\{Q_t\}$ denote the sequence of shadow prices that results from the economy with debt constraints for some given sequence of debt constraints $\{\chi_t\}$. Clearly, there exists a sequence of taste parameters $\{\psi_t\}$ that gives rise to this same sequence of shadow prices in the economy with housing. Given these shadow prices, the labor market side of the economy with housing is identical to that of the economy with debt constraints. Hence, consumption, labor allocations, and wages in the two economies coincide. Likewise, given the sequence of shadow prices that results from the economy with housing for some given sequence of taste parameters $\{\psi_t\}$, there exists a sequence of debt constraints in the economy with debt constraints that gives rise to these same shadow prices. We show these results formally in the Appendix and summarize this discussion with a proposition.

Proposition 1. *The economy with debt constraints is equivalent to the economy with housing in terms of consumption, labor allocations, and wages.*

2.2 An Economy with Illiquid Assets

Here we consider an economy with illiquid assets. Each family can save in assets that have a relatively high rate of return but are illiquid, and each can borrow at a relatively low rate. The budget constraint is

$$c_t + q_a a_{t+1} - q_b b_{t+1} = y_t + d_t + a_t - b_t - \phi(a_{t+1}, a_t), \tag{20}$$

where a_{t+1} denotes assets and b_{t+1} denotes debt. We assume that $q_a = 1/(1 + r_a) < q_b = 1/(1 + r_b)$ so that the return on assets, $1 + r_a$, is higher than the interest on debt, $1 + r_b$. We interpret r_a and r_b as after-tax interest rates. We imagine a situation in which even though the before-tax rate on debt is higher than that on assets, the after-tax rate is lower because of the tax deductibility of interest payments. For simplicity, we assume that $\beta = q_a$. The function $\phi(a_{t+1}, a_t)$ represents the cost of adjusting assets from a_t to a_{t+1} and captures the idea that assets are illiquid. Borrowing is subject to the debt constraint

$$b_{t+1} \leq \tilde{\chi}_t, \tag{21}$$

where $\{\tilde{\chi}_t\}$ is a sequence of exogenous maximal amounts of borrowing. The consumption problem of the family is to choose $\{c_t, a_{t+1}, b_{t+1}\}$ to maximize utility in (4) subject to the budget constraint (20) and debt constraint (21).

We assume that the interest rate on borrowing is sufficiently low in the illiquid asset economy that the borrowing constraint in that economy binds at the shadow prices constructed from the economy with debt constraints. That is, condition

$$q_b = \frac{1}{1 + r_b} > \frac{Q_{t+1}}{Q_t} \tag{22}$$

holds, where the right side of (22) is evaluated at the consumption allocations in the economy with debt constraints. In the Appendix, we prove an equivalence result analogous to that in Proposition 1, which we summarize in the following proposition.

Proposition 2. *Under (22), the economy with debt constraints is equivalent to the economy with illiquid assets in terms of consumption, labor allocations, and wages.*

3 Quantification

We next describe how we choose parameters for our quantitative analysis and the model's steady-state implications.

Assigned Parameters. The model is monthly. We choose utility to be $u(c) = c^{1-\alpha}/(1-\alpha)$ and present results for a range of values of α. The discount factor β is $(.96)^{1/12}$, the world bond price q is $(.98)^{1/12}$, and the survival rate ϕ is set so that consumers are in the market for 40 years on average. The bargaining weight γ is 1/2. The exogenous separation rate σ is set to 2.61% per month.[5] It turns out that the implied endogenous separation rate is negligible, only .05% per month, leading to a total separation rate of 2.66% per month, consistent with the separation rate reported in Krusell et al. (2011) for prime-age males aged 21 to 65.[6]

Home production is parameterized as $b(z) = b_0 + b_1 z$. We set the slope parameter b_1 equal to 0.25. We think of this parameter as capturing unemployment benefits, which are proportional to wages and hence to z. This value of b_1 implies a replacement rate of approximately 25%. This value is consistent with Krusell et al. (forthcoming), who argue that after taking into account unemployed workers who are ineligible or choose not to take up benefits, the relevant replacement rate is 23%. We interpret the intercept b_0 as corresponding to the value of home production. As we discuss later, given b_1, we choose b_0 in order to match the employment rate of 63% in the United States and so find b_0 to be equal to 42% of the average output produced in a match. Taken together, b_0 and b_1 imply that the ratio of the mean of home production plus benefits to the mean output produced in a match is 48%. This figure is not far from the 40% used by Shimer (2005) and is in the lower end of the 47% to 96% range estimated by Chodorow-Reich and Karabarbounis (2016).

Importantly, in the robustness exercises of Section 5, we explore the sensitivity of our results to assuming home production proportional to z ($b_0 = 0$), constant home production ($b_1 = 0$), and no home production ($b_0 = b_1 = 0$). We show that the drop in employment after a credit contraction is not very sensitive to the specification of home production. These findings make clear that our results are not driven by the intuition that arises in the Hagedorn and Manovskii (2008) recalibration of the Shimer (2005) model: if consumers are essentially indifferent between working in the market and working at home, small shocks to productivity in the market generate large increases in nonemployment. Rather, the key idea in our model is that during a credit crunch, investing in employment relationships with surplus flows that

[5]This figure is lower than the 3.6% used by Shimer (2005) because Shimer includes job-to-job transitions, whereas we focus solely on employment-to-nonemployment transitions. We also experimented with a recalibration in which we used the higher separation rate from Shimer and found very similar results. As will become evident later, employment responses in our model are determined by the duration of the benefit flows from a match, which is primarily influenced by the amount of general human capital accumulation, rather than by the length of time a worker spends in any given match.

[6]We reproduced this separation rate using the Current Population Survey (CPS) data and the seasonal adjustments and classification error corrections for reported employment status adopted by Krusell et al. (2011), who used data from 1994 to 2007. We also updated the series using data from 1978 to 2012 and, using the same corrections, replicated the total separation rate of 2.8% per month estimated by Krusell et al. (forthcoming). We note that this updated number is close to the total separation rate of 2.66% from Krusell et al. (2011) and thus close to our number as well.

have long durations is not desirable. The duration of surplus flows is primarily determined by on-the-job human capital accumulation, and we find that this duration is not affected very much by home production.

Endogenously Chosen Parameters. We jointly choose the remaining parameters, $\vartheta = (b_0, \eta, \kappa, \sigma_z, \rho, \bar{z}_e, \bar{h})$, using the method of simulated moments by requiring that the model matches 11 moments as closely as possible. We can group these parameters into two sets: $\vartheta_1 = (b_0, \eta, \kappa, \sigma_z)$ and $\vartheta_2 = (\rho, \bar{z}_e, \bar{h})$. For a given set of parameters ϑ_2, the parameters in ϑ_1 are pinned down by the following four moments: an employment-population ratio of 0.63, corresponding to the 2006 CPS estimate for people aged 16 and older; a job-finding rate of 0.45 from Shimer (2005); a vacancy cost as a fraction of monthly match output of 0.15 from Hagedorn and Manovskii (2008); and a standard deviation of wage changes in the PSID of 0.21, as in Floden and Linde (2001).

We now provide some details behind some of these numbers. Consider the average job-finding rate. In the model we define it as the average job-finding rate among consumers in active markets, namely, $\bar{\lambda}_w = \int_{z^*} \lambda_w(z) \mathrm{d}u(z) / \int_{z^*} \mathrm{d}u(z)$. Given that we target an employment-population ratio of 63%, in our model 37% of consumers are not working. As panel B of Figure 1 shows, most nonemployed consumers are inactive in that they have a zero job-finding rate: only 4% of all consumers (or about 11% of nonemployed consumers) are active, and we require that the model produces an average job-finding rate of 45% for these active consumers. Heterogeneity in job-finding rates thus allows our model to simultaneously replicate the large number of nonemployed in the data and the modest flows out of nonemployment, since, on average, less than 5% (11% times 45%) of nonemployed become employed each month.

In terms of the vacancy cost, Hagedorn and Manovskii (2008) follow the insight in Pissarides (1992) in recognizing that this cost consists of both capital and labor components. For the capital cost, they assume that when firms post vacancies, they rent the same amount of capital as when they produce output, and find that the user cost of capital for posting a vacancy is about 12% of the monthly output of a worker. For the labor cost, they follow Barron, Berger, and Black (1997) and find that the labor cost of hiring a worker is 3% of the monthly output of a worker. Thus, the total cost of hiring a worker is approximately 15% of the monthly output of a worker. Note that Hagedorn and Manovskii (2008) measure these costs in a way that applies equally well to both our model and theirs.

The parameters in ϑ_2 determine how consumers accumulate human capital and are pinned down by moments related to wage growth on the job and over the life cycle. Since on-the-job human capital accumulation is key to the model's quantitative predictions, we next describe in detail how we quantify these parameters.

Parameters of Human Capital Processes. We quantify the parameters of the human capital processes using moments and estimates derived from two sources of data on high school graduates: cross-sectional data from Elsby and Shapiro (2012) on how wages increase with experience and longitudinal data from Buchinsky et al. (2010) on how wages grow over an employment spell. Later in our robustness analysis, we show that our results change little if we rely on alternative moments and estimates, including, for example, estimates for different education groups.

The key moment we use from the cross-sectional data is the 1.21 difference in log wages between workers with 30 years of experience and those just entering the labor market. This difference corresponds to an average increase in wages of 4.1% per year of experience. This moment is calculated based on census data also used by Elsby and Shapiro (2012) on full-time workers from census years 1970, 1980, and 1990.[7]

The five moments we use from the longitudinal data relate to how wages grow during an employment spell. These moments are calculated from the parameter estimates of the wage equation in Buchinsky et al. (2010) based on PSID data from 1975 to 1992 for high school graduates. The moments are the average annual growth rate of wages during an employment spell for workers with different levels of experience. As Table 1 shows, these growth rates are equal to 10%, 7%, 6%, 6%, and 7% for workers with 1 to 10, 11 to 20, 21 to 30, 31 to 40, and 1 to 40 years of experience, respectively. For brevity, we refer to these moments as *moments from the data*.

We now turn to explaining how Buchinsky et al. (2010) estimate their wage equation and how we use their parameter estimates to quantify the parameters of our human capital processes. Buchinsky et al. (2010) estimate a structural model of worker labor market participation, mobility, and wages that allows for rich sources of heterogeneity. In particular, they estimate the parameters of an equation that relates workers' wages to their demographic characteristics and history of past employment, as well as current experience (number of years in the labor market) and tenure (number of years with a given firm). The equation describing an individual i's wages at date t is

$$\log w_{it} = z_{it}'\beta + f(\text{experience}_{it}) + g(\text{tenure}_{it}) + \epsilon_{it}, \tag{23}$$

where z_{it} captures individual characteristics as well as the history of that individual's past employment. The functions $f(\cdot)$ and $g(\cdot)$ are fourth-order polynomials in experience and tenure.

[7]Following Elsby and Shapiro (2012), we use data on labor earnings of full-time workers working 35 or more hours per week for 50 or more weeks, and interpret their estimates as corresponding to wages in our model.

We use the parameter estimates from (23) as follows. For given values for ϑ, we simulate paths for wages, experience, and tenure for a panel of individuals. Given the simulated experience and tenure profiles from our model, we compute the *annualized* wage growth predicted by (23). The last five moments we use to calibrate our model are the predicted growth rates for the five experience groups described earlier. We choose parameter values so that wages in our model grow over an employment spell at the same rates as those implied by the wage estimates in Buchinsky et al. (2010) for workers with different levels of experience.

We show later in our robustness analysis that our model's implications for the labor market responses to a credit tightening are robust to a broad range of estimates on how wages grow over the life cycle, including those of Rubinstein and Weiss (2006), and a broad range of estimates on how wages grow on the job, including those from Altonji and Shakotko (1987) and Topel (1991).

Intuition for Identification. We now provide some intuition for how we separately identify the parameters $(\rho, \bar{z}_e, \bar{h})$ governing general and firm-specific human capital accumulation. While increases in $\bar{z}_e$ and $\bar{h}$ lead to similar increases in on-the-job wage growth, an increase in $\bar{z}_e$ leads to a much larger increase in life-cycle wage growth than does an increase in $\bar{h}$. The reason is twofold: firm-specific human capital is lost after each transition into nonemployment, whereas general human capital is not, and workers typically experience multiple employment and nonemployment spells over their lifetimes. As for the serial correlation parameter ρ, note that a decrease in ρ makes the wages of young workers grow faster than the wages of older workers. This parameter is then pinned down by the moments in Table 1, reflecting how annual wage growth during an employment spell varies with experience. Hence, the combination of the cross-sectional evidence in Elsby and Shapiro (2012) and the longitudinal evidence summarized by the parameter estimates in Buchinsky et al. (2010) jointly identifies these three parameters.

Recall that the cross-sectional data imply a life-cycle annual wage growth of 4.1%, whereas the longitudinal data imply an overall on-the-job wage growth of 7%. To understand how the model can simultaneously account for these two facts, consider Figure 2, which shows the path of wages of a typical worker in our model. Such a worker experiences several employment and nonemployment spells over the life cycle and, for simplicity, no shocks to human capital. Note that wages drop after each transition into nonemployment, owing to the loss of firm-specific human capital. Thus, even though wages rise relatively rapidly on the job, they rise less rapidly over the life cycle, because of the wage declines associated with transitions into nonemployment. As we discuss later, our model's implications for the loss in wages after a transition into nonemployment are in line with the data.

Parameter Values. Panels A and B of Table 1 summarize our parameterization strategy: it shows the moments used in our calibration, the parameters that we assign, and those that we endogenously determine. As Table 1 shows, the model exactly matches the first set of four moments that pin down the parameters in ϑ_1 and closely reproduces the remaining wage growth moments that pin down the parameters in ϑ_2. Panel B of Table 1 also reports the implied parameter values. In particular, the parameter governing the drift of general human capital accumulation, $\log \bar{z}_e$, is equal to 2.44, whereas that governing firm-specific human capital accumulation, $\log \bar{h}$, is equal to 0.82.

To see what our parameter estimates imply for the relative importance of the two types of human capital, consider an alternative version of the model in which firm-specific human capital is constant at $\bar{h} = 1$ and the other parameters are unchanged. The resulting model generates an average annual on-the-job wage growth of 5%, which is about 1/3 lower than that implied by our baseline model, but it implies a difference of 1.02 between the log wages of workers with 30 years of experience and those just entering the labor market, which is only slightly lower than the 1.19 difference implied by the baseline model. Taken together, these figures imply that firm-specific human capital accumulation accounts for about 1/3 of on-the-job wage growth but for only 16% of life-cycle wage growth.

Additional Model Implications. Here we discuss additional implications of our model. We start with implications for which we have counterparts in the data. Consider first the four moments in the top half of Table 2. Based on monthly PSID data between 1987 and 1997, we compute the mean log wage difference between the first wage after a nonemployment spell and the last wage before a nonemployment spell. We find that, on average, wages drop between 4.4% (all separations), like in Krolikowski (2017), and 5.5% (separations excluding quits) after a spell of nonemployment.[8] The corresponding number in our model, 5%, is in the middle of this range. To see how these wage declines vary with tenure in the last job before nonemployment, we ran a regression of log wage differences before and after a nonemployment spell on tenure in the last job. We estimate that an additional year of tenure increases the wage drop due to nonemployment by 1.5% in the PSID data (monthly sample) and by 1.9% in the model for workers with up to 25 years of tenure in the last job before nonemployment.

Next, the standard deviation of log wages in the cross section of consumers that start a new job is equal to 0.82 in our model and 0.85 in our PSID (monthly) sample. Finally, in the model, the profit share of revenue is 6%, which matches the mean value of U.S. after-tax

[8] We also computed this wage loss in the yearly PSID sample of Buchinsky et al. (2010) and found that the wage drop after a nonemployment spell of up to one year is remarkably similar at 5.7%. See the Appendix.

corporate profits to GDP over the period 1960-2010.[9]

Now consider the model's implications for the duration of nonemployment spells. In Figure 3, we plot the fraction of consumers whose nonemployment spells last k months. Overall, the model produces a pattern that is close to the one observed in the PSID (monthly) data. The greatest discrepancy is that the model overpredicts the fraction of spells lasting one month.[10]

We now turn to discussing how our model's implications line up with those from the literature on wage losses from job displacement. This literature attempts to measure the loss in wages that a displaced worker experiences relative to an otherwise identical worker who is not displaced. (Note that measured this way, the wage loss following displacement is large when the wage growth for an otherwise identical nondisplaced worker is large, even if the actual change in wage for the displaced worker is small.) We follow Huckfeldt (2017), who implements a version of the *displaced worker regression* of Jacobson, LaLonde, and Sullivan (1993) by estimating

$$\log w_{it} = \alpha_i + \gamma_t + \beta \text{experience}_{it} + \sum_{k \geq -2}^{10} \delta_k D_{it}^k + \varphi F_{it} + \varepsilon_{it} \tag{24}$$

on data simulated from our model. Here α_i is a person-specific fixed effect, γ_t is a time effect, D_{it}^k is a dummy variable that identifies a displaced worker in the k-th year after displacement, and F_{it} is a dummy variable for workers within their first 10 years of displacement ($k = 0, \ldots, 10$). Notice that wage losses are measured for workers in the two years before displacement ($k = -2, -1$), the year of displacement ($k = 0$), and the 10 years following displacement ($k = 1, \ldots, 10$). In Figure 4, we report the sum $\delta_k + \varphi I(k \geq 0)$ for $k = 0, \ldots, 10$, which we interpret as the wage losses for displaced workers relative to nondisplaced workers. For the year of displacement, our model produces wage losses relative to nondisplaced workers that are similar to those reported in Huckfeldt (2017), both of which are about 6%. For later years, our numbers are about half as large as those in Huckfeldt (2017).

