

Agency Business Cycles*

MIKHAIL GOLOSOV

University of Chicago and NBER

GUIDO MENZIO

NYU and NBER

July 2019

Abstract

We develop a theory of endogenous and stochastic fluctuations in economic activity. Individual firms choose to randomize over firing or keeping workers who performed poorly in the past to give them an ex-ante incentive to exert effort. Different firms choose to correlate the outcome of their randomization to reduce the probability with which they fire non-performing workers. Correlated randomization leads to aggregate fluctuations. Aggregate fluctuations are endogenous—they emerge because firms choose to randomize and they choose to randomize in a correlated fashion—and they are stochastic—they are the manifestation of a randomization process. The hallmark of a theory of endogenous and stochastic fluctuations is that the stochastic process for aggregate “shocks” is an equilibrium object.

JEL Codes: D86, E24, E32.

Keywords: Endogenous and Stochastic Cycles, Coordinated Randomization, Unemployment Fluctuations.

*We are grateful to Gadi Barlevy, Paul Beaudry, Jess Benhabib, Katarina Borovickova, Veronica Guerrieri, Christian Haefke, Boyan Jovanovic, John Kennan, Narayana Kocherlakota, Ricardo Lagos, Rasmus Lentz, Igor Livschits, Paolo Martellini, Nicola Pavoni, Thijs van Rens, Guillaume Rocheteau, Tom Sargent, Karl Shell, Ben Schoefer, Rob Shimer, Mathieu Taschereau-Dumouchel and Randy Wright for comments on earlier drafts of the paper. We are also grateful to participants at seminars at the University of Wisconsin Madison, New York University, University of Chicago, University of California Berkeley, CREI, LSE, New York FRB, Minneapolis FRB, and at the Search and Matching European Conference in Aix-en-Provence, the Search and Matching Workshop at the Philadelphia FRB, the conference in honor of Chris Pissarides at Sciences Po, the Econometric Society World Congress in Montreal, the NBER EFG Meeting at the New York FRB.

1 Introduction

What causes cyclical fluctuations in economic activity? This is a central question in macroeconomics and, naturally, it has attracted a great deal of attention from both theorists and empiricists. One explanation for cyclical fluctuations is that the economy is subject to aggregate shocks to its fundamentals (see, e.g., Kydland and Prescott 1983). A second explanation for cyclical fluctuations is that the economic system admits multiple equilibria and there are switches in the equilibrium played by market participants (see, e.g., Benhabib and Farmer 1994). According to these two theories, cyclical fluctuations are exogenous and stochastic, in the sense that they are driven by an exogenously given stochastic process for the shocks to fundamentals or to the selection of equilibrium. The main criticism to these theories is that they leave the driving force of business cycles completely unexplained. A different view of cyclical fluctuations is that the economic system does not tend towards stasis—where the extent of economic activity remains constant over time—but it naturally oscillates between periods of high and low activity (see, e.g., Benhabib and Day 1982). According to this theory, cyclical fluctuations are endogenous and deterministic. The main criticism to this theory is that business cycles do not appear to follow a deterministic pattern.

In this paper, we develop a theory of endogenous and stochastic business cycles. The structure of our theory is simple: individual agents find it optimal to randomize over some choice in order to overcome a non-convexity in their decision problem, and different agents find it optimal to correlate the outcome of their randomization. Aggregate fluctuations are endogenous because they are an equilibrium outcome: individual agents choose to randomize over some decision (which endogenously creates individual uncertainty) and they choose to correlate the outcome of their randomization (which endogenously generates aggregate uncertainty from individual uncertainty). Aggregate fluctuations are stochastic because they are the manifestation of aggregate uncertainty. The distinguishing feature of our theory—and more generally the hallmark of a theory of endogenous and stochastic fluctuations—is that the stochastic process for aggregate “shocks” is endogenous. Therefore, our theory makes predictions about the economies in which shocks will be frequent or infrequent, large or small.

While the structure of our theory is fairly general, we exemplify it in the context of a search-theoretic model of the labor market in the spirit of Pissarides (1985) and Mortensen and Pissarides (1994). We consider a market populated by risk-averse workers and risk-neutral firms. Unemployed workers and vacant firms come together through a frictional search process. Once matched, a firm-worker pair bargains over the terms of

an employment contract and starts producing output. Production is subject to moral hazard—in the sense that the firm does not observe the effort of the worker but only output, which is a noisy measure of effort. The employment contract allocates the gains from trade between the firm and the worker and tries to overcome the moral hazard problem. In particular, the contract specifies the level of effort recommended to the worker, the wage paid to the worker, and the probability with which the worker is fired conditional on the realization of the worker’s output and possibly on the realization of a public sunspot.

Two features are necessary to develop our theory of endogenous and stochastic fluctuations. First, we need firms to find it optimal to use a firing lottery. In the model, we obtain this feature by assuming that the firm pays the wage before it observes the worker’s output and that the firm and the worker renegotiate the terms of the employment contract every period. Under these assumptions, the firm can provide the worker with the incentive to exert effort only by randomizing over firing or keeping him in case he produces low output. Second, we need firms to find it optimal to correlate the outcomes of their firing lotteries. In the model, we obtain this feature by assuming that the firm’s vacancy cost is convex, which, in turn, implies that the cost of losing a job to a worker is decreasing in the unemployment rate.

In the first part of the paper, we characterize the properties of the optimal employment contract. We show that the optimal contract is such that the worker is fired only when the realization of output is low and the state of the world (i.e. the realization of the sunspot) is one in which the gains from continued trade accruing to the worker are highest relative to those accruing to the firm. The result is intuitive. Firing is costly—as it destroys a valuable firm-worker relationship—but necessary—as it is the only way for the firm to give the worker an incentive to exert effort. When firing takes place, however, it is only the value of the relationship that would have accrued to the worker that provides incentives. The value that would have accrued to the firm is collateral damage. The optimal contract minimizes the collateral damage by loading the firing probability on the states of the world in which the worker’s gains from continued trade are highest relative to the firm’s. In other words, the optimal contract loads the firing probability on the states of the world where the cost to the worker from losing the job is highest relative to the cost to the firm from losing the worker.

In the second part of the paper, we characterize the equilibrium relationship between firing and relative gains from trade.¹ We show that there exists a correlated equilibrium

¹The characterization is done under the conjecture that the worker’s and firm’s gains from trade are strictly increasing in unemployment and that the equilibrium wage is strictly decreasing in unemployment.

in which all firms fire their non-performing workers for some realizations of the sunspot, and they all keep their non-performing workers for the other realizations. The measure of realizations of the sunspot for which there is firing is uniquely pinned down by the workers' incentive compatibility constraint. In this equilibrium, individual firms correlate the outcomes of their firing lottery. There is a simple logic behind this finding. Suppose that firms load up the firing probability on some states of the world. In those states of the world, the unemployment rate is higher and, because the vacancy cost is convex, the labor market tightness is lower and so is the probability with which unemployed workers find jobs. Since the job-finding probability is lower, workers have a weaker outside option when bargaining with firms and their wage is lower. Since the wage is lower, the marginal utility of income of workers (who are risk-averse) relative to the marginal utility of income of firms (who are risk-neutral) is higher and, per the Nash bargaining solution, the gains from trade accruing to the workers are high relative to those accruing to the firms. Thus, if the other firms load the firing probability on some states of the world, an individual firm finds it optimal to load the firing probability on the very same states of the world. In other words, firms want to correlate the outcome of the randomization between firing and keeping their non-performing workers, and the sunspot allows them to achieve correlation.

Alongside the correlated equilibrium described above, there exists an uncorrelated equilibrium in which firms ignore the realization of the sunspot and randomize over firing or keeping their non-performing workers in an independent fashion. The uncorrelated equilibrium, though, is not robust. The uncorrelated equilibrium only exists because firms randomize simultaneously and, hence, they must rely on an inherently meaningless signal (the sunspot) to achieve correlation. When a signal is inherently meaningless, there is always an equilibrium in which the signal is ignored. We show that, in a version of the model where firms randomize sequentially, history can always be used to achieve correlation and the uncorrelated equilibrium disappears. Specifically, firms who act later always find it optimal to correlate the outcome of their randomization to the randomization outcome of the firms who move first.²

In the correlated equilibrium, the economy goes through aggregate fluctuations, which we dub *Agency Business Cycles* or ABCs. ABCs are endogenous. They are not caused by exogenous shocks to current or future fundamentals nor by exogenous shocks to the equilibrium played by market participants. ABCs are caused by correlated randomization—i.e. individual firms find it optimal to randomize on firing or keeping their non-performing workers, and different firms find it optimal to correlate the outcomes of their randomiza-

The conjecture holds for the calibrated version of the model.

²The sequential version of the model also shows that a sunspot is not necessary for our theory of aggregate fluctuations.

tions. Correlation is achieved either through the sunspot (in the simultaneous version of the model) or through history (in the sequential version of the model). ABCs are stochastic. The economy does not follow a deterministic cycle, but a random process in which the probability of a correlated firing episode and, hence, a recession is endogenous. We show that the probability of a recession depends positively on the worker's cost of effort and negatively on the worker's cost of losing a job (which rises with unemployment).

In the last part of the paper, we calibrate the theory to the US labor market. We show that, for some parameter values, ABCs feature fluctuations in unemployment, unemployment-to-employment (UE), and employment-to-unemployment (EU) rates that have the same magnitude and pattern of comovement as in the data. Moreover, ABCs generate fluctuations in unemployment, UE and EU rates that are, as in the data, uncorrelated with fluctuations in labor productivity. However, we conclude that ABCs—at least in the formulation developed in this paper—are not a complete explanation of labor market fluctuations because they feature a correlation between unemployment and vacancies that is counterfactually positive.

The nature of recessions in ABCs is quite different than in theories of business cycles where fluctuations are driven by aggregate productivity shocks (such as Real Business Cycles, or RBCs). In RBCs, recessions are times when productivity is unusually low and so are the gains from trade in the labor market. For this reason, workers and firms have weaker incentives to trade and unemployment is high. In ABCs, recessions are states of the world in which the value of being unemployed to a worker is unusually low and, hence, the gains from trade in the labor market are high. In these states of the world, firms find it optimal to fire their non-performing workers and, hence, unemployment is high. Overall, in RBCs, the gains from trade move in the opposite direction as unemployment. In ABCs, the gains from trade move in the same direction as unemployment. We show that, at least when measured from the perspective of a worker, the gains from trade are countercyclical.

Our theory of business cycles is not a mere intellectual curiosity. There is empirical evidence consistent with the view that firms use firing as an incentive device. Cappelli and Chauvin (1991) examine the internal records of a large car manufacturing company. Exploiting geographical variation across plants, they establish a negative relationship between the plant's wage relative to the average manufacturing wage in the plant's area (which is a measure of the cost to the worker of losing the job) and the frequency of disciplinary dismissals. The finding suggests that the firm uses the threat of firing as an incentive device, that workers understand the threat, and that they adjust their effort according to the strength of the threat. Ichino and Riphahn (2005) examine days of absence per week for white-collar workers at a large Italian bank. During the first 12

weeks of their tenure, workers are in a probationary period and can be fired at will. Afterwards, workers enjoy strong employment protection. Ichino and Riphahn (2005) find that days of absence per week triple right after the end of the probationary period. That is, workers' effort (as measured by absenteeism) varies depending on whether they can or cannot be fired at will. The finding suggests that workers expect the bank to use firing as part of its incentive scheme.

There is also empirical evidence consistent with the mechanism behind our theory. Argawal and Kolev (2013) show that Fortune 500 companies tend to cluster mass layoffs within a few days of each other, even though mass layoffs are a relatively infrequent event. Interestingly, they find that an announcement of mass layoffs by one of the top 20 firms is positively related with announcements of mass layoffs by other Fortune 500 firms in the five following business days, while it is uncorrelated with mass layoffs in the five previous business days. The asymmetry suggests that firms are not being hit by a common shock. As in our theory, it could be that smaller firms use the layoff decisions of larger firms as a coordination device. Similar clustering seems to take place at the top of organizations as well. Jenter and Kanaan (2015) find that CEOs who underperform the industry average are much more likely to be fired when the industry-wide performance is poor, even though one would imagine that only a CEO's relative performance is informative about effort. The finding implies that the firing of non-performing CEOs is clustered during downturns.

The main contribution of the paper is to develop a new theory of endogenous and stochastic business cycles. We show that endogenous and stochastic cycles emerge in equilibrium when individual agents want to randomize over some economic decision and different agents find it optimal to correlate the outcome of their randomization. The hallmark of a theory of endogenous and stochastic business cycles is a stochastic process for aggregate "shocks" that is determined endogenously. Conceptually, a theory of endogenous and stochastic cycles is useful because it explains why the aggregate economy is subject to shocks, rather than simply assuming the existence of shocks. This implies that the theory has something to say about what determines the frequency of shocks, the magnitude of shocks and what policies may affect the stochastic process of shocks. Empirically, a theory of endogenous and stochastic business cycles is useful because it helps explain why economic activity seems (at first blush) to be more volatile than its fundamentals, and it does so without resorting to unobserved shocks to equilibrium selection.

