Debt Constraints and Employment*

Patrick Kehoe† Virgiliu Midrigan‡ Elena Pastorino§

June 2015

Abstract

In the Great Contraction, regions of the United States that experienced the largest declines in household debt to income also experienced the largest drops in consumption and employment. We develop a search and matching model that reproduces such patterns. Tighter debt constraints raise workers' and firms' discount rates, thus reducing match surplus, vacancy creation, and employment. Two ingredients of our model, on-the-job human capital accumulation and worker debt constraints, greatly amplify the drop in employment. On-the-job human capital accumulation implies that the returns to posting a vacancy are backloaded: the surplus from a match is thus more sensitive to changes in firm discount rates. Worker debt constraints amplify these effects further by preventing wages from falling too much. We show that the model reproduces the salient cross-sectional features of the U.S. data, including the comovement between consumption, house prices, debt-to-income as well as tradable and non-tradable employment.

Keywords: Search and Matching, Employment, Debt Constraints, Human Capital.

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

*We thank Eugenia Gonzalez Aguado, Julio Andres Blanco and Sergio Salgado Ibanez for excellent research assistance. Kehoe thanks the National Science Foundation for financial support.

†University of Minnesota and University College London, patrickjameskehoe@gmail.com.

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

§University of Minnesota and University College London, elenapastorino1@gmail.com.

1 Introduction

We develop a Diamond-Mortenson-Pissarides model with risk averse consumers, borrowing constraints, and upward sloping wage-tenure profiles and investigate its ability to account for some key features of the Great Contraction of the United States. This contraction was associated with a large fall in household debt to income, consumption, and employment. Moreover, as Mian Sufi (2009) and Midrigan and Philippon (2011) have noted, in the most recent recession, regions of the United States that experienced the largest swing in household borrowing also experienced the largest drops in output and employment. Finally, in the regions that experienced this large drop in employment, the drop was disproportionately larger in the nontraded goods sector. We show that our model can reproduce the salient features of this contraction.

A first key feature of our model is human capital accumulation through employment, which allows us to generate realistic wage-tenure profiles of the type documented by Buchinsky, Fougère, Kramarz, and Tchernis (2010). The returns to posting vacancies are *backloaded* with such upward-sloping wage and human capital profiles, and thus more sensitive to changes in debt constraints that alter an agent's discount rate. Vacancy creation and therefore employment is thus much more volatile in our framework than absent on-the-job human capital accumulation.

A second key feature of our model is that a tightening of debt constraints simultaneously changes both the firms' and workers' discount rates. Even though most of the drop in employment is accounted for by the *firms'* decision to post fewer vacancies, worker discount rates play a crucial role in equilibrium. Absent a change in worker discount rates, the fall in firm discount rates is offset by a sharp drop in equilibrium wages. This drop in wages mutes the response of vacancy creation and thus employment. In contrast, the drop in wages is much smaller when the worker's discount rate falls as well. Since debt-constrained workers value future human capital less, the wage that results from Nash bargaining between firms and workers does not fall as much and this amplifies the reduction in employment. Thus, worker-level debt constraints act as an endogenous source of wage stickiness and considerably propagate the effect of credit shocks.

As emphasized by Hall (2014), the key unresolved aspect of the Mortensen and Pissarides model is the nature of the force that depresses the payoff to job creation in recessions. In most of the papers in this tradition, the key force is a drop in productivity. As Hall also emphasizes, that explanation runs into two problems when applied to the recent recession. First, productivity did not fall much in the Great Contraction. Second, given that small fall in productivity, it is exceedingly difficult for the model to generate a large drop in output and employment. The driving force that we and Hall (2014) emphasize instead is an unexpected change of the consumer and firm discount rates. As noted, one key feature we add to the framework studied by Hall (2014) is human capital acquisition during employment. This feature greatly increases employment volatility in an otherwise standard search and matching model in which wages are continuously renegotiated.

Our work is related to a small literature that tries to link increases in financial frictions on the consumer side to economic downturns. In particular, Guerrieri and Lorenzoni (2010), Eggertson and Krugman (2011), and Midrigan and Phillipon (2011) study macroeconomic responses to a household-side credit crunch. All three of these papers find that a credit crunch has only a minor, if any, negative effect on employment unless wages are extremely sticky. Our analysis complements this work by studying a particular mechanism that prevents wages from falling in the aftermath of a consumer-side credit crunch.

Finally, our work is also related to the empirical analysis of Mian and Sufi (2011, 2013, 2014) who have documented a very robust relationship between changes in household balance sheets, consumption, employment and house prices in a cross-section of U.S. counties. Two key findings emerge from their work. First, consumption and employment positively comove from 2007 to 2009, a feature that is at odds with the predictions of standard models. Second, most of the reduction in employment associated with a tightening of household balance sheets in this period is accounted for by industries that sell non-tradable goods. The decline of tradable employment is large, but unrelated to measures of household debt or consumption. We show, in an extension of our model that allows for multiple goods, that our model successfully replicates these features of the data.

This paper is organized as follows. We first present a one-good, small open-economy version of the model. We think of the responses to a credit crunch in this model as representing those of an economy-wide shock that affects all regions in the U.S. and thus has no effect on relative prices. We then study an extension of the model with tradable and non-tradable goods in which relative prices respond to a credit crunch. We think of this extension as capturing the responses of an individual region of the U.S. and confront this version of the model with the Mian and Sufi evidence.

2 One-Good Economy

We consider a version of the Diamond-Mortenson-Pissarides model with risk-averse consumers that can borrow and save but are subject to debt constraints. A continuum of such consumers face idiosyncratic risks, the history of which we denote with s^t. The probability of any history is $\pi(s^t)$. Let $y(s^t)$ denote the income of an individual consumer. This income evolves over time both because of on-the-job human capital accumulation, as well as because of idiosyncratic productivity shocks. We consider a stochastic OLG environment in which each consumer survives from one period to another with a constant hazard ϕ. A measure $1-\phi$ of consumers replaces those that die each period.

We assume that consumers are organized into families that pool all of their members' idiosyncratic risks. Though risk-sharing is perfect within the family[1], each family faces debt constraints, changes in the severity of which affect the family's discount factor. Each member of the family maximizes the present value of the income it receives using the family's discount factor. We first describe the family's problem and then the problem of its individual members.

2.1 The Family Problem

The preferences of the family consisting of a measure of members is given by

$$\max_{c_t} \sum_{t=0} \beta^t u(c_t) \tag{1}$$

Because of full risk-sharing all members enjoy the same consumption, regardless of the idiosyncratic shocks they experience. (This type of risk sharing arrangement is familiar from the work of Merz (1995) and Andolfatto (1996).)

The budget constraint of the family is

$$c_t + qa_{t+1} = y_t + \pi_t + a_t \tag{2}$$

where q is the world price of a one period bond that pays off one unit in the next period, and a_{t+1} is the savings of the family where these and all other variables in the budget constraint are measured in traded goods. The term y_t is value of wages received by workers working in the market. The term π_t is the total profits net of vacancy posting costs that the family receives from all of the firms it owns. Here we suppose that the family owns the firms on the island on which its members reside. We also consider a scenario in which firms are owned by all the consumers in the world.

In addition to the budget constraint, the family faces an exogenous sequence of borrowing limits $\bar{d}_t$ which constrain borrowing according to

$$a_{t+1} \geq -\bar{d}_t. \tag{3}$$

[1]We also solved a version of the model similar to that in Krusell et al. (2010) in which consumers can only save with uncontingent bonds. This model has an unappealing feature. Because of human capital accumulation, the wage that solves the Nash bargaining problem is non-monotonic in individual assets. In particular, there is a region of the state space in which higher assets reduce wages. Anticipating this some consumers are deterred from savings because by doing so will decrease their future wages. This features led to odd patterns of wages, savings, and sometimes even the existence of equilibrium.

The first order conditions imply that consumption satisfies

$$u'(c_t) = \beta u'(c_{t+1})/q + \theta_t \tag{4}$$

where θ_t is the multiplier on the debt constraint. We let

$$Q_{0,t} = \frac{\beta^t u'(c_t)}{u'(c_0)} \text{and } Q_t = \frac{\beta u'(c_{t+1})}{u'(c_t)}$$

denote the effective discount factors of the family from period 0 to t and from period t to $t+1$ respectively. This will be the rate that individual members of the family use when solving their individual problems.

Here, for ease of notation, we have represented the insurance arrangements using a family construct. It should be clear that we can also represent these insurance arrangements by having appropriately defined contingent claims markets.

The objective function of a individual member of the family, henceforth worker, is

$$\sum_{t=0} \sum_{s^t} \phi^t Q_{0,t} \mu_t(s^t) \left[e_t(s^t) w_t(s^t) + (1 - e_t(s^t)) b \right].$$

where $e_t(s^t)$, which indicates the employment of an individual, is either 0 or 1 and b is the home production of a non-employed worker. Here it is important to recall that s^t denotes the history of idiosyncratic shocks of a given worker. Since the family consists of a measure 1 of such workers both the total market income and the total home production of the family is deterministic and is given by $y_t = \int \left[e_t(s^t) w_t(s^t) + (1 - e_t(s^t)) b \right] \mu_t(s^t) ds^t$.

The objective function of a individual firm owned by the family is

$$\sum_{t=0} \sum_{s^t} \phi^t Q_{f0,t} \mu_t(s^t) \pi_t(s^t). \tag{5}$$

where $\pi_t(s^t)$ is the profits net of vacancy costs of firms.

We imagine that the collection of islands that form the open economy we focus on are small in the world economy. In this world economy there are other collections of islands with symmetric economies. In each such collections of islands consumers are members of families as our small open economy with identical utility functions, inclusive of the discount factor β. Each collection of islands also have entrepreneurs with large endowments that borrow and lend to all consumers in the world and have a discount factor q.

We consider two ownership structures for firms. In one the firms on an island are owned by the families on that island so that $Q_{f0,t} = Q_{0,t}$. In the other, firms on an island are owned by all consumers in the world so that $Q_{f0,t} = \beta^t$.

An individual member of the family, henceforth worker, chooses whether to work or not, $e(s^t) \in \{0, 1\}$, in order to maximize the present value of income

$$y(s^t) = e(s^t)w(s^t) + (1 - e(s^t))b,$$

which consists of wages $w(s^t)$ when employed, and home production b when non-employed. We assume here that the amount non-employed workers earn is constant and thus independent of the worker's characteristics. We relax this assumption in a robustness section below. Since all members have a constant survival hazard ϕ, the objective of the individual member is

$$\max_{e(s^t)} \sum_{t=0} \sum_{s^t} \phi^t Q_t \pi(s^t) y(s^t). \tag{6}$$

The constraints on this problem are the search and matching frictions we describe below.

2.2 Workers and Firms

A worker is characterized by a productivity level z which gives the amount of output the worker produces if matched with a firm. We think of z as representing the individual worker's human capital. We let $w_t(z)$ denote the wage received by a worker with productivity z. The wage is the solution to a Nash bargaining problem described below. Newborn workers enter the economy endowed with a draw of z from a log-normal distribution:

$$\log(z) \sim N\left(0, \sigma_z^2/(1 - \rho_z^2)\right).$$

An employed worker's human capital evolves according to:

$$\log z' = (1 - \rho_z)\mu_z + \rho_z \log z + \sigma_z \varepsilon' \tag{7}$$

where ρ_z determines the persistence, σ_z the volatility and μ_z the mean of the process, while ε' is a Gaussian disturbance with zero mean and unit variance. A non-employed worker's human capital evolves according to a similar AR(1) process with the same persistence and volatility but with a mean equal to 0:

$$\log z' = \rho_z \log z + \sigma_z \varepsilon' \tag{8}$$

We represent below the Markov processes in (7) and (8) as $F_e(z'|z)$ and $F_u(z'|z)$.

Intuitively, a newly-born worker's initial productivity is drawn from a distribution with the lower mean and thus gradually increases over time whenever the worker is employed with a speed determined by μ_z. The employed worker's productivity thus drifts up over time. In contrast, a worker that has been employed for a number of periods and has a relatively high productivity will experience on average a loss of human capital whenever non-employed since its productivity will drift down over time. By setting appropriately the parameters of the two processes for z_t, we can match the gains from tenure and experience as well as the loss in wages a worker experiences upon separation. (Here we do not formally distinguish between unemployment and non-participation but when quantifying the model, we think of non-employment as capturing both states.)

We assume a standard aggregate matching function $M(u_t, v_t)$, which represents the measure of matches in a period with u_t non-employed workers and v_t vacancies. We assume M is given by

$$M(u, v) = Bu^{\eta}v^{1-\eta},$$

where B is the efficiency of the matching technology. Let $\theta_t = v_t/u_t$ denote the market tightness. The probability that a vacant job is filled in the current period is

$$\lambda_{f,t} = \frac{M(u_t, v_t)}{v_t} = B\left(\frac{u_t}{v_t}\right)^{\eta} = B\theta_t^{-\eta}.$$

Similarly, the probability that a non-employed worker finds a match is

$$\lambda_{w,t} = \frac{M(u_t, v_t)}{u_t} = B\left(\frac{v_t}{u_t}\right)^{1-\eta} = B\theta_t^{1-\eta}.$$

We assume that individual matches are exogenously destroyed with probability σ. In addition, since home production b is constant while z evolves over time, some matches, those that no longer have a positive surplus, are endogenously destroyed as well.

The value of an employed worker is given by its current wage as well as the discounted sum of future output:

$$W_t(z) = w_t(z) + \phi\frac{Q_{t+1}}{Q_t}(1-\sigma)\int_{z'}\max\left[W_{t+1}(z'), U_{t+1}(z')\right]dF_e(z'|z) \tag{9}$$

$$+\phi\frac{Q_{t+1}}{Q_t}\sigma\int_{z'}U_{t+1}(z')\,dF_e(z'|z).$$

Notice the max operator which reflects the decision of whether to continue a relationship, as well as the family's discount factor Q_{t+1}/Q_t with which the workers discount future output. Similarly, the value of a non-employed worker is given by the amount it produces at home, b as well as the discounted sum of future output:

$$U_t(z) = b + \phi\frac{Q_{t+1}}{Q_t}\lambda_{w,t}\int_{z'}\max\left[W_{t+1}(z'), U_{t+1}(z')\right]dF_u(z'|z) \tag{10}$$

$$+\phi\frac{Q_{t+1}}{Q_t}(1-\lambda_{w,t})\int_{z'}U_{t+1}(z')\,dF_u(z'|z)$$

Notice here that the worker may choose to turn down an offer even in the event it matches with a firm which occurs with probability $\lambda_{w,t}$.

Each firm pays a vacancy cost to form a match, produces output when matched, and pays dividends to the family. The objective of each firm is to maximize the discounted value of dividends using the discount factor Q_t of the family it belongs to. The value of a vacancy is given by the firm's current profits, output net of wages, as well as the present discounted value of future wages:

$$J_t(z) = z - w_t(z) + \phi \frac{Q_{t+1}}{Q_t}(1-\sigma)\int_{z'} \max\left[J_{t+1}(z'), 0\right] dF_e(z'|z). \tag{11}$$

Notice again that this formulation captures the possibility that a match no longer yields a positive value to the firm and is thus destroyed. The value of posting a vacancy is

$$V_t = -\kappa + \phi \frac{Q_{t+1}}{Q_t}\lambda_{f,t}\int_{z'} \max\left[J_{t+1}(z'), 0\right] dF_u(z'|z)\, d\tilde{n}_t^u(z). \tag{12}$$

There is a large number of potential entrants that can post a vacancy cost κ to create a new vacancy. Let $n_t^u(z)$ denote the measure of non-employed workers at the beginning of period t. Let $\tilde{n}_t^u(z) = n_t^u / \int dn_t^u(z)$ denote the distribution of productivity among the non-employed. Free entry drives the expected value of posting a vacancy to 0:

$$0 = -\kappa + \phi \frac{Q_{t+1}}{Q_t}\lambda_{f,t}\int_{z'} \max\left[J_{t+1}(z'), 0\right] dF_u(z'|z)\, d\tilde{n}_t^u(z). \tag{13}$$

The free entry condition pins down the vacancy-unemployment ratio θ_t and thus the flows out of non-employment.

We assume that wages are renegotiated period by period and are set by a generalized Nash bargaining protocol and thus solve

$$\max_{w_t(z)} \left[W_t(z) - U_t(z)\right]^\gamma J_t(z)^{1-\gamma}$$

or

$$\frac{\gamma}{W_t(z) - U_t(z)} = \frac{1-\gamma}{J_t(z)}.$$

Here γ represents the worker's bargaining weight. In equilibrium, firms in both sectors post vacancies until $V_t = 0$. We refer to (13) as the *free entry condition.*

Given an exogenous sequence of debt constraints $\{\bar{d}_t\}$ and world bond price q, an *equilibrium on an island in the debt constraint economy* is a set of allocations $\{c_t, a_{t+1}, b_t, y_t, \pi_t\}$ for the family, a set of discount rates for workers and firms $\{Q_t, Q_{ft}\}$, prices measures of employed workers and non-employed workers $\{n_{et}(z), n_{ut}(z)\}$,wages functions and vacancies $\{w_t(z_t), v_t\}$, matching rates for workers and firms $\{\lambda_{w,t}, \lambda_{f,t}\}$, nonemployment rates $\{u_t\}$, along with value functions of workers and firms $\{W_t(z_t), U_t(z_t), J_t(z_t), V_t\}$ such that i) given the sequences of home production, wage income, and profit income $\{b_t, y_t, \pi_t\}$ the allocations $\{c_{Nt}, c_{Tt}, a_{t+1}\}$ solve the family problem, ii) given the wage functions, the family discount factors, and the worker matching rates, the value functions of employed workers and the value functions for non-employed workers satisfy (9) and (10), iii) given the wage functions, the firm discount factors, and the firm matching rates, the value functions for firms satisfy (11) and (12) and the free entry conditions (13), iv) the family's profit income is

$$\pi_t = \int [z - w_t(z)]\, dn_{et}(z) - v_t\kappa$$

v) the income of the family from employed and nonemployed workers is $y_t = \int w_t(z)dn_{et}(z) + \int bdn_{ut}(z)$.

3 An Equivalence Result

Here we have described our economy as having debt constraints. Here we show that our economy is equivalent in terms of consumption, saving, and all labor market outcomes to one in which families own houses and a families' borrowing is subject to collateral constraints based on the value of their houses.