We do not view the underprediction of longer-term losses implied by the displaced worker regression in (24) as evidence against our mechanism. In the robustness exercises of Section 5, we show that the *high* $\bar{h}$ model, which features a higher rate of firm-specific human capital accumulation than our baseline model, matches well the estimated wage losses due to displacement and yields employment responses to a credit crunch that are essentially identical to those implied by our baseline model. See Figure 11 and panel A of Figure 12, discussed in

[9]See the U.S. Bureau of Economic Analysis, Corporate Profits After Tax (without IVA and CCAdj) [CP], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CP.

[10]This difference arises in large part because we target a job-finding rate of 45% per month used by Shimer (2005) and popular in the quantitative search literature. In the PSID sample for which we could measure both nonemployment durations and wage losses due to nonemployment at monthly frequencies, this job-finding rate is somewhat lower, about 38%.

Section 5. As Tables 4 and 5 show, the main drawback of the high $\bar{h}$ model is that it overstates both the average wage drop after a nonemployment spell and the amount of on-the-job wage growth relative to that implied by the estimates in Buchinsky et al. (2010). Note that we could have used the high $\bar{h}$ model as the baseline model with almost no effect on our results on the amplification of credit shocks. We prefer our current baseline model because the moments that inform it are at the core of much of the labor literature on the dynamics of individual wages and also pertain to a much broader cross section of the population than the particular group of workers underlying displaced worker regressions.

DMP Parameterization. Below we isolate the role of on-the-job human capital accumulation by comparing the implications of our model with those of an alternative model with no growth in either general or firm-specific human capital ($\bar{z}_e = 1$ and $\bar{h} = 1$), referred to as the *DMP model*.[11] We choose the four parameters in ϑ_1 in this version of the model to ensure that the model exactly reproduces the first four moments in panel A of Table 1.

4 Employment Response to a Credit Tightening

We illustrate how employment responds to a credit tightening and then explain the role of human capital accumulation in amplifying the employment response relative to the DMP model.

As we have argued earlier, a credit crunch reduces employment in our model because it reduces consumption and increases the family's marginal utility of consumption, making it more costly to invest resources to post vacancies. For a given drop in consumption, the lower the elasticity of intertemporal substitution, the greater the increase in the marginal utility of consumption. We maintain that the utility function is $u(c_t) = c_t^{1-\alpha}/(1-\alpha)$ and consider different values for α that reflect the range of values commonly used in the literature. Note that accounting for the low and stable short-term real interest rates in the data requires a relatively large elasticity of intertemporal substitution (EIS), around 1, a typical choice in the business cycle literature. On the other hand, estimates based on household-level data point to much lower values, in the neighborhood of 0.1. (See Guvenen (2006), Gomes and Michaelides (2005), and Best et al. (2017), and the references therein.) We therefore report results for values of α in this entire range and show that the amplification of the employment response due to human capital accumulation is essentially invariant to α, even though the level of α affects the absolute magnitude of the response.

[11] This version is not exactly the DMP model because even though it features no growth in general human capital, z, this version includes shocks to it. We found that the exact DMP model in which all consumers have the same fixed level of general human capital gave very similar results.

4.1 Impulse Responses to a Tightening of Credit

To build intuition for how the model works, we conduct an experiment in which we assume an unanticipated tightening of the debt constraint of a family. We describe the response of employment, decompose the drop in employment into the components arising from changes in the overall job-finding rate and separation rate, and then discuss why the drop in employment is persistent.

The Response of Employment and Consumption. We assume that the credit tightening is unexpected prior to the first period and that consumers have perfect foresight afterward. We choose a sequence of debt limits $\{\chi_t\}$ so that consumption drops by 5% on impact and then mean-reverts at a rate of 10% per quarter (so that $\Delta \log c_t = \varrho^t e$ where $\varrho = 0.9$ and $e = 5\%$). This mean-reversion rate is chosen so as to match the speed of postwar consumption recoveries in the data. Correspondingly, the shadow price Q_t increases proportionately to the drop in consumption and mean-reverts to its steady-state level at the same rate as consumption.

Figure 5 displays the path of employment in the baseline and DMP versions of the model for various values of α, equal to 1, 5, and 10. We also consider a version of our model with preferences given by $u(c) = \log(c - \bar{c})$, where $\bar{c}$ is the subsistence level of consumption, or *permanent habit*. We set $\bar{c}$ equal to 80% of the steady-state level of consumption, implying that the EIS is equal to $(c - \bar{c})/c = 1/5$ at the steady state but increases as the level of consumption, c, increases. With such preferences, the model can help reconcile the low EIS estimates for nonstockholders and the high EIS estimates for stockholders, assuming the latter have a greater level of consumption. (See Guvenen (2006) and the references therein.)

Figure 5 shows our main result: regardless of the value of α, the addition of human capital accumulation magnifies the drop in employment from a credit tightening by approximately a factor of 10. Of course, when α is higher, the desire to smooth consumption is stronger, the increase in the shadow price of goods is larger, and, hence, the employment drop is proportionately larger in both our baseline model and the DMP model. Thus, in both models, a higher α effectively scales the size of the shock to the shadow discount rate. For concreteness, from now on we report results only for an intermediate value of $\alpha = 5$ and note that economies with other values for α produce employment responses that are approximately scaled versions of those in the economy with $\alpha = 5$.

Note also that the drop in employment is much more persistent in our baseline model relative to the DMP model. Indeed, as we discuss later, the employment response in our baseline model has double the half-life of the response in the DMP model.

The Shimer Decomposition of Employment. We next shed light on the mechanism behind the impulse response for employment in our baseline model and the DMP model. The drop in employment is due to the combination of the drop in the job-finding rate and the increase in the endogenous component of the separation rate. We show that the first effect is by far the most important one quantitatively.

As already discussed, the job-finding rate drops because the mean-reverting increase in the shadow price of goods increases the cost of posting new vacancies relative to their benefits. To understand the increase in the separation rate, note that some consumers experience negative shocks to their human capital, which lead to a negative current flow surplus $z - b(z)$. For such consumers, there is a cutoff level of z such that it is optimal for a firm to wait for the productivity of these matches to improve above that level. During a credit crunch, since the shadow price of goods increases, the initial negative flows from such consumers are more costly to firms. Hence, the cutoff level of z for which such matches are profitable increases and more matches are dissolved.

To quantify the importance of these two factors in accounting for the employment drop after a credit tightening, we build on the approach in Shimer (2012). The transition law for total employment can be written as $e_{t+1} = (1 - s_t)e_t + f_t(1 - e_t)$, where e_t is the overall employment rate; s_t is the overall separation rate, equal to the sum of exogenous and endogenous separations; and f_t is the overall job-finding rate. We construct two counterfactual employment series in which we vary s_t and f_t, one at a time, while leaving the other at its steady-state value.

In Figure 6, we scale each of these counterfactual series by the maximal drop in employment in each model. In both our model and the DMP model, we see that the portion of the employment drop due to the change in the separation rate accounts for only a small fraction—about 10%—of the initial drop in employment. The increase in the separation rate is modest on impact because the measure of employed consumers close to the cutoffs is simply too small to have much of an effect on the separation rate, as panel A of Figure 1 makes clear.

The Persistent Drop in Employment. The employment drop in our baseline model is more persistent than that in the DMP model. One way to quantify this persistence is to use an analog to a half-life measure, namely, the number of months employment takes to recover to half of its maximal drop. In the DMP model, employment takes 29 months to recover from its trough of -0.32% to -0.16%. In contrast, in the baseline model, employment takes 57 months to recover from its trough of -2.67% to -1.33%.

Interestingly, our baseline model also generates endogenous persistence in that the drop in employment lasts much longer than the credit crunch itself. Indeed, the shadow price Q_t takes 20 months to decay to half of its maximal increase, whereas, as mentioned, employment takes 57 months to recover to half of its maximal drop.

Two forces generate this additional persistence. The quantitatively important force is from *forgone human capital accumulation*. Specifically, during the period in which the average job-finding rate falls, consumers spend less time employed and thus acquire less human capital on the job relative to what they acquire in the steady state. Several years after the credit tightening, the average skill levels of both employed and nonemployed consumers are thus lower than they were in the steady state. Hence, even after the shadow price Q_t has returned to its steady state, firms post fewer vacancies because matching with less productive consumers is less profitable.

The second force is the *hysteresis* effect due to the depreciation of human capital during nonemployment—in contrast to the forgone appreciation during employment. This depreciation occurs because nonemployed human capital drifts to a mean of $\bar{z}_u$, which is less than $\bar{z}_e$. This effect, however, is quantitatively small in our baseline model because the rate at which human capital depreciates is very low. In particular, the average drop in z during a month of nonemployment is only about 0.7%.[12] Since the average nonemployment spell lasts only two months, the drop in z during a typical nonemployment spell is only 1.4%.

We show in Section 5 that this hysteresis effect does not play an important role in amplifying the employment drop after a credit tightening in our baseline model. To do so, there we compare the response of employment in the baseline model to that in an economy with no decay in human capital during nonemployment, and hence no hysteresis, and find that employment responses are about as persistent as they are in our baseline model.

4.2 Intuition for the Amplification Effect of Human Capital

The key new feature we have introduced in an otherwise standard search and matching model is human capital accumulation on the job. Here we provide some intuition for how this feature amplifies the employment response to a credit contraction relative to that generated by the DMP model. In a nutshell, in our model with human capital accumulation, the expected flows of surplus from a match decay slowly over time. Since our mean-reverting shock implies that the initial increase in shadow prices decays over time, the present value of these benefits flows from a match increases by much less than the cost, as (16) makes clear. As a result, firms post many fewer vacancies and employment falls a lot. In contrast, in the DMP model, which features no such human capital accumulation, the high job-finding rates imply that

[12]More precisely, since $\rho = .996$, and the mean $\log z$ for the employed is 1.78, the mean percentage change in z during a month of nonemployment is $(\rho - 1)\overline{\log z} = -.004(1.78) = -0.7\%$.

the expected surplus flows from a match decay quickly over time. Thus, the present value of these benefit flows from a match increases by nearly as much as the cost. Hence, in the DMP model, in response to a credit tightening, firms post only slightly fewer vacancies and employment falls a little.

We make this point formal by building on the argument developed in Section 1.6. There we showed that the employment response to a credit crunch is larger, the greater the duration of surplus flows from forming a match. Here we show that adding human capital accumulation considerably raises these durations relative to those implied by the DMP model. Specifically, in our quantitative version of the DMP model, both the Macaulay and alternative Macaulay durations are very short, equal to 2.8 and 2.6 months, respectively. Thus, the vast bulk of the surplus flows from a match accrues very soon after a match is formed. In contrast, in our baseline model, these durations are much longer, equal to almost 120 and 24 months, respectively. Thus, the job-finding rate responds much more in our baseline model than it does in the DMP model.

To explain transparently why the DMP model implies such short durations of benefit flows from a match, we set $\bar{h} = 1$, so there is no firm-specific human capital accumulation, and $F_e(z'|z)$ equal to $F_u(z'|z)$, so there is no general human capital accumulation. When we do so, we can write the expression for match surplus recursively as

$$S_t(z) = Q_t[z - b(z)] + [1 - \sigma - \gamma\lambda_{wt}(z)]\,\phi \int \max\left[S_{t+1}(z'), 0\right] \mathrm{d}F(z'|z), \tag{25}$$

where $F(z'|z)$ is the common law of motion of human capital for employed and nonemployed consumers. This expression, familiar from the DMP model, is simply the discounted sum of the difference between the output of a consumer when employed, z, and when not employed, $b(z)$. In the DMP model, most of the present value of the return from a match accrues very early in the match because, as the second term in (25) makes clear, the effective decay rate of these flows, $\sigma + \gamma\lambda_{wt}(z)$, is high due to the relatively high average job-finding probability in the data. Intuitively, the fact that this probability is high implies that an employed consumer produces more output than a nonemployed consumer for only a few months on average because the nonemployed consumer quickly finds a job. Hence, the decay rate of the expected benefit flows is very high.

Consider next why our model with human capital accumulation implies a much slower decay rate of such benefit flows and thus a higher duration. In this case, match surplus can be written recursively as

$$S_t(z,h) = Q_t\left[zh - b\left(z\right)\right] + \left[1-\sigma-\gamma\lambda_{wt}\left(z\right)\right]\phi\int \max\left[S_{t+1}\left(z',h'\right),0\right]\mathrm{d}F_e\left(z'|z\right) \tag{26}$$
$$+\,\phi\int U_{t+1}\left(z'\right)\left[\mathrm{d}F_e\left(z'|z\right)-\mathrm{d}F_u\left(z'|z\right)\right]$$
$$+\,\gamma\lambda_{wt}\left(z\right)\phi\left[\int \max\left[S_{t+1}\left(z',h'\right),0\right]\mathrm{d}F_e\left(z'|z\right)-\int \max\left[S_{t+1}\left(z',1\right),0\right]\mathrm{d}F_u\left(z'|z\right)\right],$$

where the value of a nonemployed consumer is

$$U_t\left(z\right) = Q_t b\left(z\right) + \gamma\lambda_{wt}\left(z\right)\phi\int \max\left[S_{t+1}\left(z',1\right),0\right]\mathrm{d}F_u\left(z'|z\right) + \phi\int U_{t+1}\left(z'\right)\mathrm{d}F_u\left(z'|z\right). \tag{27}$$

Equation (26) shows that the surplus from a match combines three components. The first component, given by the right side in the first line of (26), is analogous to the DMP component just discussed. The second component is due to general human capital accumulation. Since general human capital grows faster when the consumer is employed, as the second and third lines of (26) show, the surplus includes values weighted by the laws of motion of general human capital for employed and nonemployed consumers, $\mathrm{d}F_e(z'|z)$ and $\mathrm{d}F_u(z'|z)$. The third component is due to the difference in firm-specific human capital between continuing matches, h', and new matches, $h=1$, and appears in the third line of (26) in the sense that even if $\mathrm{d}F_e = \mathrm{d}F_u$, the terms in brackets in this line would not be zero.

Since general human capital is transferable across matches, the general human capital acquired in a current match will be used by a consumer in all future matches. Hence, this component of match surplus decays slowly over time. To see that it is general human capital and not firm-specific human capital that increases the duration of surplus flows the most, note that eliminating general human capital accumulation by setting $F_e(z'|z) = F_u(z'|z)$ reduces the alternative Macaulay duration by a large amount, from 24 months in our baseline model to 7 months. In contrast, eliminating firm-specific human capital accumulation by setting $\bar{h} = 1$ reduces the alternative Macaulay duration only a little, from 24 months in our baseline model to 23 months.

At a technical level, (26) and (27) define a nonlinear vector dynamical system that is quite involved because of uncertainty about the evolution of z and allows us only to provide broad intuition about the forces determining the duration of surplus flows. We next consider a simplified version of our model without shocks to human capital, for which this system reduces to a second-order linear difference equation in S, so that we can derive closed-form expressions that allow us to sharpen the intuition developed here.

4.3 Simple Examples

Here we study simplified versions of our baseline model. In our *general human capital* model, we assume that general human capital grows at a constant rate g when the consumer is employed but we abstract from firm-specific human capital. In our *firm-specific human capital* model, we assume that firm-specific human capital grows at a constant rate g_h when the consumer is employed but we abstract from general human capital.