Intellectually, a theory of endogenous and stochastic cycles adds to the class of theories that we can use to understand macroeconomic fluctuations. Some of the existing theories of aggregate fluctuations are based on exogenous shocks to fundamentals. These can be shocks to the current value of economy-wide fundamentals (e.g., Kydland and Prescott

1982 or Mortensen and Pissarides 1994), to the future value of fundamentals (e.g., Beaudry and Portier 2004 or Jaimovich and Rebelo 2009), to the stochastic process of fundamentals (e.g., Bloom 2009), or to higher-order beliefs (e.g., Angeletos and La’O 2013). Relatedly, there are granular theories of business cycles, in which aggregate fluctuations are driven by shocks to the fundamentals of individual agents who are large enough or connected enough to others to cause aggregate swings in economic activity (e.g., Jovanovic 1987 or Gabaix 2011). Other theories of business cycles are driven by exogenous shocks to equilibrium selection (e.g., Heller 1986, Cooper and John 1988, Benhabib and Farmer 1994, Kaplan and Menzio 2016). Finally, there are theories of endogenous and deterministic aggregate fluctuations, where the economy converges to a limit cycle (e.g., Benhabib and Nishimura 1979, Diamond 1982, Diamond and Fudenberg 1989, Benhabib and Rustichini 1990, Mortensen 1999 or Beaudry, Galizia and Portier 2015) or follows chaotic dynamics (e.g., Benhabib and Day 1982, Boldrin and Montrucchio 1986 or Boldrin and Woodford 1990). The only other theory of endogenous and stochastic business cycles of which we are aware is Benhabib, Wang and Wen (2015). Their theory shares with ours the fact that the probability distribution of aggregate “shocks” is an equilibrium object. However, the mechanism leading to endogenous, stochastic cycles is different from ours, as it builds on a signal extraction problem.

The particular illustration of our theory contributes to the literature on labor market fluctuations. Shimer (2005) showed that the basic search-theoretic model of the labor market implies very small fluctuations in unemployment in response to the observed fluctuations in labor productivity. Building on this observation, many papers have identified channels through which labor productivity shocks can lead to sizeable movements in unemployment (e.g., wage rigidity in Hall 2005, Menzio 2005, Kennan 2010, Menzio and Moen 2010, small gap between home productivity and market productivity in Hagedorn and Manovskii 2008, match heterogeneity in Menzio and Shi 2011). These papers typically generate a perfect negative correlation between labor productivity and unemployment. Yet, since 1984, this correlation has vanished. Recent work has thus focused on identifying different sources of unemployment fluctuations (e.g., Farmer 2013, Galí and Van Rens 2014, Kaplan and Menzio 2016, Beaudry, Galizia and Portier 2016, Hall 2016). Our model offers a novel explanation for why unemployment is so volatile and why its volatility is uncorrelated with productivity. A distinguishing feature of our explanation relative to others is that it implies a positive correlation between the net value of employment and unemployment. We find evidence of this positive correlation in the data.

Finally, let us briefly relate our paper to Shapiro and Stiglitz (1984). Shapiro and Stiglitz (1984) showed that the existence of a moral hazard problem between firms and

workers generates unemployment in a frictionless labor market. Our paper also features a moral hazard problem between firms and workers, but the focus of our paper is not on explaining the existence of unemployment—which is caused by search frictions—but the volatility of unemployment.

2 Environment and Equilibrium

In this section, we describe the physical and contractual environment of the model and derive the conditions for a Recursive Equilibrium.

2.1 Environment

Time is discrete and continues forever. The economy is populated by a measure 1 of identical workers. Every worker has preferences described by the expected sum of current and future periodical utilities discounted at the factor $\beta \in (0, 1)$. When a worker is unemployed, his periodical utility is given by $v(b) + \zeta$, where $v(\cdot)$ is a strictly increasing, strictly concave function of consumption, b is the worker's unemployment income, and ζ is the worker's utility from leisure. When a worker is employed, his periodical utility is given by $v(w_t) - \psi e_t$, where w_t is the worker's labor income³ and ψe_t is the worker's disutility from putting effort on the job, where $\psi > 0$ and $e_t \in \{0, 1\}$.

The economy is also populated by a measure 1 of identical firms. Every firm has preferences described by the expected sum of current and future periodical profits, discounted at the factor β . Every firm operates a constant returns to scale production technology that transforms one unit of labor (i.e. one employee) into y_t units of output, where y_t is a random variable that depends on the employee's effort e_t . In particular, y_t takes the value y_h with probability $p_h(e)$ and the value y_ℓ with probability $p_\ell(e) = 1 - p_h(e)$, with $y_h > y_\ell \geq 0$ and $0 < p_h(0) < p_h(1) < 1$. Production suffers from moral hazard, in the sense that the firm does not directly observe the effort of its employee, but only the output.

Every period t is divided into five stages: sunspot, separation, matching, bargaining and production. At the first stage, a random variable, z_t , is drawn from a uniform distribution with support $[0, 1]$.⁴ The random variable is aggregate, in the sense that it

³As the reader can infer from the notation, we assume that workers are hand-to-mouth, in the sense that they consume their income in every period. We discuss the robustness of our theory to this assumption in the conclusions.

⁴Assuming that the sunspot is an i.i.d. draw from a uniform with support $[0, 1]$ is without loss of generality. In fact, as the sunspot does not directly affect preferences or technology, an equilibrium of a model where the sunspot is drawn from some arbitrary CDF can always be represented as an equilibrium of the model where the sunspot is drawn from a uniform distribution. Moreover, in any equilibrium in

is publicly observed by all market participants. The random variable is a sunspot, in the sense that it does not directly affect technology, preferences or any other fundamentals, although it may help correlate the outcome of the lotteries played by different market participants.

At the separation stage, some employed workers become unemployed. An employed worker becomes unemployed for exogenous reasons with probability $\delta \in (0, 1)$. In addition, an employed worker becomes unemployed because he is fired with probability $s(y_{t-1}, z_t)$, where $s(y_{t-1}, z_t)$ is determined by the worker's employment contract and it is allowed to depend on the output of the worker in the previous period, y_{t-1} , and on realization of the sunspot in the current period, z_t . For the sake of simplicity, we assume that a worker who becomes unemployed in period t can search for a new job only starting in period $t + 1$.

At the matching stage, some unemployed workers become employed. Firms decide how many vacancies v_t to create at the unit cost $k(v_t)$, where $k(\cdot)$ is a strictly increasing function such that $k(0) = 0$. Then, the u_{t-1} workers who were unemployed at the beginning of the period search for the v_t vacancies created by firms. The outcome of the search process is described by a constant return to scale matching function, $M(u_{t-1}, v_t)$, which gives the measure of bilateral matches formed between unemployed workers and vacancies. Hence, the probability that an unemployed worker meets a vacancy is $\lambda(\theta_t) \equiv M(1, \theta_t)$, where $\theta_t \equiv v_t/u_{t-1}$ is the tightness of the labor market and $\lambda(\cdot)$ is a strictly increasing and concave function such that $\lambda(0) = 0$. The probability that a vacancy meets an unemployment worker is $\eta(\theta_t) \equiv M(1/\theta_t, 1)$, where $\eta(\cdot)$ is a strictly decreasing function such that $\eta(\theta) = \lambda(\theta)/\theta$.

At the bargaining stage, new and continuing firm-worker pairs negotiate the terms of a one-period employment contract x_t . The contract x_t specifies the effort e_t recommended to the worker in the current period, the wage w_t paid by the firm to the worker in the current period, and the probability $s(y_t, z_{t+1})$ with which the firm fires the worker at the next separation stage, conditional on the output of the worker in the current period and on the realization of the sunspot at the beginning of next period. We assume that the outcome of the bargain between the firm and the worker is the Axiomatic Nash Bargaining Solution.

At the production stage, an unemployed worker home-produces and consumes b units of output. An employed worker chooses an effort level, e_t , and consumes w_t units of output. Then, the output of the worker, y_t , is realized and observed by both the firm and

which firms perfectly correlate their randomization (which is the “robust” equilibrium of the model), the probability of the different aggregate events is uniquely pinned down. Therefore, the same distribution of aggregates would emerge if the sunspot was i.i.d. or autocorrelated.

the worker.

A few comments about the environment are in order. We assume that employment contracts are short-term—in the sense that they can only specify effort, wages, and separation probabilities for one period before being renegotiated—and incomplete—in the sense that they cannot specify a wage that depends on the contemporaneous realization of output. Short-term, incomplete contracts imply that firms must use firing lotteries to provide incentives to their workers.⁵ We assume that vacancy costs are convex, in the sense that the cost of an additional vacancy is increasing in the number of vacancies opened by the firm. Convex vacancy costs imply that the value of unemployment to an individual worker is decreasing in aggregate unemployment.⁶

2.2 Definition of Equilibrium

We now derive the conditions for an equilibrium in our model economy.⁷ Let u denote the measure of unemployed workers at the production stage. Let $W_0(u)$ denote the lifetime utility of a worker who is unemployed at the production stage. Let $W_1(x, u)$ denote the lifetime utility of a worker who is employed under the contract x at the production stage. Let $W(x, u)$ denote the difference between $W_1(x, u)$ and $W_0(u)$. Let $F(x, u)$ denote the present value of profits for a firm that, at the production stage, employs a worker under the contract x . Let $x^*(u)$ denote the equilibrium employment contract between a firm and a worker. Finally, let $\theta(u, \hat{z})$ denote the tightness of the labor market at the matching stage of next period, given that next period's sunspot is \hat{z} . Similarly, let $h(u, \hat{z})$ denote the unemployment at the production stage of next period, given that next period's sunspot is \hat{z} .

The lifetime utility $W_0(u)$ of an unemployed worker is such that

$$W_0(u) = v(b) + \zeta + \beta E_{\hat{z}} [W_0(h(u, \hat{z})) + \lambda(\theta(u, \hat{z}))W(x^*(h(u, \hat{z})), h(u, \hat{z}))]. \quad (2.1)$$

In the current period, the worker home-produces and consumes b units of output. At the matching stage of next period, the worker finds a job with probability $\lambda(\theta(u, \hat{z}))$. In this case, the worker's continuation lifetime utility is $W_0(h(u, \hat{z})) + W(x^*(h(u, \hat{z})), h(u, \hat{z}))$.

⁵In the conclusions, we discuss the robustness of our theory to relaxing the assumption of short-term, incomplete employment contracts.

⁶In the textbook search-theoretic model of the labor market (see, e.g., Pissarides 1985 or Mortensen and Pissardies 1994), the vacancy cost is assumed to be linear. For this reason, the equilibrium is such that the value of unemployment to an individual worker is independent from aggregate unemployment. Empirical studies of the hiring behavior of firms (see, e.g., Gavazza, Mongey and Violante 2017) suggest that vacancy costs are convex.

⁷We restrict attention to Recursive Equilibria. That is, we restrict attention to equilibria in which the value and policy functions depend only on the state of the economy, which in our model is the measure of unemployed workers.

With probability $1 - \lambda(\theta(u, \hat{z}))$, the worker does not find a job and his continuation lifetime utility is $W_0(h(u, \hat{z}))$.

The lifetime utility $W_1(x, u)$ of a worker employed under the contract $x = (e, w, s)$ is such that

$$W_1(x, u) = v(w) - \psi e + \beta E_{y, \hat{z}} [W_0(h(u, \hat{z})) + (1 - \delta)(1 - s(y, \hat{z}))W(x^*(h(u, \hat{z})), h(u, \hat{z})) | e]. \quad (2.2)$$

In the current period, the worker consumes w units of output and exerts effort e . At the separation stage of next period, the worker keeps his job with probability $(1 - \delta)(1 - s(y, \hat{z}))$. In this case, the worker's continuation lifetime utility is $W_0(h(u, \hat{z})) + W(x^*(h(u, \hat{z})), h(u, \hat{z}))$. With probability $1 - (1 - \delta)(1 - s(y, \hat{z}))$, the worker loses his job and his continuation lifetime utility is $W_0(h(u, \hat{z}))$.

The difference $W(x, u)$ between $W_1(x, u)$ and $W_0(u)$ represents the gains from trade to a worker employed under the contract x . From (2.1) and (2.2), it follows that $W(x, u)$ is such that

$$W(x, u) = v(w) - v(b) - \psi e - \zeta + \beta E_{y, \hat{z}} \{ [(1 - \delta)(1 - s(y, \hat{z})) - \lambda(\theta(\hat{z}, u))] W(x^*(h(u, \hat{z})), h(u, \hat{z})) | e \}. \quad (2.3)$$

We find it useful to denote as $V(u)$ the gains from trade for a worker employed under the equilibrium contract $x^*(u)$, i.e. $V(u) = W(x^*(u), u)$.

The present value of profits $F(x, u)$ for a firm that employs a worker under the contract $x = (e, w, s)$ is such that

$$F(x, u) = E_y[y|e] - w + \beta E_{y, \hat{z}} [(1 - \delta)(1 - s(y, \hat{z}))F(x^*(h(u, \hat{z})), h(u, \hat{z})) | e]. \quad (2.4)$$

In the current period, the firm enjoys a profit equal to the expected output of the worker net of the wage. At the separation stage of next period, the firm retains the worker with probability $(1 - \delta)(1 - s(y, \hat{z}))$. In this case, the firm's continuation present value of profits is $F(x^*(h(u, \hat{z})), h(u, \hat{z}))$. With probability $1 - (1 - \delta)(1 - s(y, \hat{z}))$, the firm loses the worker, in which case the firm's continuation present value of profits is zero. We find it useful to denote as $J(u)$ the present value of profits for a firm that employs a worker at the equilibrium contract $x^*(u)$, i.e. $J(u) = F(x^*(u), u)$.