The preferences of the family are

$$\max_{c_t} \sum_{t=0} \beta^t[u(c_t) + \psi_t v(h_t)] \tag{14}$$

where c_t is the consumption of any of its members and h_t is the amount of housing consumed. Because of full risk-sharing all members enjoy the same amount of housing and consumption, regardless of the idiosyncratic shocks they experience. (This type of risk sharing arrangement is familiar from the work of Merz (1995) and Andolfatto (1996).)

The family faces a budget constraint

$$c_t + a_{t+1} + p_t h_{t+1} = (1 + r)a_t + p_t h_t + \int y_{it} di + T_t \tag{15}$$

where a_{t+1} are the family's end-of-period financial assets, p_t is the price of housing, r is the exogenously-given interest rate at which the family saves with the rest of the world. The term $\int y_{it}di$ represents the total income of all the members of the family, which includes wages for those members that are employed and home production for those that are not employed. Finally, T_t are the profits net of vacancy posting costs the family receives from its ownership of firms in the economy. As we show below, since firms are owned by families, they discount future flows using the same discount rate that workers use to discount future wages.

Finally, the family faces a collateral constraint that limits the maximum amount it can borrow to a fraction χ of the value of the family's home:

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

where χ is the maximum loan to value ratio.

Our model of houses is very simple. There is a fixed supply of houses, normalized to 1, and each unit of a house delivers one unit of housing services each period. The housing stock does not depreciate. The parameter ψ_t in the utility function (14) 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) the amount of debt the family can borrow.[2]

Given an exogenous sequence of taste parameters $\{\psi_t\}$ and world bond price q, an *equilibrium on an island in the collateral constraint economy* is a set of allocations $\{c_t, , h_t, a_{t+1}, b_t, y_t, \pi_t\}$ for the family, a set of discount rates for workers and firms $\{Q_t, Q_{ft}\}$, prices for houses $\{p_t, \}$, measures of workers in the market and in non-employment $\{n_{et}(z), n_{ut}(z)\}$,wages functions and vacancies $\{w_t(z_t), v_t\}$, matching rates for workers and firms $\{\lambda_{w,t}, \lambda_{f,t}\}$, nonemployment rates $\{u_t\}$, along with value functions of workers and firms $\{W_t(z_t), U_t(z_t), J_t(z_t), V_t\}$ such that *i*) given the sequences of home production, wage income, and profit income $\{b_t, y_t, \pi_t\}$ and the prices $\{p_t\}$ the allocations $\{c_t, c_t, h_t, a_{t+1}\}$ solve the family problem, *ii*) given the wage functions, the family discount factors, and the worker matching rates, the value functions of employed workers and the value functions for non-employed workers satisfy (9) and (10), *iii*) given the wage functions, the firm discount factors, and the firm matching rates, the value functions for firms satisfy (11) and (12) and the free entry conditions (13), *iv*) the home production constraint

$$b_t = \int b dn_{ut}(z)$$

v) market clearing in houses $h_t = 1, vi$) the family's profit income is

$$\pi_t = \int [z - w_t(z)]\, dn_{et}(z) - v_t \kappa$$

vi) the income from market work and home production of the family is $y_t = \int w_t(z) dn_{et}(z) + \int b dn_{ut}(z)$.

We claim that the debt constraint economy is equivalent to the collateral constraint economy. We first show that given a debt constraint economy with a sequence of borrowing limits $\{\bar{d}_t\}$ we can construct a sequence of taste parameters $\{\psi_t\}$ for the collateral constraint economy such that, neglecting the price of housing and the allocation of housing, the two

[2] We have also studied a version of the model in which shocks to the loan-to-value ration χ are the source of credit shocks and have obtained virtually identical results.

equilibria coincide. We then show that given a collateral constraint economy with taste parameters $\{\psi_t\}$, we can construct a sequence of borrowing limits $\{\bar{d}_t\}$ such that, neglecting the price of housing and the allocation of housing, the two equilibria coincide.

Let us start with a debt constraint economy and note that the solution to the family problem can be summarized by the first order conditions

$$qu'(c_t) = \beta u'(c_{t+1}) + q\theta_t \tag{16}$$

$$\theta_t \left[a_{t+1} + \bar{d}_t\right] = 0 \tag{17}$$

the budget constraint (2) and the debt constraint , where θ_t is the multiplier on the debt constraint. In the collateral constraint economy the solution to the family problem evaluated at the equilibrium value of $h_t = 1$ is summarized by the first order conditions

$$qu'(c_t) = \beta u'(c_{t+1}) + q\theta_t \tag{18}$$

$$\theta_t \left[a_{t+1} + \chi p_t\right] = 0 \tag{19}$$

$$\beta\psi_{t+1}v'(1) + \beta p_{t+1} + \theta_t \chi p_t = u'(c_t)p_t. \tag{20}$$

and the budget constraint (15). Note first that with $h_t = 1$ the budget constraints in the two economy coincide. Next, note that given the allocations and multipliers from the debt constraint economy, if we set the price of houses $p_t = \bar{d}_t/\chi$ and set ψ_{t+1} according to

$$\psi_{t+1} = \frac{1}{\beta v'(1)}\left[u'(c_t)p_t - \beta u'(c_{t+1})p_{t+1} - \theta_t \chi p_t\right]$$

then if $\{c_t, b_t, a_{t+1}, \theta_t\}$ satisfy (16) and (17) then these same variables along with p_t satisfy (18)-(20) and the collateral constraint.

Now for the rest of the construction, since c_t coincide in the two economies so do the discount factors $\{Q_{t,t+1}, Q_{ft,t+1}\}$. A moment's reflection then makes clear that if the consumer and firm discount factors coincide then so do the allocations and wages in the labor market.

For the converse, we start with a collateral constraint economy with taste parameters $\{\psi_t\}$ and construct a sequence of borrowing limits $\{\bar{d}_t\}$ such that, neglecting the price of housing and the allocation of housing, the two equilibria coincide. Here we simply set $\bar{d}_t = \chi p_t$ and note that the budget constraints and first order conditions for the two equilibria coincide. Thus, the discount factors also coincide and hence so do the allocations and wages in the labor market.

We summarize this discussion with a propostion.

Proposition 1. *The debt constraint economy is equivalent to the collateral constraint economy.*

Note that in the collateral constraint economy we have time-varying taste parameters $\{\psi_t\}$ but a constant maximal loan-to-value ratio χ. It is immediate that a similar equivalence proposition holds if we consider a collateral constraint economy with a constant taste parameter ψ and time-varying maximal loan-to-value ratios $\{\chi_t\}$.

Next, note there is no aggregate uncertainty in this economy: we study the effect of a one-time, unanticipated, sequence of changes in ψ_t that lead to a drop in house prices, credit and consumption of the magnitude that we have seen in the U.S. data.

We assume that the rate of time preference is sufficiently small relative to the interest rate, $\beta(1+r) < 1$, so that the debt constraint binds both in steady state and all of our experiments. Given a particular path for ψ_t, the Euler equation (20) together with the borrowing and budget constraints, as well as the equilibrium condition $h_{t+1} = 1$, pin down a path of consumption c_t and therefore the family's marginal valuation of date t output

$$Q_t = \beta^t \frac{u'(c_t)}{u'(c_0)}.$$

and therefore the one-period discount factor Q_{t+1}/Q_t. Clearly, a tightening of credit constraints that leads to a transitory decline of consumption in period t reduces the one-period discount factor.

Here, for ease of notation, we have represented the insurance arrangements using a family construct. It should be clear that we can also represent these insurance arrangements by having appropriately defined contingent claims markets.

4 Quantification and Results

We next discuss how we have chosen parameters for our model, the model's steady state implications, and the responses of the economy to a credit crunch. We show that employment falls sizably in response to a credit crunch in our model. We then argue that two key ingredients account for the employment drops: on-the-job human capital accumulation and the fact that not only firms but also workers are subject to the credit shock. Absent any of these two ingredients, the employment responses are very small. Finally, we study a robustness check on our model in which we assume that home production is proportional to a worker's productivity.

4.1 Quantification

The model period is one quarter. We set the discount factor equal to $\beta = 0.94^{1/4}$ and the interest rate equal to $1 + r = 0.96^{1/4}$, thus corresponding to a 4% annual rate. We set the survival rate ϕ so that $1 - \phi = 1/160$ so that households are in the market for 40 years on average. The probability of separation $\sigma = .1$ is set so that the average employment spell is about 2 and 1/2 years as in Shimer (2005). The bargaining weight γ is set to 1/2 and we set the elasticity of the matching function $\eta = \gamma$.

The utility function is

$$u(c_t) = \frac{c_t^{1-\alpha}}{1-\alpha}$$

A large literature (see Guvenen 2006 and the references therein) finds that the elasticity of intertemporal substitution (EIS) is very low when estimated using data on households, on the order of 0.1 to 0.2. We follow this literature and set this elasticity, $1/\alpha$, equal to 0.2. We have found that the responses of employment in our model to a credit shock are approximately linear in the size of the shock: doubling (halving) the value of α would lead to doubling (halving) the employment responses.

We have six additional parameters that are jointly chosen so that the model matches exactly six statistics in the data. These parameters are: κ, the fixed cost of posting a vacancy; B, the efficiency of the matching function; ρ_z, the persistence of productivity shocks; σ_z, the standard deviation of productivity shocks; μ_z, the parameter governing the evolution of z for an employed worker and thus the returns to employment; and b, the home production parameter. The targets we use to pin down these moments are: a vacancy-unemployment ratio of 1 (as in Shimer (2005), this simply reflects a choice of units and is otherwise inconsequential), an employment-population ratio of 80%, a ratio of home production to mean wages of 40%, following Shimer (2005), a standard deviation of the log of initial wages (in the first year of a given employment spell) of 0.94, which we compute using the PSID data, and a standard deviation of changes in log wages of 0.21 per year, as computed by Floden and Linden (2001). The final statistic we require the model to match are the average returns to employment in the data which we calculate using an indirect inference approach that we discuss in detail next.

We calculate average returns to employment using the parameter estimates of Buchinsky, Fougère, Kramarz, and Tchernis (2010) who estimate a structural model of worker job participation and mobility decisions allowing for rich sources of heterogeneity among workers in the PSID. The structural model allows these researchers to identify the parameters of a Mincerian wage equation that relates a worker's wages to its demographic characteristics, education, as well as the history of past employment, including its tenure (number of years with a given firm) and experience (number of years in the labor market) at the current and its previous jobs:

$$\log(w_{it}) = c_i + x_{it}'\beta + f(\text{experience}_{it}) + g(\text{tenure}_{it}) + J_{it} + \varepsilon_{it}$$

where

$$J_{it} = \sum_{l=1}^{M_{it}} \sum_{k=1}^{4} \left(\phi_k^0 + \phi_k^s \text{tenure}_i^l + \phi_k^e \text{experience}_i^l \right) d_{ki}^l$$

summarizes the history of the worker's past employment. We use these estimates, together with the profiles for experience and tenure implied by *our* model, to evaluate a predicted wage

$$\widehat{\log(w_{it})} = f(\text{experience}_{it}) + g(\text{tenure}_{it}) + \widehat{J}_{it}$$

for all individuals in our model's simulation. We use the average growth rate of these predicted wages as our sixth target in the calibration and thus choose parameters so as to ensure that the growth rate of actual wages in our model is equal to the growth rate of the predicted wages, about 5.2% per year in our benchmark parameterization.

Intuitively, some parameters in the model have relatively more importance for some statistics in the data. Roughly, the fixed cost of posting vacancies pins down the steady state vacancy-to-unemployment ratio, the efficiency of the matching function pins down how often non-employed agents are matched with firms and thus the employment-population ratio, the parameter governing the law of motion of z for employed workers, μ_z, pins down the average growth rate of wages, while the persistence and standard deviations of shocks to z determines the unconditional dispersion and the volatility of changes in wages.

Table 1 summarizes our parameterization strategy and shows the moments used in our calibration, the parameters that we assigned and those that we endogenously chose. As the table shows, the model matches all 6 parameters exactly given that the model is exactly identified.

Figure 1 illustrates our model's implications for returns to employment by showing the path of wages of a typical worker that experiences several non-employment spells over the life-time, and contrasts the model's wage paths with those predicted by the BFKT (2010) estimates. We initialize the simulation in Figure 1 for an agent with zero experience, a mean productivity equal to the mean of newborn workers' productivity and assume no shocks to productivity over the lifetime. Notice that our model's wage path tracks that implied by the empirical estimates fairly closely.

4.2 Steady State Implications

Table 1 also shows several additional implications of the model. Notice in the lower-left panel of Table 1 that 18% of workers have a level of productivity below the amount they would produce at home. This reflects the backloaded nature of returns to working in this environment. Also notice that very few, 0.2% of matches are voluntarily destroyed, reflecting the fact that a worker's productivity typically drifts up. However, since worker's productivity is very dispersed, a substantial fraction, 0.278 of matches yield negative surplus and are thus not formed.

Our model also predicts , since productivity of the non-employed gradually drifts down over time, that wages typically fall after a non-employment spell. In our model, the typical wage drop after a non-employment spell is about 2%, reflecting the fact that non-employed workers do not stay non-employed too long. In contrast, workers that are non-employed for a year (two years) experience much larger drops of 6.1% (8.8%). While the wages drops predicted by the model are somewhat smaller than those in studies of the earnings losses of displaced workers, they are consistent with micro-economic estimates that study all employment to non-employment to employment transitions. For example, as Figure 1 shows, the estimates of Buchinsky, Fougère, Kramarz, and Tchernis (2010) which are based on all such transitions, are similar with those predicted by our model.

Consider next the four panels of Figure 2 which display the policy and value functions. The upper-left panel shows that when a firm hires a worker the profits $z - \omega(z)$ are negative for a range of human capital z. The bottom left panel shows that the value of a newly filled vacancy $J(z)$ increases as the productivity of the worker in that vacancy increases. Comparing the top and the bottom left hand panels show that a firm creates matches with consumers even for levels of productivity in which profits are initially negative. The reason again is the backloading of returns to employment. In this sense, there is an additional investment aspect to the relationship between firms and consumers in a match in addition to the cost of posting the vacancy.

The upper-right panel shows the value $\omega(z)$ of wages for different values of human capital z . The bottom right panel shows the value of working relative to that of being unemployed. As the graph shows, an unemployed consumer accepts jobs even when its wages are below the current value of home production b, so that "worker profits" are negative. As discussed above, the consumer finds it optimal to do so because as z increases with time in the job the value of wages will tend to increase above b. Here also, in contrast to standard search models, there is also an investment component to the consumer from forming a match with the firm.

Figure 3 displays the measures of employed workers $n_e(z)$ and unemployed workers $n_u(z)$ in the steady state. The top panel shows that reservation human capital level $\bar{z}$ is less than one. The lower panel shows that there are two humps in the measure of human capital among non-employed workers. Any non-employed worker with human capital above $\bar{z}$ will accept a job. Such workers are typically ones that have held a job relatively recently. The non-employed workers with human capital less than $\bar{z}$ will not accept a job. Such workers have relatively longer periods of time before they find jobs and end up with rather long durations of unemployment.

4.3 A Tightening of Credit

We now turn to our experiment. In it we suppose that there is an unanticipated drop in the taste for houses ψ_t. This drop is unexpected prior to the first period. Consumers have perfect foresight afterwards. We choose the size of the drop in preferences for housing so that in equilibrium the tighter credit leads to an immediate consumption drop of 5% after which it reverts to its original steady state at a rate of 10% per quarter. Our choice of a mean-reversion rate of 10% is chosen so as to match the speed of postwar consumption recoveries in the data.

Figure 4 displays the resulting path of consumption c_t, house prices, the one-period discount factor $\beta Q_{t+1}/Q_t$, as well as the implied annual discount rate. We note that house prices fall by about 20% on impact and recover much more gradually than consumption. The quarterly discount factor falls from 0.985 to about 0.960, implying that the annualized discount rate increases on impact from about 6 to 15%. The size of the discount rate shock we consider is thus similar to that studied by Hall (2014) who considers the impact of an increase in the discount rate from 10 to 20% in a standard search and matching model. Hall (2014) argues that such large swings in discount rates are necessary to rationalize the behavior of the stock market in recent years.

Figure 5 displays the path for employment relative to that of consumption. The 5% fall in consumption is accompanied by a 2% fall in employment. Note that the employment drop is more persistent than the consumption drop. A simple way to measure the relative persistence is to compare the relative cumulative impulse responses (the area under the impulse responses) for consumption and employment at long horizons. Overall, the cumulative impulse response for employment is 92% of the cumulative impulse response for consumption.

The next two figures help shed light on the mechanism behind the impulse response for employment. The upper panels of Figure 6 shows that the credit tightening leads to a 15% drop in the vacancy-unemployment ratio on impact, which reduces a worker's matching probability by about 8% initially. Both of these recover very gradually. The lower panels of the figure show that the credit tightening is also associated with a transitory spike in the number of job separations, from about 10% per quarter, to 10.5%. Finally, the fraction of profitable matches reduces somewhat as well, from about 72.5% to 70% on impact.

We next quantify the importance of each of these three factors in accounting for the employment drop. We do so using an approach that builds on Shimer (2012). The transition law for total employment can be written

$$E_{t+1} = (1 - s_t)E_t + \lambda_{wt} a_t (1 - E_t)$$

where the separation rate s_t is the sum of exogenous and endogenous separations, λ_{wt} is the probability of any given worker finding a match with a firm, and a_t is the acceptance rate. We construct three counterfactual series in which we vary one of s_t, λ_{wt}, and a_t leaving the others at the steady state value. In Figure 7 we see that the decline in the matching rate λ_{wt} account for the lion share of the movements in employment while, except for the first couple of periods, the increase in the separation rate accounts for little. Interestingly, the decline in the acceptance rate accounts for only a modest share of the decline in employment early on but is much more persistent and thus accounts for about half of the employment decline about 10 years after the shock.

4.4 Two Key Ingredients of Our Model

We have introduced two ingredients in an otherwise standard search and matching model, on-the-job human capital accumulation and worker debt constraints. We argue next that each of these ingredients is critical in generating sizable employment responses to a credit crunch.

On-the-Job Human Capital Accumulation. Consider first the role of on-the-job human capital accumulation. We compare the implications of our model with those of an otherwise identical model in which z is constant over time and the returns to work are thus equal to zero. Figure 8 shows that employment responds very little to the credit crunch: it falls by only about 0.4%, thus about one-fifth of the drop in our Benchmark model with returns to work.