We derive closed-form expressions for the surplus flows accruing at each date and show that the implied alternative Macaulay durations are much greater in the general human capital model than they are in the firm-specific human capital model. The key force driving this difference is that in the firm-specific human capital model, all the human capital accumulated during a match is lost upon its dissolution, whereas in the general human capital model, the human capital accumulated during the match is transferred to all subsequent matches. For both models, we write surplus in the form

$$S = \sum_{k=0}^{\infty} \beta^k s_{k+1}, \tag{28}$$

where, for brevity, we slightly abuse notation and denote the expected surplus flow $\mathbb{E}_0 s_{k+1}$ as simply s_{k+1} and, to keep the algebra simple, we express (13) in date $t+1$ rather than date 0 goods.[13] Given closed-form solutions for s_{k+1}, we can compute closed-form solutions for durations. For simplicity, we let the cost of posting vacancies and home production be proportional to human capital so that job- and worker-finding rates are independent of human capital, as we will show.

4.3.1 General Human Capital

Suppose that employed consumers' human capital grows deterministically at rate $g > 0$, nonemployed consumers' human capital stays constant at its level at the end of the last match, and $h = 1$. For convenience, let β denote the effective discount factor, namely, the product of a family's discount factor and the consumer survival probability, ϕ. The value functions of employed consumers, nonemployed consumers, and firms, in units of current consumption goods, satisfy

$$W(z) = w(z) + \beta[(1-\sigma) W((1+g) z) + \sigma U((1+g) z)], \tag{29}$$

$$U(z) = bz + \beta[\lambda_w W(z) + (1-\lambda_w) U(z)], \tag{30}$$

$$J(z) = z - w(z) + \beta(1-\sigma) J((1+g) z), \tag{31}$$

[13] To obtain (28) from (13), divide by Q_t and evaluate Q_{t+k}/Q_t at its steady-state value of β^k.

where $w(z)$ is the wage received by the consumer and bz is home production, assumed proportional to z. We can then express match surplus as

$$S(z) = \frac{1 - b + \frac{\beta g}{1-\beta} b}{1 - \beta \left[(1-\sigma)(1+g) - \gamma \lambda_w + \frac{\beta g}{1-\beta} \gamma \lambda_w \right]} z = Sz, \tag{32}$$

with the constant S defined by the second equality in (32). The free-entry condition is

$$\kappa z = \beta \lambda_f (1-\gamma) S z, \tag{33}$$

where we have imposed that vacancy costs κz are proportional to z. Dividing both sides of (33) by z proves that the firm's worker-finding rate is independent of z and, hence, so is the consumer's job-finding rate.

We consider a special case of this model, namely the DMP model, with $g = 0$. To derive the representation in (28), we manipulate (29)–(31) to obtain the first-order difference equation

$$S_t = 1 - b + \beta \delta_{DMP} S_{t+1},$$

which, iterating forward, gives (28) with $s_{k+1} = (\delta_{DMP})^k (1-b)$, where $\delta_{DMP} = 1 - \sigma - \gamma\lambda_w$ is the *DMP root*. The weight ω_{k+1} used in the duration calculation is the fraction of the surplus accrued in the $k+1$-th period of the match, namely, $\omega_{k+1} = \beta^k s_{k+1}/S = (\beta\delta_{DMP})^k / \sum_{k=0}^{\infty} (\beta\delta_{DMP})^k$ where $S = (1-b)/(1-\beta\delta_{DMP})$. The Macaulay duration of these flows is

$$\sum_{k=0}^{\infty} (k+1) \omega_{k+1} = \frac{1}{1 - \beta\delta_{DMP}},$$

whereas the alternative Macaulay duration, defined in (17), is

$$\sum_{k=0}^{\infty} \frac{1 - \varrho^{k+1}}{1 - \varrho} \omega_{k+1} = \frac{1}{1 - \beta\varrho\delta_{DMP}}.$$

When $g > 0$, we can also express the surplus as the discounted value of expected surplus flows as in (28). Here, though, the presence of general human capital turns the relevant difference equation for the surplus flows into a second-order difference equation with two roots, so that the surplus flow has the form

$$s_{k+1} = c_s \delta_s^k + c_\ell \delta_\ell^k, \tag{34}$$

where δ_s denotes the *small root* and δ_ℓ the *large root*. We establish the following proposition in the Appendix.

Proposition 3. *In the general human capital economy with $g > 0$, the surplus flow has the form in (28) with roots*

$$\delta_s = \alpha - \frac{1}{2} \left(\sqrt{(1-\alpha)^2 + 4\gamma\lambda_w g} - \sqrt{(1-\alpha)^2} \right) < 1, \tag{35}$$

$$\delta_\ell = \alpha + \frac{1}{2}\left(\sqrt{(1-\alpha)^2 + 4\gamma\lambda_w g} - \sqrt{(1-\alpha)^2}\right) > 1, \tag{36}$$

a Macaulay duration of

$$\frac{1}{S}\left[\frac{c_s\beta\delta_s}{(1-\beta\delta_s)^2} + \frac{c_\ell\beta\delta_\ell}{(1-\beta\delta_\ell)^2}\right], \tag{37}$$

and a alternative Macaulay duration of

$$\frac{1}{1-\varrho}\left[1 - \frac{1}{S}\frac{c_s\varrho}{1-\beta\varrho\delta_s} - \frac{1}{S}\frac{c_\ell\varrho}{1-\beta\varrho\delta_\ell}\right], \tag{38}$$

and implies a surplus of

$$S = \frac{c_s}{1-\beta\delta_s} + \frac{c_\ell}{1-\beta\delta_\ell}, \tag{39}$$

where $\alpha = (1-\sigma)(1+g) - \gamma\lambda_w < 1$, $c_\ell = [(1-b)(\alpha-\delta_s) + bg]/(\delta_\ell - \delta_s)$, *and* $c_s = 1 - b - c_\ell$.

To gain some intuition about the size of the small and large roots, note that a first-order Taylor approximation around $g = 0$ yields

$$\delta_s \approx (1-\sigma-\gamma\lambda_w)\left(1 + \frac{\sigma}{\sigma+\gamma\lambda_w}g\right) \text{ and } \delta_\ell \approx 1 + \frac{\gamma\lambda_w}{\sigma+\gamma\lambda_w}g. \tag{40}$$

Now, at our parameter values, $\gamma\lambda_w$ is an order of magnitude larger than σ, so $\delta_s \approx \delta_{DMP}$ and $\delta_\ell \approx 1 + g$, which is actually very accurate in the neighborhood of our parameters.[14]

Figure 7 shows how these two durations measures vary with the growth rate of general human capital. Since the small root is approximately equal to the DMP root and that root, by itself, produces flows with short durations, the large root indeed accounts for the high duration of the surplus flows. Formally, the presence of general human capital adds a persistent component to the standard DMP surplus flows, and this component generates a much longer duration of flows. Technically, this result arises because the presence of general human capital makes the system governing the evolution of surplus flows a second-order system with one root that is larger than one. This system imparts much greater persistence to the surplus flows relative to the DMP model, which, as discussed earlier, follows a first-order system with a small root that dies out quickly.[15]

[14] Intuitively, this approximation is very accurate when $\sigma/\gamma\lambda_w$ is close to zero. For our parameters and, say, $g = .3\%$ per month, exact solutions are $r_s = .7491$ and $r_\ell = 1.0027$; (40) gives $r_s = .7489$ and $r_\ell = 1.003$.

[15] Note that in our quantitative model, the durations are concave in the growth rate of human capital, whereas in these simple examples, they are not. In our quantitative model, this concavity arises because the growth in human capital tapers off with both tenure and experience. Observe also that when $g = 0$, these formulas reduce to the DMP case studied above: $\delta_s = \delta_{DMP} = \alpha$ and, thus, c_ℓ equals zero.

4.3.2 Firm-Specific Human Capital

Suppose now that consumers only accumulate firm-specific human capital and $z = 1$. Firm-specific human capital equals $h_0 = 1$ for all newly hired consumers, grows at rate $1+g_h$ during an employment spell, and is lost upon separation, so that the evolution of firm-specific human capital is $h_{t+1} = (1+g_h)h_t$ for employed consumers and $h_{t+1} = h_0$ for not employed consumers at t. As in the general human capital case, β denotes the effective discount factor, namely, the product of a family's discount factor and the consumer survival probability, ϕ. Since all nonemployed consumers are identical, home production and vacancy costs are simply denoted by b and κ.

The value functions of employed consumers, nonemployed consumers, and firms satisfy

$$\begin{aligned} W(h) &= w(h) + \beta\left[(1-\sigma)\, W((1+g_h)h) + \sigma U\right], \\ U &= b + \beta\left[\lambda_w W(h_0) + (1-\lambda_w)U\right], \\ J\,(h) &= h - w(h) + \beta(1-\sigma)J((1+g_h)h), \end{aligned}$$

where we have expressed these values in units of current consumption goods. Since all non-employed consumers have firm-specific human capital h_0, the relevant surplus for vacancy creation is

$$S(h_0) = \frac{1}{1-\beta(1-\sigma-\gamma\lambda_w)}\left[\frac{1-\beta(1-\sigma)}{1-\beta(1-\sigma)(1+g_h)}h_0 - b\right]. \tag{41}$$

We can decompose the surplus (41) into flows accruing at each date k after a match is formed and write it in the same form as in the general human capital model, namely

$$s_{k+1} = c_s\delta_s^k + c_\ell\delta_\ell^k. \tag{42}$$

We establish the following proposition in the Appendix.

Proposition 4. *In the firm-specific human capital economy, the surplus flow has the form in (42) with roots $\delta_s = \delta_{DMP}$ and $\delta_\ell = (1-\sigma)(1+g_h)$ and constants $c_s = 1-b-c_\ell$ and $c_\ell = (1-\sigma)g_h/\left[(1-\sigma)g_h + \gamma\lambda_w\right]$. The Macaulay duration, the alternative Macaulay duration, and the surplus have the forms in (37), (38), and (39).*

The presence of firm-specific human capital also adds a second root to the surplus equation in addition to the DMP root, which, as discussed, converges quickly to 0. The key difference between the firm-specific and general human capital case is that the large root is smaller than one for reasonable parameterizations in the firm-specific human capital case, rather than larger than one, as in the general human capital case. To see why $\delta_\ell = (1-\sigma)(1+g_h)$ is less than one, recall that $\sigma = 2.6\%$ per month, whereas, for any reasonable estimate of wage growth per month, the monthly growth rate of firm-specific human capital g_h is much lower than that. In sum, the presence of firm-specific human capital does not increase the duration of surplus flows very much relative to the DMP model.

From Figure 7 we see that the Macaulay duration and the alternative Macaulay duration increase with the growth rate of firm-specific human capital, g_h, at a much smaller rate than they do in the general human capital model. The key insight here is that since firm-specific human capital acquired in a match is destroyed upon separation, rather than being transferred to subsequent matches, reasonable separation rates imply a low duration of resulting surplus flows.

5 Robustness

Here we discuss the robustness of our results to a range of parameter values. We focus mostly on the human capital processes, which are the key ingredients of our model. We also investigate robustness with respect to the parameterization of home production and vacancy costs.

5.1 Human Capital Accumulation

We start by showing that our results are robust to alternative estimates of wage growth. We then show that our results are not driven by hysteresis resulting from the decay in human capital during nonemployment.

5.1.1 Alternative Estimates of Wage Growth

Consider first how wages increase with experience in the cross section. Recall that Elsby and Shapiro (2012) document that the wage differential between workers who have high school degrees with 30 years of experience and those entering the labor market is 1.21 log points. In contrast, Rubinstein and Weiss (2006) report a range of values for this wage differential centered on 0.80 log points. In the two robustness experiments that follow, we target this lower wage differential.

Consider next the estimates of on-the-job wage growth. Although the estimates of the wage equation in (23) from Buchinsky et al. (2010) imply an average on-the-job wage growth of 7% per year, the earlier literature finds smaller numbers. Accordingly, we consider two alternative parameterizations derived from the estimated wage equations in the influential studies of Altonji and Shakotko (1987) and Topel (1991).

As earlier, for a given parameterization of our model, we simulate paths for wages, experience, and tenure for a panel of individuals. Given the simulated experience and tenure profiles from our model, we compute the annualized wage growth predicted by the estimated wage equations in Altonji and Shakotko (1987) and Topel (1991). We then choose parameters in our model to ensure that our model's implications for wage growth at different levels of experience are consistent with the implications of these estimates, which we refer to as moments from the data as before.

Table 3A shows the moments from the data and our model for these alternative sets of estimates. For simplicity, we refer to the first robustness experiment that uses estimates from Altonji and Shakotko (1987) as the *Altonji-Shakotko* experiment, and the experiment that uses estimates from Topel (1991) as the *Topel* experiment. The first five moments are common to both experiments. In particular, we use the cross-sectional log-wage differential from Rubinstein and Weiss (2006, Figure 3a, page 9), as well as the employment rate, job-finding rate, vacancy posting cost, and standard deviation of wage changes from our baseline model. Notice that we match all of these moments almost perfectly. Importantly, on-the-job wage growth rates are now substantially smaller than in our baseline model. For example, the average on-the-job wage growth for workers with 1 to 40 years of experience is 7% in our baseline model but only 3% in the Altonji-Shakotko experiment and 5% in the Topel experiment.

Table 3B shows the resulting parameter values and compares them with those in the baseline model. Notice that the parameters governing the rate of human capital accumulation $\bar{z}_e$ and $\bar{h}$ are substantially smaller in these two robustness experiments than in our baseline model. Nevertheless, both of these alternatives imply a sizable amount of general human capital accumulation, as indicated by $\log \bar{z}_e$ being substantially greater than 0. In these experiments, all other parameters are suitably adjusted so that the model replicates the remaining targeted moments.

Figure 8 shows how employment responds in the Altonji-Shakotko and Topel experiments to a credit tightening that leads to the same 5% drop in consumption as in our baseline model. Note that the maximal drop in employment is slightly larger than in our baseline model: it is 2.7% in our baseline and about 3% in both of these alternative experiments. Also note that both of these alternative experiments predict that the employment drop is at least as persistent as in the baseline model. In particular, the cumulative impulse response of employment relative to that of consumption after four years is 0.91 in the Altonji-Shakotko experiment and 0.98 in the Topel experiment, which are both larger than the corresponding 0.9 statistic in the baseline model.

5.1.2 Rates of Human Capital Accumulation when Employed

Here we show that as long as we choose the drift in general human capital so that it generates life-cycle wage growth of at least 1% per year, the model's implications for employment are virtually identical.

In particular, in Figure 9 we vary the parameters governing the rate of general human capital accumulation, $\log \bar{z}_e$, and of firm-specific human capital accumulation, $\log \bar{h}$, to levels well outside those typically estimated and report the model's implications for wage growth and the response of employment to a credit tightening. As earlier, when we vary these parameters, all other parameters are suitably adjusted so that the model replicates all other statistics used in our baseline parameterization.

Panels A and B of Figure 9 report the model's implications for wage growth on the job and over the life cycle in these experiments. The lines marked with triangles set $\bar{h}$ equal to our baseline value, and the lines marked with circles eliminate all firm-specific human capital accumulation by setting $\bar{h}$ equal to 1. As we vary $\bar{z}_e$, the model produces a wide range of rates of wage growth on the job and over the life cycle. As panel B makes clear, the cross-sectional wage-experience profiles vary little with the amount of firm-specific human capital accumulation.

Consider next the models' implications for the employment response to a credit tightening. We summarize these implications in panels C and D with two statistics: the maximal employment drop after a 5% drop in consumption and the cumulative drop in employment relative to consumption in the first four years after the credit tightening.

Notice from these figures that as long as the drift parameter of general human capital, $\log \bar{z}_e$, is greater than 1, our model's key predictions for employment responses are remarkably similar. In particular, the maximal employment drop is about 2.7% to 3%, whereas the cumulative drop is about 0.9 that of consumption in the first four years after the credit tightening. Indeed, Figure 9 shows that we would have obtained similar employment responses even if we had set $\log \bar{z}_e$ slightly less than 1 and $\log \bar{h}$ to zero, which imply rates of general and firm-specific human capital acquisition that yield an on-the-job wage growth as low as 1.5% per year (panel A) and a wage differential of 0.3 log points between workers with 30 years of experience and those with no experience (panel B). For these parameter values, life-cycle wage growth is equal to $1\% = 0.3/30$ per year.

We next provide some intuition for why the employment responses are nonlinear in the rate of general human capital accumulation. Part of this intuition can be obtained by realizing that the duration of benefit flows is nonlinear in the rate of human capital accumulation. Panels A and B of Figure 10 show how the Macaulay duration and the alternative Macaulay duration vary as we vary the drift parameter $\log \bar{z}_e$. This figure makes clear that both durations are highly concave in the drift parameter: as this parameter increases, the marginal effect on duration decreases. These durations asymptote to about 130 months and 24 months, respectively.