The equilibrium contract $x^*(u)$ is the Axiomatic Nash Solution to the bargaining problem between the firm and the worker. That is, $x^*(u)$ is such that

$$\max_{x=(e,w,s)} W(x, u)F(x, u), \quad (2.5)$$

subject to the logical constraints

$$e \in \{0, 1\} \text{ and } s(y, \hat{z}) \in [0, 1],$$

and the worker's incentive compatibility constraints

$$\begin{aligned} \psi &\leq \beta(p_h(1) - p_h(0))E_{\hat{z}} [(1 - \delta)(s(y_\ell, \hat{z}) - s(y_h, \hat{z}))V(h(u, \hat{z}))], \quad \text{if } e = 1, \\ \psi &\geq \beta(p_h(1) - p_h(0))E_{\hat{z}} [(1 - \delta)(s(y_\ell, \hat{z}) - s(y_h, \hat{z}))V(h(u, \hat{z}))], \quad \text{if } e = 0. \end{aligned}$$

In words, the equilibrium contract $x^*(u)$ maximizes the product between the gains from trade accruing to the worker, $W(x, u)$, and the gains from trade accruing to the firm, $F(x, u)$, among all contracts x that satisfy the worker's incentive compatibility constraints. The first incentive compatibility constraint states that, if the contract specifies $e = 1$, the cost to the worker from exerting effort must be smaller than the benefit. The second constraint states that, if the contract specifies $e = 0$, the cost to the worker from exerting effort must be greater than the benefit. The cost of effort is ψ . The benefit of effort is given by the effect of effort on the probability that the realization of output is high, $p_h(1) - p_h(0)$, times the effect of a high realization of output on the probability of keeping the job, $(1 - \delta)(s(y_\ell, \hat{z}) - s(y_h, \hat{z}))$, times the value of the job to the worker, $\beta V(h(u, \hat{z}))$.

The equilibrium market tightness $\theta(u, \hat{z})$ must be consistent with the firms' incentives to create vacancies. The cost to the firm from creating an additional vacancy is $k(u\theta(u, \hat{z}))$. The benefit to the firm from creating an additional vacancy is given by the job-filling probability, $\eta(\theta(u, \hat{z}))$, times the value to the firm of filling a vacancy, $J(h(u, \hat{z}))$. The market tightness is consistent with the firm's incentives to create vacancies if

$$k(u\theta(u, \hat{z})) = \eta(\theta(u, \hat{z}), u)J(h(u, \hat{z})). \quad (2.6)$$

The equilibrium law of motion for unemployment, $h(u, \hat{z})$, must be consistent with the equilibrium firing probability $s^*(y, \hat{z}, u)$ and with the job-finding probability $\lambda(\theta(u, \hat{z}), u)$. Specifically, $h(u, \hat{z})$ must be such that

$$h(u, \hat{z}) = u - u\lambda(\bar{\theta}(J(h(u, \hat{z})), u)) + (1 - u)E_y [\delta + (1 - \delta)s^*(y, \hat{z}, u)], \quad (2.7)$$

where $\bar{\theta}(J, u)$ denotes the unique solution with respect to θ of (2.6) given that the value of a worker to a firm is J and the measure of unemployed workers is u . The first term on the right-hand side of (2.7) is unemployment at the production stage in the current period. The second term is the measure of unemployed workers who become employed during the matching stage of next period, which is given by unemployment u times the probability $\lambda(\bar{\theta}(J(h(u, \hat{z})), u))$ that an unemployed worker becomes employed. The last

term is the measure of employed workers who become unemployed during the separation stage of next period. The sum of the three terms on the right-hand side of (2.7) is the unemployment at the production stage of next period.

We are now in the position to define a recursive equilibrium.

Definition 1: *A Recursive Equilibrium is a tuple (W, F, V, J, x^*, h) such that: (i) The gains from trade accruing to the worker, $W(x, u)$, and to the firm, $F(x, u)$, satisfy (2.3) and (2.4) and $V(u) = W(x^*(u), u)$, $J(u) = F(x^*(u), u)$; (ii) The employment contract $x^*(u)$ satisfies (2.5); (iii) The law of motion $h(u, \hat{z})$ satisfies (2.7).*

Over the next two sections, we characterize the properties of equilibrium. We carry out the characterization under the maintained assumptions that the gains from trade are strictly positive, i.e. $J(u) > 0$ and $V(u) > 0$, and that the employment contract requires the worker to exert effort, i.e. $e^*(u) = 1$. The first assumption guarantees that firms and workers want to trade, and the second assumption guarantees that firms and workers want to solve the moral hazard problem.

3 Optimal Contract

In this section, we characterize the properties of the employment contract that maximizes the product of the gains from trade accruing to the worker and the gains from trade accruing to the firm subject to the worker's incentive compatibility constraint. We refer to this employment contract as the optimal contract. We first show that the optimal contract is such that the firing probability is loaded onto the states of the world (i.e. realizations of the sunspot) in which the worker's cost of losing the job is largest relative to the firm's cost of losing a worker. We then show that these states of the world are those in which the wage is lowest.

To lighten up the notation, and without risk of confusion, we drop the dependence of the gains from trade to the worker, W , and to the firm, F , as well as the dependence of the optimal contract, x^* , on unemployment u . We also drop the dependence of the continuation gains from trade to the worker and to the firm on unemployment and write $V(h(u, \hat{z}))$ as $V(\hat{z})$ and $J(h(u, \hat{z}))$ as $J(\hat{z})$.

Lemma 1. *(Firing Probabilities). Any optimal contract x^* is such that:*

(i) *The worker's incentive compatibility constraint holds with equality:*

$$\psi = \beta(p_h(1) - p_h(0))E_{\hat{z}} [(1 - \delta)(s^*(y_\ell, \hat{z}) - s^*(y_h, \hat{z}))V(\hat{z})]. \quad (3.1)$$

(ii) If the realization of output is high, the worker is fired with probability 0:

$$s^*(y_h, \hat{z}) = 0, \forall \hat{z} \in [0, 1]. \quad (3.2)$$

(iii) If the realization of output is low, there exists a ϕ^* such that the worker is fired with probability 0 if $\phi(\hat{z}) < \phi^*$, and the worker is fired with probability 1 if $\phi(\hat{z}) > \phi^*$ with $\phi(\hat{z}) \equiv V(\hat{z})/J(\hat{z})$:

$$s^*(y_\ell, \hat{z}) = \begin{cases} 1, & \forall \hat{z} \in [0, 1] \text{ s.t. } \phi(\hat{z}) > \phi^*, \\ 0, & \forall \hat{z} \in [0, 1] \text{ s.t. } \phi(\hat{z}) < \phi^*. \end{cases} \quad (3.3)$$

Proof. In Appendix A. ■

Parts (i) and (ii) of Lemma 1 are standard properties of optimal contracts under moral hazard and are easy to understand. To understand part (i), consider a contract x such that the worker's incentive compatibility constraint is lax. Since the incentive compatibility constraint is lax, the contract x is such that the worker is fired with positive probability when the realization of output is low, i.e. $s(y_\ell, \hat{z}_0) > 0$ for some \hat{z}_0 . Now, consider modifying the contract x by lowering the firing probability $s(y_\ell, \hat{z}_0)$ by some amount that is small enough to still satisfy the worker's incentive compatibility constraint. The survival probability of the match is higher in the modified contract than in the original one. Since the continuation value of the match is strictly positive for both the worker and the firm, a higher survival probability of the match implies higher gains from trade W for the worker, higher gains from trade F to the firm, as well as a higher product $W \cdot F$. Thus, the modified contract is an improvement over the original one.

To understand part (ii) of Lemma 1, consider a contract x such that the worker is fired with positive probability when the realization of output is high, i.e. $s(y_h, \hat{z}_0) > 0$ for some \hat{z}_0 . Again, consider modifying the contract x by lowering the firing probability $s(y_h, \hat{z}_0)$. The modified contract satisfies the worker's incentive compatibility constraint. Moreover, the survival probability of the match is higher in the modified contract than in the original one. Hence, in the modified contract, the gains from trade W accruing to the worker, the gains from trade F accruing to the firm, and the product $W \cdot F$ are higher than in the original contract.

Part (iii) of Lemma 1 is more interesting and is one of the key insights of the paper. It states that any optimal contract is such that, when the realization of output is low, the worker is fired with probability 1 in states of the world \hat{z} where the continuation gains from trade to the worker, $V(\hat{z})$, relative to the continuation gains from trade to the firm, $J(\hat{z})$, are greater than some cutoff ϕ^* . The worker is fired with probability 0 in states of the world \hat{z} where the relative continuation gains from trade to the worker

are smaller than the cutoff ϕ^* . There is a simple intuition behind this property. Firing is costly—as it destroys a valuable relationship—but also necessary—as it is the only tool to provide the worker with an incentive to exert effort. However, only the value of the destroyed relationship that would have accrued to the worker serves the purpose of providing incentives. The value of the destroyed relationship that would have accrued to the firm is collateral damage. The optimal contract minimizes the collateral damage by loading the firing probability on states of the world in which the value of the relationship to the worker would have been highest relative to the value of the relationship to the firm. In other words, the optimal contract minimizes the collateral damage by loading the firing probability on states of the world in which the cost to the worker from losing the job, $V(\hat{z})$, is highest relative to the cost to the firm from losing the worker, $J(\hat{z})$.

Next, we want to show that the cutoff ϕ^* is unique and, hence, so are the firing probabilities. To this aim, note that the cutoff ϕ^* must solve

$$\psi = \beta(p_h(1) - p_h(0))(1 - \delta) \left[\int_{\phi(\hat{z}) > \phi^*} V(\hat{z}) d\hat{z} + \int_{\phi(\hat{z}) = \phi^*} s^*(y_\ell, \hat{z}) V(\hat{z}) d\hat{z} \right], \quad (3.4)$$

The equation above is the worker's incentive compatibility constraint (3.1), which we have rewritten using the fact that $s^*(y_h, \hat{z})$ and $s^*(y_\ell, \hat{z})$ are respectively given by (3.2) and (3.3). Figure 1 plots the right-hand side of (3.4) as a function of ϕ^* . On any interval $[\phi_0, \phi_1]$ where the distribution of the random variable $\phi(\hat{z})$ has positive density, the right-hand side of (3.4) is strictly decreasing in ϕ^* . On any interval $[\phi_0, \phi_1]$ where the distribution of $\phi(\hat{z})$ has no density, the right-hand side of (3.4) is constant. At any value ϕ^* where the distribution of $\phi(\hat{z})$ has a mass point, the right-hand side of (3.4) can take on an interval of values, as the firing probability $s^*(y_\ell, \hat{z})$ for \hat{z} such that $\phi(\hat{z}) = \phi^*$ may take any value between 0 and 1.

The cutoff ϕ^* must be such that the right-hand side of (3.4) equals ψ . There are three possible cases to consider. The first case is illustrated by $\psi = \psi_1$ in Figure 1. In this case, there exists a unique ϕ^* that solves equation (3.4) and the random variable $\phi(\hat{z})$ has positive density around ϕ^* and no mass point at ϕ^* . In this case, the firing probability $s^*(y_\ell, \hat{z})$ is either 0 (for all \hat{z} such that $\phi(\hat{z}) < \phi^*$) or 1 (for all \hat{z} such that $\phi(\hat{z}) > \phi^*$). The second case is illustrated by $\psi = \psi_2$. In this case, there also exists a unique ϕ^* that solves equation (3.4) but the random variable $\phi(\hat{z})$ has a mass point at ϕ^* . In this case, the firing probability $s^*(y_\ell, \hat{z})$ is 0 for all \hat{z} such that $\phi(\hat{z}) < \phi^*$, 1 for all \hat{z} such that $\phi(\hat{z}) > \phi^*$, and between 0 and 1 for all \hat{z} such that $\phi(\hat{z}) = \phi^*$. The third and last case is illustrated by $\psi = \psi_3$. In this case, there exists an interval of ϕ^* 's that solve equation (3.4). However, the choice of ϕ^* is immaterial as the random variable $\phi(\hat{z})$ has no mass




Figure 1: Equilibrium Cutoff ϕ^*

over this interval. In this case, the firing probability $s^*(y_\ell, \hat{z})$ is either 0 (for all $\phi(\hat{z})$ to the left of the interval) or 1 (for all $\phi(\hat{z})$ to the right of the interval).

Next, we want to characterize the worker's wage in an optimal contract.

Lemma 2. (*Wage*) Any optimal contract x^* is such that the wage w^* satisfies

$$\frac{v'(w^*)}{1} = \frac{W(x^*)}{F(x^*)}. \quad (3.5)$$

Proof. In Appendix B. ■

Lemma 2 states that an optimal contract prescribes a wage w^* such that the marginal utility of a higher wage to the worker relative to the marginal cost of a higher wage to the firm, $v'(w^*)/1$, is equal to the worker's gains from trade relative to the firm's gains from trade, $W(x^*)/F(x^*) = V/J$. Lemma 2 is not surprising. The optimality condition for w^* is nothing more than the standard condition for maximizing the Nash product. Lemma 2 is, however, important for our business cycle theory. The lemma implies that the states of the world in which the worker's relative gains from trade are highest are the states of the world in which the wage is lowest. Hence, in light of Lemma 1, the states of the world in which the firm is more likely to fire the worker are the states of the world in which the worker's wage would have been low (i.e. recessions).