We note that these results are similar to those of Hall (2014) who studies the effects of a permanent doubling of the discount rate (from 10% to 20%) in a standard search and matching framework and finds that steady-state unemployment only increases from 5.8 to 5.88%. Discount rate shocks thus have very little effect on vacancy creation and employment in standard search models.

To see why this is the case, consider a simplified continuous time version of our model.[3] Assume that employed workers' productivity increases at a rate g so that $\mathrm{d}z = gz\mathrm{d}t$ if a worker is employed and zero otherwise. The values of the employed, non-employed and firms

[3]The expressions we derive below convey a similar intuition but are much more complicated in a discrete-time setting.

are characterized by

$$\begin{aligned}
\rho W(z) &= \omega(z) - \sigma\left(W(z) - U(z)\right) + zgW'(z) \\
\rho U(z) &= b + \lambda_w\left(W(z) - U(z)\right) \\
\rho J(z) &= z - w(z) - \sigma J(z) + zgJ'(z)
\end{aligned}$$

where ρ is the discount rate of both firms and workers, σ is the separation rate and λ_w is the worker's job finding rate. The terms $zgW'(z)$ and $zgJ'(z)$ reflect the on-the-job human capital accumulation of those in an existing match. Solving for the surplus of the match, $S(z) = W(z) - U(z) + J(z)$, under the assumption that the bargaining share γ is equal to $1/2$ gives

$$S(z) = \frac{z-b}{\tilde{\rho}} + \frac{\tilde{g}z}{(\tilde{\rho}-\tilde{g})\tilde{\rho}},$$

where

$$\tilde{\rho} = \rho + \sigma + \frac{1}{2}\lambda_w$$

is the *effective* rate at which agents discount the match surplus and

$$\tilde{g} = g\left(1 + \frac{\lambda_w}{2\rho}\right)$$

depends both on the rate at which productivity grows on the job g as well as the discount rate ρ.

Consider first an economy in which g is equal to 0 so that there is no on-the-job human capital accumulation. The surplus of the match does fall when ρ increases, but not very much. Since ρ is much smaller than σ and λ_w (empirical separation and job finding rates are much greater than discount rates), a doubling of ρ does not change $\tilde{\rho}$ and thus the surplus from a relationship very much. The free entry condition then implies that market tightness and thus employment changes little with changes in discount rates. Intuitively, if separation and job finding rates are high, the surplus from a match is very transitory: an employed workers is very likely to lose its job and a non-employed worker is very likely to find a match. Since the surplus is transitory, its valuation does not depend much on the rate at which agent discount future flows.

In contrast, when g is positive, the discount rate ρ also affects the effective rate at which the agents value the future increases in productivity arising from employment, $\tilde{g}$. Notice that when λ_w is sufficiently large relative to ρ, a doubling of the discount factor approximately halves $\tilde{g}$ and thus the surplus from a match, leading in equilibrium to a large drop in employment.

A numerical example may clarify this intuition further. Suppose $z = 1$, $b = 0.4$, $g = \rho = 0.05$, $\sigma = 0.25$ and $\lambda_w = 2$ as it is approximately the case in an annual calibration of our model. Consider the effect of a doubling of the discount rate to $\rho = 0.10$. The effective discount rate $\tilde{\rho}$ barely changes in this experiment, from 1.3 to 1.35. Hence, if g were equal to zero, the surplus from the match would only fall by several percentage points. In contrast, doubling the discount rate approximately halves the effective growth rate $\tilde{g}$ from 1.05 to 0.55, thus reducing the match surplus four-fold, from 3.69 to 0.95.

Clearly, allowing for on-the-job human capital accumulation is critical for credit shocks to have a sizable impact on employment.

Worker Credit Constraints.

A second key ingredient of our model, in contrast to a sizable literature in macroeconomics that focuses solely on the effect of firm-level credit frictions, is that both firms and workers are affected by discount rate shocks in our environment. We show below that credit frictions on the workers prevent wages from falling too much in response to firm-level shocks and are thus a source of endogenous wage stickiness.

To see this, we consider next a version of the model in which only firms are subject to the change in discount rates. Figure 9 shows that employment responds very little, by only about 0.5% in this version of the model, thus about one-fourth the drop in our Benchmark model. Figure 10 shows that the reason for the much smaller employment drop is that wages fall a lot, by about 5%, when firms only experience a change in discount rates. Intuitively, absent a wage change firms are the ones bearing the entire loss of surplus from the credit shock. Nash bargaining requires, however, that the surplus of the firms and workers is equated and this requires a large drop in wages. The drop in wages makes hiring workers relatively more attractive and vacancy creation and employment does not fall much, despite the firm-level credit shock.

Indeed, one simple way to amplify the volatility of employment in models with firm-only frictions is to assume that wages are exogenously sticky. As Figure 11 shows, doing so (we simply impose that $w_t(z) = w_0(z)$ during the transition) triples the drop in employment in the model with firm-level shocks alone. Recall, however, from Figure 9, that wages respond much less to the credit crunch in our model with both firm and worker discount rate changes, despite the fact that they are continuously renegotiated. Intuitively, since both workers and firms are now affected by the credit friction and face backloaded returns to employment, both the worker and firm surplus falls a lot even if wages do not change. Since workers are also adversely affected by the shock, there is not much room for wages to fall under Nash bargaining and firms cannot pass on the losses from the credit shock to their workers. Worker credit shocks thus act as a source of endogenous wage rigidity and amplify the effect of firm-level credit shocks.

Another way to see why wages are sticky in our Benchmark model is to notice that the credit crunch has two offsetting effects on the equilibrium wage. From the worker's viewpoint, at the higher discount rates the original upward sloping path of wages is less valuable. This force tends to push up wages. From the firm's standpoint, at the higher discount rate the original upward sloping path of profits from attracting the worker are lower than before. This force tends to push down wages. The net result of these offsetting movements is that the wage barely moves.

4.5 Robustness: Economy with Proportional Benefits

As discussed above, the drop in employment in our model is accounted for by several forces: an initial spike in job separations, a reduction in the acceptance rate of new matches as well as the drop in vacancy creations. One may argue that the first two forces are unappealing since they imply that workers willingly turn down job offers during recessions which is perhaps counterfactual. We show next that our results are very similar in an environment in which we shut down fluctuations in the separation and acceptance rate.

We formalize this argument by studying an alternative version of our model in which the amount of output a non-employed worker produces is proportional to its productivity. That is, we now assume that $b(z) = \lambda z$, where $\lambda < 1$ and is chosen to ensure that home production is equal to 40% of the average wage. Clearly, in this version of the model all matches produce positive surplus and are accepted. Similarly, all employed workers choose to stay employed.

Figure 12 compares the response of employment in this model to that in our Benchmark model. Clearly, the responses are not too dissimilar: employment drops by about three quarters as much as in our Benchmark model, and by almost the same amount as our Benchmark model would predict if we were to shut down fluctuations in the acceptance and separation rate (Figure 6). This suggests that general equilibrium effects arising from fluctuations in the acceptance and separation rates, (e.g., firms not hiring since they anticipate to be turned down by workers), do not explain much of the drop in vacancy creation in our Benchmark model.

5 An Economy with Traded and Nontraded goods

We consider an economy that consists of a continuum of islands each of which produces a nontraded good that can only be consumed on the island and a traded good that is consumed in all islands. Each island also has a fixed stock of houses. Labor is immobile across islands but can switch sectors. We first discuss the setup of the model and then results from our quantitative experiments.

5.1 Model Setup

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 ingredients that we introduce. Preferences are defined now over the consumption of a composite good that aggregates purchases of non-traded (N) and traded (T) goods:

$$c_t = \left[\tau^{\frac{1}{\sigma}} (c_{Nt})^{\frac{\mu-1}{\mu}} + (1-\tau)^{\frac{1}{\sigma}} (c_{Tt})^{\frac{\mu-1}{\mu}}\right]^{\frac{\mu}{\mu-1}}.$$

Traded goods are imported from the rest of the world. Since an individual island is infinitesimal and we consider the effect of a shock on one island in isolation, the price of traded goods is constant over time and we normalize it to 1. Let p_{Nt} denote the price of nontraded goods.

The demand for nontradable goods on the island is

$$c_{Nt} = \tau \left(\frac{p_{Nt}}{P_t}\right)^{-\mu} c_t \tag{21}$$

where

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

is the composite price index on the island.

We consider a credit crunch that is identical to that we have studied in the one-good model: a reduction in the preference for housing that causes a tightening of the credit limit and thus a reduction of consumption of 5% on impact. Importantly, we assume that firms are owned by households on *all* islands. Firm discount rates are thus unaffected by credit shocks on an individual island, only worker discount rates are.

There are two additional forces that shape the response of employment to a credit crunch in this model. First, as (21) makes it clear, a drop in consumption on the island leads to a drop in nontraded goods prices (since output is initially determined by the measure of agents that are employed in the nontradable sector and is thus predetermined). The drop in nontraded goods prices thus leads to a flow of employment from the nontraded sector to the traded sector. Second, since the island's aggregate price index falls as well, the discount factor changes by more then in the one-good model, since it now reflects the anticipation of an increase in the price level in future periods. That is, the one-period discount factor is now $\beta Q_{t+1}/Q_t$, where $Q_t = u'(c_t)/P_t$.

The matching technology is the same as in the one-good model. There are two sectors now and firms post vacancies in each of the two sectors. Non-employed workers search in both sectors but receive an offer from at most one firm. The matching functions in the two sectors are therefore

$$M_{Tt} = B_T(u_t)^{\eta} (v_{Tt})^{1-\eta} \quad \text{and} \quad M_{Nt} = B_N(u_t)^{\eta} (v_{Nt})^{1-\eta},$$

where M_{Tt} and M_{Nt} are the number of matches in each sector and v_{Tt} and v_{Nt} are the number of vacancies post by firms in each sector. The number of matches then determines the flows into each of these sectors. For example, the probability that a non-employed worker matches with a firm in each of the two sectors is:

$$\lambda_{wt}^{T} = B_T \left(\frac{v_{Tt}}{u_t} \right)^{1-\eta} = B_T \left(\theta_{Tt} \right)^{1-\eta}$$

and

$$\lambda_{wt}^{N} = B_N \left(\frac{v_{Nt}}{u_t} \right)^{1-\eta} = B_N \left(\theta_{Nt} \right)^{1-\eta}$$

Consider next the value functions that characterize the worker's problem. Let $w_{Nt}(z)$ and $w_{Tt}(z)$ denote the wages received by workers in the two sectors. We assume that non-employed workers produce b units of the composite good, so the value of home production is $P_t b$. The value of a worker employed in the traded sector is

$$\begin{aligned} W_{Tt}(z) &= \omega_{Tt}(z) + \phi \frac{Q_{t+1}}{Q_t} (1-\sigma) \int \max \left[W_{Tt+1}(z'), U_{t+1}(z') \right] dF_e(z'|z) \\ &+ \phi \frac{Q_{t+1}}{Q_t} \sigma \int U_{t+1}(z') \, dF_e(z'|z) \, . \end{aligned}$$

The value of a worker in the nontraded sector is

$$\begin{aligned} W_{Nt}(z) &= \omega_{Nt}(z) + \phi \frac{Q_{t+1}}{Q_t} (1-\sigma) \int \max \left[W_{Nt+1}(z'), U_{t+1}(z') \right] dF_e(z'|z) \\ &+ \phi \frac{Q_{t+1}}{Q_t} \sigma \int U_{t+1}(z') \, dF_e(z'|z) \, . \end{aligned}$$

Finally, the value of an unemployed worker reflects the possibility of joining any of the two sectors:

$$\begin{aligned} U_t(z) &= P_t b + \phi \frac{Q_{t+1}}{Q_t} \lambda_{T,t}^{w} \int \max \left[W_{Tt+1}(z'), U_{t+1}(z') \right] dF_u(z'|z) \\ &+ \phi \frac{Q_{t+1}}{Q_t} \lambda_{N,t}^{w} \int \max \left[W_{Nt+1}(z'), U_{t+1}(z') \right] dF_u(z'|z) \\ &+ \phi \frac{Q_{t+1}}{Q_t} \left(1 - \lambda_{Tw,t} - \lambda_{w,t}^{n} \right) \int U_{t+1}(z') \, dF_u(z'|z) \, . \end{aligned}$$

The production technology is the same as earlier: a worker with productivity z produces z units of the nontradable good if matched in the nontradable sector and z units of the tradable good if matched in the tradable goods sector. One key difference between this model and the one-good model is that we now assume that firms are owned by consumers on all islands and thus discount future flows at a constant rate $1/\beta - 1$. Nevertheless, firms operating on an island that experiences the credit shock are indirectly affected by two forces: i) the drop in the relative price of nontraded goods makes posting vacancies in the nontraded sector less profitable and ii) the drop in worker's surplus from a given relationship prevents wages from falling as much as they would have done otherwise.

The value of a matched firm in the tradable sector is

$$J_{Tt}(z) = z - \omega_{Tt}(z) + \phi\beta(1-\sigma)\int \max\left(J_{Tt+1}(z'),0\right) dF_e(z'|z)$$

and that in the nontradable sector is:

$$J_{Nt}(z) = p_{Nt}z - \omega_{Nt}(z) + \phi\beta(1-\sigma)\int \max\left(J_{Nt+1}(z'),0\right) dF_e(z'|z)$$

The free-entry conditions into the traded sector is

$$0 = -\kappa_T + \phi\beta\lambda^f_{T,t}\int \max\left[J_{Tt+1}(z'),0\right] \mathrm{d}F_u(z'|z)\, d\tilde{n}_{ut}(z)$$

while in the nontraded sector is

$$0 = -\kappa_N + \phi\beta\lambda^f_{N,t}\int \max\left[J_{Nt+1}(1,z'),0\right] \mathrm{d}F_u(z'|z)\, \mathrm{d}\tilde{n}_{ut}(z).$$

Notice that we are implicitly assuming that the cost of posting vacancies is denominated in units of the tradable good.

5.2 Quantitative Findings

There are several additional parameters in this model. We set the elasticity of substitution μ across traded and non-traded goods equal to 4, consistent with estimates in the trade literature.[4] We set τ, the preference weight on non-traded goods equal so that 2/3 of employment on the island is in the non-traded sector. This is consistent with the evidence in Mian and Sufi (2013) for the U.S. The costs of posting vacancies κ_T and κ_N are chosen so as to ensure that market tightness θ_T and θ_N is the same in both sectors, and equal to 1. As earlier, the choice of 1 is a normalization. Finally, we choose the efficiency of the matching function in the two sectors, B_T and B_N to ensure i) an employment-population ratio of 80% and ii) the steady state price of non-traded goods, p_N is equal to that of traded goods, 1. This latter restriction implies that workers in both sectors receive identical wages, $w_T(z) = w_N(z)$ in the ergodic steady state of our model. Given these choices, workers are indifferent between which sectors to join and the steady-state predictions of the multi-good model are identical to those of the one-good model.

[4] The macro literature sets a much lower elasticity, of up to 0.5. Using that elasticity would imply a much larger drop in prices in the island hit by a credit shock and amplify the employment responses. We conjecture that such large drops in prices are counterfactual.

Figure 13 shows the response of overall employment to the same credit shock we have considered in the one-good model. The maximal drop in employment is about 1.75%, slightly lower than the 2% in the benchmark model, reflecting the fact that firms now are owned by all islands and are thus not subject to the credit shock. This makes it then clear that having shocks on the firm side is not critical in our environment: worker-level credit shocks account for most of the employment drops. Intuitively, since workers keep most of their accumulated human capital when transiting to a new firm, their horizon is much longer than that of firms (who lose a separated worker and thus its human capital for good). Workers are thus much more sensitive to discount rate shocks and fluctuations in the worker's discount rate account for most of the employment losses in our model.

We also see in Figure 13 that the drop in employment in the model with no on-the-job human capital accumulation are much lower than in our model, only about 0.4%. Once again, backloaded returns to employment are critical for credit shocks to have a quantitatively significant impact. Figure 14 shows the behavior of wages in the economy with and without returns to work. Absent returns to work, the standard DMP model predicts a sharp decline in wages in the aftermath of the credit shock, reflecting the drop in prices on the island and the fact that matches are now less profitable. But the drop in wages insulates firm profits from the credit shock and hence employment does not react much. Our model with returns to work produces, in contrast, a form of wage rigidity: wages fall much less now and this strongly reduces firms' incentive to post vacancies.

Figures 15 and 16 show the response of nontradable and tradable employment. Notice that both models with and without returns to work can generate a sizable drop in non-traded employment, albeit the model with returns to work predicts a drop that is about twice larger. The reason non-traded employment falls even in the standard model is the drop in the price of non-traded goods which makes posting vacancies in that sector particularly undesirable. Notice, however, in Figure 16, that these relative price effects imply, in the standard model without returns to work, that employment flows from the non-traded to the traded sector: tradable employment expands by about 1.5% in response to the credit shock. This prediction of the model is, however, counterfactual, as documented by Mian and Sufi (2013) and as we also show below. In contrast, the model with returns to work predicts, if anything, a slight drop in tradable employment.

5.3 Comparison with Cross-Section from the U.S.

Mian and Sufi (2011, 2013, 2014) have documented a very robust relationship between changes in consumption, employment and house prices in a cross-section of U.S. counties. Two key findings emerge from their work. First, consumption and employment positively comove from 2007 to 2009, a feature that is at odds with the predictions of standard models. Second, most of the reduction in employment associated with a tightening of household balance sheets in this period is accounted for by industries that sell non-tradable goods. The decline of tradable employment is large, but unrelated to measures of household debt or consumption. We argue next that our model successfully replicates these features of the data.

We do so by engineering a fall in the taste for houses in 20 different islands, each of which is of measure 0 and thus does not affect the price of tradable goods or interest rates. The shock is chosen so that consumption drops by 1%, 2% ... 20% in each of the 20 islands 2 years after the shock. Figure 17 plots the response of employment against the drop in consumption in each of these islands (right panel) and contrasts it with the U.S. evidence (left panel). Notice that the model implies a linear relationship between the drop in employment and that in house prices, with a slope of about 0.56. A similar relationship, with a slope of about 0.5, is apparent in the U.S. data as well. The model thus successfully reproduces the consumption-employment comovement in the data.