In our baseline model, we focused on estimates of life-cycle and on-the-job wage growth for high school graduates. We also repeated our exercise for analogous estimates from Elsby and Shapiro (2012) and Buchinsky et al. (2010) for individuals with less than a high school degree and those with a college degree. Both sets of estimates imply life-cycle wage growth well above our 1% cutoff, and hence, from Figure 9, panels C and D, lead to maximal employment drops and cumulative employment drops that are similar to those implied by our baseline model.

Consider finally the implications of increasing the rate of firm-specific human capital accumulation for the wage losses after displacement. In the *high* $\bar{h}$ model, we double $\log \bar{h}$ relative to our baseline level. Figure 11 shows that such a model generates a pattern of wage losses from displacement that are much closer to those implied by the *displaced worker* regression of Huckfeldt (2017) than those generated by our baseline model. Panel A of Figure 12 further shows that this model generates impulse responses to a credit crunch that are essentially identical to those generated by our baseline model. The main drawback of the high $\bar{h}$ model is that it generates on-the-job wage growth of 9% per year, which is greater than the 7% implied by the data—see the last column in Table 4.

5.1.3 Rates of Decay in Human Capital When Nonemployed

So far we have imposed that the serial correlation of the processes for general human capital accumulation during employment and nonemployment are equal. Here we allow these serial correlations to differ and denote them by ρ_e and ρ_u, respectively, for the rates during employment and nonemployment.

We first demonstrate that hysteresis does not play an important role in our baseline model. To see this, we compare the response of employment in the baseline model to a *low decay* (*high* ρ_u) model in which human capital does not depreciate during nonemployment and hence, by construction, does not give rise to hysteresis. Specifically, we set the persistence of the human capital process during nonemployment, ρ_u, to 1, keeping the persistence of the human capital process during employment, ρ_e, at its baseline value. Here, as well as in all robustness exercises that follow, we recalibrate all parameters in order to match the same set of statistics used in our baseline parameterization.

Panel B of Figure 12 shows that the employment drop in the low decay model is larger initially but somewhat less persistent than in our baseline model. In the low decay model, the half-life of the employment drop is equal to 47 months, compared to 57 months in our baseline model. Hence, hysteresis plays a minimal role in generating either the magnitude or the persistence of the employment drop in our baseline model.

Of course, we can increase the importance of the hysteresis effect by increasing the decay rate of general human capital during nonemployment. To illustrate this point, we consider a *high decay* (*low* ρ_u) model in which we set the persistence of the human capital process during nonemployment, ρ_u, equal to 0.95 while keeping the persistence of the human capital process during employment, ρ_e, at its baseline value. In this case, the average drop in productivity during a month of nonemployment is nearly 9% as opposed to 0.7% per month as in our baseline model.[16] Panel B of Figure 12 further shows that the employment drop is larger and more persistent in the high decay model than in our baseline model. In the high decay model, employment reaches half of its maximal drop after 147 months, whereas it reaches half of its maximal drop after only 57 months in our baseline model.

Two additional points are worth noting about this high decay model. First, as Table 5 shows, the high decay (low ρ_u) model implies an average wage drop of 9% upon displacement which is much higher than the 5% drop in the data. Second, as Figure 11 shows, the implications for the displaced worker regression in the high decay model are nearly identical to those in our baseline model in Figure 4.

5.2 Home Production and Vacancy Costs

Here we discuss the sensitivity of our results to the specification of home production and vacancy costs.

Home Production. Consider first the specification of home production. In our baseline model, a nonemployed consumer produces $b_0 + b_1 z$. We consider three extremes: in the *constant b* model, we set $b_1 = 0$ so that home production is independent of z; in the *proportional b* model, we set $b_0 = 0$ so that home production is proportional to z; and in the *no home production* model, we set $b_0 = b_1 = 0$ so that consumers do not produce when nonemployed. In all three exercises, we recalibrate the rest of the parameters to match the targeted moments in the data from panel A of Table 1, including an employment-to-population ratio of 0.63. One exception is that when $b_0 = 0$, the job-finding rate is essentially constant in z and there are no parameter values that can simultaneously match the low employment rate of 63% and the high job-finding rate of 45%.[17]

In panel C of Figure 12, we see that the impulse response of the constant b model is nearly identical to that in our baseline model. The employment drop in the proportional b model is both smaller and less persistent than in our baseline model. This occurs because as the average level of general human capital falls, home production $b(z) = b_1 z$ falls and

[16] Since $\rho_u = .95$ and the mean $\log z$ for the employed is 1.78, the mean percentage change in z during a month of nonemployment is $(\rho_u - 1)\overline{\log z} = -.05(1.78) = -9\%$.

[17] To see why, note that with an approximately constant job-finding rate, the employment transition equation, $e_{t+1} = (1 - \sigma)e_t + \lambda(1 - e_t)$, approximately yields $e = \lambda/(\lambda + \sigma)$ in steady state. With $\lambda = 0.45$ and $\sigma = 0.0261$, e is 94.5% rather than 63%, as in the data.

so do the returns to being nonemployed, relative to the case in which either $b(z) = b_0$ or $b(z) = b_0 + b_1 z$. Figure 13 shows that, unlike both our baseline model and the constant b model, the proportional b model yields counterfactual implications for the distribution of nonemployment spell durations.

Finally, consider the no home production model. As panel D of Figure 12 shows, this model actually produces a maximal decline in employment that is over twice as large as that produced by the baseline model. As noted earlier, we include this model to make the point that our results are not driven by the intuition that arises in the recalibration by Hagedorn and Manovskii (2008) of the model of Shimer (2005): when all consumers are essentially indifferent between working in the market and working at home, small shocks to productivity in the market generate a large increase in nonemployment. Instead, in our model, the key idea is that during a credit crunch, investing in employment relationships with surplus flows that have long durations is not desirable. We emphasize that we do not think of this model as a serious alternative to the baseline model, because for it to generate a 37% nonemployment rate, vacancy costs need to be extremely high (650% rather than 15% of a worker's monthly output; see Table 6).

Vacancy Costs. In the baseline model, we assume that vacancy creation costs do not vary with z. In our *proportional* κ model, we assume that the cost of posting a vacancy in market z is κz. Panel C of Figure 12 shows that this specification produces an employment response that is nearly identical to that produced by our baseline model, although Figure 13 shows that it generates nonemployment spell durations that match the data less well than those generated by our baseline model.

6 An Economy with Tradable and Nontradable Goods

Now we turn to evaluating the ability of our mechanism to account for the cross-state evidence in the United States during the Great Recession on the comovement of consumption, employment, and wages. To do so, we embed the labor market structure of the one-good model considered so far into a richer model with tradable and nontradable goods. Importantly, this richer model's steady-state implications for labor market variables and human capital accumulation are identical to those of the one-good model and can thus match both cross-sectional and longitudinal evidence on how wages grow over the life cycle and during an employment spell, as discussed earlier.

In the following, we discuss evidence on the cross section of U.S. states during the Great Recession, develop a multi-good version of our model, and then present our main quantitative findings. We show that our richer model reproduces well all key cross-state patterns.

6.1 Motivating Evidence from U.S. States

Our work is motivated by several patterns that are closely related to those documented by Mian and Sufi (2014) and Beraja, Hurst, and Ospina (2016) for a cross section of U.S. regions.

The first pattern is that the regions of the United States that experienced the largest declines in consumption also saw the largest declines in employment, especially in the nontradable goods sector during the Great Recession. The second pattern is that regions that experienced the largest employment declines also saw the largest declines in real wages relative to trend.

Here we illustrate the first pattern by using annual data on employment and consumption from the Bureau of Economic Analysis (BEA). We provide a brief description of the data and provide more detail in the Appendix. Employment is measured as total state-level private nonfarm employment, excluding construction, relative to the total state-level working-age population. (We exclude construction since our model abstracts from housing investment.) We follow the BEA classification of sectors to break down overall employment into nontradable and tradable employment. We measure consumption as per capita consumption expenditure in each state deflated by the aggregate CPI. In the spirit of the model, we isolate changes in consumption triggered by changes in households' ability to borrow—or more generally in credit conditions—as proxied by changes in house prices, by projecting state-level consumption growth on the corresponding growth in state-level (Zillow) house prices. We use the resulting series for consumption growth in our analysis. See Charles, Hurst, and Notowidigdo (2015) for a similar approach.

In panel A of Figure 14, we plot state-level employment growth between 2007 and 2009 against state-level consumption growth over this same period. The figure shows that the elasticity of employment to consumption is 0.38: a 10% decline in consumption is associated with a 3.8% decline in employment.

Panels B and C show that consumption declines are associated with relatively large declines in nontradable employment and essentially no changes in tradable employment across states: a 10% decline in consumption across states is associated with a 5.5% decline in nontradable employment and a negligible (and statistically insignificant) 0.3% increase in tradable employment. As the large negative intercept in panel C shows, the decline in tradable employment is sizable in all states but unrelated to changes in state-level consumption across states.

Beraja, Hurst, and Ospina (2016) also interpret the U.S. cross-regional variation in employment during the Great Recession as arising from discount rate shocks that lead to differential consumption declines across U.S. states. Moreover, they document that states that experienced the largest employment declines were also characterized by the largest decline in real wages relative to trend.

Here we reproduce a version of their findings. For wages, we use census data from the Integrated Public Use Microdata Series, and we control for observable differences in workforce composition both across states and within a state over time by closely following the approach of Beraja, Hurst, and Ospina (2016). We show in panel D of Figure 14 that a decline in employment of 10% across states is associated with a decline in wages of 7.8%. As Beraja, Hurst, and Ospina (2016) argue, in this sense, wages are moderately flexible in the cross section.

In sum, state-level data show that consumption, employment, and wages all strongly positively comove. We summarize these comovements in the first column of Table 7.

6.2 The Richer Model with Tradable and Nontradable Goods

Here we extend our economy to one that can address the cross-state evidence just discussed. We first present the setup of the model and then the results from our quantitative experiments. Most of the details of the model are identical to those of the one-good model and are omitted for brevity. We only discuss the additional features that we introduce.

The economy consists of a continuum of islands, each of which produces intermediate goods, which are combined to make nontradable goods that are only consumed on the island, and a differentiated variety of tradable goods that is consumed everywhere. Consumers receive utility from a composite consumption good that is either purchased in the market or produced at home. Each consumer is endowed with one of two types of skills, which are used in different intensities in the nontradable and tradable goods sectors. Labor is immobile across islands but can switch sectors. We let s index an individual island.

In our experiments, we consider shocks to only a subset of islands that, taken together, are small in the world economy and borrow from the rest of the world at a constant bond price $q > \beta$. In our simple interpretation, this subset of islands is a net borrower from the rest of the world. Given our earlier equivalence results, though, these experiments admit alternative interpretations, such as the illiquid asset interpretation, in which the subset of the islands we consider is not a net borrower, but rather a net saver, from the rest of the world.

Preferences and Demand. The composite consumption good on island s is produced from nontradable goods on island s and tradable goods according to

$$x_t(s) = \left[\tau^{\frac{1}{\mu}} x_{Nt}(s)^{1-\frac{1}{\mu}} + (1-\tau)^{\frac{1}{\mu}} x_{Mt}(s)^{1-\frac{1}{\mu}}\right]^{\frac{\mu}{\mu-1}}.$$

The demand for nontradable goods and tradable goods on island s is given by

$$x_{Nt}(s) = \tau \left(\frac{p_{Nt}(s)}{p_t(s)}\right)^{-\mu} x_t(s) \text{ and } x_{Mt}(s) = (1-\tau)\left(\frac{p_{Mt}}{p_t(s)}\right)^{-\mu} x_t(s),$$

where $p_{Nt}(s)$ is the price of nontradable goods, p_{Mt} is the world price of tradable goods, and

$$p_t(s) = \left[\tau p_{Nt}(s)^{1-\mu} + (1-\tau)\, p_{Mt}^{1-\mu}\right]^{\frac{1}{1-\mu}}$$

is the price of the composite consumption good on island s.

The tradable good itself is a composite of varieties of differentiated goods produced in all other islands s', given by

$$x_{Mt}(s) = \left(\int x_{Mt}(s,s')^{\frac{\mu_X-1}{\mu_X}} \mathrm{d}s'\right)^{\frac{\mu_X}{\mu_X-1}},$$

where $x_{Mt}(s,s')$ is the amount of the variety of tradable goods produced on island s' and consumed on island s, μ_X is the elasticity of substitution between varieties produced on different islands. Let $p_{Xt}(s')$ be the price of the traded variety produced on island s'. We assume that there are no costs of shipping goods from one island to another, so that the law of one price holds and all islands purchase a variety s at the common price $p_{Xt}(s)$. The price of the composite tradable good is common to all islands and given by

$$p_{Mt} = \left(\int p_{Xt}(s)^{1-\mu_X} \mathrm{d}s\right)^{\frac{1}{1-\mu_X}}.$$

The demand on island s' for a tradable variety produced on s is therefore

$$x_{Mt}(s',s) = \left(\frac{p_{Xt}(s)}{p_{Mt}}\right)^{-\mu_X} x_{Mt}(s'),$$

so that the world demand for tradable goods produced by an island s is given by

$$y_{Xt}(s) = \int x_{Mt}(s',s)\,\mathrm{d}s' = \left(\frac{p_{Xt}(s)}{p_{Mt}}\right)^{-\mu_X} y_{Wt}, \tag{43}$$

where $y_{Wt} = \int x_{Mt}(s')\mathrm{d}s'$.

Since any individual island is of measure zero, shocks to an individual island do not affect either the world aggregate price, p_{Mt}, or world demand, y_{Wt}. Given that we consider shocks to only a subset of islands that, taken together, are small in the world economy, world aggregate quantities and prices are also constant with respect to these shocks. We normalize the constant world price of the composite tradable good, p_{Mt}, to 1 so that the composite tradable good is the numeraire.

A Family's Problem. Consider the problem of a family on a given island s. Since from now on we focus on one island, for simplicity we suppress the dependence on s in what follows. The preferences of a family are described by $\sum_{t=0}^{\infty} \beta^t u(c_t)$, where the family's consumption $c_t = x_t + b_t$ is the sum of the goods purchased in the market, x_t, and home produced, b_t, which can only be consumed by that family. The budget constraint is

$$p_t x_t + q a_{t+1} = y_t + d_t + a_t,$$

where p_t is the price of composite consumption good on the island, a_t are the family's assets, y_t is the income of the family's workers in the form of wages, and d_t are the profits from the firms the family owns on island s. The family's debt constraint on island s is, as before,

$$a_{t+1} \geq -\chi_t.$$

Note that the consumption problem of the family is almost identical to that in the one-good model. The one difference is that the shadow price of one unit of composite tradable good at date t in units of the date 0 composite tradable good is $Q_t = \beta^t u'(c_t)/p_t$, where for simplicity we choose p_0 so that $u'(c_0)/p_0$ is 1.

Technology. Nontradable and tradable goods are produced with locally produced intermediate goods. These intermediate goods are used by the nontradable and tradable sectors in different proportions. This setup effectively introduces costs of sectoral reallocations of workers because it implies a curved production possibility frontier between nontradable and tradable goods.