We are now in the position to summarize the characterization of the optimal employment contract.

Theorem 3. (*Optimal Contract*). Given the distribution of the random variable $\phi(\hat{z})$, the optimal contract x^* is unique. (i) The wage w^* is given by (3.5); (ii) The firing probability

$s^*(y_h, \hat{z})$ is given by (3.2); (iii) The firing probability $s^*(y_\ell, \hat{z})$ is given by (3.3); (iv) The cutoff ϕ^* is uniquely determined by the worker's incentive compatibility constraint (3.4).

4 Properties of Equilibrium

In this section, we characterize the equilibrium. In the first part, we characterize the equilibrium relationship between the realization of the sunspot, the workers' relative gains from trade, and the firing strategy of firms. We show that there is an equilibrium in which firms use the sunspot to perfectly correlate the outcomes of their randomization over firing and keeping non-performing workers. There is also an equilibrium in which firms ignore the sunspot and randomize over firing and keeping non-performing workers independently of one another. The correlated equilibrium exists because firms find it optimal to correlate the randomization to economize on agency costs. The uncorrelated equilibrium exists because the only instrument that firms have to achieve correlation is an inherently meaningless sunspot, that has meaning to an individual firm only to the extent that it has meaning to others. In the second part of this section, we argue that the only robust equilibrium is the correlated equilibrium. Specifically, in a version of the model in which firms fire sequentially (and thus history can act as a correlation device), the unique equilibrium is one with perfect correlation. In the last part of this section, we describe the key properties of the equilibrium dynamics.

4.1 Stage Equilibrium

In equilibrium, the firms' firing probability $s^*(y_\ell, \hat{z})$ and the workers' relative gains from trade $\phi(\hat{z})$ must simultaneously satisfy two conditions. For any realization \hat{z} of the sunspot, the firing probability $s(\hat{z}) = s^*(y_\ell, \hat{z})$ must be part of the optimal employment contract given that the workers' relative gains from trade are $\phi(\hat{z})$ and the firing cutoff is ϕ^* -cutoff which depends on the whole distribution of the workers' relative gains from trade across realizations of the sunspot. Moreover, for any realization \hat{z} of the sunspot, the workers' relative gains from trade $\phi(\hat{z})$ must be consistent with the optimal employment contracts that are negotiated next period given the firing probability $s(\hat{z})$. Formally, in any equilibrium, the functions $\phi(\hat{z})$ and $s(\hat{z})$ must be a fixed-point of the mapping we just described. Borrowing language from game theory, we refer to such a fixed-point as the *Stage Equilibrium*, as it describes the key equilibrium outcomes between the bargaining stage of the current period and the bargaining stage of the next period.

We characterize the Stage Equilibrium under the conjectures that the optimal wage $w^*(u)$ is strictly decreasing in unemployment, and that the worker's and firm's gains

from trade $V(u)$ and $J(u)$ are strictly increasing in unemployment. These conjectures are natural. If unemployment is higher, the convexity of the vacancy cost implies that the job-finding probability of unemployed workers is lower, and so is the lifetime utility of unemployed workers. In turn, this implies that the gains from trade between workers and firms are higher and, through Nash bargaining, so are the gains from trade accruing to the workers and the firms, V and J . Since the output of a firm-worker match is independent of unemployment, a higher J requires a lower wage. In Section 5, we verify that the conjectures hold for the calibrated version of the model.

We can now characterize the effect of the firms' firing probability $s(\hat{z})$ on the worker's relative gains from trade $\phi(\hat{z})$. Given that the unemployment at the bargaining stage of the current period is u and that the firing probability at the separation stage of next period is $s(\hat{z})$ for a realization of the sunspot \hat{z} , it follows that unemployment at the bargaining stage of the next period is $\hat{u}(s(\hat{z}))$ given by

$$\hat{u}(s(\hat{z})) = u - u\lambda(\bar{\theta}(J(\hat{u}(s(\hat{z})), u)) + (1 - u)(\delta + (1 - \delta)p_\ell(1)s(\hat{z})). \quad (4.1)$$

Under the conjecture that J is a strictly increasing function, there exists a unique $\hat{u}(s(\hat{z}))$ that solves (4.1) and $\hat{u}(s(\hat{z}))$ is strictly increasing in $s(\hat{z})$. At the bargaining stage of next period, the optimal employment contract signed by workers and firms is such that

$$\phi(\hat{z}) = v'(w^*(\hat{u}(s(\hat{z}))))). \quad (4.2)$$

Under the conjecture that w^* is a strictly decreasing function, the workers' relative gains from trade $\phi(\hat{z})$ are strictly increasing in $\hat{u}(s(\hat{z}))$ and, hence, strictly increasing in $s(\hat{z})$. The solid red line in Figure 2 illustrates the effect of the firing probability $s(\hat{z})$ on the worker's relative gains from trade $\phi(\hat{z})$.

Next, we characterize the effect of the workers' relative gains from trade $\phi(\hat{z})$ on the firms' firing probability $s(\hat{z})$. From the characterization of the optimal contract in Theorem 3, it follows that $s(\hat{z})$ is such that

$$s(\hat{z}) = \begin{cases} 0, & \forall \hat{z} \in [0, 1] \text{ s.t. } \phi(\hat{z}) < \phi^*, \\ \in [0, 1], & \forall \hat{z} \in [0, 1] \text{ s.t. } \phi(\hat{z}) = \phi^*, \\ 1, & \forall \hat{z} \in [0, 1] \text{ s.t. } \phi(\hat{z}) > \phi^*, \end{cases} \quad (4.3)$$

where ϕ^* is implicitly defined by the worker's incentive compatibility constraint (3.4). The dashed green line in Figure 2 illustrates the effect of the worker's relative gains from trade $\phi(\hat{z})$ on the firing probability $s(\hat{z})$.

The two equilibrium conditions relating the firms' firing probability $s(\hat{z})$ and the workers' relative gains from trade $\phi(\hat{z})$ are intuitive. The higher is $s(\hat{z})$, the higher is unem-




Figure 2: Stage Equilibrium

ployment at the next bargaining stage, the lower is the workers' outside option when bargaining, and the lower is the wage. Since Nash bargaining equates the workers' relative gains from trade to the worker's relative marginal utility of the wage and workers are risk averse, it follows that a higher $s(\hat{z})$ implies a higher $\phi(\hat{z})$. Conversely, the higher is $\phi(\hat{z})$, the lower is the firms' cost of firing workers in state of the world \hat{z} . Hence, a higher $\phi(\hat{z})$ implies a higher $s(\hat{z})$.

For any realization \hat{z} of the sunspot, the firms' firing probability $s(\hat{z})$ must be optimal given the workers' relative gains from trade $\phi(\hat{z})$ (i.e. we must be on the dashed green line) and the workers' relative gains from trade must be consistent with the firms' firing probability (i.e. we must be on the solid red line). As it is clear from Figure 2, for any realization of \hat{z} , only three outcomes are possible: points A, B and C. The first outcome, point A, is such that the firms' firing probability $s(\hat{z})$ is 0 and the workers' relative gains from trade $\phi(\hat{z})$ are smaller than ϕ^* . The second outcome, point B, is such that the firms' firing probability $s(\hat{z})$ is interior and the workers' relative gains from trade $\phi(\hat{z})$ are equal to ϕ^* . The third outcome, point C, is such that the firms' firing probability $s(\hat{z})$ is 1 and the workers' relative gains from trade $\phi(\hat{z})$ are greater than ϕ^* .

Let Z_A denote the realizations of the sunspot for which firms fire non-performing workers with probability 0, and let π_A denote the measure of Z_A . Let Z_B denote the realizations of the sunspot for which firms fire non-performing workers with a probability $s^*(y_\ell, \hat{z}) = s_B \in (0, 1)$, and let π_B denote the measure of Z_B . Similarly, let Z_C denote the realizations of the sunspot for which firms fire non-performing workers with probability 1, and let π_C denote the measure of Z_C .

Depending on π_B , we can identify three qualitatively different types of equilibria. If

$\pi_B = 1$, we have an *Uncorrelated Equilibrium*, in which firms fire their non-performing workers with probability $s_B \in (0, 1)$ for all $\hat{z} \in [0, 1]$. That is, in an Uncorrelated Equilibrium, firms ignore the realization of the sunspot and randomize over keeping or firing their non-performing workers independently of each other. The firms' firing probability s_B is such that the workers' incentive compatibility constraint holds with equality, i.e.

$$s_B = \frac{\psi}{\beta(p_h(1) - p_h(0))(1 - \delta)V(\hat{u}(s_B))}. \quad (4.4)$$

If $\pi_B = 0$, we have a *Correlated Equilibrium*. In this equilibrium, every firm fires its non-performing workers with probability 0 whenever the realization of the sunspot is $\hat{z} \in Z_A$, and every firm fires its non-performing workers with probability 1 whenever the realization of the sunspot is $\hat{z} \in Z_C$, with $Z_A \cup Z_C = [0, 1]$. That is, in a Correlated Equilibrium, firms use the sunspot to randomize over firing or keeping their non-performing workers in a correlated fashion. The measures of Z_A and Z_C are not free, but must be such that the workers' incentive compatibility constraint holds with equality. Specifically, $\pi_A = 1 - \pi_C$ and

$$\pi_C = \frac{\psi}{\beta(p_h(1) - p_h(0))(1 - \delta)V(\hat{u}(1))}. \quad (4.5)$$

If $\pi_B \in (0, 1)$, we have a *Partially Correlated Equilibrium*. In this equilibrium, firms fire workers with probability $s_B \in (0, 1)$ for all $\hat{z} \in Z_B$. That is, when $\hat{z} \in Z_B$, firms randomize over firing or keeping their non-performing workers independently from each other. However, if $\hat{z} \notin Z_B$, firms fire their workers with probability 0 if $\hat{z} \in Z_A$ and with probability 1 if $\hat{z} \in Z_C$. That is, when $\hat{z} \notin Z_B$, firms use the sunspot to randomize over firing or keeping their non-performing workers in a correlated fashion. A Partially Correlated Equilibrium is a mixture of an Uncorrelated and a Correlated Equilibrium. Given π_B and π_C , the firms' firing probability s_B is such that the workers' incentive compatibility constraint holds with equality, i.e.

$$s_B = \frac{\psi - \beta(p_h(1) - p_h(0))(1 - \delta)\pi_C V(\hat{u}(1))}{\beta(p_h(1) - p_h(0))(1 - \delta)\pi_B V(\hat{u}(s_B))}. \quad (4.6)$$

The above results are summarized in Theorem 4.

Theorem 4. (*Stage Equilibrium*). *There exist 3 types of Stage Equilibria:*

- (i) *A unique Uncorrelated Equilibrium: $s(\hat{z}) = s_B$ for all $\hat{z} \in Z_B$, with $\pi_B = 1$ and s_B given by (4.4);*
- (ii) *A unique Correlated Equilibrium: $s(\hat{z}) = 0$ for all $\hat{z} \in Z_A$ and $s(\hat{z}) = 1$ for all $\hat{z} \in Z_C$, with $\pi_A = 1 - \pi_C$ and π_C given by (4.5);*
- (iii) *A double continuum of Partially Correlated Equilibrium: $s(\hat{z}) = 0$ for all $\hat{z} \in Z_A$,*

$s(\hat{z}) = s_B$ for all $\hat{z} \in Z_B$, and $s(\hat{z}) = 1$ for all $\hat{z} \in Z_C$, with $\pi_A = 1 - \pi_B - \pi_C$, $\pi_B \in (0, 1)$, $\pi_C \in (0, 1 - \pi_B]$ and s_B given by (4.6).

4.2 Stage Equilibrium Refinement

The Correlated Equilibrium exists because firms want to correlate the outcomes of the randomization over keeping and firing non-performing workers because doing so allows them to minimize the agency cost of moral hazard. Moreover, firms are able to correlate the outcome of the randomization because they can all observe the sunspot. The Uncorrelated Equilibrium (as well as the Partially Correlated Equilibria) exists because the sunspot is inherently meaningless and, hence, there is always an equilibrium in which firms ignore it. If firms do not need to rely on an inherently meaningless signal, the Uncorrelated Equilibrium (as well as the Partially Correlated Equilibria) should disappear. In this subsection, we show that—if firms fire sequentially and, hence, can use history as a correlation device—then the unique equilibrium of the stage game is the correlated one.

Here is a formal description of the stage game. Let $1 - u$ denote the measure of employed workers at the bargaining stage of the current period. The measure of employed workers is equally divided into a large number $N \cdot K$ firms, each employing 1 worker of “measure” $(1 - u)/NK$. Firms are clustered into a large number K of groups, each comprising a large number N of firms. Firms and workers bargain over the terms of the one-period employment contract knowing the group to which they belong. At the separation stage of next period, firm-worker pairs in different groups decide to break up or stay together sequentially. First, the firm-worker pairs in group 1 decide whether to separate or not. After observing the outcomes of group 1, the firm-worker pairs in group 2 decide whether to separate or not. The process continues until the firm-worker pairs in group K decide to separate or not.