Consider next the response of non-tradable employment. In Figure 18 we plot the drop in non-tradable employment against the change in house prices in the model (right panel) as well as in the data (left panel). Notice that in the data the drop in non-tradable employment is positively related to that in house prices, with a slope of about 0.13. Thus, for every 10% drop in house prices, nontradable employment fell by about 1.3%. As the right panel of Figure 18 shows, our model implies a very similar relationship, with an identical slope.

Consider next the response of tradable employment. As the left panel of Figure 19 shows, there is essentially no relationship in the data between changes in tradable employment and house prices. Our model replicates this feature of the data as well. In contrast, a standard DMP model without returns to work would predict a negative relationship between changes in house prices and tradable employment, due to the flow of labor from the nontraded to tradable goods sector.

6 Conclusions

The key insight we have explored in this paper is that since returns to employment are backloaded in the data, employment is sensitive to changes in worker debt constraints. We have shown, using a standard DMP model, that this force makes wages endogenously sticky and amplifies the drop in employment in response to a credit shock. We have argued that this mechanism is a quantitatively promising avenue for understanding both the evolution of employment in the U.S. aggregate as well as for accounting for the cross-regional U.S.

evidence documented by Mian and Sufi (2013).

References

[1] Andolfatto, David. 1996. “Business Cycles and Labor-Market Search, ” *American Economic Review*, 86(1), 1112-132.

[2] Chari, V.V., P. Kehoe, and E. McGrattan. 2005. “Sudden Stops and Output Drops,” *American Economic Review*, 95(2), 381-387.

[3] Eggertson, G. and P. Krugman. 2012. “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach.” *Quarterly Journal of Economics*, 1459-1513.

[4] Guerrieri, V, and G. Lorenzoni. 2010. “Credit Crises, Precautionary Savings, and the Liquidity Trap, ” Chicago Booth mimeo.

[5] Hall, R.E. 2014 “High Discounts and High Unemployment, ” NBER working paper 19871.

[6] Krusell, P., T. Mukoyama, and A. Sahin. 2010. “Labor-Market Matching with Precautionary Savings and Aggregate Fluctuations,” *Review of Economic Studies*, 77(4), 1477-1507.

[7] Mendoza, E. 2010. “Sudden Stops, Financial Crises, and Leverage, ” *American Economic Review*, 100(5), 1941-1946.

[8] Merz, M. 1995. “Search in the Labor Market and the Real Business Cycle, ” *Journal of Monetary Economics*, 36 (2), 269–300.

[9] Mian, A., and A. Sufi. 2011. “House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis,” *American Economic Review*, 101(5), 2132-2156.

[10] Mian, A., K. Rao and A. Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 128(4), 1687-1726.

[11] Mian, A., and A. Sufi. 2014. “What Explains the 2007-2009 Drop in Employment,” Chicago Booth Research Paper No. 13-43.

[12] Midrigan, V., and T. Philippon. 2011. “Household Leverage and the Recession,” NYU mimeo.

[13] Shimer, R. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95(1), 25-49.

[14] Shimer, R. 2012. “Reassessing the Ins and Outs of Unemployment,” *Review of Economic Dynamics*, 15(2), 127-148.

Table 1: Parameterization

Panel A: *Moments*

	Data	Model
Moments used in calibration		
Market tightness	1.00	1.00
Fraction employed	0.80	0.80
Home production / mean wage	0.40	0.40
Std. dev. wage changes	0.21	0.21
Std. dev. initial wages	0.94	0.94
Mean wage growth of employed	0.052	0.052
Additional model predictions		
Fraction workers with $w < b$		0.181
Fraction voluntary separations		0.002
Probability worker matches		0.595
Fraction rejected matches		0.278
Wage drop after non-employment spell		1.9%
Wage drop if non-employed 1 year		6.1%
Wage drop if non-employed 2 years		8.8%

Panel B: *Parameters*

Endogenously chosen	
B, efficiency matching function	0.595
κ, vacancy cost (rel. mean output)	0.90
ρ_z, persistence	$0.952^{1/4}$
σ_z, volatility	0.112
μ_z, mean efficiency employed	2.82
b, home production (rel. mean output)	0.27
Assigned	
period length	1 quarter
β, discount factor	$0.94^{1/4}$
$1 + r$, interest rate	$0.96^{-1/4}$
ϕ, survival probability	1-1/160
σ, probability of separation	0.10
α, inverse IES	5
η, elasticity matching function	0.50
γ, worker's bargaining share	0.50

Figure 1: Wage paths: Model and BFKT (2010) Estimates

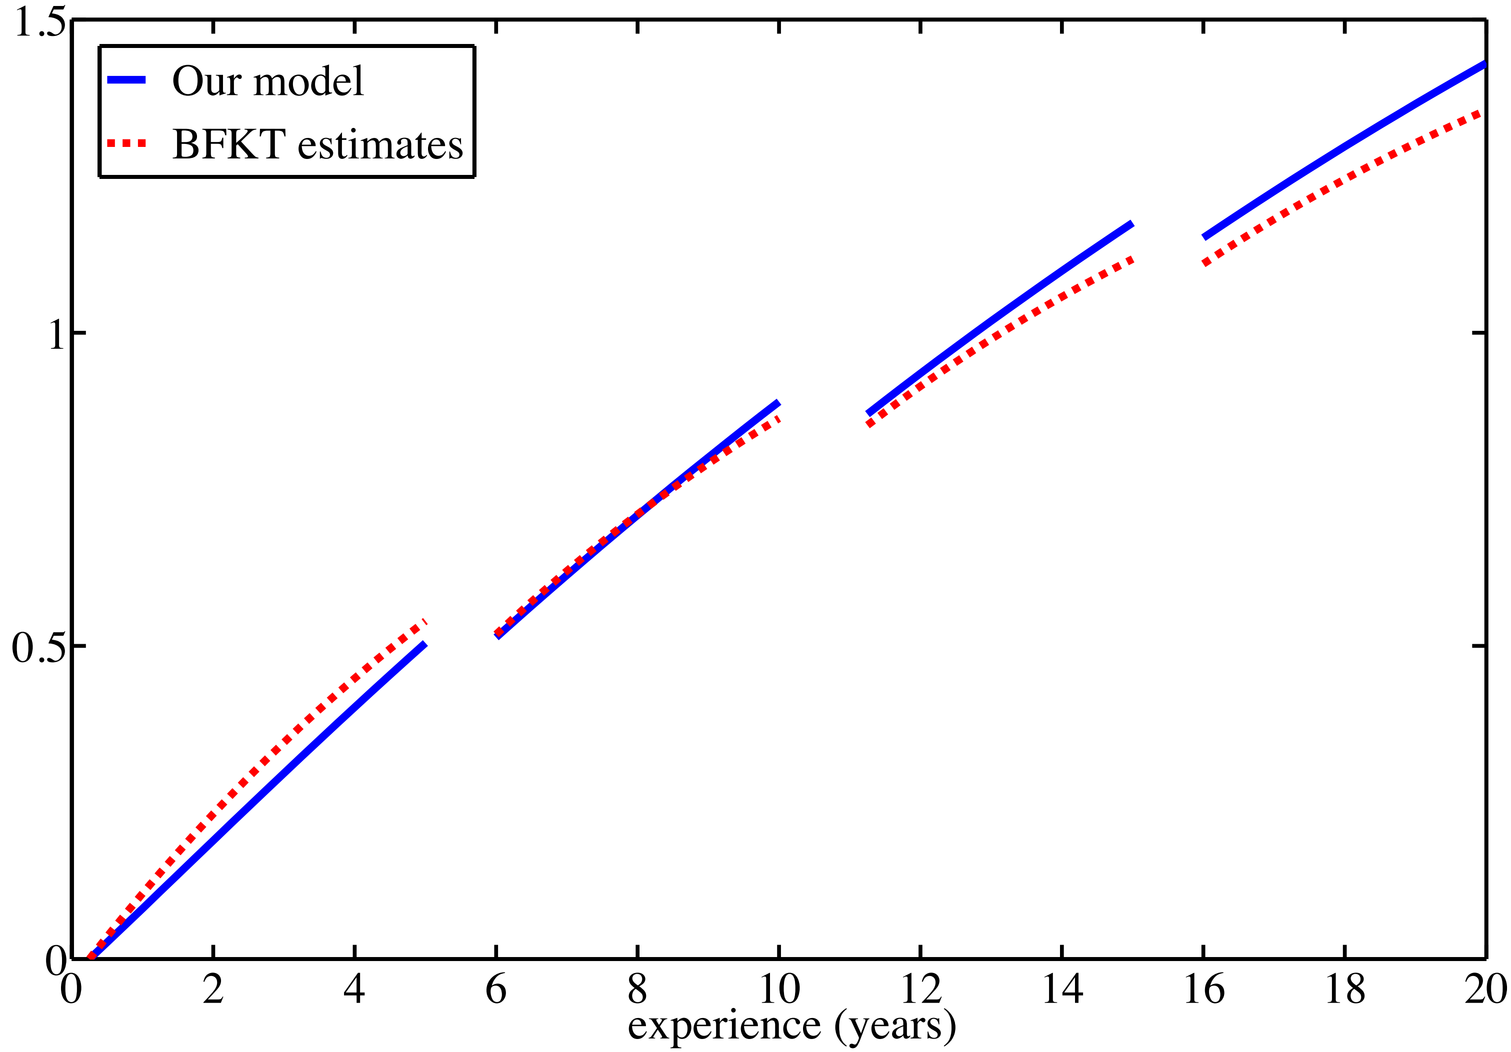

We initialize the simulation for an agent with zero experience, a mean productivity equal to the mean of newborn workers' productivity and assume no shocks to productivity over the lifetime. The figure shows the evolution of actual wages in the model as well as those predicted by the BKFT estimates given the experience and tenure history of the worker.