Specifically, this economy has two types of intermediate goods: type $\mathcal{N}$ and type $\mathcal{X}$ goods. The technology for producing nontradable goods disproportionately uses type $\mathcal{N}$ goods, whereas the technology for producing tradable goods disproportionately uses type $\mathcal{X}$ goods according to the production technologies

$$y_{Nt} = A\left(y_{Nt}^{\mathcal{N}}\right)^{\nu}\left(y_{Nt}^{\mathcal{X}}\right)^{1-\nu} \quad \text{and} \quad y_{Xt} = A\left(y_{Xt}^{\mathcal{N}}\right)^{1-\nu}\left(y_{Xt}^{\mathcal{X}}\right)^{\nu}, \tag{44}$$

with $\nu \geq 1/2$. Here $y_{Nt}^{\mathcal{N}}$ and $y_{Xt}^{\mathcal{N}}$ denote the use of intermediate inputs of type $\mathcal{N}$ in the production of nontradable and tradable goods, whereas $y_{Nt}^{\mathcal{X}}$ and $y_{Xt}^{\mathcal{X}}$ denote the use of intermediate inputs of type $\mathcal{X}$ in the production of nontradable and tradable goods. Both nontradable goods producers and tradable goods producers are competitive and take as given the price of their goods, p_{Nt} and p_{Xt}. The demands for intermediate inputs in the nontradable goods sector are given by

$$y_{Nt}^{\mathcal{N}} = \nu\left(\frac{p_t^{\mathcal{X}}}{p_t^{\mathcal{N}}}\right)^{1-\nu} y_{Nt} \quad \text{and} \quad y_{Nt}^{\mathcal{X}} = (1-\nu)\left(\frac{p_t^{\mathcal{N}}}{p_t^{\mathcal{X}}}\right)^{\nu} y_{Nt},$$

where $p_t^{\mathcal{X}}$ and $p_t^{\mathcal{N}}$ are the prices of the intermediate goods $\mathcal{N}$ and $\mathcal{X}$, and we have used the convenient normalization $A = \nu^{-\nu}(1-\nu)^{-(1-\nu)}$. Likewise, the demands for intermediate inputs in the tradable goods sector are

$$y_{Xt}^{\mathcal{N}} = (1-\nu)\left(\frac{p_t^{\mathcal{X}}}{p_t^{\mathcal{N}}}\right)^{\nu} y_{Xt} \quad \text{and} \quad y_{Xt}^{\mathcal{X}} = \nu\left(\frac{p_t^{\mathcal{N}}}{p_t^{\mathcal{X}}}\right)^{1-\nu} y_{Xt}.$$

The zero profit conditions in the nontradable and tradable goods sectors imply

$$p_{Nt} = \left(p_t^{\mathcal{N}}\right)^{\nu}\left(p_t^{\mathcal{X}}\right)^{1-\nu} \quad \text{and} \quad p_{Xt} = \left(p_t^{\mathcal{N}}\right)^{1-\nu}\left(p_t^{\mathcal{X}}\right)^{\nu}.$$

We assume that there are measures of consumers, $\pi^{\mathcal{N}}$ and $\pi^{\mathcal{X}} = 1 - \pi^{\mathcal{N}}$, who supply labor to produce the two types of intermediate goods, $\mathcal{N}$ and $\mathcal{X}$. We refer to these consumers as being in *occupations* $\mathcal{N}$ and $\mathcal{X}$. Consumers in occupation $\mathcal{N}$ can produce good $\mathcal{N}$, and consumers in occupation $\mathcal{X}$ can produce good $\mathcal{X}$. Consumers are hired by intermediate goods firms that produce intermediate goods of either type $\mathcal{N}$ or type $\mathcal{X}$. These goods are then sold at competitive prices $p_t^{\mathcal{N}}$ and $p_t^{\mathcal{X}}$ to firms in the nontradable and tradable goods sectors. Of course, it is equivalent to think that the consumers in each occupation work in the sector that purchases the goods they produce. Under this interpretation, we can think of consumers in occupation $\mathcal{X}$ as employed in sectors N and X and consumers in occupation $\mathcal{N}$ as also employed in sectors N and X in different proportions: sector N employs consumers in occupation $\mathcal{N}$ relatively intensively, whereas sector X employs consumers in occupation $\mathcal{X}$ relatively intensively.

This setup captures in a simple way the idea that switching sectors is relatively easy, whereas switching occupations is difficult. Here any individual consumer faces no cost of switching sectors, but if a positive measure of consumers moves from one sector to another, it reduces the marginal revenue products of those workers and, thus, their wages. This reduction in marginal revenue products acts like a switching cost in the aggregate.

Labor Market. Firms that produce intermediate good $i \in \{\mathcal{N}, \mathcal{X}\}$ post vacancies for consumers in occupation i with general human capital z and firm-specific human capital h, who produce intermediate good i when matched. We assume that consumers cannot switch occupations, so the measure of consumers in each occupation is fixed. The values of consumers in occupation i with general human capital z and firm-specific human capital h are similar to those in our one-good model and are given by

$$W_t^i(z,h) = Q_t\omega_t^i(z,h) + \phi(1-\sigma)\int \max\left[W_{t+1}^i(z',h'), U_{t+1}^i(z')\right] \mathrm{d}F_e(z'|z) \tag{45}$$
$$+ \phi\sigma \int U_{t+1}^i(z')\,\mathrm{d}F_e(z'|z)$$

for employed consumers and

$$U_t^i(z) = Q_t p_t b(z) + \phi \lambda_{wt}^i(z) \int \max\left[W_{t+1}^i(z',1), U_{t+1}^i(z')\right] \mathrm{d}F_u(z'|z) \qquad (46)$$
$$+ \phi[1-\lambda_{wt}^i(z)] \int U_{t+1}^i(z') \, \mathrm{d}F_u(z'|z)$$

for nonemployed consumers, where $\omega_t^i(z,h)$ is the wage received by a consumer in occupation i as a function of human capital and $\lambda_{wt}^i(z)$ is the job-finding probability of a consumer in occupation i.

The value of a firm producing intermediate good i matched with a consumer in occupation i with productivity (z,h) is

$$J_t^i(z,h) = Q_t[p_t^i zh - \omega_t^i(z,h)] + (1-\sigma)\,\phi \int \max\left[J_{t+1}^i(z',h'), 0\right] \mathrm{d}F_e(z'|z). \qquad (47)$$

That is, at date t a consumer in occupation i with human capital (z,h) matched with a firm in intermediate good sector i produces zh units of good i, which sell for $p_t^i zh$, and the firm pays the consumer $\omega_t^i(z,h)$. The cost of posting a vacancy is κ units of the composite tradable good. The free-entry condition in sector i is analogous to (10).

The technology for firms producing intermediate good i is the same as in the one-good model. The matches of firms that produce intermediate good i with consumers with general human capital z are given by $m_t^i(z) = u_t^i(z)\, v_t^i(z) / [u_t^i(z)^\eta + v_t^i(z)^\eta]^{\frac{1}{\eta}}$, where $u_t^i(z)$ is the measure of nonemployed consumers and $v_t^i(z)$ is the measure of vacancies directed at such consumers. The associated worker job-finding rate $\lambda_{wt}^i(z)$ and firm worker-finding rate $\lambda_{ft}^i(z)$ then follow as before. The determination of wages by Nash bargaining is exactly analogous to that in the one-good model.

Consider next the market clearing conditions. Market clearing for the two types of intermediate goods requires

$$\int_{z,h} zh \, \mathrm{d}e_t^i(z,h) = y_{Nt}^i + y_{Xt}^i \text{ for } i \in \{\mathcal{N}, \mathcal{X}\}. \qquad (48)$$

The left side of this equation is the total amount of intermediate goods of type i produced by the measure of employed consumers, $e_t^i(z,h)$, and the right side is the total amount of these intermediate goods used by firms in the nontradable and tradable sectors.

Market clearing for nontradable goods can be written as $x_{Nt} = y_{Nt}$. Market clearing for tradable goods requires that the demand for these goods in (43) equals the supply of these goods in (44).

6.3 The Workings of the Richer Model

Consider how employment responds to a credit tightening in this version of the model. In contrast to the one-good model, here consumers can reallocate across sectors. The cost of this reallocation is governed by the curvature of the production possibility frontier between nontradable and tradable goods: the more curved this frontier (the higher is ν), the higher the cost of reallocation. Mechanically, as ν increases, a given flow of consumers into a sector leads to a greater fall in the marginal product and, thus, wages in that sector.

In this environment, a credit tightening has two effects on employment. The first, the *investment* effect, is similar to that in the one-good model: the cost of posting vacancies increases by more than the surplus from a match, leading firms in both sectors to post fewer vacancies and, hence, leading to a drop in overall employment.

The second, the *relative demand* effect, is due to the reduction in the demand for nontradable goods produced on the island and thus their relative price. This drop in prices amplifies the drop in employment in the nontradable sector relative to that in the tradable sector. When the cost of sectoral reallocation is small, a large flow of labor from the nontradable sector to the tradable sector ensues. This reallocation can be so large that even though overall employment declines, employment in the tradable sector increases. In contrast, when the cost of sectoral reallocation is large, the flow of labor from the nontradable sector to the tradable sector is small, so that employment falls in both sectors.

Thus, in response to a credit tightening, nontradable employment falls unambiguously because of the combination of the investment and relative demand effects. The response of tradable employment is, instead, ambiguous. If sectoral reallocation is costly (high ν), then tradable employment falls because the investment effect dominates the relative demand effect. If, in contrast, sectoral reallocation is not too costly (low ν), then tradable employment increases because the relative demand effect more than offsets the investment effect.

To see these effects, consider first the extreme case in which the two sectors employ consumers from the two occupations equally intensively, that is, an economy with $\nu = 1/2$, in which the costs of sectoral reallocation are relatively low and the remaining parameter values are equal to those in the quantitative model discussed later. Figure 15 shows the response to a credit tightening that generates a 5% drop in consumption. As the figure shows, in this case tradable employment expands because of the inflow of consumers from the nontradable sector.

Consider the second extreme case in which consumers cannot switch sectors, that is, an economy with $\nu = 1$. As Figure 16 shows, in this case a similar credit tightening leads both tradable and nontradable employment to fall. The drop in tradable employment is somewhat smaller since it is driven solely by the investment effect. Because the price of nontradable goods falls more than that of tradable goods, firms find it even less attractive to post vacancies in that sector.

For ν in between these two extremes, overall employment falls, employment in the nontradable sector falls, and employment in the tradable sector can either rise or fall depending on ν. As we discuss next, we discipline our choice of ν by confronting the patterns of nontradable and tradable employment during the Great Recession.

6.4 Comparison with Cross-Sectional Data

We begin by discussing how we set parameters in this extended version of the model. There are five new parameters, in addition to those in the one-good model, namely $(\tau, \mu_X, \pi^{\mathcal{N}}, \nu, \mu)$. We choose the parameter τ so that the share of spending on nontradable goods is 2/3 and set the trade elasticity, μ_X, at 4. Both of these numbers are fairly standard in the trade literature. We choose the fraction of consumers in the two occupations to ensure that in the steady state, wages in the two occupations are the same for a given level of human capital. Given this choice, the steady-state implications of this richer model are identical to those of the one-good model reported in panel A of Table 1, and so we choose the rest of the parameters as we did before in panel B of Table 1.

As discussed earlier, one key parameter in this richer version of the model is the parameter ν governing the curvature of the production possibility frontier between tradable and nontradable goods. This parameter allows us to capture, in a parsimonious yet flexible way, the cost of reallocating consumers across the two sectors. We choose this parameter by requiring our model to reproduce the Mian and Sufi (2014) observation that declines in consumption across states were essentially unrelated to changes in tradable employment. Specifically, we choose ν so as to generate an elasticity of tradable employment to consumption of -0.03, as observed in the data. The resulting share is $\nu = 0.87$.

A second key parameter is μ, the elasticity of substitution between tradable and nontradable goods. The lower μ is, the more the relative price of nontradable goods falls following a credit tightening and, thus, the more wages fall. We choose this parameter to reproduce the observation of Beraja, Hurst, and Ospina (2016) that in the cross section of U.S. states, wages were moderately flexible during the Great Recession. Specifically, we choose μ so as to generate an elasticity of wages to employment of 0.78. The resulting elasticity of substitution is $\mu = 2.5$.

We next describe the experiments we conduct. For each state, we choose a sequence of shocks to the debt limit so that the model exactly reproduces the predicted consumption series for each state in the data. Each of these shocks is unanticipated. As in our one-good model, the path for shocks is such that agents believe that consumption will revert to its steady state at a rate of 10% per quarter. Given these paths for shocks, we calculate the evolution of state-level variables. We then compare the summary statistics in our model to the corresponding statistics in the data summarized in Table 7.

As discussed, we have chosen the parameters ν and μ so that the model reproduces the observed elasticities of tradable employment to consumption and of wages to employment. We now evaluate the extent to which the model can account for how overall employment and nontradable employment fell as consumption fell during the Great Recession.

Recall that in the data, a fall in consumption across states of 10% is associated with a fall in nontradable employment of 5.5% and in overall employment of 3.8% between 2007 and 2009. As Table 7 shows, our model implies that such a fall in consumption is associated with a fall in nontradable employment of 5.7% and a fall in overall employment of 3.3%. Thus, our model successfully accounts for the comovements of consumption and both nontradable employment and overall employment.

Table 8 shows how these elasticities change as we vary the cost of reallocating consumers from one extreme case with a low cost of reallocation ($\nu = 0.5$) to the other extreme case with a prohibitively high cost of reallocation ($\nu = 1$). When reallocation costs are low, the relative demand effect dominates the investment effect, leading to a counterfactually sharp decline in nontradable employment and a sharp increase in tradable employment. The model's implication for the comovement of employment and wages is also grossly at odds with the data: wages comove too little with employment compared to the data. When reallocation costs are high, the investment effect dominates the relative demand effect, leading, counterfactually, to similarly sized declines in nontradable and tradable employment following a credit tightening.

Table 8 also shows how these elasticities change as we vary the elasticity of substitution between tradable and nontradable goods from a relatively low elasticity of $\mu = 1$ to a relatively high elasticity of $\mu = 5$. When this elasticity is low, nontradable goods prices fall a lot, leading to a counterfactually large wage drop, whereas when this elasticity is high, nontradable goods prices fall a little, leading to a counterfactually small wage drop.

Table 9 shows the effects of eliminating the growth in human capital accumulation. In our baseline model, we chose the values of ν and μ so that the model reproduces the features that wages are moderately flexible and that tradable employment does not comove with consumption in the cross section. Table 9 shows that, regardless of the values of these parameters, the model without human capital is unable to come anywhere close to reproducing these features of the data. Moreover, consistent with our results from the one-good model, for all such parameter values, the overall employment responses are much smaller than those in both the baseline model and the data.

7 Conclusion

In a search and matching model, hiring a worker is an investment activity that requires paying a cost today in return for a sequence of expected surplus flows in the future. As is well known, the present value of a sequence of flows is sensitive to changes in discount rates if and only if those flows have long durations. We demonstrated that the standard DMP model implies a short duration of surplus flows from a match between a worker and a firm, with an alternative Macaulay duration of only 2.6 months. Our search model with general human capital, consistent with both cross-sectional and longitudinal evidence on wage growth, implies a duration that is nearly 10 times as long. Hence, our model amplifies the negative effect of a mean-reverting increase in shadow interest rates on employment by a factor of 10, compared to the DMP model. For a stripped-down version of our economy for which we obtain closed-form solutions, we prove that human capital accumulation amplifies the employment response to a credit tightening relative to the DMP model, when these surplus flows have long durations. Our extensive robustness exercises show that this link between general human capital accumulation and employment amplification survives alternative parameterizations of our model.

Even though it is purposely sparsely parameterized, the model successfully reproduces salient features of the micro-data with respect to the dynamics of wages on the job, over the life cycle, and after spells of nonemployment. An extended version of our model also closely matches the cross-regional patterns of consumption, employment, and wages observed in the United States during the Great Recession.

References

Altonji, Joseph G., and Robert A. Shakotko. 1987. Do wages rise with job seniority? *Review of Economic Studies*, 54 (3), 437-459.

Andolfatto, David. 1996. Business cycles and labor-market search. *American Economic Review*, 86 (1), 112-132.

Barron, John M., Mark C. Berger, and Dan A. Black. 1997. On-the-job training. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2016. The aggregate implications of regional business cycles. National Bureau of Economic Research Working Paper 21956.

Bernanke, Ben, and Mark Gertler. 1989. Agency costs, net worth, and business fluctuations. *American Economic Review,* 79 (1), 14-31.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. 1999. The financial accelerator in a quantitative business cycle framework. In *Handbook of Macroeconomics*, vol. 1C, ed. John

B. Taylor and Michael Woodford, 1341-1393. Amsterdam: North-Holland.

Best, Michael Carlos, James Cloyne, Ethan Ilzetzki, and Henrik Kleven. 2017. Estimating the elasticity of intertemporal substitution using mortgage notches. Mimeo, Stanford University.

Buchinsky, Moshe, Denis Fougère, Francis Kramarz, and Rusty Tchernis. 2010. Interfirm mobility, wages and the returns to seniority and experience in the United States. *Review of Economic Studies*, 77 (3), 972-1001.

Chodorow-Reich, Gabriel, and Loukas Karabarbounis. 2016. The cyclicality of the opportunity cost of employment. *Journal of Political Economy*, 124 (6), 1563-1618.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2015. Housing booms and busts, labor market opportunities, and college attendance. National Bureau of Economic Research Working Paper 21587.

Davis, Morris A., and Stijn Van Nieuwerburgh. 2014. Housing, finance and the macroeconomy. National Bureau of Economic Research Working Paper 20287.

Davis, Stephen J., and Till von Wachter. 2011. Recessions and the cost of job loss. *Brookings Papers on Economic Activity*, 43 (Fall), 1-72.

den Haan, Wouter J., Garey Ramey, and Joel Watson. 2000. Job destruction and the experiences of displaced workers. In *Carnegie-Rochester Conference Series on Public Policy*, vol. 52, ed. Bennett T. McCallum and Charles I. Plosser, 87-128. Amsterdam: North-Holland.

Eggertsson, Gauti B., and Paul Krugman. 2012. Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *Quarterly Journal of Economics*, 127 (3), 1469-1513.

Eisfeldt, Andrea L., and Adriano A. Rampini. 2007. New or used? Investment with credit constraints. *Journal of Monetary Economics*, 54 (8), 2656-2681.