Let T_i denote the measure of workers who separate from firms in groups 1 through i . We assume that each firm in group i takes as given the probability distribution of T_i conditional on T_{i-1} , which we denote as $P_i(T_i|T_{i-1})$. The assumption implies that each firm views itself as small compared to its group, which is reasonable for N large. We also assume that $P_i(T_i|T_{i-1})$ is strictly increasing—in the sense of first-order stochastic dominance—with respect to T_{i-1} . The assumption implies that firms in group i view themselves as small compared to the whole economy, which is reasonable for K large. To keep the analysis simple, we approximate the worker’s gains from trade, $V(u)$, and the firm’s gains from trade, $J(u)$, with linear functions. The approximation implies that the worker’s expected gains from trade relative to the firm’s are equal to the worker’s

relative gains from trade evaluated at the expectation of next period's unemployment, $E[V(\hat{u})]/E[J(\hat{u})] = \phi(E[\hat{u}])$.

We now characterize the optimal firing probability for firms in different groups. For firms in groups $i = 2, 3, \dots, K$, the firing probability $s_i(y, T_{i-1})$ depends on the realization of the worker's output y and on the measure T_{i-1} of workers separating from firms in groups 1 through $i - 1$. As in Section 3, we can show that the optimal firing probability is such that: (i) the worker's incentive compatibility constraint holds with equality; (ii) when the realization of output is high, the worker is fired with probability 0, i.e. $s_i(y_h, T_{i-1}) = 0$ for all T_{i-1} ; (iii) when the realization of output is low, the worker is fired with probability 0 if the relative gains from trade are below a cutoff ϕ_i^* , and with probability 1 if they are above the cutoff, i.e. $s_i(y_\ell, T_{i-1}) = 0$ if $\phi(E[\hat{u}|T_{i-1}]) < \phi_i^*$, and $s_i(y_\ell, T_{i-1}) = 1$ if $\phi(E[\hat{u}|T_{i-1}]) > \phi_i^*$. Under the same conjectures about V , J and w made in Section 4.1, $\phi(E[\hat{u}|T_{i-1}])$ is strictly increasing in $E[\hat{u}|T_{i-1}]$. For firms in group 1, the firing probability can only depend on the realization of the worker's output y . Hence, the optimal contract is such that $s_1(y_h) = 0$ and $s_1(y_\ell) = s_1$, where s_1 is such that the worker's incentive compatibility constraint holds with equality.

Firms in group K understand that $E[\hat{u}|T_{K-1}]$ is strictly increasing in T_{K-1} , as $P_K(T_K|T_{K-1})$ is strictly increasing in T_{K-1} and unemployment is strictly increasing in T_K . Therefore, there exists a cutoff T_{K-1}^* such that firms in group K fire their non-performing workers with probability 0 if $T_{K-1} < T_{K-1}^*$ and with probability 1 if $T_{K-1} > T_{K-1}^*$. Firms in group $K - 1$ understand that $E[\hat{u}|T_{K-2}]$ is strictly increasing in T_{K-2} , as $P_K(T_{K-1}|T_{K-2})$ is strictly increasing in T_{K-2} and the firing probability of firms in group K is increasing in T_{K-1} . Hence, there exists a cutoff T_{K-2}^* such that firms in group $K - 1$ fire their non-performing workers with probability 0 if $T_{K-2} < T_{K-2}^*$ and with probability 1 if $T_{K-2} > T_{K-2}^*$. The same reasoning implies that there exists a cutoff T_{i-1}^* for firms in all groups $i = 2, 3, \dots, K$.

Next, we compute the probability distribution of the measure t_i of workers who separate from firms in group i conditional on T_{i-1} . A worker separates from a firm in group 1 with probability $\tau_1 = \delta + (1 - \delta)p_\ell(1)s_1$. Since N is large, we can apply the Central Limit Theorem to approximate t_1 as a Normal with mean $\tau_1(1 - u)/K$ and variance $\tau_1(1 - \tau_1)[(i - u)/K]^2/N$. Conditional on T_{i-1} , workers separate from firms in group $i = 2, 3, \dots, K$ with probability $\tau_i = \delta$ if $T_{i-1} < T_{i-1}^*$ and with probability $\tau_i = \delta + (1 - \delta)p_\ell(1)$ if $T_{i-1} > T_{i-1}^*$. Again, we can approximate t_i as a Normal with mean $\tau_i(1 - u)/K$ and variance $\tau_i(1 - \tau_i)[(i - u)/K]^2/N$. We find it convenient to define t_ℓ as $\delta(1 - u)/K$, and t_h as $[\delta + (1 - \delta)p_\ell(1)](1 - u)/K$.

Lastly, we compare the incentive compatibility constraint for workers employed by

firms in groups 2 and 3 to characterize the equilibrium for $N \rightarrow \infty$. The incentive compatibility constraint for workers in firms of group 2 is

$$\psi = \beta(1 - \delta)(p_h(1) - p_\ell(1)) \Pr(T_1 > T_1^*)V(E[\hat{u}|T_1 > T_1^*]). \quad (4.7)$$

The incentive compatibility constraint for workers in firms of group 3 is

$$\psi = \beta(1 - \delta)(p_h(1) - p_\ell(1)) \Pr(T_2 > T_2^*)V(E[\hat{u}|T_2 > T_2^*]). \quad (4.8)$$

For $N \rightarrow \infty$, T_2 is approximately equal to $\tau_1(1 - u)/K + t_\ell$ if $T_1 < T_1^*$ and approximately equal to $\tau_1(1 - u)/K + t_h$ if $T_1 > T_1^*$.

Suppose that $\Pr(T_2 > T_2^*) > \Pr(T_1 > T_1^*)$. By the reservation property of the firing probability of firms in group 3, $\Pr(T_2 > T_2^*) > \Pr(T_1 > T_1^*)$ implies that $s_3(y_\ell, T_2) = 1$ for all realizations of T_1 such that $T_1 > T_1^*$, realizations that induce $T_2 = \tau_1(1 - u)/K + t_h$, and also for some of the realizations of T_1 such that $T_1 < T_1^*$, realizations that induce $T_2 = \tau_1(1 - u)/K + t_\ell$. In light of these observations, we can rewrite (4.8) as

$$\begin{aligned} \psi &= \beta(1 - \delta)(p_h(1) - p_h(0)) \cdot \\ &\cdot \{\Pr(T_1 > T_1^*)V(E[\hat{u}|T_2 = \tau_1(1 - u)/K + t_h]) \\ &+ \Pr(T_1 < T_1^* \vee T_2 > T_2^*)V(E[\hat{u}|T_2 = \tau_1(1 - u)/K + t_\ell])\} \end{aligned} \quad (4.9)$$

However, the first term on the right-hand side of (4.9) is precisely the benefit of exerting effort to a worker employed by a firm in group 2 and, by (4.7), it must equal to the cost ψ . Therefore, (4.7) and (4.9) can simultaneously hold only if $\Pr(T_2 > T_2^*) \leq \Pr(T_1 > T_1^*)$.

Now, suppose that $\Pr(T_2 > T_2^*) < \Pr(T_1 > T_1^*)$. By the reservation property of the firing probability of firms in group 3, $\Pr(T_2 > T_2^*) < \Pr(T_1 > T_1^*)$ implies that $s_3(y_\ell, T_2) = 0$ for all realizations of T_1 such that $T_1 < T_1^*$, realizations that induce $T_2 = \tau_1(1 - u)/K + t_\ell$, and also for some realizations of T_1 such that $T_1 > T_1^*$, realizations that induce $T_2 = \tau_1(1 - u)/K + t_h$. In light of these observations, we can rewrite (4.7) as

$$\begin{aligned} \psi &= \beta(1 - \delta)(p_h(1) - p_h(0)) \cdot \\ &\cdot \{\Pr(T_1 > T_1^* \vee T_2 > T_2^*)V(E[\hat{u}|T_2 = \tau_1(1 - u)/K + t_h]) \\ &+ \Pr(T_1 > T_1^* \vee T_2 < T_2^*)V(E[\hat{u}|T_2 = \tau_1(1 - u)/K + t_\ell])\}. \end{aligned} \quad (4.10)$$

However, the first term on the right-hand side of (4.10) is precisely the benefit of exerting effort to a worker employed by a firm in group 3 and, by (4.8), it must equal to the cost ψ . Therefore, (4.7) and (4.8) can simultaneously hold only if $\Pr(T_2 > T_2^*) \geq \Pr(T_1 > T_1^*)$.

Combining the above observations, it follows that any equilibrium is such that $\Pr(T_2 > T_2^*) = \Pr(T_1 > T_1^*)$. And, by repeating the same argument, it follows that any equilibrium is such that $\Pr(T_i > T_i^*) = \Pr(T_1 > T_1^*)$ for $i = 2, 3, \dots, K$. Hence, in any equilibrium,

if $T_1 > T_1^*$, then all firms in the following groups fire their non-performing workers with probability 1. If $T_1 < T_1^*$, then all firms in the following groups fire with probability 0. It is then straightforward to show that such an equilibrium does exist.

We have thus established the following.

Theorem 5. (*Equilibrium Refinement*) *For $N \rightarrow \infty$ and $K \rightarrow \infty$, the unique equilibrium of the stage game with K groups of N firms firing sequentially is the Correlated Equilibrium.*

Theorem 5 is relevant for three reasons. First, as we have already discussed, the theorem shows that—if firms can use history and not only an inherently meaningless signal to coordinate their behavior—the unique equilibrium of the stage game is the correlated equilibrium. Second, the theorem shows that our theory does not rely on the existence of a sunspot (in contrast with business cycle theories based on equilibrium indeterminacy). Third, the theorem exemplifies how firms may coordinate their behavior in practice. While a sunspot is a convenient theoretical construct to allow coordination, it does not have a clear empirical counterpart. What is the sunspot that firms use in the real world? The theorem suggests that, in practice, firms may achieve coordination by looking at the firing decisions of focal firms (e.g., industry leaders, large firms, etc. . .).

4.3 Agency Business Cycles

Having established that the unique robust equilibrium of the stage game is the Correlated Equilibrium, we can now proceed to characterize some of the key properties of the dynamic equilibrium. A Recursive Equilibrium features aggregate uncertainty.⁸ Given a current unemployment rate of u , firms fire their non-performing workers for all realizations of the sunspot $\hat{z} \in Z_C$, an event which occurs with probability π_C given by (4.5), and firms keep their non-performing workers for realizations of the sunspot $\hat{z} \in Z_A$, an event which occurs with probability $\pi_A = 1 - \pi_C$. In the first case, the fraction of employed workers who become unemployed is $\delta + (1 - \delta)p_\ell(1)$ and the unemployment rate goes to $h(u, Z_C)$. In the second case, the fraction of employed workers who become unemployed is δ and the unemployment rate goes to $h(u, Z_A)$. Since $h(u, Z_C) > h(u, Z_A)$, the equilibrium features aggregate uncertainty.

⁸In an earlier version of the paper (Goloso and Menzio, 2015), we established the existence of a Recursive Equilibrium that satisfies the conjectures that we made in Section 4.1. The sufficient conditions for existence were quite stringent and not particularly informative and, for these reasons, here we simply check the existence of a Recursive Equilibrium and verify the conjectures for the calibrated version of the model.

Aggregate uncertainty causes aggregate unemployment fluctuations, which we dub *Agency Business Cycles* (or ABCs). These unemployment cycles are illustrated in Figure 3. Imagine an economy in which the unemployment rate is u_0 . As long as the realization of the sunspot falls in Z_A , firms keep their non-performing workers, the rate at which employed workers lose their job is δ , and the unemployment rate falls towards u^* . When the realization of the sunspot falls in Z_C , firms fire their non-performing workers, the rate at which employed workers lose their job is $\delta + (1 - \delta)p_\ell(1)$, and the unemployment rate goes back up. The unemployment rate starts falling again towards u^* when the realization of the sunspot returns in Z_A .

Agency Business Cycles are endogenous. ABCs are not caused by changes in fundamentals, which remain constant over the cycle. ABCs are not caused by changes in expectations about fundamentals in the future, as these expectations remain constant over the cycle. ABCs are not caused by changes in the selection of equilibrium. In fact, the same stage equilibrium is played throughout the cycle. Moreover, while the unique stage equilibrium features randomization over two possible outcomes (e.g., firing and keeping non-performing workers), these two outcomes are not equilibria on their own but only if properly mixed. Indeed, keeping non-performing workers with probability 0 is not an equilibrium as this violates the workers' incentive compatibility constraint. Firing non-performing workers with probability 1 is not an equilibrium as doing so would be suboptimal. ABCs are endogenous because the unique equilibrium of the stage game features correlated randomization.