Figure 2: Decision Rules and Value Functions

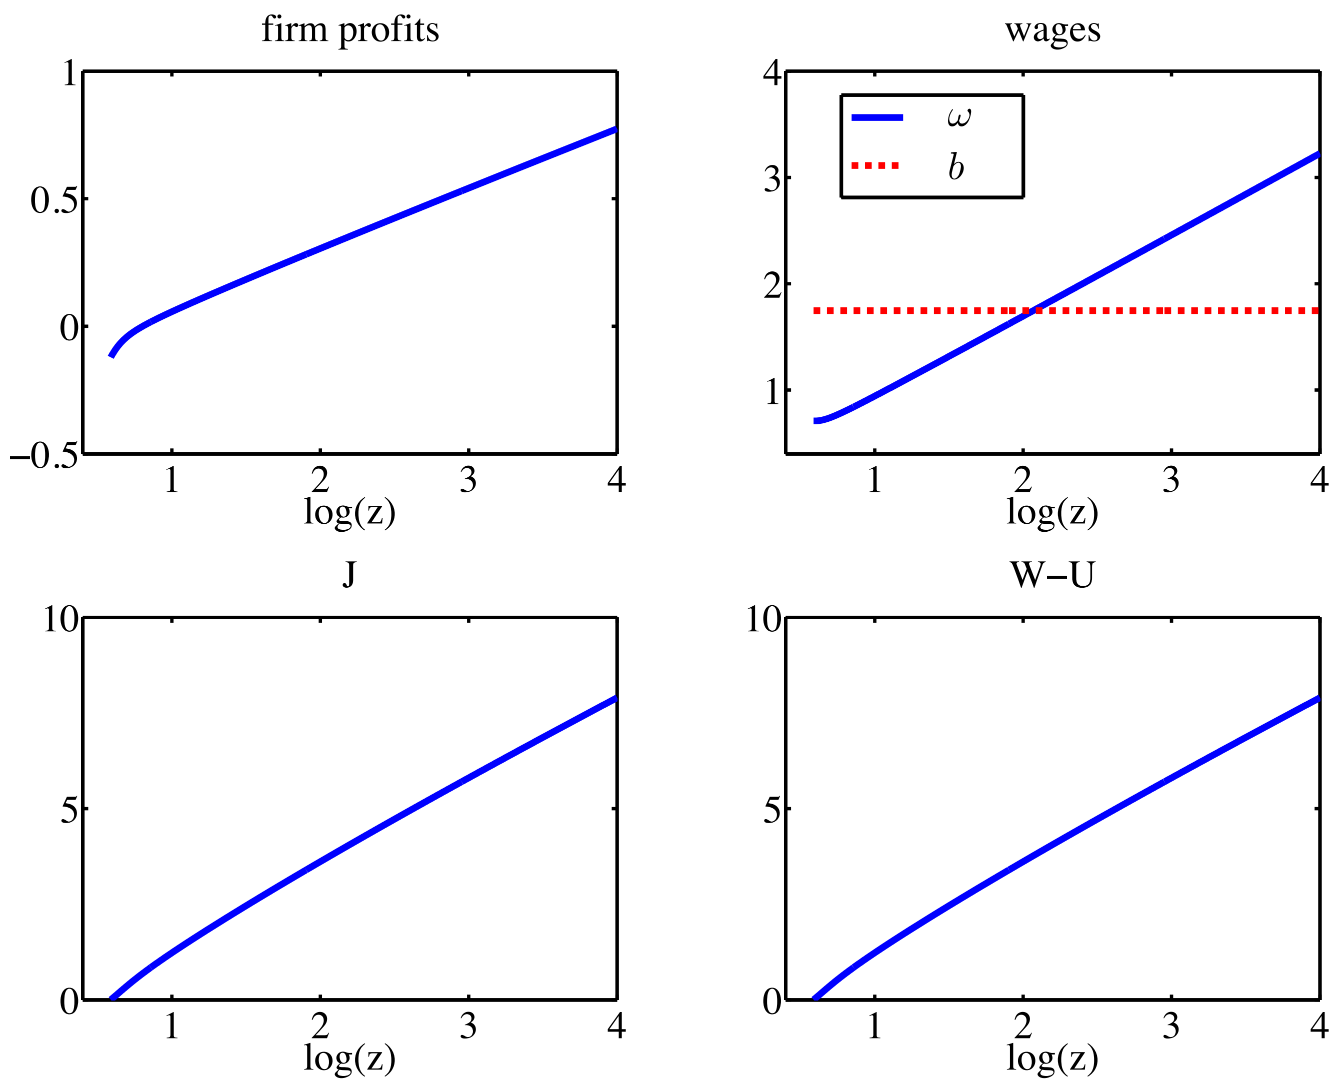

Figure 3: Steady State Measures

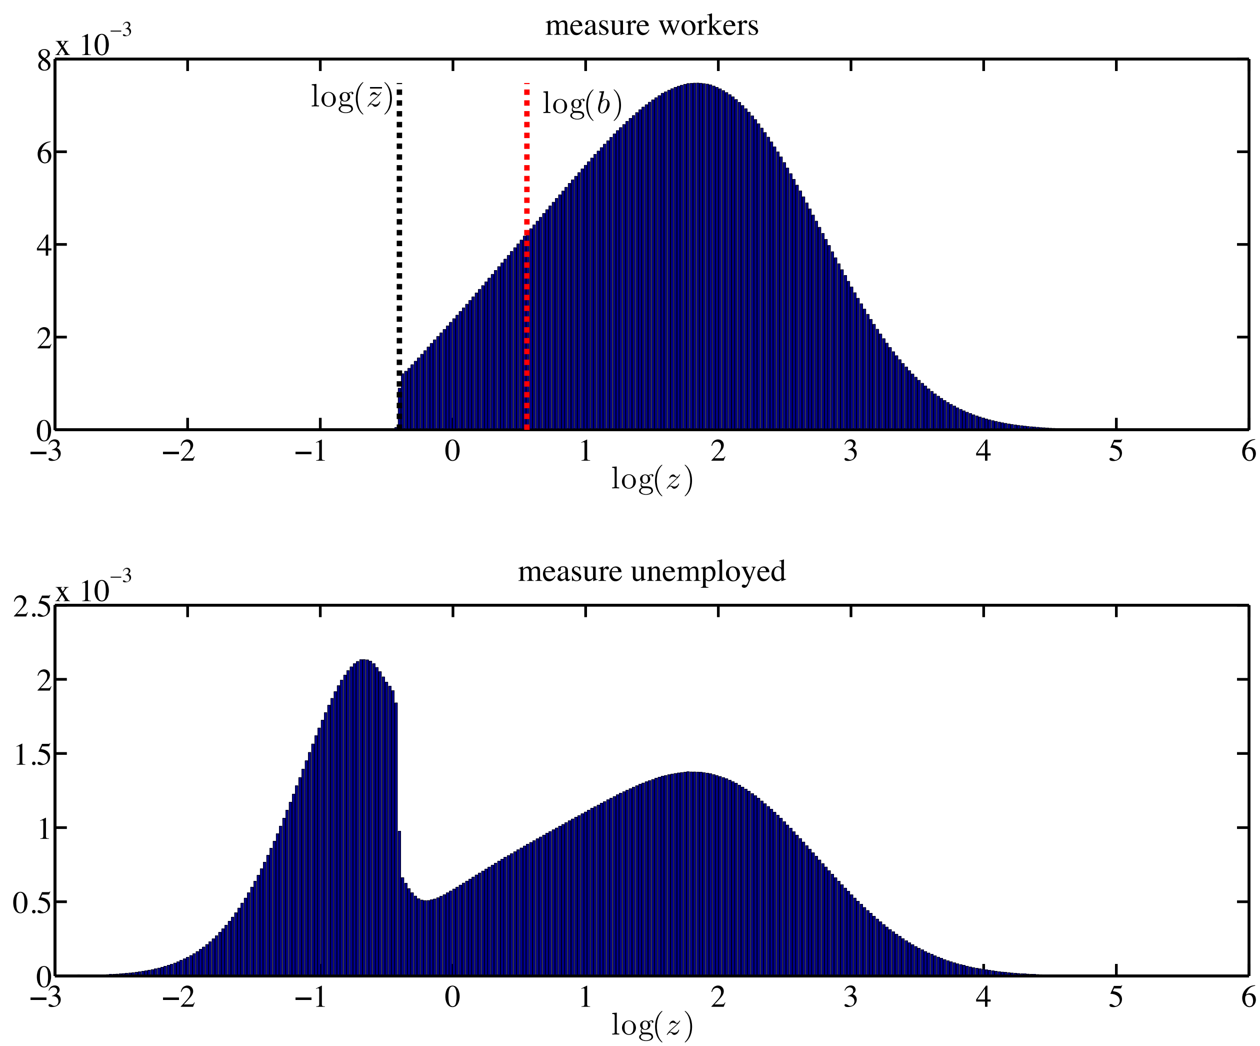

Figure 4: Credit Tightening Experiment

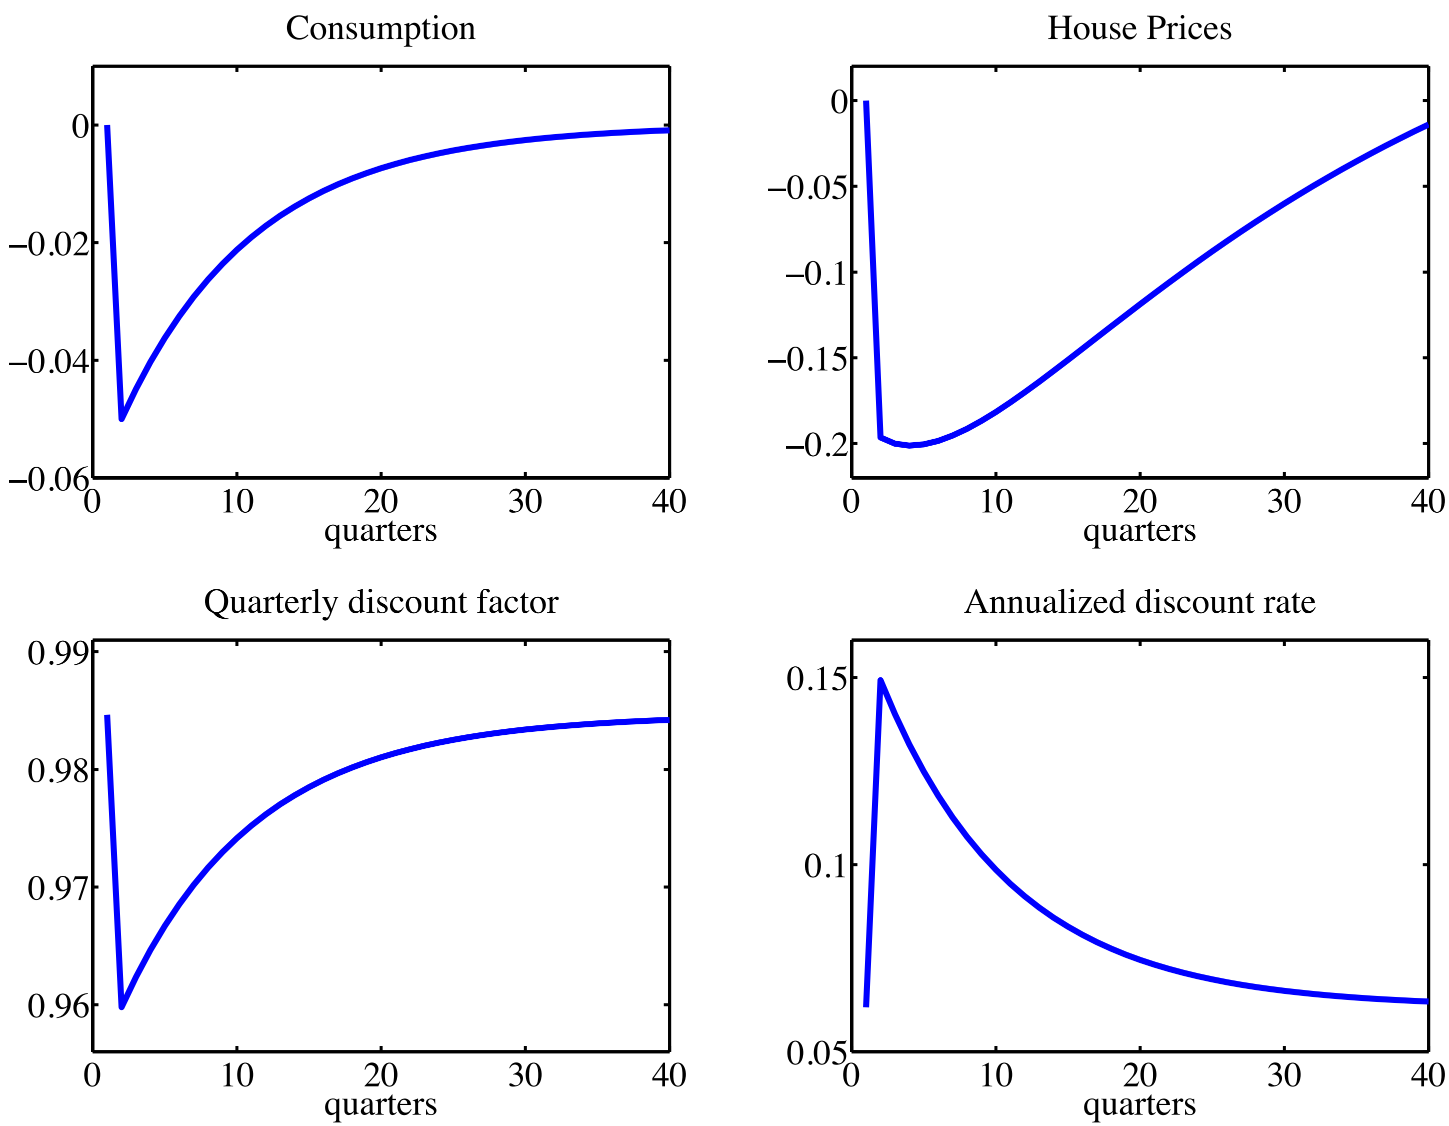

Figure 5: Employment Response

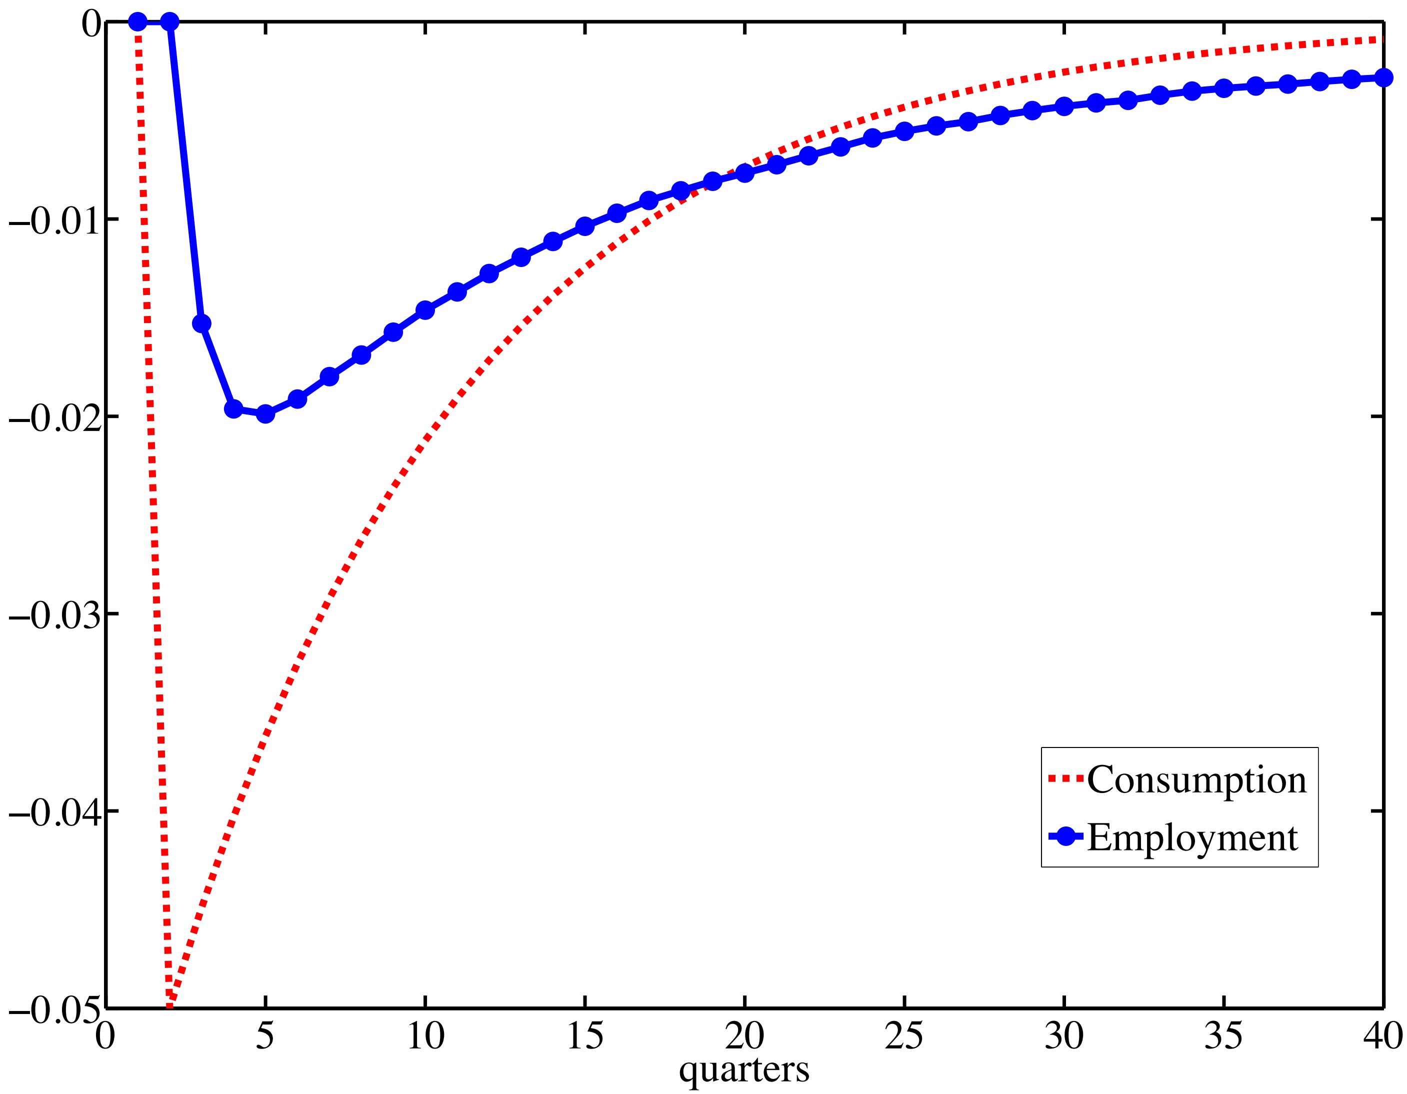

Figure 6: Sources of Employment Drop

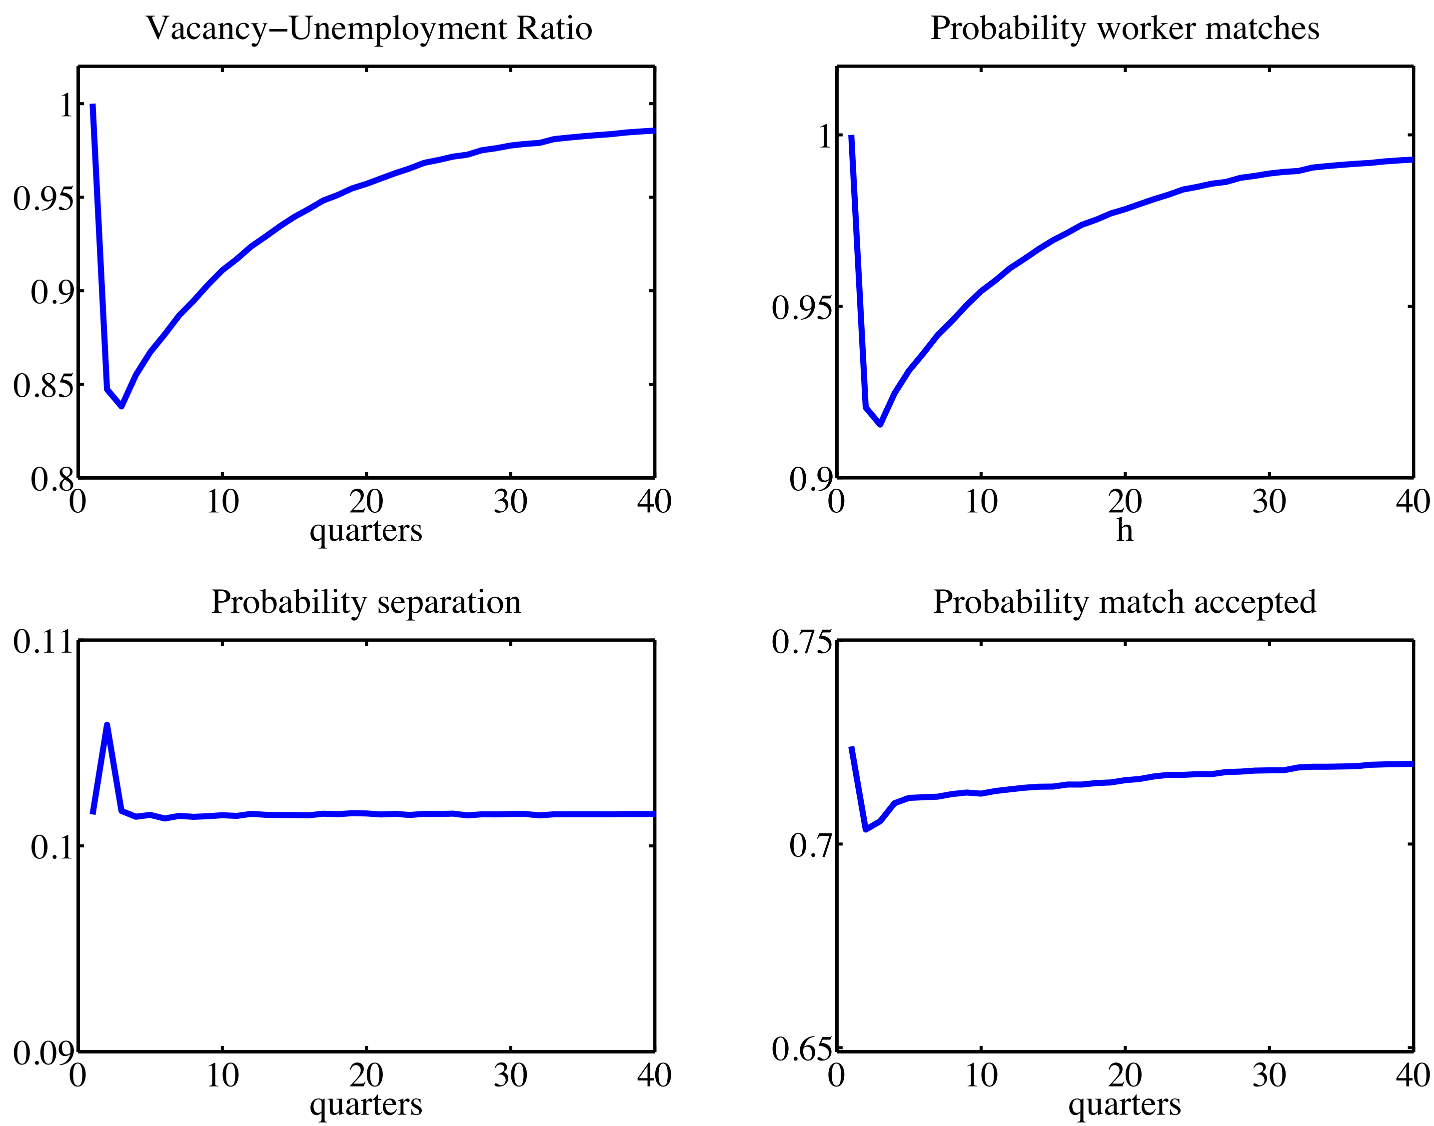