Elsby, Michael W.L., and Matthew D. Shapiro. 2012. Why does trend growth affect equilibrium employment? A new explanation of an old puzzle. *American Economic Review*, 102 (4), 1378-1413.

Floden, Martin, and Jesper Lindé. 2001. Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance? *Review of Economic Dynamics,* 4 (2), 406-437.

Gertler, Mark, and Peter Karadi. 2011. A model of unconventional monetary policy. *Journal of Monetary Economics*, 58 (1), 17-34.

Gertler, Mark, and Nobuhiro Kiyotaki. 2010. Financial intermediation and credit policy in business cycle analysis. In *Handbook of Monetary Economics*, vol. 3, ed. Benjamin M. Friedman and Michael Woodford, 547-599. Amsterdam: North-Holland.

Gilchrist, Simon, and Egon Zakrajšek. 2012. Credit spreads and business cycle fluctuations. *American Economic Review*, 102 (4), 1692-1720.

Gomes, Francisco, and Alexander Michaelides. 2005. Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance*, 60 (2), 869-904.

Guerrieri, Veronica, and Guido Lorenzoni. 2015. Credit crises, precautionary savings, and the liquidity trap. Mimeo, Chicago Booth.

Guvenen, Fatih. 2006. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics*, 53 (7), 1451-1472.

Hagedorn, Marcus, and Iourii Manovskii. 2008. The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review*, 98 (4), 1692-1706.

Hall, Robert E. 2017. High discounts and high unemployment. *American Economic Review*, 107 (2), 305-330.

Hall, Robert E., and Paul R. Milgrom. 2008. The limited influence of unemployment on the wage bargain. *American Economic Review*, 98 (4), 1653-1674.

Head, Allen, and Huw Lloyd-Ellis. 2012. Housing liquidity, mobility, and the labour market. *Review of Economic Studies*, 79 (4), 1559-1589.

Huckfeldt, Christopher. 2014. The scarring effect of recessions: A quantitative analysis. Mimeo, New York University.

Itskhoki, Oleg, and Elhanan Helpman. 2015. Trade liberalization and labor market dynamics with heterogeneous firms. Mimeo, Princeton University.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. Earnings losses of displaced workers. *American Economic Review,* 83 (4), 685-709.

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

Kiyotaki, Nobuhiro, and John Moore. 1997. Credit cycles. *Journal of Political Economy*, 105 (2), 211-248.

Krolikowski, Pawel. 2017. Job ladders and earnings of displaced workers. *American Economic Journal: Macroeconomics*, 9 (2), 1-31.

Krusell, Per, Toshihiko Mukoyama, and Ayşegül Şahin. 2010. Labour-market matching with precautionary savings and aggregate fluctuations. *Review of Economic Studies*, 77 (4), 1477-1507.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2011. A three state model of worker flows in general equilibrium. *Journal of Economic Theory*, 146 (3), 1107-1133.

————(forthcoming) Gross worker flows over the business cycle. *American Economic Review*.

Ljungqvist, Lars, and Thomas J. Sargent. 1998. The European unemployment dilemma. *Journal of Political Economy*, 106 (3), 514-550.

————2008. Two questions about European unemployment. *Econometrica*, 76 (1), 1-29.

————

McCall, John J. 1970. Economics of information and job search. *Quarterly Journal of Economics*, 84 (1), 113-126.

Mendoza, E. 2010. Sudden stops, financial crises, and leverage. *American Economic Review*, 100 (5), 1941-1946.

Merz, Monika. 1995. Search in the labor market and the real business cycle. *Journal of Monetary Economics*, 36 (2), 269–300.

Mian, Atif, and Amir Sufi. 2011. House prices, home equity-based borrowing, and the U.S. household leverage crisis. *American Economic Review*, 101 (5), 2132-2156.

————2014. What explains the 2007–2009 drop in employment? *Econometrica*, 82 (6), 2197-2223.

Midrigan, Virgiliu, and Thomas Philippon. 2016. Household leverage and the recession. Mimeo, New York University.

Ohanian, Lee E. 2010. The economic crisis from a neoclassical perspective. *Journal of Economic Perspectives*, 24 (4), 45-66.

Pinheiro, Roberto, and Ludo Visschers. 2015. Unemployment risk and wage differentials. *Journal of Economic Theory*, 157 (3), 397-424.

Pissarides, Christopher A. 1992. Loss of skill during unemployment and the persistence of employment shocks. *Quarterly Journal of Economics*, 107 (4), 1371-1391.

Rubinstein, Yona, and Yoram Weiss. 2006. Post schooling wage growth: Investment, search and learning. In *Handbook of the Economics of Education*, vol. 1, ed. Eric A. Hanushek and Finis Welch, 1-67. Amsterdam: North Holland.

Shimer, Robert. 2005. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95 (1), 25-49.

————2012. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15 (2), 127-148.

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. Forthcoming. Firming up inequality. In *Firms and the Distribution of Income: The Roles of Productivity and Luck*. National Bureau of Economic Research.

Topel, Robert. 1991. Specific capital, mobility, and wages: Wages rise with job seniority. *Journal of Political Economy*, 99 (1), 145-176.

Table 1: Parameterization

Panel A: Moments

	Data	Model
Employment rate	0.63	0.63
Job-finding rate	0.45	0.45
Vacancy cost (% output)	15.0	15.0
Standard deviation of wage changes	0.21	0.21
Cross-sectional difference in log wages:		
30 to 1 years of experience	1.21	1.19
Annual wage growth during employment spell:		
1-10 years of experience	0.10	0.10
11-20 years of experience	0.07	0.08
21-30 years of experience	0.06	0.06
31-40 years of experience	0.06	0.05
1-40 years of experience	0.07	0.07

Panel B: Parameters

Endogenously chosen	
b_0, home production (rel. to mean output)	0.42
η, matching function elasticity	0.61
κ, vacancy cost (rel. to mean output)	0.15
σ_z, standard deviation of shocks	0.06
ρ, convergence rate	$0.95^{1/12}$
$\log \overline{z}_e$, general human capital drift	2.44
$\log \overline{h}$, firm-specific human capital drift	0.82
Assigned	
Period length	1 month
β, discount factor	$0.96^{1/12}$
q, bond price	$0.98^{1/12}$
ϕ, survival probability	1-1/480
σ, probability of separation	0.026
ϱ_c, mean-reversion rate of consumption	0.10
b_1, replacement rate	0.25
γ, worker's bargaining share	0.50

Table 2: Additional Model Implications

	Data	Model
External Validation		
Mean wage drop after nonemployment spell	0.044–0.055	0.05
Sensitivity of wage loss to additional tenure year, %	1.54	1.95
Std. deviation of initial log wages	0.85	0.82
Profit share of revenue	0.06	0.06
Other Implications		
Probability of endogenous separation	-	0.001
Mean home production to mean wage	-	0.48
Fraction nonemployed with positive match probability	-	0.04

Table 3A: Moments Targeted in Robustness to Alternative Wage Estimates*

	Altonji-Shakotko		Topel	
	Data	Model	Data	Model
Employment rate	0.63	0.63	0.63	0.63
Job-finding rate	0.45	0.44	0.45	0.45
Vacancy cost (% output)	0.15	0.15	0.15	0.15
Std. deviation of wage changes	0.21	0.21	0.21	0.21
Cross-sectional $\Delta \log w$ (30 to 1 years)	0.80	0.78	0.80	0.81
On-the-job wage growth (1-10 years)	0.05	0.05	0.08	0.08
On-the-job wage growth (11-20 years)	0.04	0.04	0.05	0.06
On-the-job wage growth (21-30 years)	0.03	0.03	0.04	0.05
On-the-job wage growth (31-40 years)	0.02	0.02	0.04	0.04
On-the-job wage growth (1-40 years)	0.03	0.03	0.05	0.05

*Both calibrations use cross-sectional wage growth from Rubinstein and Weiss (2006).

Table 3B: Parameters Used in Robustness to Alternative Wage Estimates

	Baseline	Altonji-Shakotko	Topel
b_0, home production (rel. to mean output)	0.42	0.39	0.51
η, matching function elasticity	0.61	0.77	0.72
κ, vacancy cost (rel. to mean output)	0.15	0.15	0.15
σ_z, std. deviation of shocks	0.06	0.06	0.06
ρ, convergence rate	$0.95^{1/12}$	$0.96^{1/12}$	$0.93^{1/12}$
$\log \bar{z}_e$, general human capital drift	2.44	1.68	1.56
$\log \bar{h}$, firm-specific human capital drift	0.82	0.00	0.57

Table 4: Moments Targeted in Robustness Checks

	Data	High ρ_u	Low ρ_u	Prop. κ	Prop. b	Const. b	$b=0$	High $\bar{h}$
Employment rate	0.63	0.63	0.63	0.63	0.63	0.63	0.63	0.63
Job-finding rate	0.45	0.46	0.48	0.46	0.05	0.46	0.05	0.45
Vacancy cost (% output)	15.0	15.0	14.9	15.0	15.1	15.0	6.5	14.9
Std. deviation of wage changes	0.21	0.21	0.21	0.21	0.21	0.21	0.21	0.21
Cross-section $\Delta \log w$	1.21	1.18	1.15	1.17	1.21	1.19	1.04	1.22
On-the-job $\Delta \log w$ (annual)	0.07	0.07	0.08	0.07	0.09	0.07	0.08	0.09

*High ρ_u and Low ρ_u are the no decay and high decay models. Prop. κ, Prop. b, Const. b, and $b=0$ are the proportional vacancy model, the proportional b model, the constant b model, and the no home production model, respectively. High $\bar{h}$ is the high $\bar{h}$ model.

Table 5: Additional Moments for Robustness Checks

	Data	High ρ_u	Low ρ_u	Prop. κ	Prop. b	Const. b	$b=0$	High $\bar{h}$
External Validation								
Mean wage drop after nonemp.	0.05	0.04	0.09	0.05	0.05	0.05	0.03	0.10
Sens. wage loss to tenure, %	1.54	1.99	0.60	1.96	0.70	1.94	0.39	3.82
Std. deviation initial log w	0.85	0.79	0.74	0.80	0.93	0.82	0.91	0.81
Profit share of revenue	0.06	0.03	0.28	0.07	0.29	0.07	0.31	0.07
Other Implications								
Prob. endogenous separation	-	0.03	0.10	0.12	0	0.06	0	0.04
Mean home prod. to mean w	-	0.46	0.61	0.50	0.48	0.47	0	0.48
Frac. nonemployed & $\lambda_w > 0$	-	0.04	0.04	0.04	0.37	0.04	0.33	0.04

*High ρ_u and Low ρ_u are the no decay and high decay models. Prop. κ, Prop. b, Const. b, and $b=0$ are the proportional vacancy model, the proportional b model, the constant b model, and the no home production model, respectively. High $\bar{h}$ is the high $\bar{h}$ model.

Table 6: Parameters Used in Robustness Checks

	High ρ_u	Low ρ_u	Prop. κ	Prop. b	Const. b	$b=0$	High $\bar{h}$
b_0, home production to mean y	0.41	0.41	0.43	0	0.44	0	0.42
η, matching function elasticity	0.77	0.48	0.51	0.20	0.61	0.61	0.61
κ, vacancy cost to mean y	0.15	0.15	0.15	0.15	0.15	6.5	0.15
σ_z, std. deviation of shocks	0.06	0.06	0.06	0.05	0.06	0.04	0.06
$\log \bar{z}_e$, general human capital drift	2.08	3.65	2.44	2.00	2.44	1.55	2.44
$\log \bar{h}$, firm-specific human capital drift	0.82	0	0.82	0.10	0.82	0	1.65

*High ρ_u and Low ρ_u are the no decay and high decay models. Prop. κ, Prop. b, Const. b, and $b=0$ are the proportional vacancy model, the proportional b model, the constant b model, and the no home production model, respectively. High $\bar{h}$ is the high $\bar{h}$ model.

Table 7: Cross-State Elasticities

	Data	Baseline Model
Elasticity Δe vs. Δc	0.38	0.33
Elasticity Δe_N vs. Δc	0.55	0.57
Elasticity Δe_T vs. Δc	-0.03	-0.03
Elasticity Δw vs. Δe	0.78	0.78

Note: Δe, Δe_{NT}, and Δe_T denote changes in overall, non-tradable, and tradable employment; Δc denotes changes in predicted consumption; Δw denotes changes in wages.

Table 8: Alternative Parameterizations

	Baseline	Varying ν		Varying μ	
	Model	$\nu = 0.5$	$\nu = 1$	$\mu = 1$	$\mu = 5$
Elasticity Δe vs. Δc	0.33	0.49	0.27	0.25	0.40
Elasticity Δe_N vs. Δc	0.57	1.42	0.30	0.57	0.58
Elasticity Δe_T vs. Δc	-0.03	-0.85	0.23	-0.21	0.13
Elasticity Δw vs. Δe	0.78	0.03	1.21	1.70	0.22

Table 9: Model without Human Capital Accumulation

	$\nu = 0.87$	Varying ν		Varying μ	
	$\mu = 2.5$	$\nu = 0.5$	$\nu = 1$	$\mu = 1$	$\mu = 5$
Elasticity Δe vs. Δc	0.07	0.14	0.05	0.08	0.07
Elasticity Δe_N vs. Δc	0.41	0.90	0.22	0.51	0.31
Elasticity Δe_T vs. Δc	-0.47	-1.07	-0.23	-0.61	-0.33
Elasticity Δw vs. Δe	5.72	2.07	9.89	6.71	4.57