Agency Business Cycles are stochastic. ABCs are driven by correlated firing episodes, which occur randomly. Since ABCs are endogenous, the probability of a correlated firing episode is not determined outside of the model by some free parameter, but it is determined in equilibrium. Specifically, the probability of a firing episode is given by

$$\pi_C = \frac{\psi}{\beta(1 - \delta)(p_h(1) - p_h(0))V(h(u, Z_C))}. \quad (4.11)$$

The probability π_C of a firing episode (and, hence, of a recession) depends on the worker's cost of exerting effort, ψ , the difference between the probability of a positive realization of output when the worker does and does not exert effort, $p_h(1) - p_h(0)$, and the worker's cost of losing a job, $V(h(u, Z_C))$. Since the worker's cost of losing a job is increasing in unemployment, it follows that the probability π_C of a firing episode is higher when the unemployment rate is lower. Hence, going back to Figure 3, as the unemployment rate falls from u_0 to u^* , the probability of a firing episode gets larger and larger. When the firing episode eventually takes place, the unemployment rate increases and the probability of another firing episode falls.




Figure 3: Agency Business Cycles

The endogeneity of the probability distribution of aggregate “shocks” is a genuinely distinctive feature of our business cycle theory and, more generally, it would be a distinctive feature of any theory of endogenous and stochastic fluctuations. Existing theories of business cycles take the probability distribution of shocks as exogenous and, hence, have nothing to say about the magnitude, persistence, and determinants of the distribution of shocks. In contrast, our theory has something to say about the relationship between the structure of aggregate shocks and the fundamentals of the economy. For instance, it says that negative shocks become more likely as the unemployment rate falls. Similarly, it says that negative shocks are less frequent in economies or sectors where agency problems are less severe because either the cost of unobserved effort ψ is lower or the ability to detect low effort, as captured by $p_h(1) - p_h(0)$, is higher.

Lastly, we wish to point out that the nature of recessions in ABCs is very different from the nature of recessions in business cycle theories driven by aggregate productivity shocks (e.g., the Real Business Cycles (RBCs) of Kydland and Prescott 1983 and Mortensen and Pissarides 1994). In RBCs, recessions are times when productivity is unusually low and so are the gains from trade between workers and firms in the labor market. As the gains from trade are small, workers and firms search less intensely and the unemployment rate is higher. In ABCs, recessions are states of the world where the workers’ value of unemployment is unusually low and, hence, the gains from trade between workers and firms are large. For this reason, recessions are states of the world in which firms find it optimal to fire their non-performing workers, in which unemployment is high and, because

of convex vacancy costs, in which the workers' value of unemployment is low. Thus, the correlation between unemployment and gains from trade is negative in RBCs and positive in ABCs. To paint a picture, in RBCs, recessions are times when it is raining on the marketplace. In ABCs, recessions are times when the TV set at home is broken.

The above observations are summarized in the following theorem.

Theorem 6. (*Recursive Equilibrium*) *A Recursive Equilibrium features:*

- (i) *Aggregate uncertainty: For any $u \in [0, 1]$, next period's unemployment is $h(u, Z_C)$ with probability π_C and $h(u, Z_A)$ with probability $1 - \pi_C$, with $h(u, Z_C) > h(u, Z_A)$ as long as $u < 1$.*
- (ii) *Endogenous probability of a recession: The probability of a correlated firing episode π_C is given by (4.11), and it is increasing in ψ and decreasing in $p_h(1) - p_h(0)$ and in u .*
- (iii) *Countercyclical gains from trade: $V(u)$ and $J(u)$ are increasing in u .*

5 Quantifying the Theory

In this section, we calibrate our theory of endogenous and stochastic fluctuations in order to assess quantitatively the features of Agency Business Cycles and compare them with the data. We find that there are parameter values for which ABCs display the same volatility and the same pattern of comovement of unemployment, UE and EU rates as in the data. Moreover, ABCs display fluctuations in unemployment, UE and EU rates that are, as in the data, large relative to and uncorrelated with fluctuations in labor productivity. We find that the main counterfactual prediction of ABCs is that the correlation between unemployment and vacancies is positive rather than negative. Lastly, we show that, in the data, the gains from trade in the labor market are countercyclical, as predicted by ABCs, and not procyclical, as predicted by RBCs.

5.1 Calibration

We calibrate the model to the US labor market between 1951 and 2014. We want to understand whether our theory of endogenous and stochastic cycles can possibly match the cyclical features of the US labor market. For this reason, we use as targets in the calibration not only average values of key labor market variables, but also their cyclical volatility.

Let us start by reviewing the parameters of the model. Preferences are described by the discount factor β and by the worker's periodical utility. When the worker is unemployed,

his periodical utility is given by $v(b) + \zeta$, where $v(\cdot)$ is the utility of consumption, b is unemployment income and ζ is the value of leisure. When the worker is employed, his periodical utility is given by $v(w_t) - \psi e_t$. We specialize the worker's utility function $v(c)$ to be $\log(c)$. The production process is described by the possible realizations of the worker's output, y_h and y_ℓ , by the probability distribution over realizations of the worker's output conditional on effort, $p_h(1)$ and $p_h(0)$, and by the probability of exogenous job destruction, δ . The search and matching process is described by the vacancy cost function, $k(v)$, and by the matching function, $M(u, v)$. We specialize the vacancy cost function to be of the form $k(v) = k_0 v^\rho$, where $k_0 > 0$ is a scale parameter and $\rho > 0$ is the elasticity of the vacancy cost with respect to vacancies. We specialize the matching function to be of the form $M(u, v) = uv(u^\gamma + v^\gamma)^{-1/\gamma}$, which is a constant returns to scale function with an elasticity of substitution γ between u and v .

We calibrate the model to the US labor market between 1951 and 2014. The calibration strategy for some of the parameters of the model is standard. We choose the model period to be 1 month. We set the discount factor, β , so that the annual real interest rate is 5 percent. We set the scale coefficient in the vacancy cost function, k_0 , and the probability of exogenous job destruction, δ , so that the average unemployment rate and the average EU rate are the same in the model as in the data.⁹ We normalize the expected labor productivity of a worker to 1, i.e. $\bar{y} \equiv p_h(1)y_h + p_\ell(1)y_\ell = 1$. We set the unemployment income b to be 40% of expected labor productivity, so as to reflect the typical replacement rate of US unemployment benefits.¹⁰ We set the value of leisure ζ so that the flow value of unemployment expressed in units of output, $b + \zeta/v'(b)$, is 70% of expected labor productivity, which Hall and Milgrom (2008) argue is a reasonable estimate for the US economy. We choose the elasticity of substitution γ between unemployment and vacancies in the matching function to 1.24, which is the value estimated by Menzio and Shi (2011).

We calibrate the remaining parameters, which are novel to our theory, to match some key cyclical properties of the US labor market. First, note that the probability that the realization of the worker's output is low, $p_\ell(1) = 1 - p_h(1)$, determines the rate at which employed workers become unemployed during a firing episode and, hence, the cyclical volatility of the EU rate. Second, note that the elasticity ρ of the vacancy cost with respect to vacancies determines the extent to which the UE rate falls with unemployment and, hence, in conjunction with $p_\ell(1)$, determines the cyclical volatility of unemployment.

⁹We measure the UE and the EU rates using the civilian unemployment and short-term unemployment rates from the CPS, following the methodology in Shimer (2005). We measure labor productivity as output per-worker in the non-farm sector.

¹⁰To be more precise, the replacement ratio is a ratio between unemployment benefits and wages and not between unemployment benefits and output. However, because the UE rate is so high, the average wage is close to the average output.

TABLE 1: CALIBRATION

	Value	Target	Data	Model
y_h	1.04	normalization	1	1
δ	.025	average EU rate	2.57%	2.57%
k_0	8.80	mean u rate	5.87%	5.87%
$p_h(1)$.963	std EU rate	9.74	9.58
ρ	1.3	std u rate	19.65	20.51
ψ	.001	recession probability	1.4%	1.4%
γ	1.24	elast. subst. in $M(u, v)$	1.24	1.24
b	.40	u benefit replacement	.40	.40
ζ	.75	u relative payoff	.70	.70

Notes: Calibrated values of parameters, empirical targets and their model-generated counterpart. In the calibration $y_\ell = 0$ and $p_h(1) - p_h(0) = 0.5$.

Third, note that the worker's disutility of effort ψ determines the probability of a firing episode. Based on these observations, we choose $p_h(1)$ so that the cyclical volatility of the EU rate is the same as in the data and ρ so that the cyclical volatility of unemployment is the same as in the data. We choose ψ so that a firing episode (and, hence, a recession) occurs once every 6 years.

Lastly, note that the parameters y_ℓ and $p_h(0)$ cannot be uniquely identified. Given that expected output is 1, the realizations y_h and y_ℓ do not affect the equilibrium, as long as it is optimal for firms to require effort from their workers. Similarly, the probability $p_h(0)$ only affects the equilibrium through the ratio $\psi/(p_h(1) - p_h(0))$. Therefore, we choose some arbitrary values for y_ℓ and $p_h(0)$ such that firms find it optimal to require effort.

Table 1 contains the calibrated parameter values, the empirical moments used to calibrate the parameters and their model-generated counterpart. Three of the calibrated parameter values are worth discussing. The calibrated value of $p_h(1)$ is 96.4%, which is the value required by the model to match the empirical volatility of the EU rate. The value of $p_h(1)$ means that, in each month, a worker who exerts effort has a 3.6% probability of generating low output. The calibrated value of ρ is 1.3, which is the value required by the model to match the empirical volatility of the unemployment rate. The value of ρ means that, if the measure of vacancies increases by 10%, the marginal cost of a vacancy increases by 13%. The calibrated value of ψ is 0.1%, which is the value required by the model to generate a recession once every 6 years. The value of ψ means that the consumption-equivalent disutility of effort is approximately 0.1% of the consumption of an employed worker.

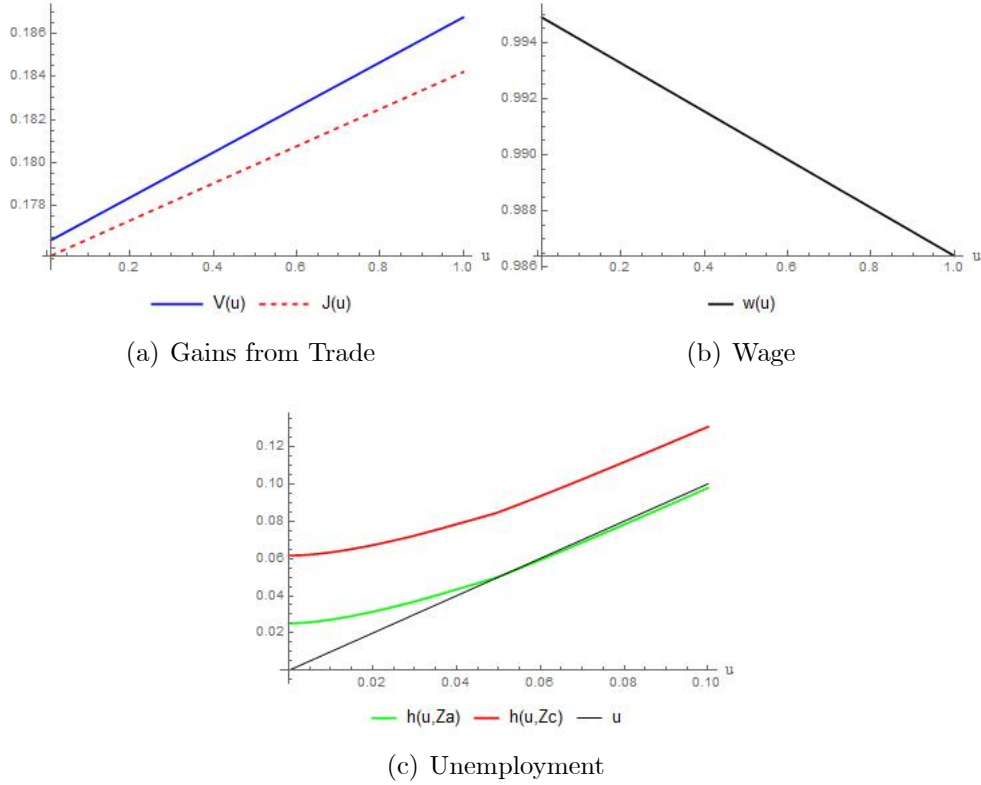


Figure 4: Value and Policy Functions

5.2 Quantitative Properties of ABCs

We now examine the properties of equilibrium and the cyclical features of ABCs. The left panel of Figure 4 is the plot of the worker's gains from trade, $V(u)$, and the firm's gains from trade, $J(u)$. The right panel is the plot of the worker's wage, $w^*(u)$. The bottom panel is the plot of the law of motion for unemployment when the realization of the sunspot is $\hat{z} \in Z_A$ and when the realization of the sunspot is $\hat{z} \in Z_C$. The worker's and firm's gains from trade are strictly increasing in the unemployment rate u and the worker's wage is strictly decreasing in u . Therefore, the conjectures made in Section 3 are verified, and Theorems 4, 5 and 6 apply to the calibrated model. Moreover, for the calibrated model, the Recursive Equilibrium exists, as we find a solution to the system of equations laid out in Section 2.