Figure 7: Ins an Outs Decomposition

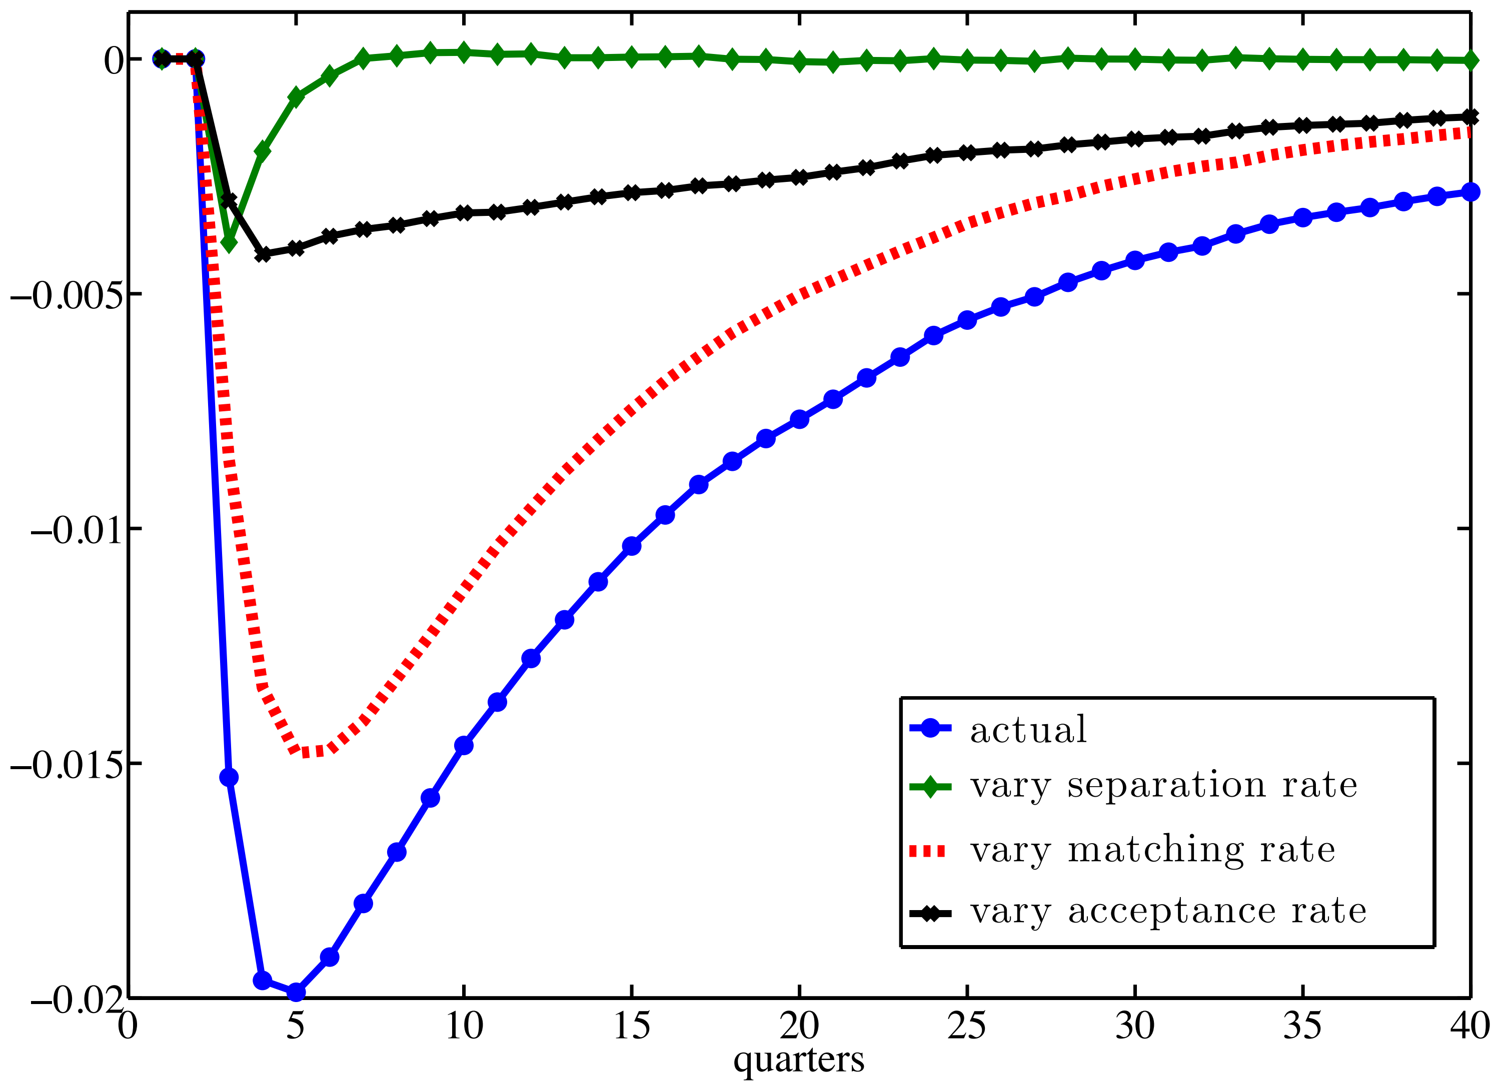

Figure 8: Employment Responses Absent Returns to Work

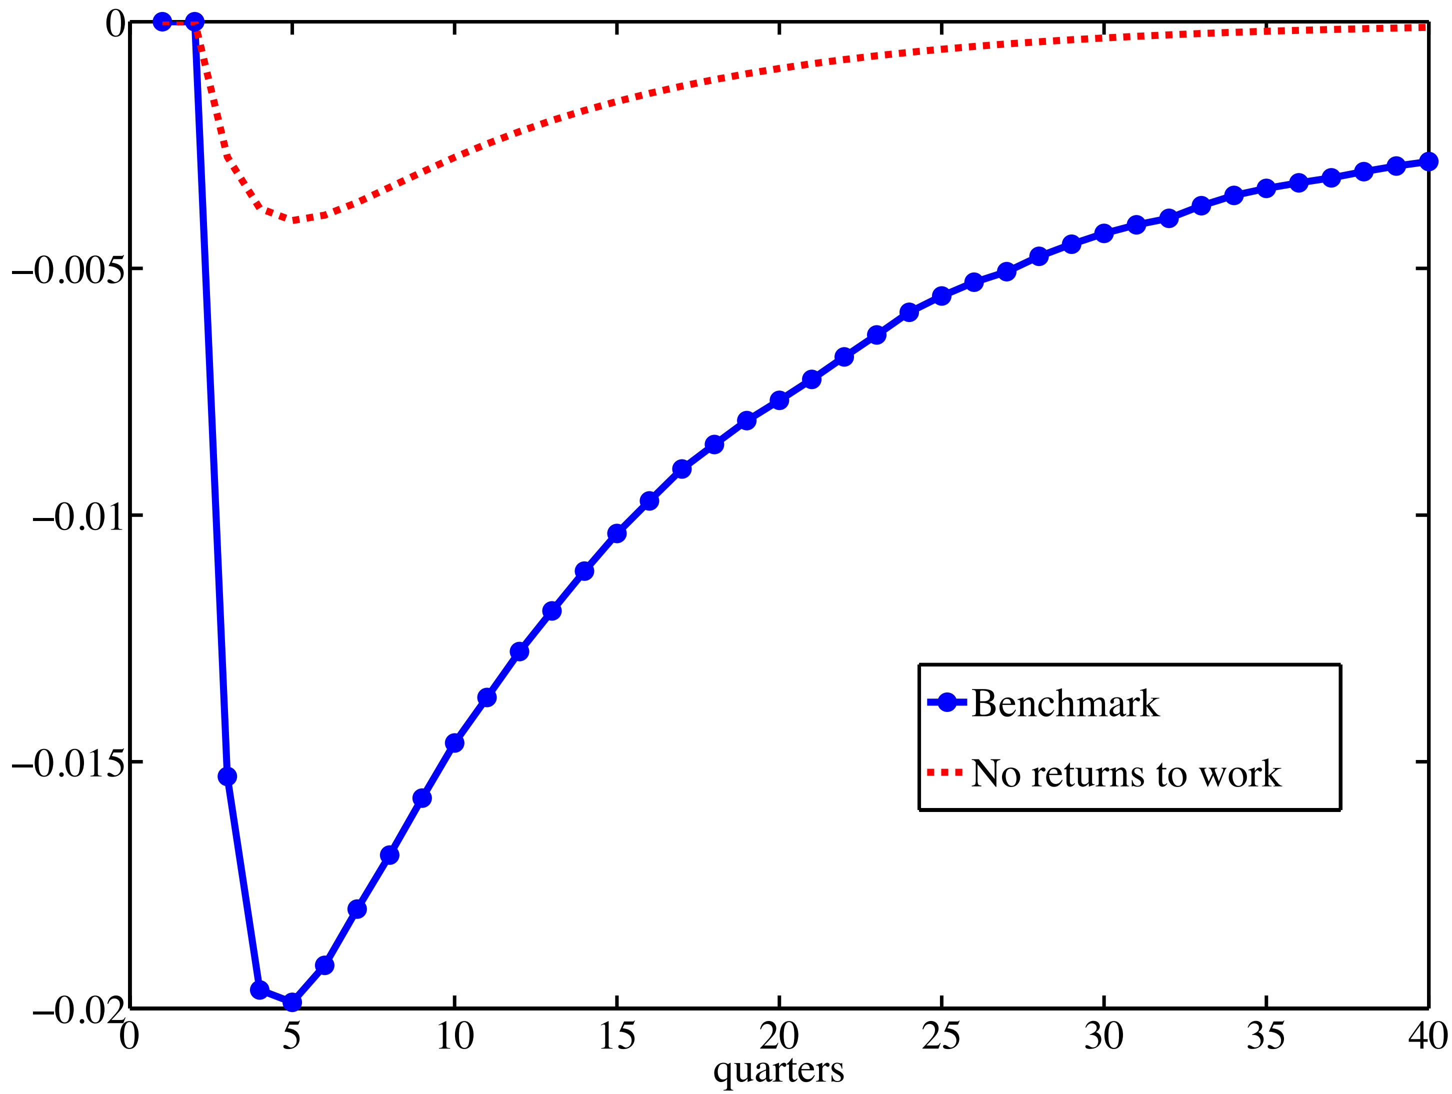

Figure 9: Employment Responses with Firm Discount Change

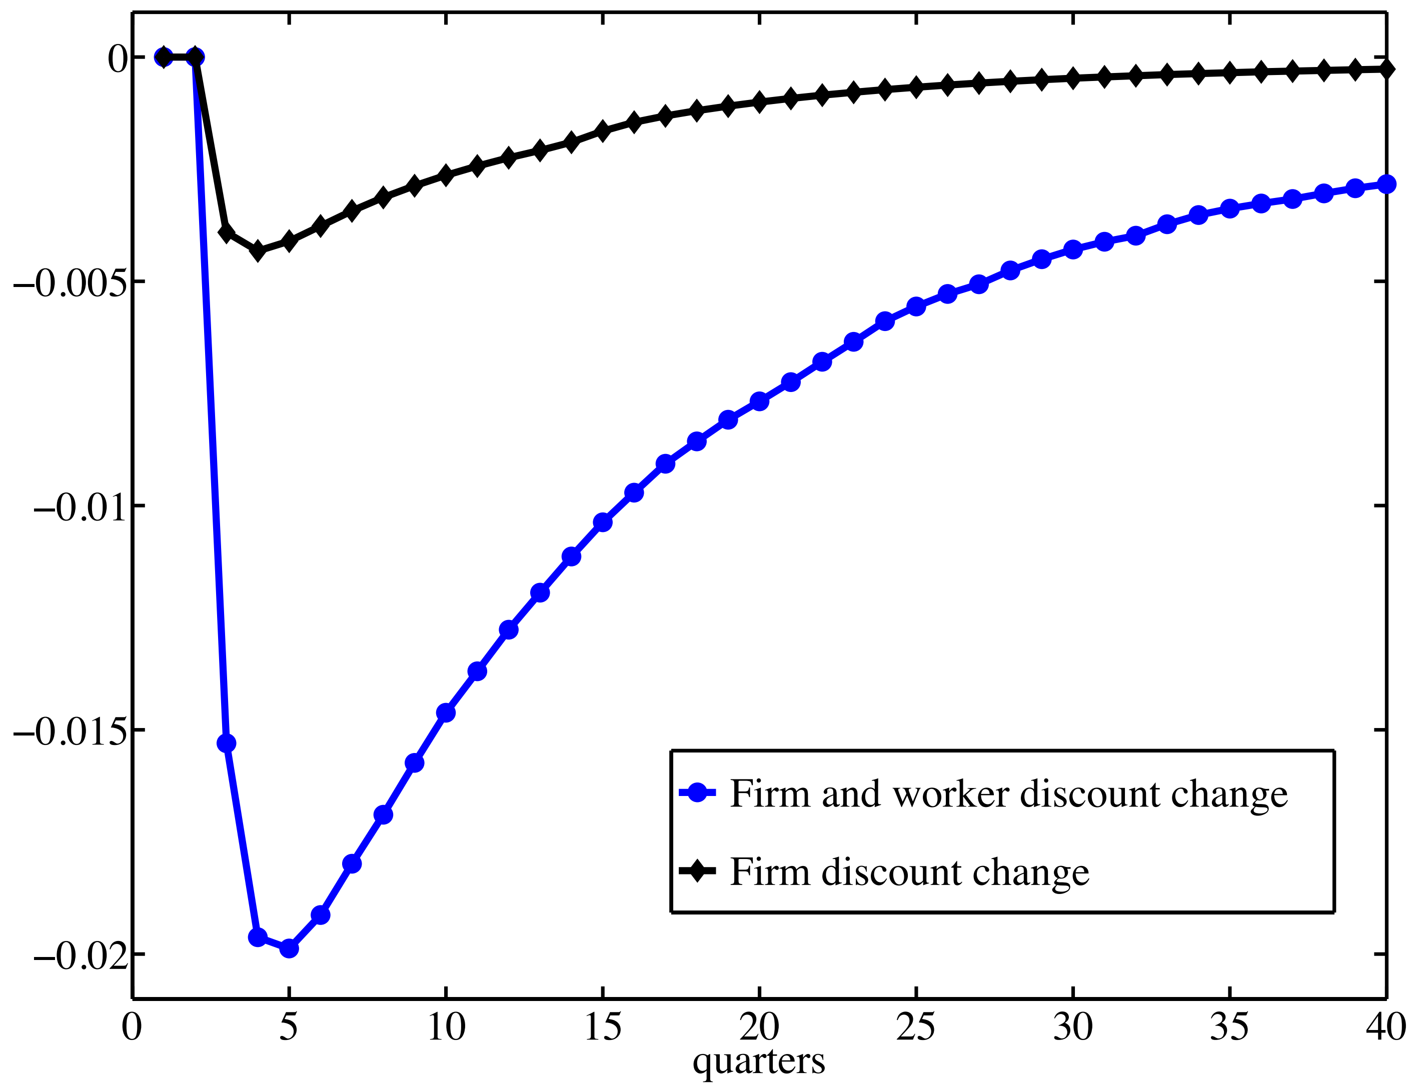

Figure 10: Wage Responses with Firm Discount Change

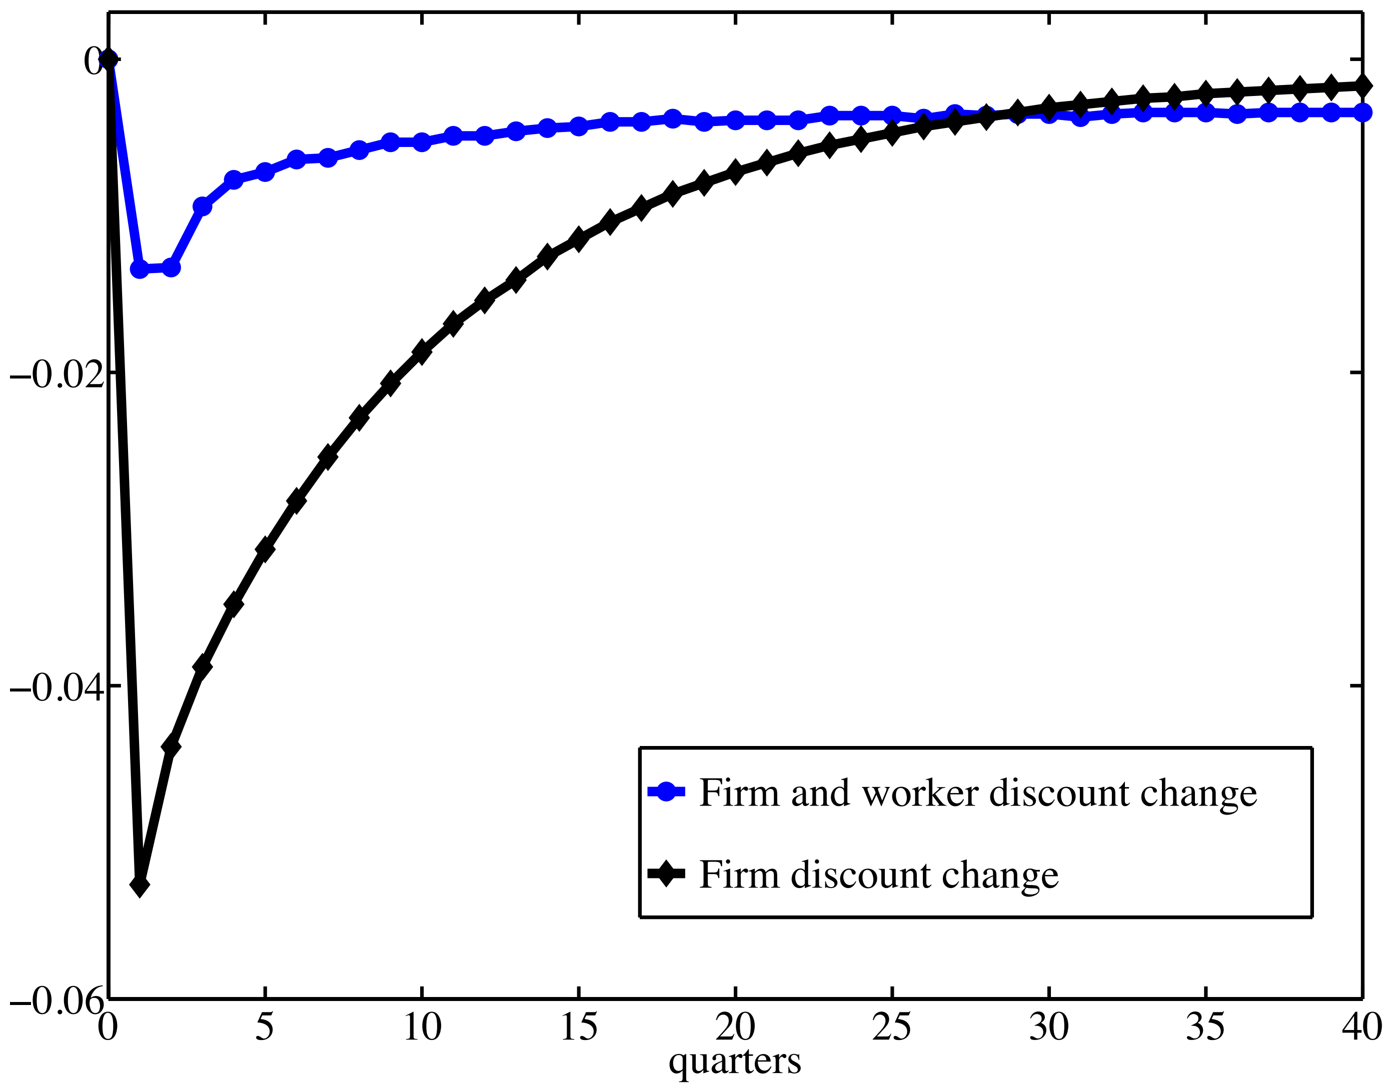

Figure 11: Employment Responses with Firm Discount Change and Sticky Wages