Figure 1: Steady-State Measures and Matching Rates

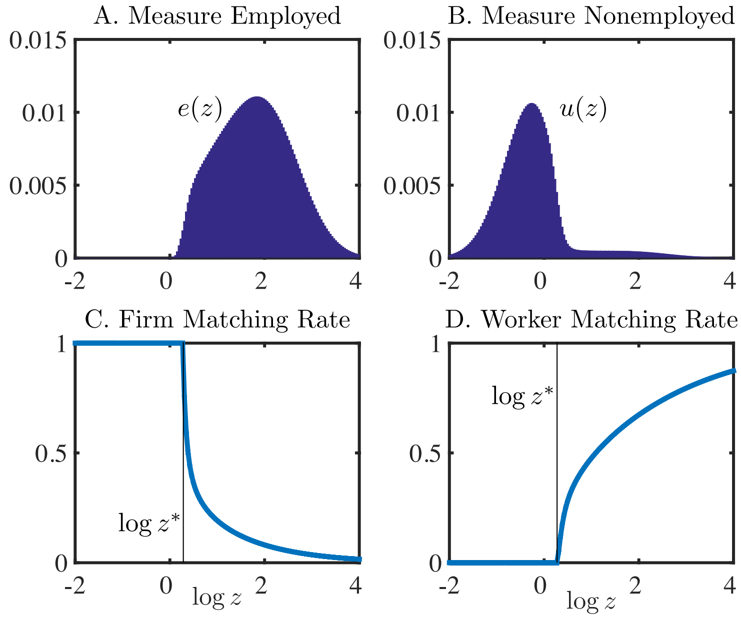

Figure 2: Example of Individual Wage Path

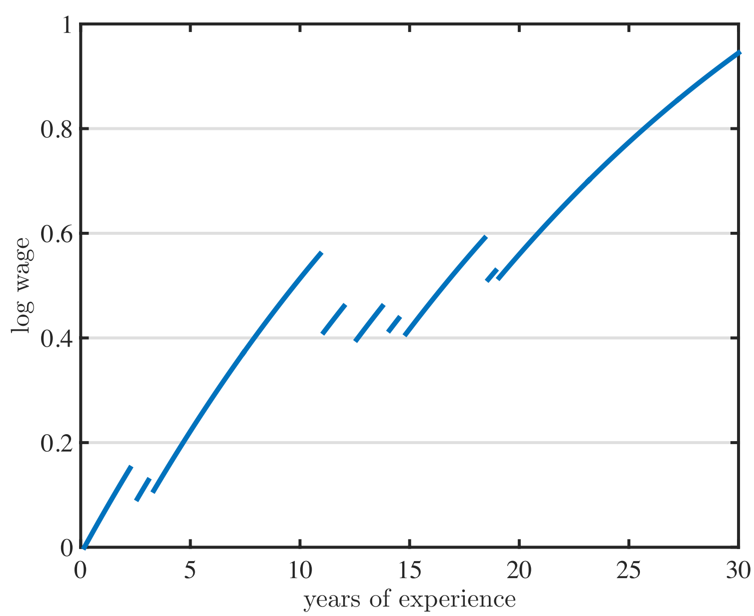

Figure 3: Distribution of Nonemployment Durations

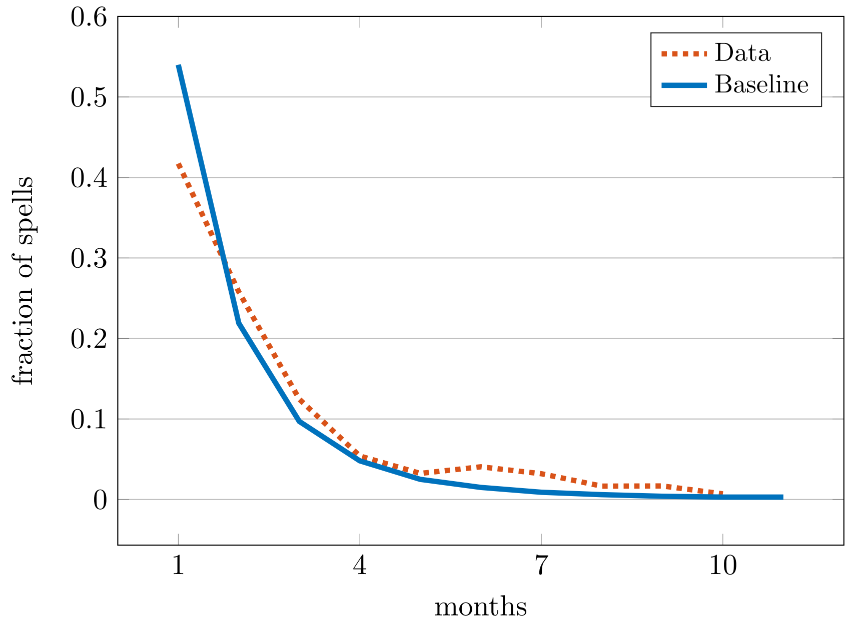

Figure 4: Displaced Worker Regressions

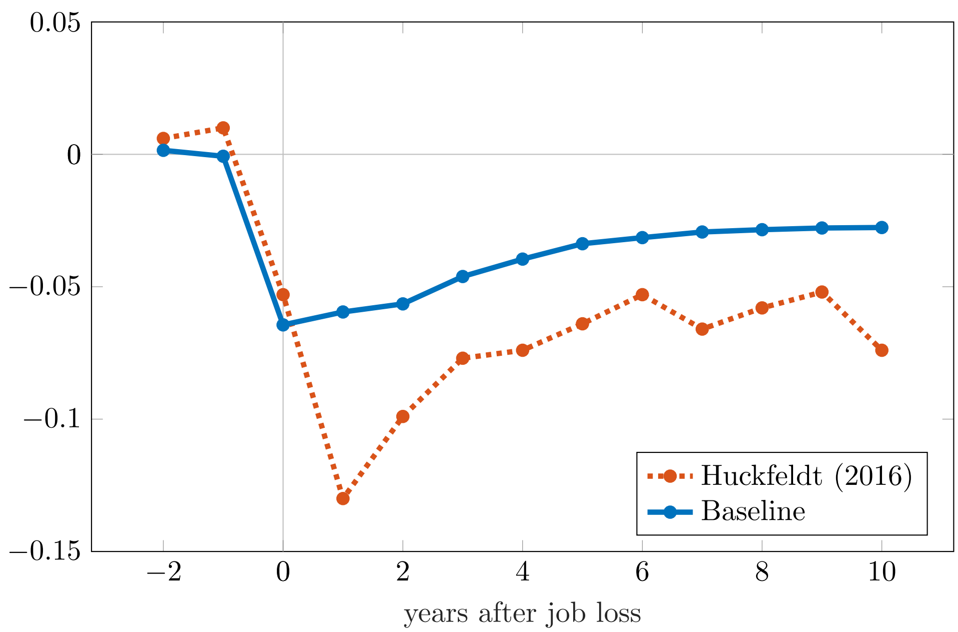

Figure 5: Employment Response after Credit Tightening

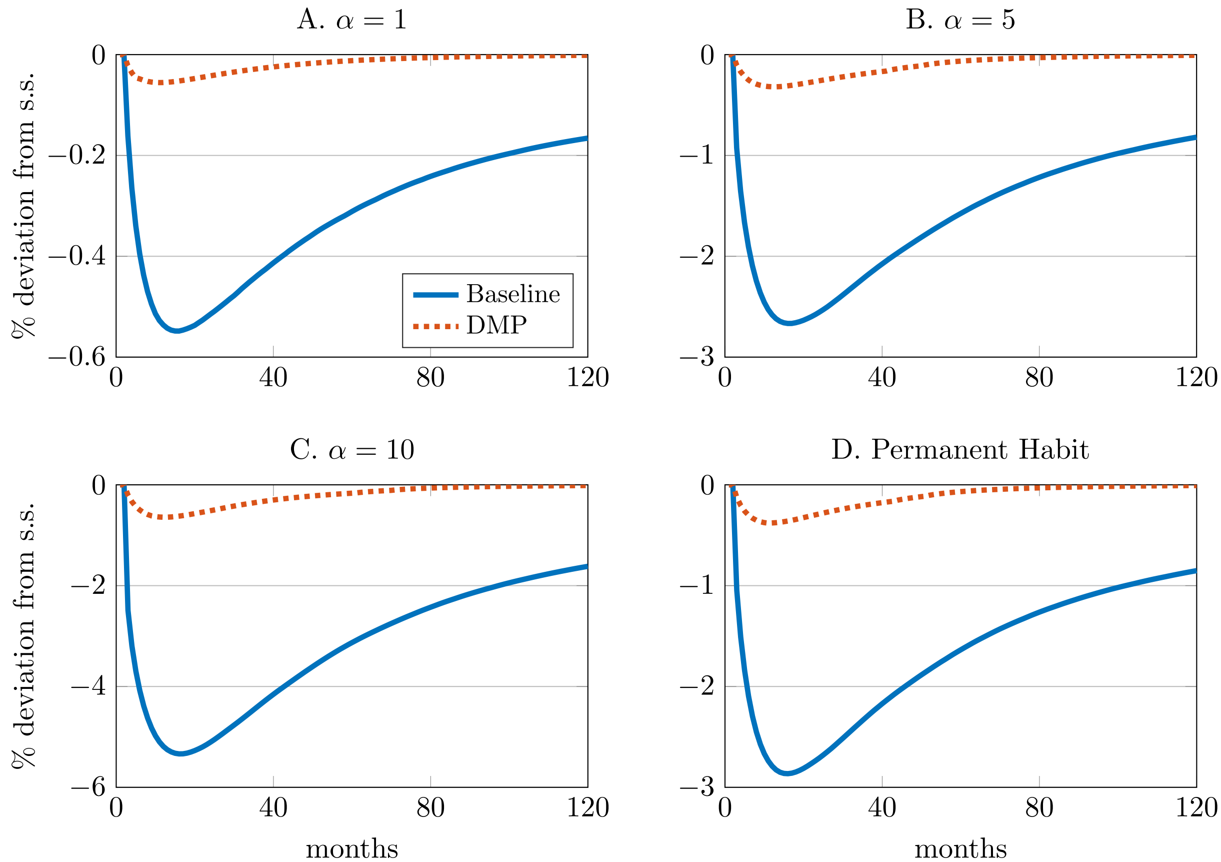

Figure 6: Shimer Employment Decomposition

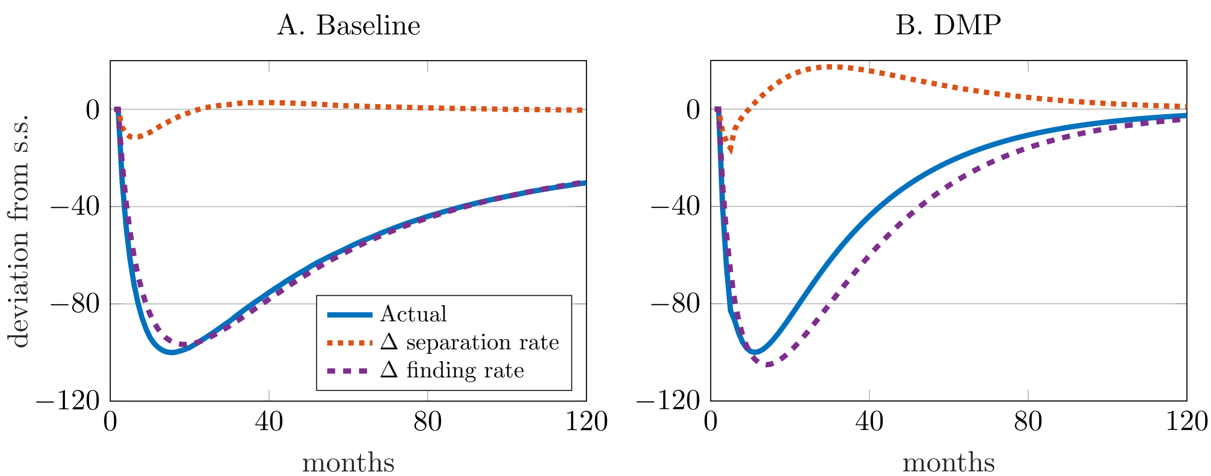

Figure 7: Duration of Surplus Flows in Simple Models

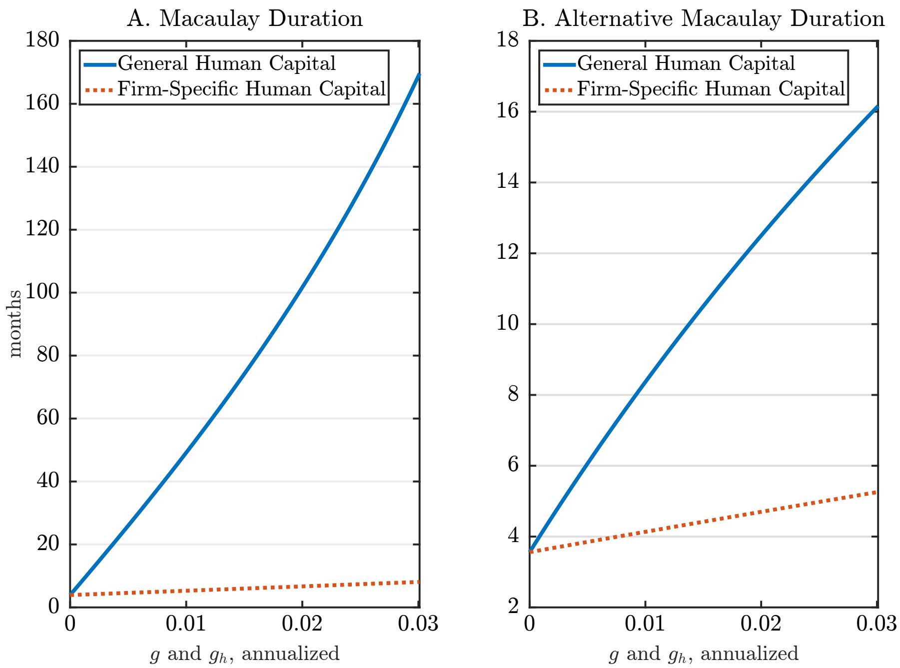

Figure 8: Employment Responses under Alternative Parameterizations

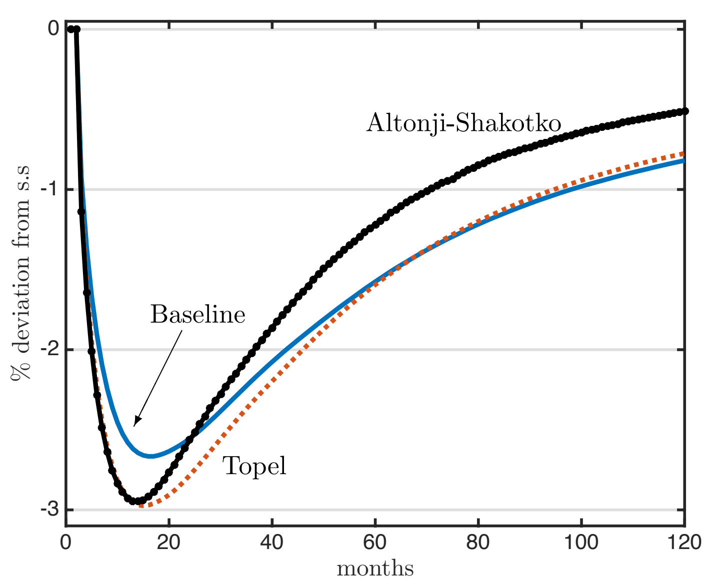

Figure 9: Robustness to Alternative Parameterization of Human Capital Process

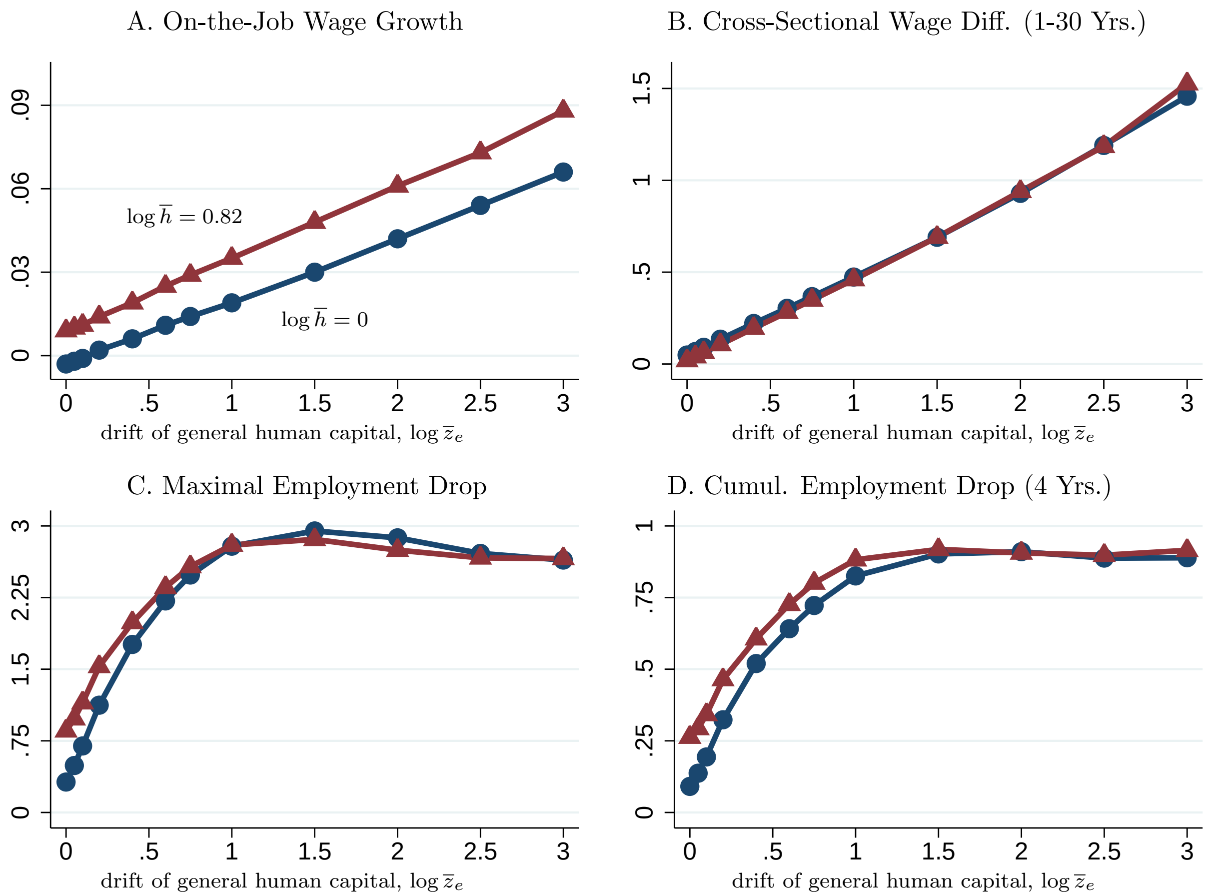

Figure 10: Duration of Surplus Flows

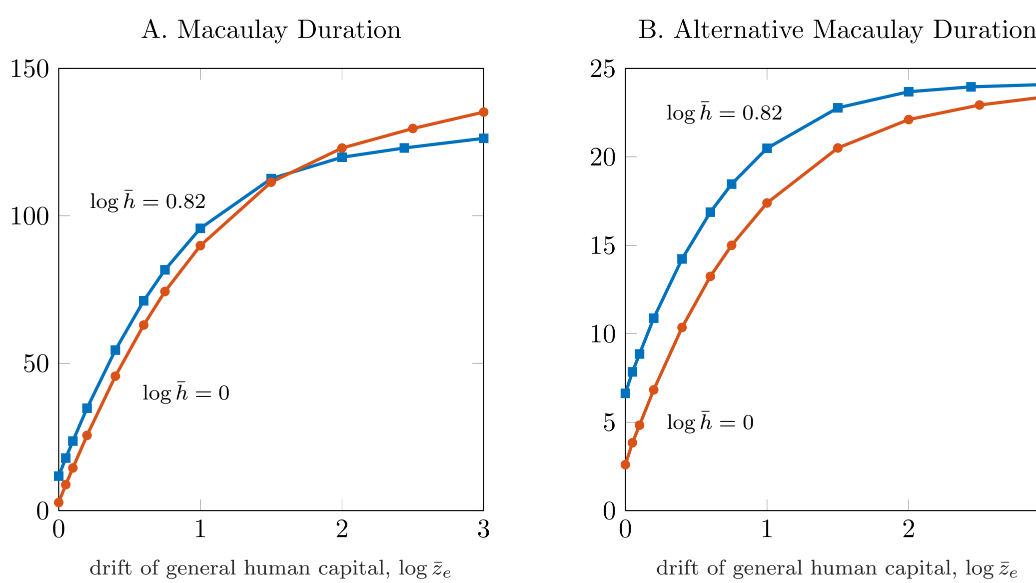

Figure 11: Displaced Worker Regressions (Robustness Checks)

