Welfare Cost of Business Cycles When Markets Are Incomplete

June, 2004

Tom Krebs[*]
Brown University[†]

Abstract

This paper analyzes the welfare costs of business cycles when workers face uninsurable idiosyncratic labor income risk. In accordance with the previous literature, this paper decomposes labor income risk into an aggregate and an idiosyncratic component, but in contrast to the previous literature, this paper allows for multiple sources of idiosyncratic labor income risk. Using the multi-dimensional approach to idiosyncratic risk, this paper provides a general characterization of the welfare cost of business cycles when preferences and the (marginal) process of individual labor income in the economy with business cycles are given. The general analysis shows that the introduction of multiple sources of idiosyncratic risk never decreases the welfare cost of business cycles, and strictly increases it if there are cyclical fluctuations across the different sources of risk. Finally, this paper also provides a quantitative analysis of multi-dimensional labor income risk based on a version of the model that is calibrated to match U.S. labor market data. The quantitative analysis suggests that realistic variations across two particular dimensions of idiosyncratic labor income risk increase the welfare cost of business cycles by a substantial amount.

I. Introduction

In a highly influential contribution, Lucas (1987) argues that the welfare costs of business cycles are likely to be very small, and that therefore the potential gains from counter-cyclical stabilization policy are negligible. His argument is based on a representative-agent model with no production and "standard" preferences. In other words, Lucas (1987) assumes that i) there is no uninsurable risk (complete markets), ii) there is no link between business cycles and economic growth, and iii) preferences allow for a time-additive expected utility representation with moderate degree of (relative) risk aversion. In principle, any one of these three assumptions could be questioned, and it is therefore important to know how relaxing any one of them will change the welfare conclusions.[1] This paper analyzes to what extent the introduction of uninsurable idiosyncratic labor income risk (incomplete markets) increases the welfare cost of business cycles. In order to disentangle the complete-markets assumption from the other two assumptions made in Lucas (1987), the analysis presented in this paper is based on a model with standard preferences and no link between business cycles and economic growth.

There is strong empirical evidence that idiosyncratic labor income risk varies over the business cycle.[2] If risk-averse agents dislike variations in idiosyncratic risk, then this observation suggests a channel through which market incompleteness might increase the welfare cost of business cycles.[3] A growing body of work has studied the quantitative importance

---

[1]For example, Lucas (2003) argues that this type of analysis constitutes one of the main macroeconomic priorities.

[2]Clearly, the job displacement rate and the unemployment duration increase during recessions (Hall, 1995). Furthermore, the magnitude of permanent income losses of displaced workers depends on cyclical conditions (Jacobson, LaLonde, and Sullivan, 1993), and the business cycle has an effect on the placement success of initial labor market entrants. Finally, Storesletten, Telmer, and Yaron (2004), and to a certain extent also Meghir and Pistaferri (2004), find that the variance of persistent income shocks is highly sensitive to business cycle conditions.

[3]Another possible channel is the interaction between aggregate productivity shocks and idiosyncratic

of this channel finding mostly modest effects.[4] In his recent review of the literature Lucas (2003) therefore concludes: *"But I argue in the end that, based what we know now, it is unrealistic to hope for gains larger than a tenth of a percent from better countercyclical policy"*. In this paper, we introduce multiple sources (dimensions) of idiosyncratic labor income risk, and show that this innovation has the potential to generate welfare cost of business cycles that are substantially larger than one tenth of a percent of lifetime consumption.

The previous literature has commonly decomposed individual labor income risk into an aggregate and an idiosyncratic component, and has further assumed that the idiosyncratic component is represented by a one-dimensional variable. In this paper, we also decompose individual labor income risk into an aggregate and an idiosyncratic component, but in contrast to the previous literature we allow for multiple sources of idiosyncratic labor income risk. Using the multi-dimensional approach to idiosyncratic risk, this paper first provides a general characterization of the welfare cost of business cycles when preferences and the (marginal) process of individual income in the economy with business cycles are given. The general analysis shows that the introduction of multiple sources of idiosyncratic risk never decreases the welfare cost of business cycles, and strictly increases it if there are cyclical fluctuations across the different dimensions of idiosyncratic risk. In other words, this paper derives lower and upper bounds on the welfare cost of business cycles, and shows that the case of one-dimensional idiosyncratic risk leads to the minimum welfare cost of business

---

shocks due to the non-linearity of the utility function (risk-aversion). However, Krebs (2003a) finds that this interaction effect tends to be relatively small. Storesletten, Telmer, and Yaron (2001) find a larger interaction effect, but they use a higher degree of risk aversion.