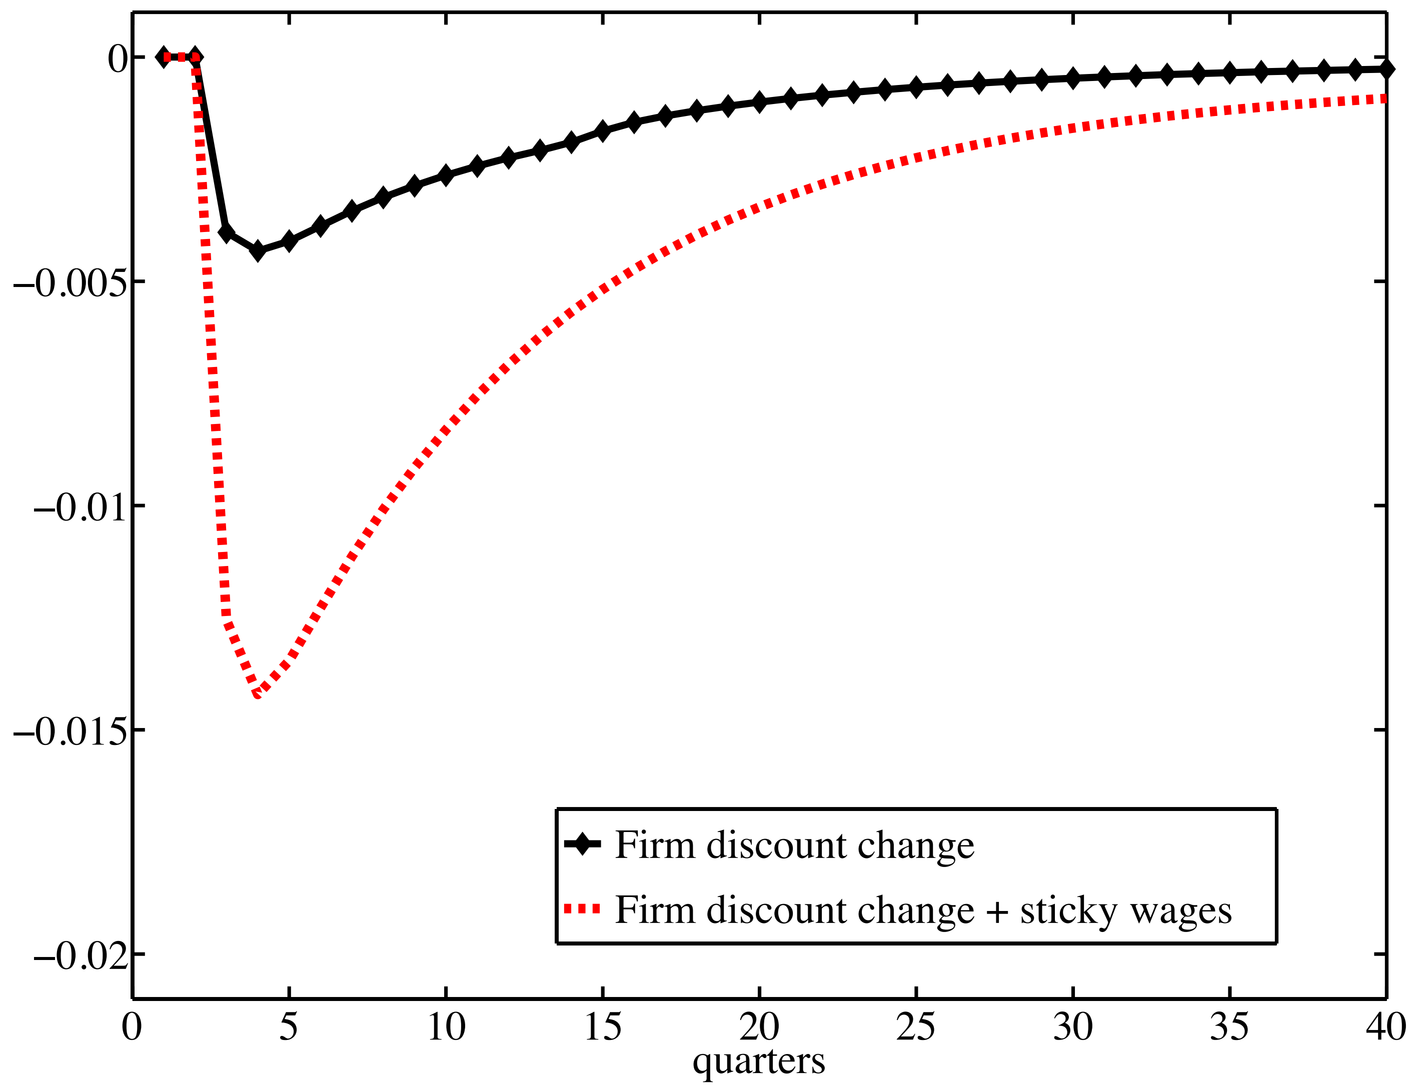

Figure 12: Employment Responses with Proportional Benefits

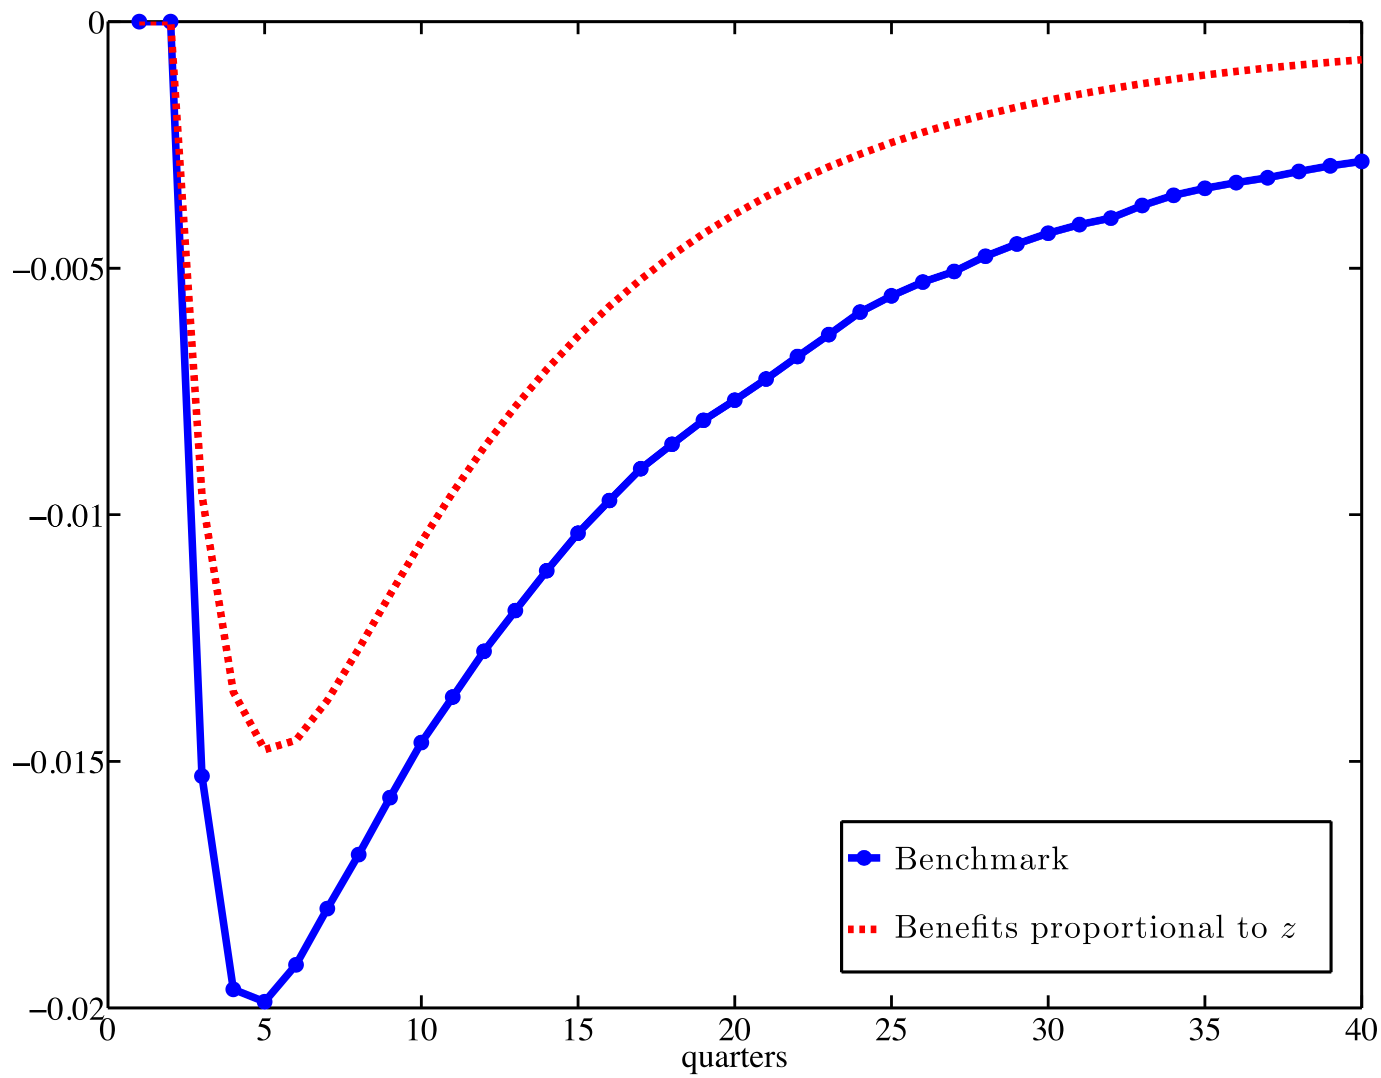

Figure 13: Employment Responses in Multi-Good Model

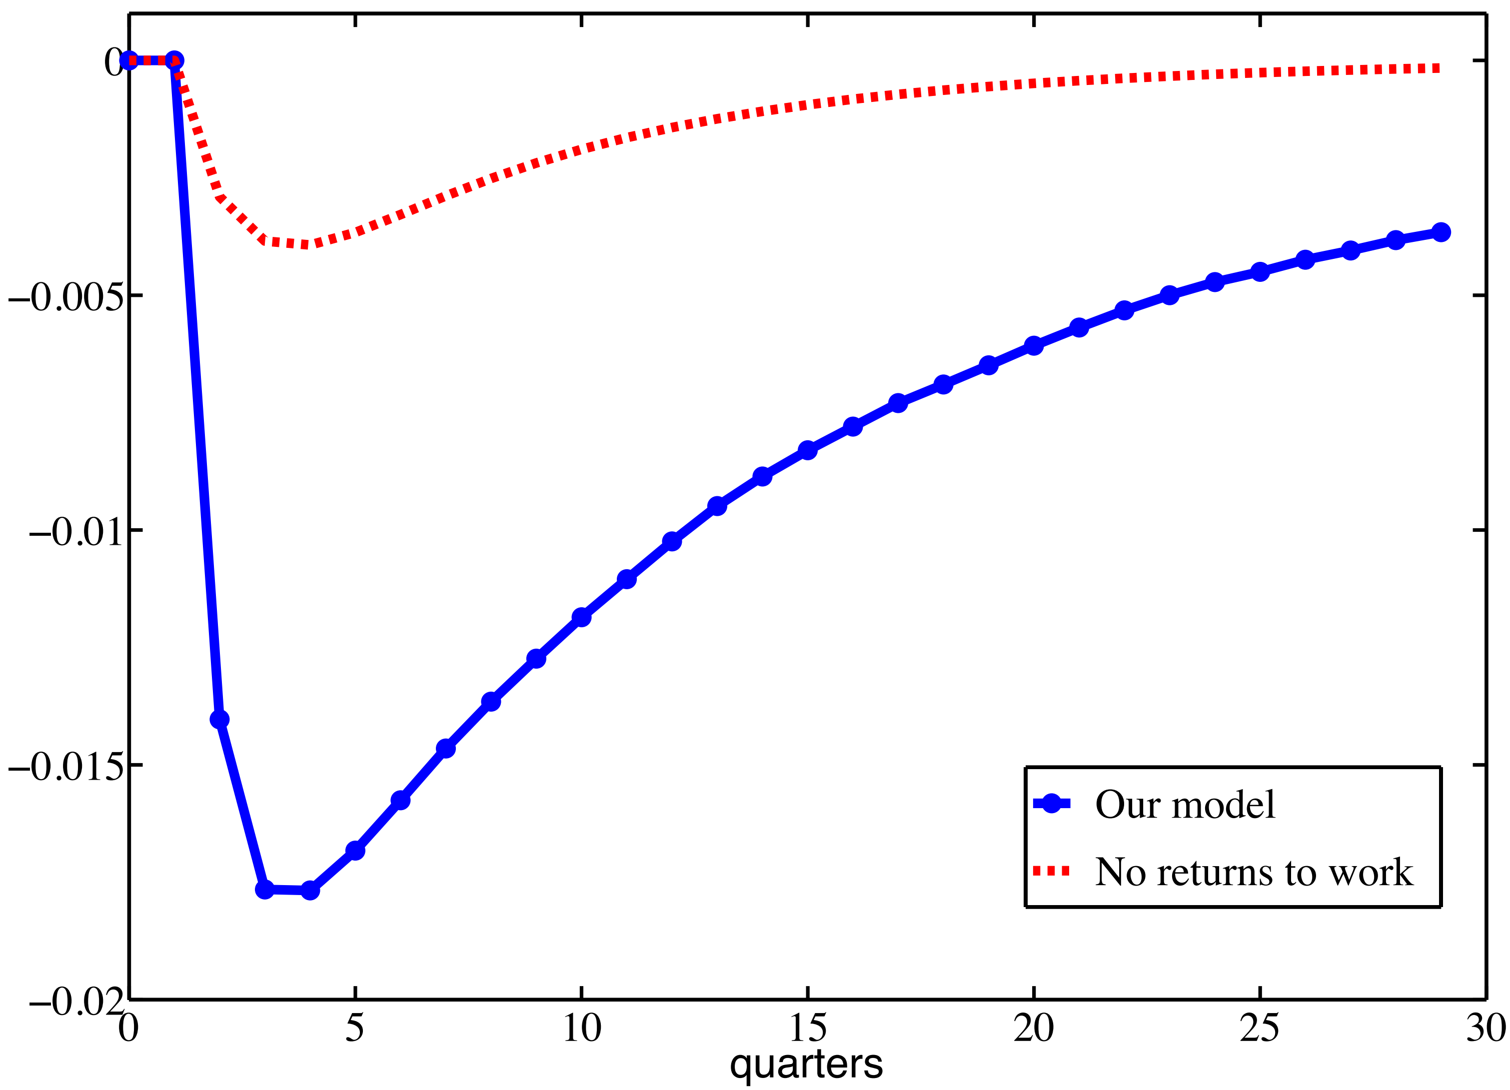

Figure 14: Wage Responses in Multi-Good Model

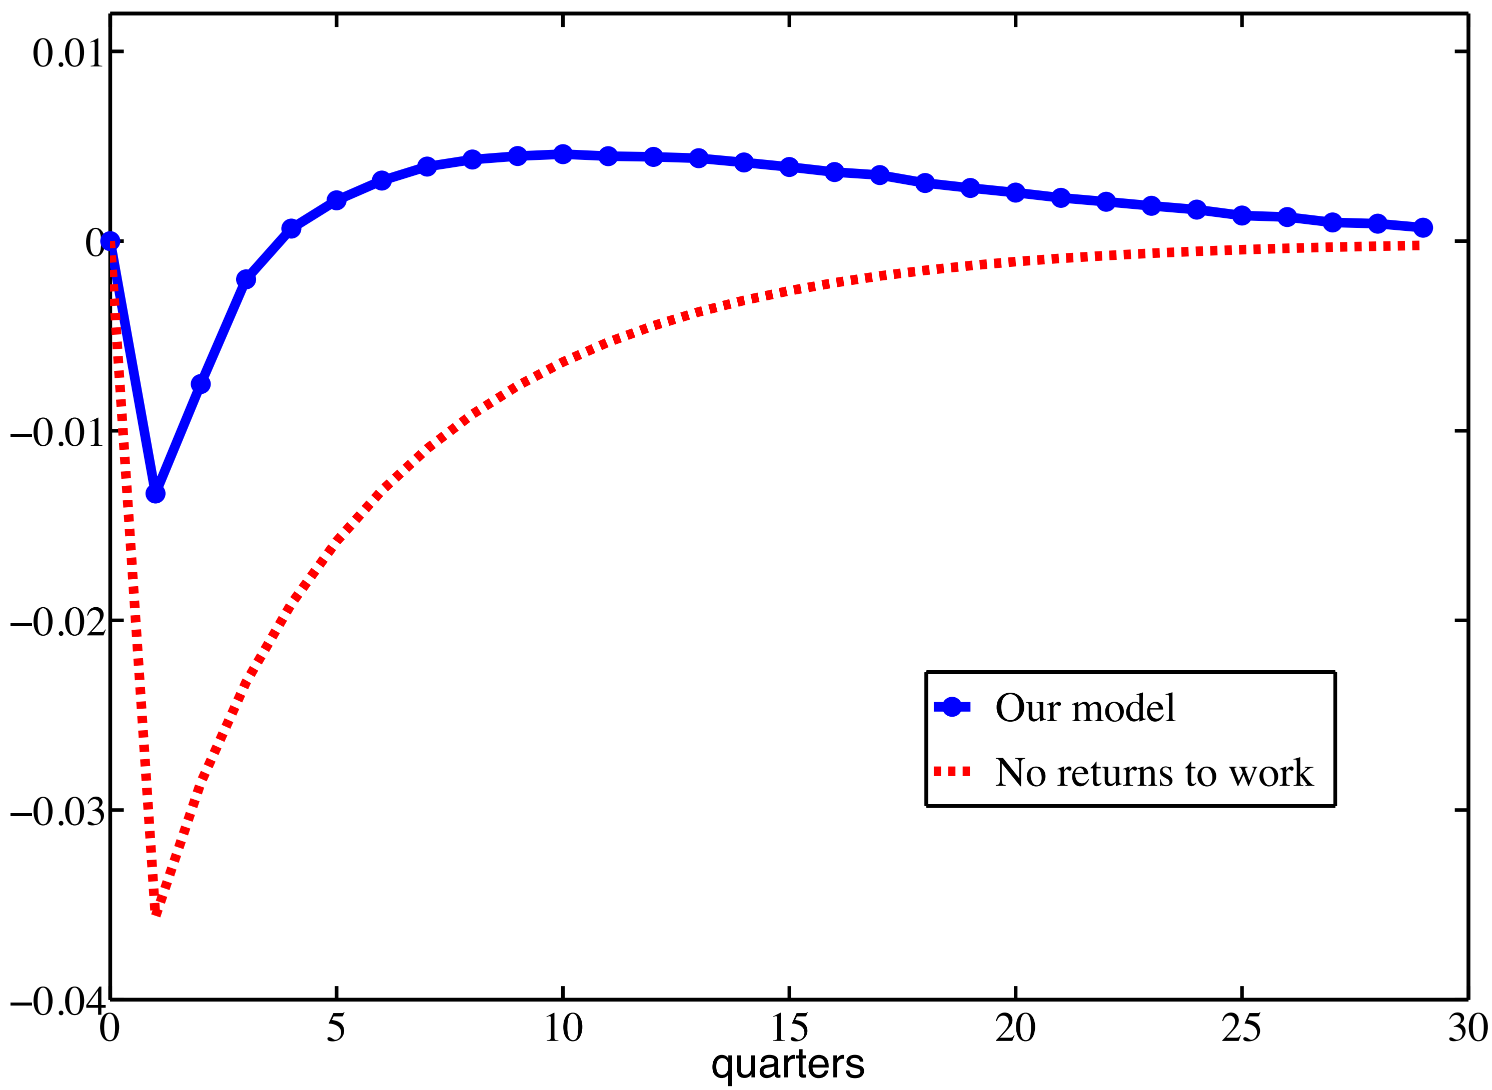

Figure 15: Non-traded Employment Responses in Multi-Good Model