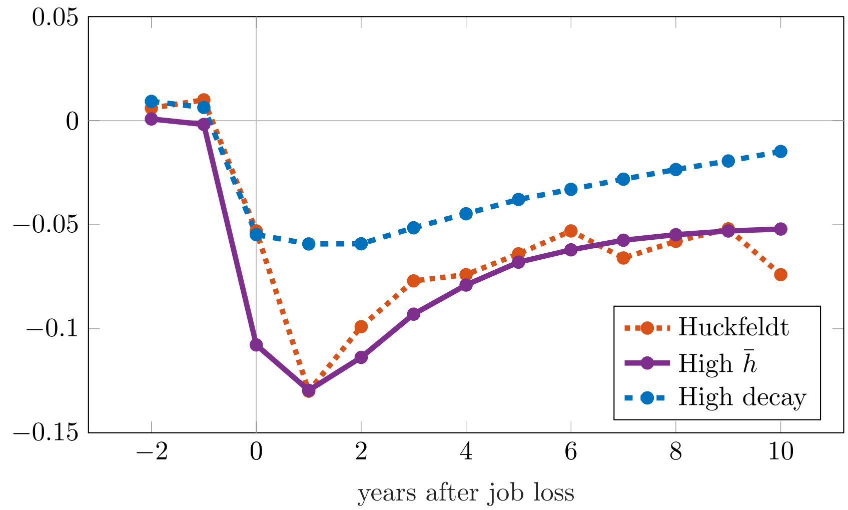

Figure 12: Employment Responses (Robustness Checks)

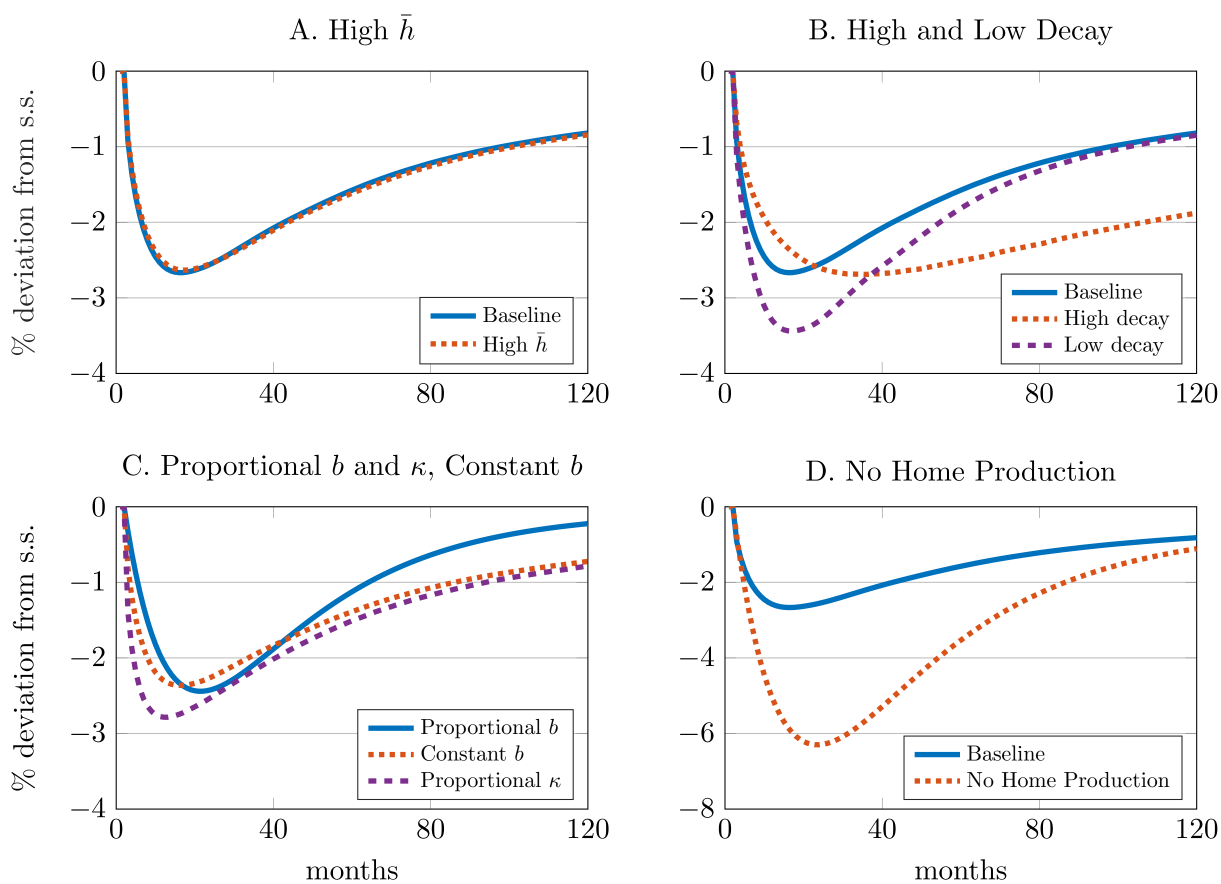

Figure 13: Distribution of Nonemployment Durations (Robustness Checks)

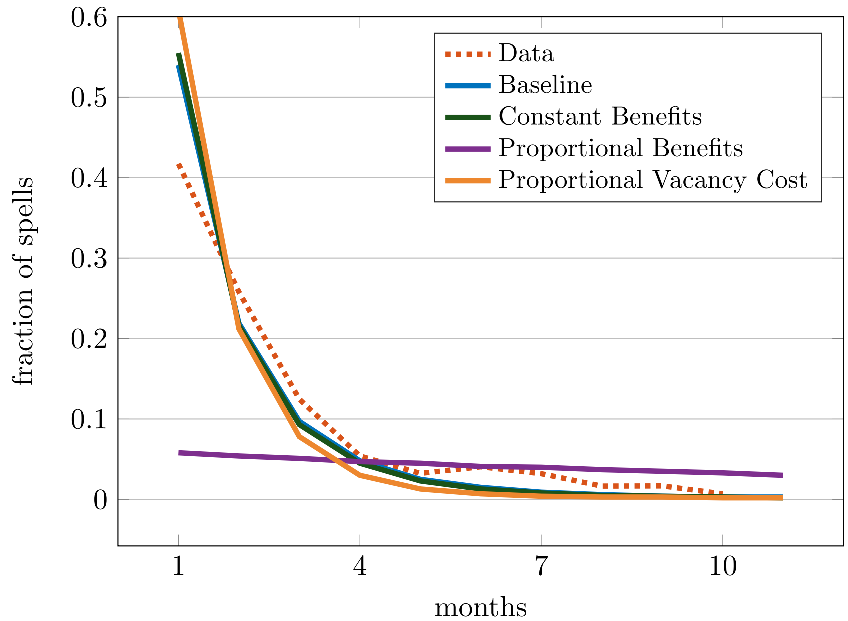

Figure 14: Employment, Consumption, and Wages

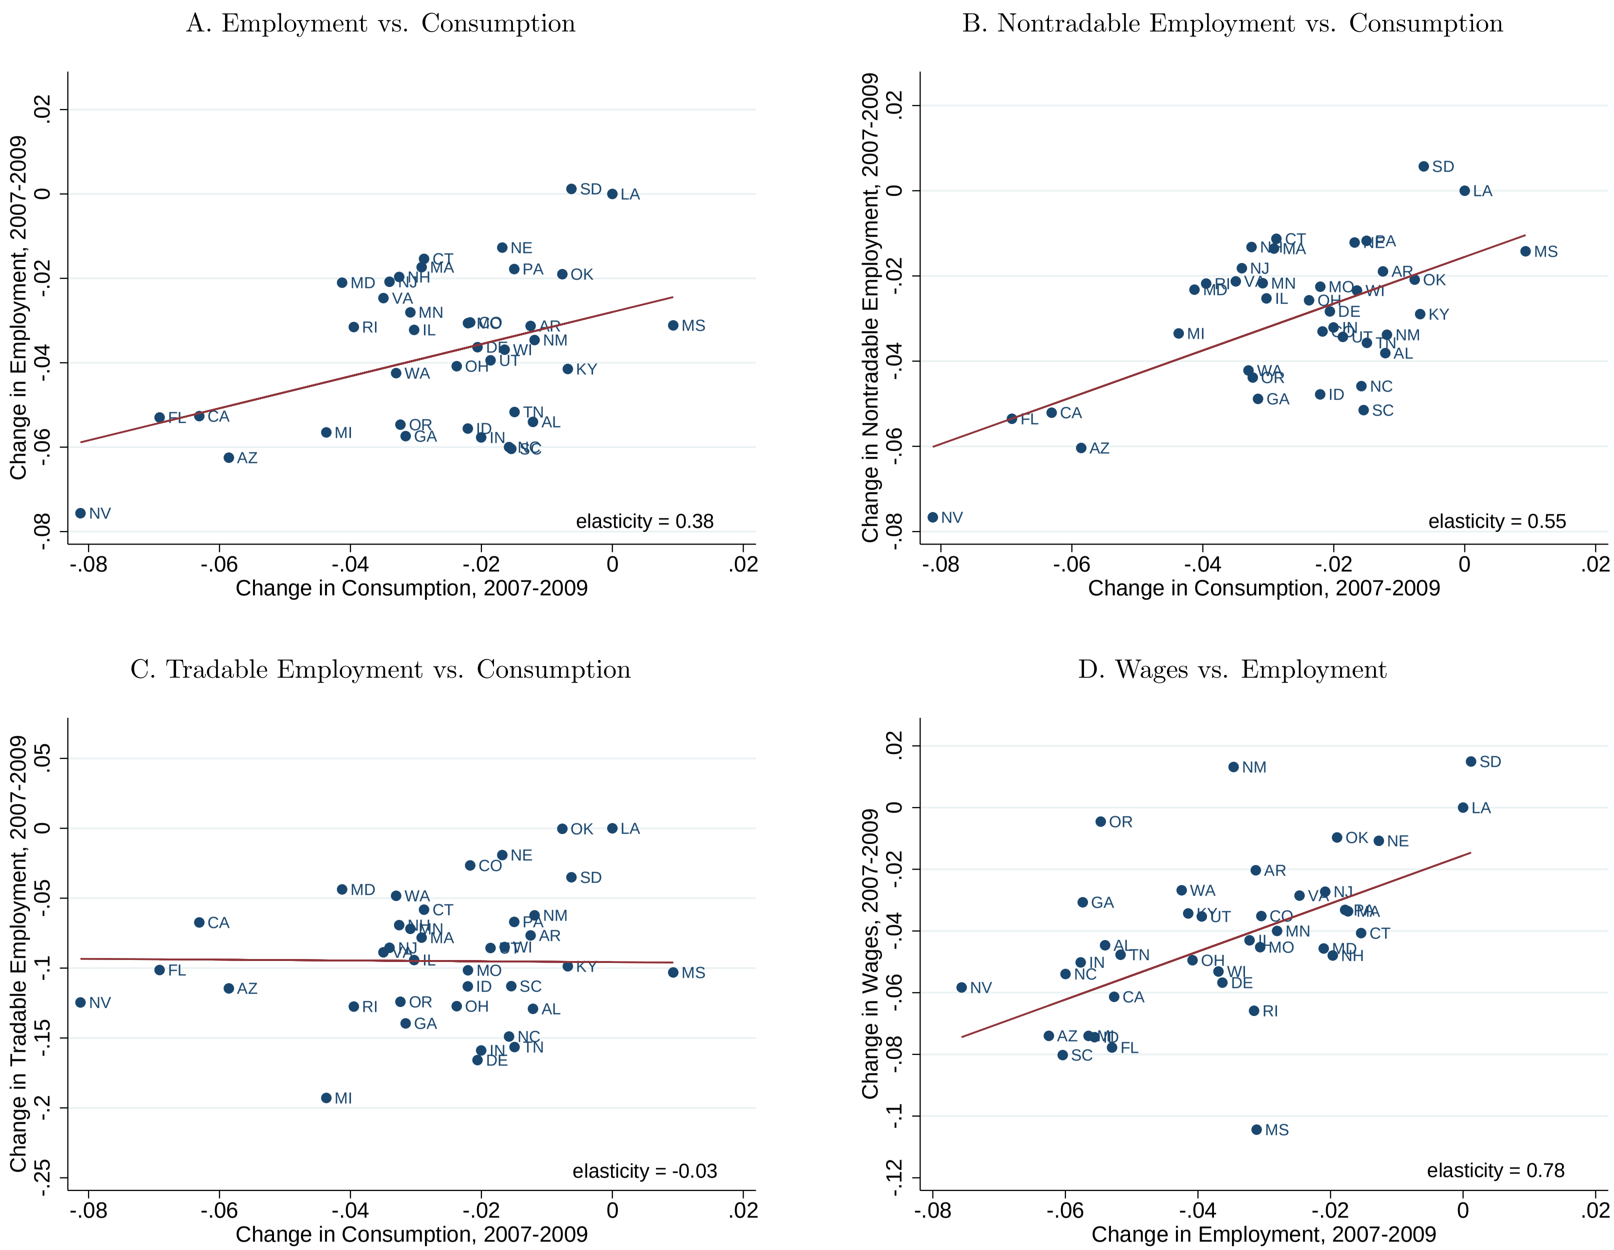

Figure 15: Low Cost of Sectoral Reallocation ($\nu = 0.5$)

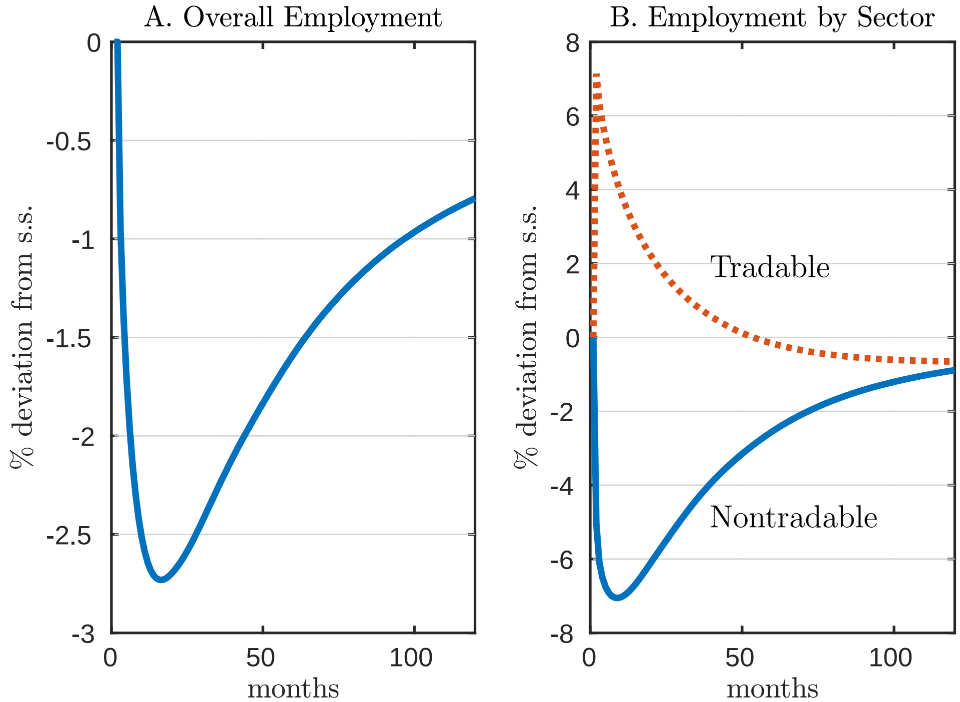

Figure 16: High Cost of Sectoral Reallocation ($\nu = 1$)