Table 2 compares the cyclical volatility of the unemployment, UE and EU rates, tightness, vacancies and labor productivity in the model and in the data. Given the calibrated parameter values, ABCs display approximately the same volatility of unemployment, UE and EU rates as in the data. Moreover, ABCs display the same correlation between fluctuations in these variables as in the data. Specifically, the unemployment rate is positively

TABLE 2: AGENCY BUSINESS CYCLES


		u	UE	EU	v	θ	LP
Model	std	20.51	13.81	9.58	6.36	10.35	0
	cor	1	-.98	.18	.99	-.98	0
Data: 1951-2014	std	19.65	12.98	9.74	19.52	38.59	1.98
	cor	1	-.94	.81	-.91	-.97	-.37
Data: 1984-2014	std	17.35	13.88	6.92	17.01	33.96	1.38
	cor	1	-.96	.70	-.88	-.95	.09

Notes: For each variable, we construct quarterly time-series by taking 3-month averages. We then compute the cyclical component of each variable as the percentage deviation of its quarterly value from a Hodrick-Prescott trend constructed using a smoothing parameter of 100,000.

correlated with the EU rate and negatively correlated with the UE rate.

Figure 5 takes a closer look at the structure of leads and lags of unemployment, UE and EU rates. In ABCs, a recession starts with an episode of correlated firing which increases the EU rate. The increase in the EU rate leads to an increase in unemployment and, because of the convexity of the vacancy cost, to a decline in the UE rate. The decline in the UE rate slows down the return of unemployment to its pre-recession level. This chain of events can be seen in Figure 5. In the model, the EU rate leads unemployment by a quarter (i.e. the highest correlation is between unemployment at t and EU at $t - 1$) and it returns to trend after one quarter. The UE rate is coincidental to the unemployment rate (i.e. the highest correlation is between unemployment at t and UE at t) and it returns to trend more slowly than the EU rate. The same qualitative pattern of leads and lags can be seen in the data.

Overall, our analysis reveals that—for some parameter values—ABCs generate fluctuations in unemployment, UE and EU rates of the same magnitude and with the same structure of comovement as in the data. We wish to make two additional observations. First, in ABCs, UE and EU fluctuations contribute in equal measure to the volatility of unemployment but all labor market fluctuations are caused by endogenous shocks to the EU rate. This observation implies that decomposing unemployment fluctuations into the contribution of UE and EU rates (as in Shimer 2005 and subsequent papers) is informative about the morphology of labor market fluctuations, but not about their root cause. Second, in ABCs, unemployment, UE and EU rates are as volatile as in the data but the standard deviation of labor productivity is zero, and so is the correlation between unemployment, UE and EU rates and labor productivity. This observation is important because, as it is well-known, the empirical volatility of labor productivity is an order of magnitude smaller than the volatility of unemployment. Moreover, as shown in Table 2, the empirical correlation between unemployment and labor productivity is very small



Notes: Correlation between u , UE and EU rates in quarter $x+t$ and unemployment in quarter x . Data are HP-filtered with a smoothing coefficient of 10^5 .

Figure 5: Leads and Lags

and, since 1984, basically zero. Our theory identifies a mechanism that can explain why unemployment is so volatile in the face of small, uncorrelated fluctuations in technology.

ABCs, however, do not look like US business cycles along one key dimension. Table 2 shows that ABCs feature a positive correlation between vacancies and unemployment, while the correlation is negative in the data. In the model, a correlated firing episode leads to an increase in unemployment and to an increase in the measure of vacancies (a less than proportional increase because convex vacancy costs, but an increase nonetheless). In the data, recessions feature high unemployment and low vacancies. From this observation, we conclude that our simple theory of endogenous, stochastic fluctuations cannot provide a complete explanation of the fluctuations of the US labor market. Other forces must be at work to make vacancies move together with the business cycle. Demand externalities, for example, would lower the firm's benefit from creating vacancies in recessions and could flip the sign of the correlation between unemployment and vacancies. Wage rigidities, procyclical productivity or countercyclical discount rates could do the same. In any case, as shown by Menzio and Shi (2011), the theory would have to introduce on-the-job search to make credible predictions about the cyclical behavior of vacancies.

5.3 Countercyclical Gains from Trade

In Section 4, we pointed out that one of the differences between recessions in ABCs and RBCs is the cyclicity of the gains from trade in the labor market. In RBCs, recessions are times when the gains from trade in the labor market are unusually low because the productivity of labor is unusually low. In ABCs, recessions are times when the gains from

trade in the labor market are unusually large because the value of being unemployed is unusually low. Given this striking difference between RBCs and ABCs, it is natural to ask whether the gains from trade in the labor market are pro or countercyclical.

In order to answer the question, we construct a rudimentary time-series of the workers' gains from trade in the labor market.¹¹ We define the value of employment to a worker, $W_{1,t}$, and the value of unemployment to a worker, $W_{0,t}$, as


$$\begin{aligned} W_{1,t} &= w_t + \beta [h_{t+1}^{EU} W_{0,t+1} + (1 - h_{t+1}^{EU}) W_{1,t+1}], \\ W_{0,t} &= b_t + \beta [h_{t+1}^{UE} W_{1,t+1} + (1 - h_{t+1}^{UE}) W_{0,t+1}], \end{aligned} \tag{5.1}$$

where w_t is the real wage in month t , b_t the unemployment benefit/value of leisure in month t , h_{t+1}^{EU} is the EU rate in month $t + 1$, h_{t+1}^{UE} the UE rate in month $t + 1$, and $W_{1,t+1}$ and $W_{0,t+1}$ are respectively the value of employment and unemployment in month $t + 1$. We define the gains from trade to a worker, V_t , as the difference between $W_{1,t}$ and $W_{0,t}$.

We measure w_t using the time-series for the hourly wage that have been constructed by Haefke, Sonntag and van Rens (2013). We consider two alternative time-series: the average hourly wage in the cross-section of all employed workers, and the average hourly wage in the cross-section of newly hired workers after controlling for the composition of new hires. The first time-series may be more appropriate if we want to interpret V_t as the cost of losing a job to a worker, the second-time series may be more appropriate if we want to interpret V_t as the benefit of finding a job to a worker. As in Shimer (2005), we measure h_t^{UE} and h_t^{EU} using, respectively, the values for the EU rate and UE rates implied by the time-series for unemployment and short-term unemployment. In order to make the time-series for w_t , h_t^{UE} and h_t^{EU} stationary, we construct their Hodrick-Prescott trend using a smoothing parameter of 10^5 . We then take the difference between the value of each variable and its trend and add this difference to the time-series average for that variable. As b_t is not directly observable, we follow Hall and Milgrom (2008) and set it equal to 70% of the average of the detrended wage.

We now have all the inputs to construct the time-series for the value of employment to a worker, $W_{1,t}$, and the value of unemployment to a worker, $W_{0,t}$, over the period going from January 1979 to December 2014. We compute the values for W_1 and W_0 in December 2014 by assuming that, from January 2015 onwards, w , h^{UE} , and h^{EU} are equal to their historical averages. Given the values for W_1 and W_0 in December 2014, we compute the

¹¹We only attempt to measure the gains from trade accruing to the workers. Measuring the gains from trade accruing to the firms would require constructing a time-series for profit per employee, a task that is beyond the scope of this paper. For what it is worth, if we measure profit per employee as labor productivity net of wages, we find that the gains from trade accruing to the firms are countercyclical as well.



Notes: Time series of unemployment (thick black solid line), worker’s gains from trade using average wages (thick dashed blue line) and using the wage of new hires (thin dashed red line).

Figure 6: Gains from Trade to a Worker

values for W_1 and W_0 from November 2014 back to January 1979 by using equation (5.1) and the time-series for w , b , h^{UE} , and h^{EU} . Notice that the values of W_1 and W_0 thus computed differ from their theoretical counterpart because they are constructed using the realizations rather than the expectations of future w , h^{UE} , and h^{EU} .

Figure 6 presents the result of our calculations. The figure displays the time-series for the gains from trade to a worker, V_t , computed using the average wage of all employed workers (thick dashed line) and the average wage of newly hired workers (thin dashed line). The figure also displays the time-series for the detrended unemployment rate (solid line). The figure clearly shows that V_t is countercyclical, in the sense that V_t moves together with the unemployment rate. This is true whether we measure V_t using the average wage of all employed workers—in which case the correlation between V_t and u_t is 80%—or whether we measure V_t using the average wage of newly hired workers—in which case the correlation between V_t and u_t is 71%. Mechanically, V_t is countercyclical because, in recessions, the decline in the value of unemployment caused by the persistent decline in the UE rate is larger than the decline in the value of employment caused by the small decline in wages and by the transitory increase in the EU rate.¹²

The finding that V_t is countercyclical means that recessions are times when unemployed workers find it especially valuable to find a job, and when employed workers find it

¹²Our calculation of V_t does not allow for the possibility that the EU rate for newly hired workers might systematically differ from the EU rate for all employed workers. If the EU rate for newly hired workers was more countercyclical than the EU rate for all employed workers, our calculations would overstate the countercyclicity of V_t .

especially costly to lose a job. The finding is in stark contrast with the view of recessions as “days of rain in the marketplace” advanced by Kydland and Prescott (1982) or by Mortensen and Pissarides (1994). In contrast, the finding is supportive of the view of recessions as “days of no TV at home” advanced by our theory. Notice that our finding that V_t is countercyclical should not be entirely surprising. Using a different, more sophisticated approach and richer data, Davis and von Wachter (2011) show that the lifetime earning cost of losing a job is much higher in recessions than in expansions. Even if one is skeptical about our theory of recessions, the finding that V_t is countercyclical represents a challenge for many existing theories of business cycles.

6 Conclusions

We conclude by discussing the robustness of our theory of endogenous and stochastic fluctuations to some stark modelling assumptions.

First, we assumed that employment contracts are incomplete—in the sense that the wage cannot depend on the current realization of output—and short-term—in the sense that future wages cannot depend on the current realization of output because they are renegotiated in every period. These assumptions guarantee that employment contracts feature a firing lottery. However, employment contracts may also feature a firing lottery when they are complete and long-term. Consider an employment contract that can specify current and future wages as a function of the entire history of the worker’s output. Wages, however, are constrained from below by a minimum wage. Suppose, for now, that the contract is such that the worker keeps his job with probability 1. When the value of the contract to the worker is close to the discounted value of the minimum wage, the worker has no incentive to exert any effort. The value to the firm of keeping a worker who exerts no effort and is paid the minimum wage may be negative. Hence, the value of the contract to the firm as a function of the value of the contract to the worker (i.e. the Pareto frontier) is hump-shaped. If the value of breaking the match lies above the Pareto frontier, the optimal contract should involve—for low values of the contract to the worker—a firing lottery. Clementi and Hopenhayn (2006) show, in the context of a contract between a lender and a borrower, that the lottery is played after a sufficiently long sequence of low realizations of output. We were able to show the same in a simple two-period version of the model.

Second, we assumed that workers are hand-to-mouth, in the sense that they consume their income in every period. The assumption guarantees that the workers’ relative gains from trade are decreasing in the equilibrium wage. Note that, because of Nash bargaining,

the workers' relative gains from trade are decreasing in the equilibrium wage as long as the workers' value function is concave with respect to cash-on-hand (i.e., wealth plus wage). Even when workers are allowed to borrow and save, the value function is typically concave in cash-on-hand.

Third, we assumed that the vacancy costs are convex. The assumption guarantees a negative relationship between the value of unemployment to a worker and the unemployment rate. This negative relationship implies that the wage is decreasing and, in turn, the workers' relative gains from trade are increasing in the unemployment rate. This is why firms want to correlate the outcomes of their firing lotteries. There are several alternative assumptions that make the value of unemployment decreasing in the unemployment rate. In an earlier version of the paper, we assumed that the matching function has decreasing returns to scale. Inspired by Chodorow-Reich and Karabarbounis (2016), we also considered a version of the model in which the unemployment income b is decreasing with the unemployment rate.

Finally, we assumed that the economy is not subject to aggregate shocks to fundamentals. We made this assumption in order to develop—as cleanly as possible—our theory of endogenous and stochastic aggregate fluctuations. In reality, though, there may be aggregate shocks, as the growth rate of technological progress, the generosity of unemployment benefits, and the tax code may vary unexpectedly over time. In a version of our model with aggregate shocks, firms may use the realization of these shocks to correlate the outcome of their firing lotteries. Consequently, correlated firing would act as a mechanism that amplifies exogenous shocks, rather than replacing them as the root cause of aggregate fluctuations. Such a version of our model would not only be more realistic than the one presented in this paper, but also, as explained in Section 5, would probably provide a better fit of the cyclical behavior of the US labor market.