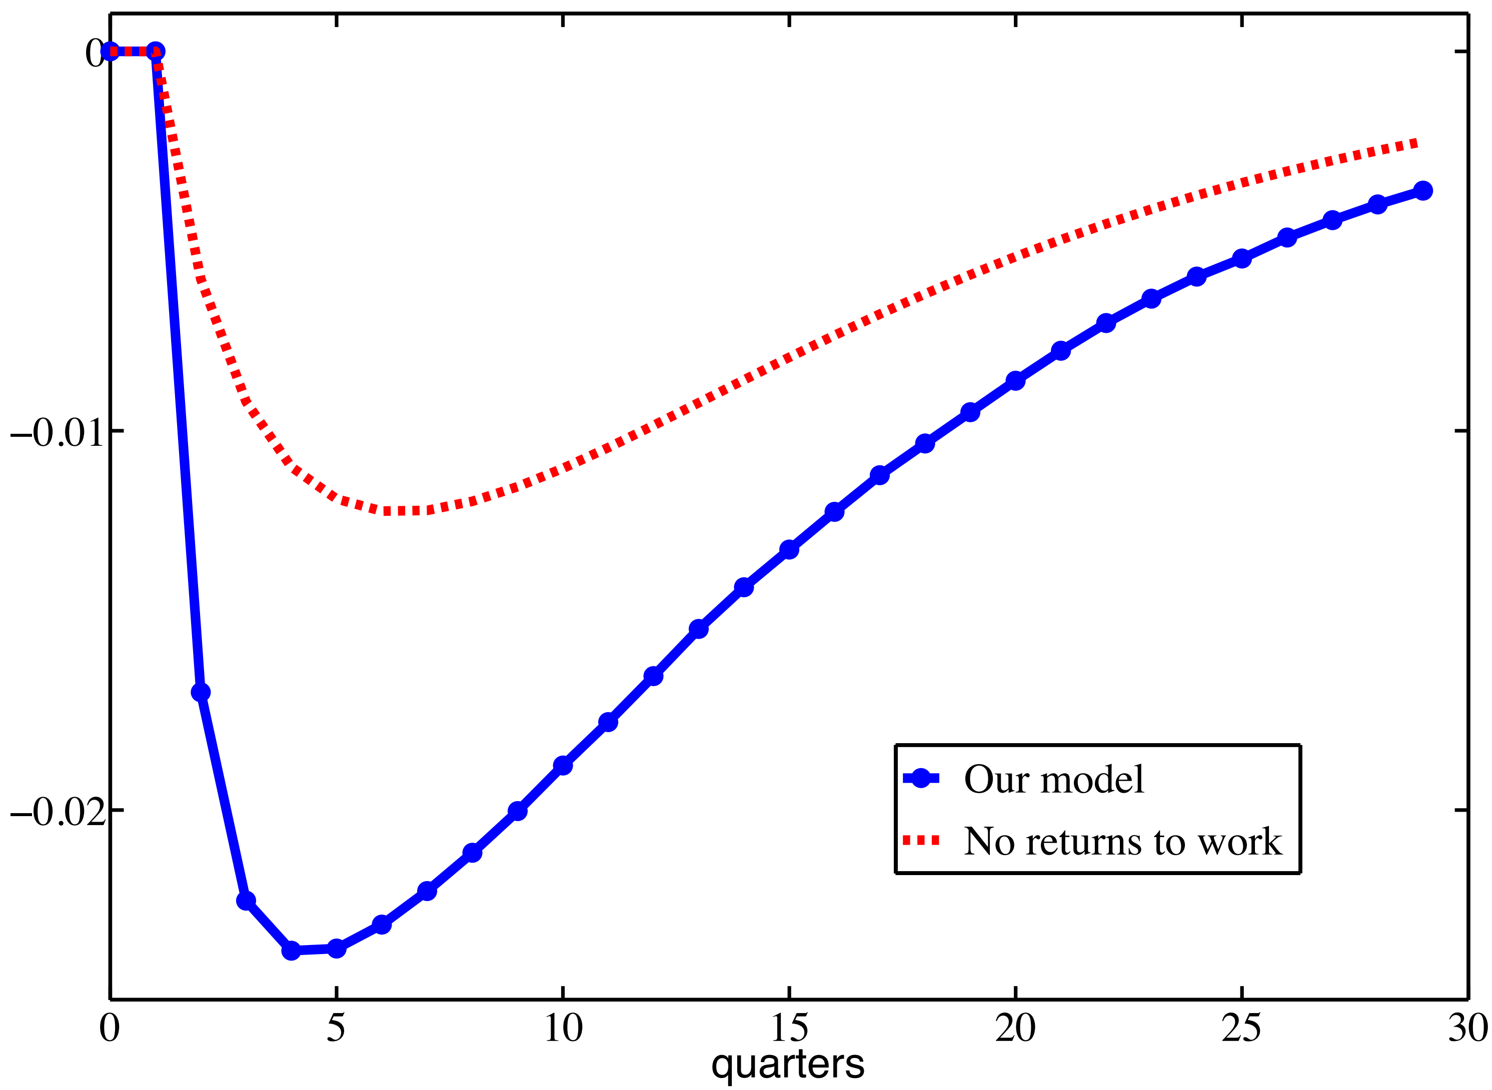

Figure 16: Tradable Employment Responses in Multi-Good Model

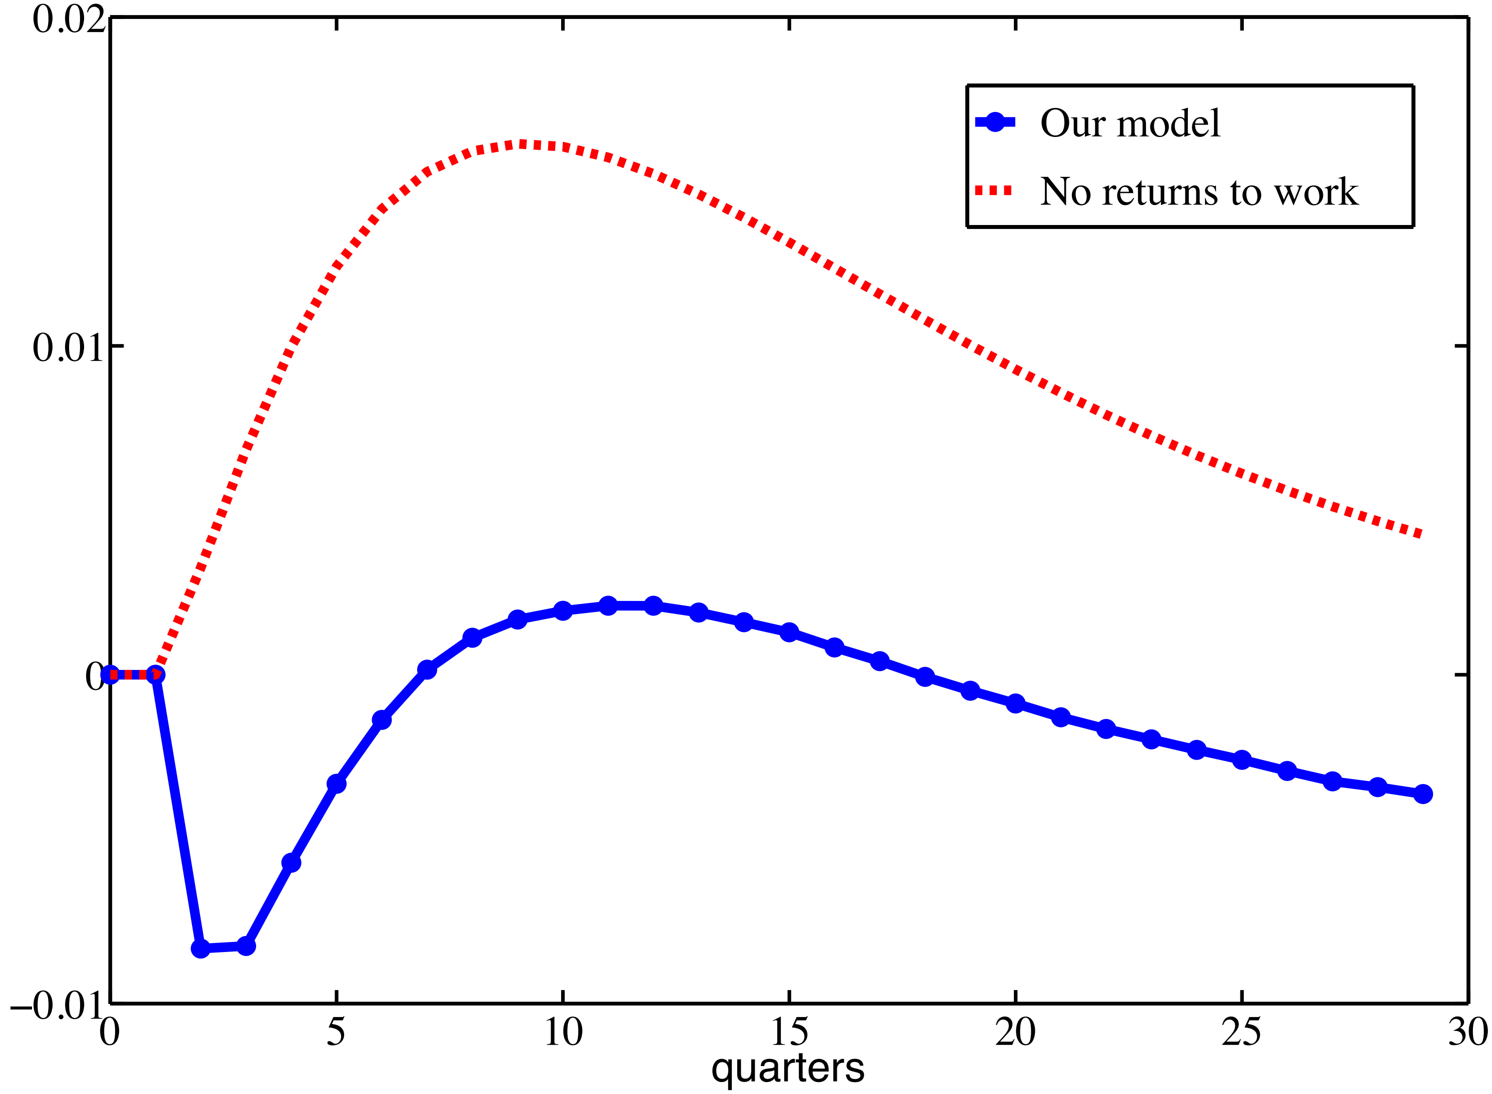

Figure 17: Employment vs. Consumption in Model and Data

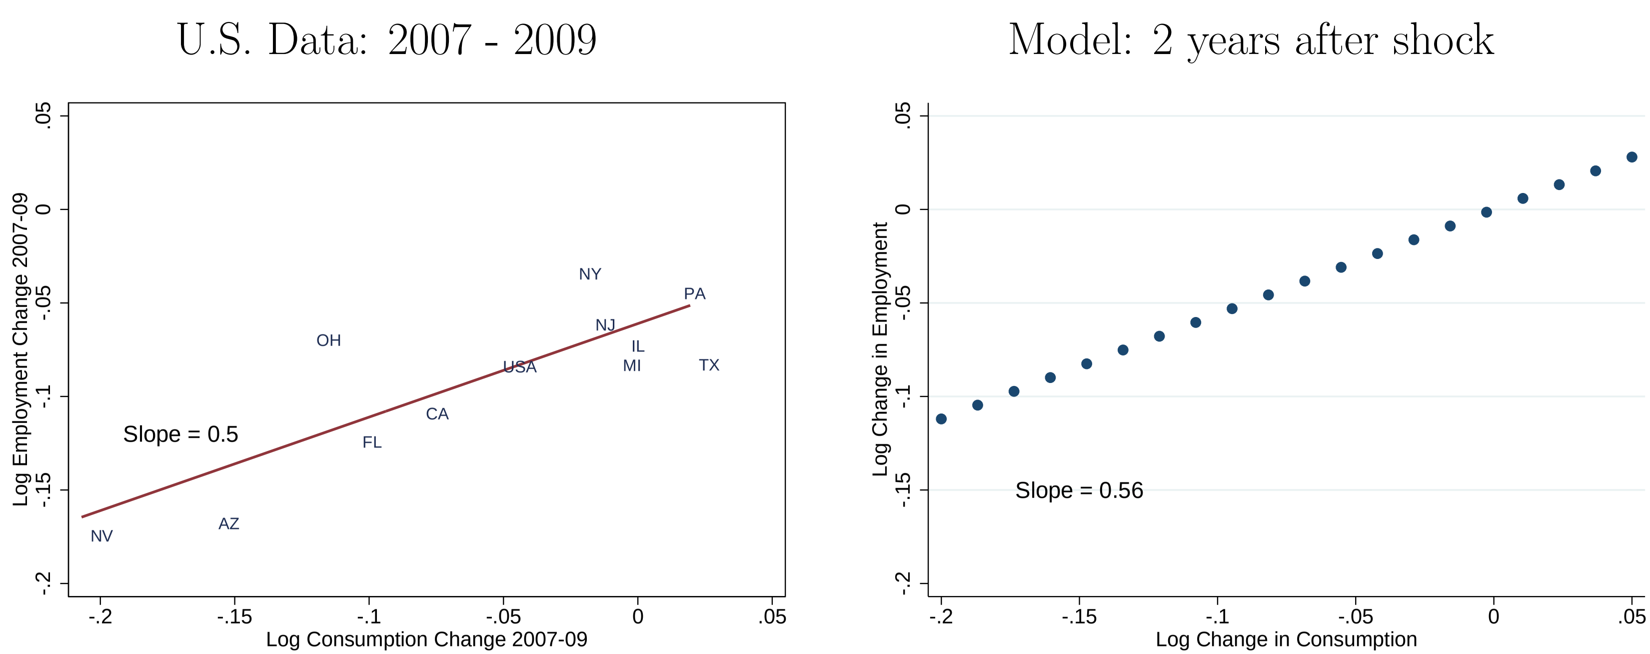

Figure 18: Non-Traded Employment vs. House Prices in Model and Data

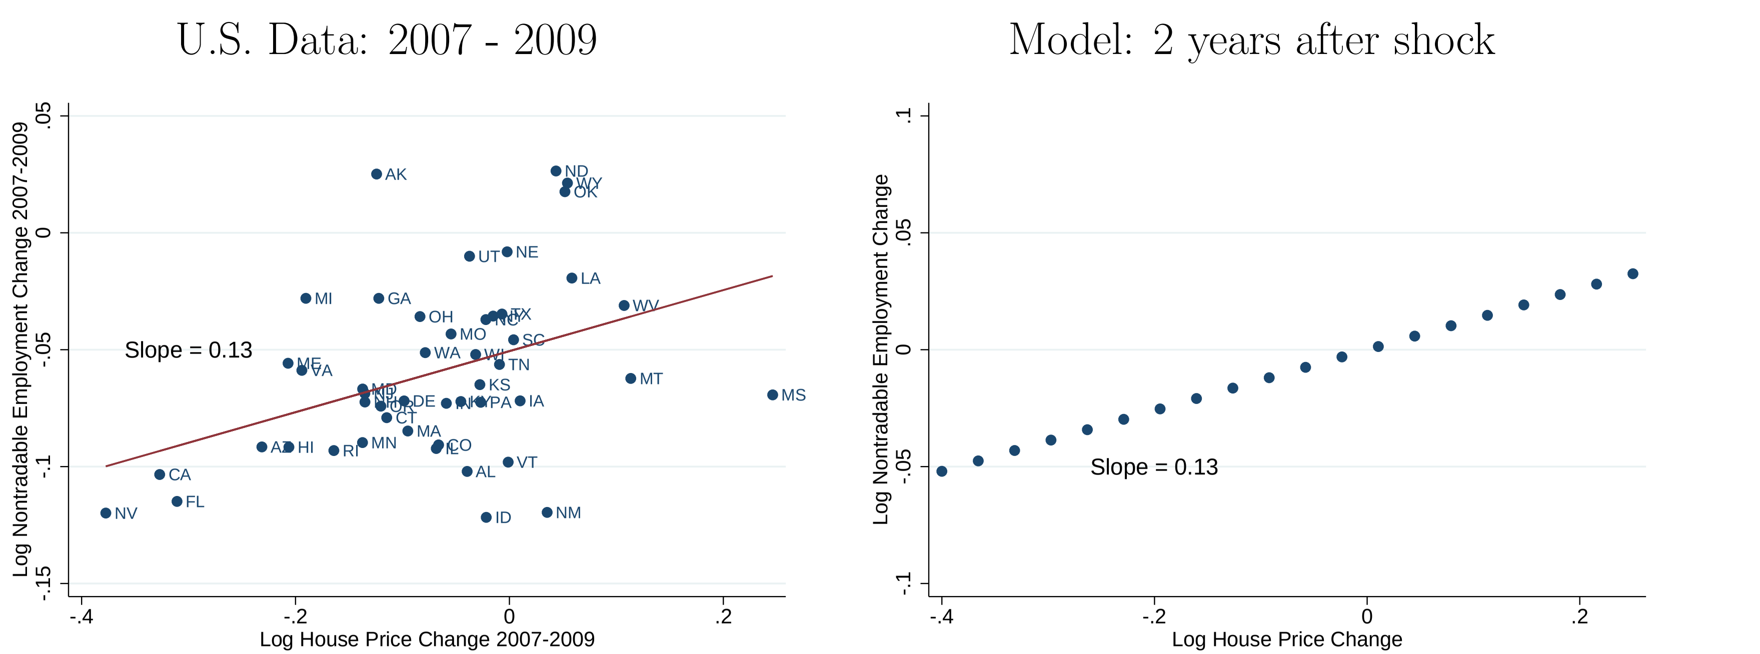

Figure 19: Tradable Employment vs. House Prices in Model and Data