[4]For example, Atkeson and Phelan (1994) suggest that the welfare cost of business cycles might be nil, Krusell and Smith (1999) find only negligible effects, and Imrohoroglu (1989) and Storesletten, Telmer, and Yaron (2001) find non-negligible, but relatively modest effects (at least for degrees of risk aversion around one). Further, Gomes, Greenwood, and Rebelo (2001) suggest that business cycles might have positive welfare effects. More recently, Krebs (2003a) provides an example where the welfare cost of business cycles is large, and Krusell and Smith (2002) find substantial costs for some groups of agents, even though the net effect on welfare for the entire population is small. See also Lucas (2003) for a recent survey of the literature.

cycles.[5]

There is a simple economic intuition why additional dimensions of idiosyncratic risk increase the welfare costs of business cycles when there are cyclical fluctuations across the different sources of idiosyncratic risk. Consider, for example, the case in which the labor income of an individual worker is determined by the random returns to his endowment of two different skill (human capital) types.[6] Suppose further that the first skill is only useful during economic booms and the second skill is only useful during recessions. In this case, an individual household working in the economy with business cycles receives a labor income that at each point in time is determined by the random return to one skill only, whereas after the elimination of business cycles the same household receives a labor income that at each point in time is determined by a weighted average of two random returns.[7] Clearly, in this case the elimination of business cycles leads to a reduction in the amount of labor income risk the individual household has to bear, and this risk-reduction effect is absent in any one-dimensional model of labor income risk. Thus, introducing multiple dimensions of idiosyncratic risk and allowing for cyclical fluctuations across dimensions of risk leads to the conclusion that the welfare costs of business cycles are likely to be higher than suggested by

---

[5]Notice that welfare of workers depends on preferences and the marginal process of individual income only, and that throughout the entire analysis we fix preferences and the marginal process of individual income in the economy with business cycles. Thus, the introduction of multi-dimensional idiosyncratic risk changes the welfare cost of business cycles because it changes welfare in the economy without business cycles (keeping welfare in the economy with business cycles constant). This result holds even though business cycles are eliminated in accordance with the integration principle (Krusell and Smith, 1999, and equation 10).

[6]Note that any Mincer-type wage regression (Mincer, 1974) implicitly makes the assumption that the labor income of a worker is composed of individual components that reflect the return to different dimensions of human capital. Notice also that a similar argument applies when skill returns are constant, but skill-levels evolve stochastically over time (for example, loss of specific human capital in the event of job displacement).

[7]This is an application of the integration principle as first proposed by Krusell and Smith (1999). Atkeson and Phelan (1994) suggest that the elimination of business cycles amounts to eliminating any correlation across individual income shocks. If the independence assumption is satisfied (equation 15), then the application of the integration principle removes any correlation across individual income shocks (for details, see the discussion following proposition 2).

the previous literature.

In addition to the theoretical analysis, this paper also provides a numerical analysis based on a version of the model calibrated to U.S. data. In this quantitative part, we compute the welfare cost of business cycles as a function of one parameter, which measures the degree to which the economy with business cycles exhibits cyclical fluctuations across the different dimensions of idiosyncratic labor market risk. Depending on the particular value of this one parameter, the welfare cost of business cycles could be anything between 0.3 (the one-dimensional case) and 9 percent of lifetime consumption (for log-utility preferences). That is, the introduction of multi-dimensional idiosyncratic risk has the potential to generate welfare cost of business cycles that are quite large indeed.

The preceding discussion highlights that it is not multi-dimensionality *per se* that increases the welfare cost of business cycles, but the cyclical fluctuations across different dimensions of risk. Thus, a natural question to ask is whether the U.S. labor market exhibits cyclical fluctuations of this type, and whether these fluctuations are substantial enough to generate sizable welfare cost of business cycles. Clearly, there are in principle many approaches to investigating this issue empirically, but one approach that seems particularly promising is the following. First, identify particular dimensions of human capital of individual workers and their empirical counterparts, say specific human capital (proxied by tenure at a firm) and general human capital (proxied by education). Second, estimate that part of total labor income risk that can be attributed to each individual component of human capital. Third, use these estimates to find the cyclical fluctuations across the individual dimensions of human capital risk. In this paper, we do not attempt to conduct a full-blown empirical analysis along these lines, but instead provide some preliminary evidence that is based on recently obtained empirical results on permanent labor income risk (Meghir and Pistaferri, 2004) and insights derived from the literature on Mincer-type wage regressions

4

(Mincer, 1974). This preliminary analysis suggests that the U.S. labor market exhibits cyclical variations across two particular dimensions of human capital risk (specific and general human capital) that lead to welfare cost of business cycles around 1.5 percent of lifetime consumption (for log-utility preferences). Clearly, this particular quantitative result is highly tentative, and much more empirical work on the issue of cyclical fluctuations across dimensions of human capital risk is required to arrive at firm conclusions. However, the result indicates that the multi-dimensional view of labor market risk is likely to generate welfare cost of business cycles that are significantly larger than what has been suggested by the previous literature.