References

- [1] Agarwal, Ruchir, and Julian Kolev. 2013 “Strategic Corporate Layoffs.” Manuscript, IMF.
- [2] Angeletos, Marios, and Jennifer La’O. 2013. “Sentiments.” *Econometrica*, 81: 739-780.
- [3] Beaudry, Paul, Dana Galizia, and Franck Portier. 2016 “Reviving the Limit Cycle View of Macroeconomic Fluctuations,” NBER Working Paper 21241.
- [4] Beaudry, Paul, and Franck Portier. 2004. “Exploring Pigou’s Theory of Cycles.” *Journal of Monetary Economics*, 51: 1183-1216.
- [5] Benhabib, Jess, and Richard Day. 1982. “A Characterization of Erratic Trajectories in the Overlapping Generations Model.” *Journal of Economic Dynamics and Control*, 4: 37–55.
- [6] Benhabib, Jess, and Roger Farmer. 1994. “Indeterminacy and Increasing Returns.” *Journal of Economic Theory*, 63: 19–41.
- [7] Benhabib, Jess, and Kazuo Nishimura. 1979. “The Hopf Bifurcation and Existence and Stability of Closed Orbits in Multisector Models of Optimal Economic Growth.” *Journal of Economic Theory*, 21: 421–444.
- [8] Benhabib, Jess, and Aldo Rustichini. 1990. “Equilibrium Cycling with Small Discounting.” *Journal of Economic Theory*, 52: 423–432.
- [9] Benhabib, Jess, Pengfei Wang and Yi Wen. 2015. “Sentiments and Aggregate Demand Fluctuations.” *Econometrica* 83: 549–585
- [10] Bloom, Nicholas. 2009. “The Impact of Uncertainty Shocks.” *Econometrica*, 77: 623-685.
- [11] Boldrin, Michele, and Luigi Montrucchio. 1986. “On the Indeterminacy of Capital Accumulation Paths.” *Journal of Economic Theory*, 40: 26-39.
- [12] Boldrin, Michele, and Michael Woodford. 1990. “Equilibrium Models Displaying Endogenous Fluctuations and Chaos.” *Journal of Monetary Economics*, 25: 189-222.
- [13] Cappelli, P., and K. Chauvin. 1991. “An Interplant Test of the Efficiency Wage Hypothesis.” *Quarterly Journal of Economics*, 106: 769-787

- [14] Chodorow-Reich, Gabriel, and Loukas Karabarbounis. 2016. “The Cyclicity of the Opportunity Cost of Employment.” Forthcoming, *Journal of Political Economy*.
- [15] Clementi, Gianluca, and Hugo Hopenhayn. 2006. “A Theory of Financing Constraints and Firm Dynamics.” *Quarterly Journal of Economics*, 121: 229-65.
- [16] Cooper, Russell, and Andrew John. 1988. “Coordinating Coordination Failures in Keynesian Models.” *Quarterly Journal of Economics*, 103: 441–63.
- [17] Davis, Steve, and Till von Wachter. 2011. “Recessions and the Costs of Job Loss.” *Brookings Papers on Economic Activity*.
- [18] Diamond, Peter. 1982. “Aggregate Demand Management in Search Equilibrium.” *Journal of Political Economy*, 90: 881–94.
- [19] Diamond, Peter, and Drew Fudenberg. 1989. “Rational Expectations Business Cycles in Search Equilibrium.” *Journal of Political Economy*, 97: 606–19.
- [20] Farmer, Roger. 2013 “Animal Spirits, Financial Crises and Persistent Unemployment”, *Economic Journal*, 123: 317-340.
- [21] Fujita, Shigeru, and Garey Ramey. 2009. “The Cyclicity of Separation and Job Finding Rates.” *International Economic Review*, 50: 415–30.
- [22] Gabaix, Xavier. 2011. “The Granular Origins of Aggregate Fluctuations.” *Econometrica* 79: 733-772.
- [23] Galí, Jordi, and Thijs van Rens. 2014. “The Vanishing Procyclicality of Labor Productivity.” Manuscript, University of Warwick.
- [24] Gavazza, Alessandro, Simon Mongey, and Gianluca Violante. 2018. “Aggregate Recruiting Intensity.” *American Economic Review*, Forthcoming.
- [25] Gertler, Mark, and Antonella Trigari. 2009. “Unemployment Dynamics with Staggered Nash Wage Bargaining.” *Journal of Political Economy*, 117: 38-86.
- [26] Golosov, Mikahil, and Guido Menzio. 2015. “Agency Business Cycles.” NBER Working Paper 21743.
- [27] Haefke, Christian, M. Sonntag, and Thijs van Rens, 2013. “Wage Rigidity and Job Creation.” *Journal of Monetary Economics*.
- [28] Hagedorn, Markus, and Iourii Manovskii. 2008. “The Cyclical Behavior of Unemployment and Vacancies Revisited.” *American Economic Review*, 98: 1692–706.

- [29] Hall, Robert, 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review*, 95: 50-65.
- [30] Hall, Robert 2016. “High Discounts and High Unemployment,” *American Economic Review*, Forthcoming.
- [31] Hall, Robert, and Paul Milgrom. 2008. “The Limited Influence of Unemployment of the Wage Bargain.” *American Economic Review*, 98: 1653–74.
- [32] Heller, Walter. 1986. “Coordination Failure Under Complete Markets with Applications to Effective Demand.” *Equilibrium analysis: Essays in Honor of Kenneth J. Arrow*, 2: 155–175.
- [33] Ichino, Andrea, and Regina Riphahn. 2005. “The Effect of Employment Protection on Worker Effort: Absenteeism during and after Probation.” *Journal of the European Economic Association*, 3: 120-143.
- [34] Jaimovich, Nir, and Sergio Rebelo. 2009. “Can News About the Future Drive Business Cycles?” *American Economic Review*, 99: 1097-1118.
- [35] Jenter, Dirk, and Fadi Kanaan. 2015. “CEO Turnover and Relative Performance Evaluation.” *Journal of Finance* 50: 2155-2183.
- [36] Jovanovic, Boyan. 1987. “Micro Shocks and Aggregate Risk.” *Quarterly Journal of Economics*, 102: 395-409.
- [37] Kaplan, Greg, and Guido Menzio. 2016. “Shopping Externalities and Self-Fulfilling Unemployment Fluctuations.” *Journal of Political Economy*, 124: 771-825.
- [38] Kennan, John 2010. “Private Information, Wage Bargaining and Employment Fluctuations.” *Review of Economic Studies*, 77: 633-64.
- [39] Kydland, Finn, and Edward Prescott. 1982. “Time to Build and Aggregate Fluctuations.” *Econometrica*, 50: 1345-70.
- [40] Menzio, Guido 2005. “High Frequency Wage Rigidity.” Manuscript, Northwestern University.
- [41] Menzio, Guido, and Espen Moen. 2010. “Worker Replacement.” *Journal of Monetary Economics*, 57: 623-36.
- [42] Menzio, Guido, and Shouyong Shi. 2011. “Efficient Search on the Job and the Business Cycle.” *Journal of Political Economy*, 119: 468–510.

- [43] Mortensen, Dale 1999. “Equilibrium Unemployment Dynamics.” *International Economic Review*, 40: 889–914.
- [44] Mortensen, Dale, and Christopher Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies*, 61: 397–415.
- [45] Pissarides, Christopher 1985. “Short-Run Equilibrium Dynamics of Unemployment, Vacancies and Real Wages.” *American Economic Review*, 75: 676–90.
- [46] Shapiro, Carl, and Joseph Stiglitz. 1984. “Equilibrium unemployment as a worker discipline device.” *American Economic Review*, 74: 433-443.
- [47] Shimer, Robert. 2005. “The Cyclical Behavior of Unemployment and Vacancies.” *American Economic Review*, 95: 25–49.

Appendix

A Proof of Lemma 1

(i) Let $\rho \geq 0$ denote the Lagrange multiplier on the worker's incentive compatibility constraint, let $\bar{\nu}(y, \hat{z}) \geq 0$ denote the multiplier on the constraint $1 - s(y, \hat{z}) \geq 0$ and let $\underline{\nu}(y, \hat{z})$ denote the multiplier on the constraint $s(y, \hat{z}) \geq 0$.

The first order condition with respect to the firing probability $s(y_\ell, \hat{z})$ is given by

$$(1 - \delta) [F(x)V(\hat{z}) + W(x)J(\hat{z})] = \rho\beta(1 - \delta)(p_h(1) - p_h(0))V(\hat{z}) + \underline{\nu}(y_\ell, \hat{z}) - \bar{\nu}(y_\ell, \hat{z}), \quad (\text{A1})$$

together with the complementary slackness conditions $\bar{\nu}(y_\ell, \hat{z}) \cdot (1 - s(y_\ell, \hat{z})) = 0$ and $\underline{\nu}(y_\ell, \hat{z}) \cdot s(y_\ell, \hat{z}) = 0$. The left-hand side of (A1) is the marginal cost of increasing $s(y_\ell, \hat{z})$. This cost is given by the decline in the product of the worker's and firm's gains from trade caused by a marginal increase in the firing probability $s(y_\ell, \hat{z})$. The right-hand side of (A1) is the marginal benefit of increasing $s(y_\ell, \hat{z})$. This benefit is given by the value of relaxing the worker's incentive compatibility and the $s(y_\ell, \hat{z}) \geq 0$ constraints net of the cost of tightening the $s(y_\ell, \hat{z}) \leq 1$ constraint by marginally increasing the firing probability $s(y_\ell, \hat{z})$.

Similarly, the first order condition with respect to the firing probability $s(y_h, \hat{z})$ is given by

$$(1 - \delta) [F(x)V(\hat{z}) + W(x)J(\hat{z}) + \rho\beta(p_h(1) - p_h(0))V(\hat{z})] = \underline{\nu}(y_h, \hat{z}) - \bar{\nu}(y_h, \hat{z}), \quad (\text{A2})$$

together with the complementary slackness conditions $\bar{\nu}(y_h, \hat{z}) \cdot (1 - s(y_h, \hat{z})) = 0$ and $\underline{\nu}(y_h, \hat{z}) \cdot s(y_h, \hat{z}) = 0$. The left-hand side of (A2) represents the marginal cost of increasing $s(y_h, \hat{z})$. The right-hand side of (A2) represents the marginal benefit of increasing $s(y_h, \hat{z})$. Notice that increasing the firing probability $s(y_h, \hat{z})$ tightens the worker's incentive compatibility constraint and, hence, the term in ρ is now on the left-hand side of (A2).

Suppose $\rho = 0$. First, notice that the left-hand side of (A1) is strictly positive as $V(\hat{z}) > 0$, $J(\hat{z}) > 0$ by assumption, and $W(x) > 0$, $F(x) > 0$ at the optimum x^* . The right-hand side of (A1) is strictly positive only if $\underline{\nu}(y_\ell, \hat{z}) > 0$. Hence, if $\rho = 0$, the only solution to the first order condition with respect to the firing probability $s(y_\ell, \hat{z})$ is 0. Next, notice that the left-hand side of (A2) is strictly positive and the right-hand side is strictly positive only if $\underline{\nu}(y_h, \hat{z}) > 0$. Hence, if $\rho = 0$, the only solution to the first order condition with respect to the firing probability $s(y_h, \hat{z})$ is 0. However, if $s(y_\ell, \hat{z}) = s(y_h, \hat{z}) = 0$, the worker's incentive compatibility constraint is violated. Therefore, $\rho > 0$ and the worker's

incentive compatibility constraint holds with equality.

(ii) The first order condition with respect to $s(y_h, \hat{z})$ is given by (A2) together with the complementary slackness conditions $\bar{\nu}(y_h, \hat{z})(1-s(y_h, \hat{z})) = 0$ and $\underline{\nu}(y_h, \hat{z})s(y_h, \hat{z}) = 0$. The left-hand side of (A2) is strictly positive. The right-hand side of (A2) is strictly positive only if $\underline{\nu}(y_h, \hat{z}) > 0$. Therefore, the first order condition is satisfied only if $\underline{\nu}(y_h, \hat{z}) > 0$ and, hence, only if $s(y_h, \hat{z}) = 0$.

(iii) Using the definition of $\phi(\hat{z})$, we can rewrite the first order condition with respect to the firing probability $s(y_\ell, \hat{z})$ as

$$(1 - \delta)V(\hat{z}) [F(x) + W(x)/\phi(\hat{z}) - \rho\beta(p_h(1) - p_h(0))] = \underline{\nu}(y_\ell, \hat{z}) - \bar{\nu}(y_\ell, \hat{z}), \quad (\text{A3})$$

together with $\bar{\nu}(y_\ell, \hat{z}) \cdot (1 - s(y_\ell, \hat{z})) = 0$ and $\underline{\nu}(y_\ell, \hat{z}) \cdot s(y_\ell, \hat{z}) = 0$. The left-hand side of (A3) is strictly decreasing in $\phi(\hat{z})$. The right-hand side of (A3) is strictly positive if $\underline{\nu}(y_\ell, \hat{z})$ is strictly positive and it is strictly negative if $\bar{\nu}(y_\ell, \hat{z})$ is strictly positive. Therefore, there exists a ϕ^* such that if $\phi(\hat{z}) > \phi^*$, the left-hand side is strictly negative and the solution to (A3) requires $\bar{\nu}(y_\ell, \hat{z}) > 0$. In this case, the solution to the first order condition for $s(y_\ell, \hat{z})$ is 1. If $\phi(\hat{z}) < \phi^*$, the left-hand side is strictly positive and the solution to (A3) requires $\underline{\nu}(y_\ell, \hat{z}) > 0$. In this case, the solution to the first order condition for $s(y_\ell, \hat{z})$ is 0.

B Proof of Lemma 2

The first order condition with respect to the wage w is given by

$$F(x)v'(w) - W(x) = 0. \quad (\text{A4})$$

The left-hand side of (A4) is the increase in the product of the worker's and firm's gains caused by a marginal increase in the worker's wage w . A marginal increase in w , increases the worker's gains from trade by $v'(w)$ and decreases the firm's gains from trade by 1. Therefore, a marginal increase in w , increases the product of the worker's and firm's gains from trade by $F(x)v'(w) - W(x)$. The first order condition for w states that the effect of a marginal increase in w is zero.