At this stage, it is worth pointing out that the results reported here do not imply that macroeconomic stabilization policy necessarily leads to substantial welfare gains. More specifically, this paper does not explicitly model how counter-cyclical fiscal or monetary policy might affect the business cycle (the cyclical nature of idiosyncratic labor market risk), and it is therefore not well-suited for studying the question whether policy makers can design macroeconomic policies that stabilize the business cycle.[8] Moreover, this paper only focuses on one channel through which the business cycle might affect welfare, and the welfare calculation presented here are therefore necessarily incomplete. For example, there are many ways in which business cycles affect aggregate output (growth), and this paper rules out any such output-effect by assumption. Thus, the welfare results derived here have to be interpreted with caution keeping these limitations in mind. In short, this paper does not show that macroeconomic stabilization policy is desirable, but it does show that our current knowledge of idiosyncratic labor market risk is not sufficient to rule out large welfare costs of business cycles.

Finally, a comment regarding modeling strategy is in order. In the spirit of Lucas (1987),

---

[8]This type of "black-box" approach is common in the literature on the welfare cost of business cycles.

this paper attempts to analyze the welfare cost of business cycles keeping the mean of aggregate output fixed.[9] Given this restriction, one approach would be to use preferences and estimates of the process of individual consumption risk to derive the welfare cost of business cycles. One problem with this approach is that the empirical literature has mainly focused on cyclical variations in individual income risk.[10] Put differently, we know much more about cyclical variations in individual income risk than we know about cyclical variations in individual consumption risk. Thus, for the quantitative application, it is necessary to have a model that translates income shocks into consumption fluctuations. The model used in this paper is an exchange economy with only permanent income shocks as in Constantinides and Duffie (1996) and Krebs (2004). This model has the advantage that it is highly tractable, but still rich enough to allow for a tight link between the theoretical model and the empirical literature on labor income risk.

The paper is organized as follows. Section II develops the model that is used to discuss the effect of multi-dimensional risk on the welfare cost of business cycles. Section III discusses how to eliminate business cycles in economies with incomplete markets and derives a closed-form expression for the welfare cost of business cycles. Section IV contains the main results. More specifically, section IV.A provides a general characterization of the welfare cost of business cycles, and shows that the introduction of multi-dimensional idiosyncratic risk increases the welfare cost of business cycles if there are cyclical variations across the different

---

[9]This is not to say that an analysis of the relationship between business cycles and long-run output growth is unimportant. However, keeping the mean of the level and growth rate of aggregate output fixed allows us to focus on the issue of multi-dimensional risk in a simple and transparent manner. Moreover, several links between business cycles and the mean of aggregate output (growth) have been suggested in the literature (see Lucas, 2003, for a survey), and no clear consensus has yet emerged as to which mechanism is the most important one.

[10]There is, however, some recent empirical work that discusses the asset pricing implications of cyclical variations in individual consumption risk (for example, Brav, Constantinides, and Gescy, 2002). One well-known problem with the consumption approach is that the currently available individual consumption data are either limited to food consumption (PSID) or lack a sufficient panel dimension (CEX).

dimensions of idiosyncratic risk. Section IV.B conducts the quantitative analysis based on a calibrated version of the model. Section IV.C provides some preliminary evidence suggesting that the U.S. labor market is characterized by substantial cyclical variations across two particular dimensions of idiosyncratic labor market risk. Section V concludes.

II. Model

As mentioned before, the model is similar to the model considered in Constantinides and Duffie (1996) and Krebs (2004). It features ex-ante identical, long-lived households (workers) with homothetic preferences that make consumption/saving choices in the face of uninsurable income shocks. Income shocks are permanent, which implies that self-insurance is an ineffective means to smooth out income fluctuations. Indeed, the economy is set up in a way so that in equilibrium households will not self-insure at all. That is, income shocks translate one-to-one into consumption changes (proposition 1). Notice that this result does not depend on the assumption that aggregate saving is zero, even though we will make it to simplify the analysis. For example, Krebs (2003a,2003b) considers a production economy with only permanent income shocks (log-income follows a random walk) and ex-ante identical households, and shows again that self-insurance is highly ineffective.[11]

---

[11]In the current model, there is no production and no asset in positive net supply, so that in equilibrium all households end up with zero wealth. Thus, a stationary cross-sectional distribution of wealth exists, even though it is degenerate. Once we introduce assets in positive net supply (Constantinides and Duffie, 1996, and Krebs, 2004) or physical capital and production (Krebs 2003a,b), the cross-sectional distribution of wealth will diverge since individual wealth (approximately) follows a logarithmic random walk. However, Constantinides and Duffie (1996) show how to modify the model by introducing death probabilities so that a stationary distribution of wealth always exists. Indeed, they show that by choosing the death probabilities appropriately, the model can match any cross-sectional distribution of wealth. Huggett (1996) discusses the wealth distribution in life-cycle economies, and Castanedas, Diaz-Gimenez, and Rios-Rull (2003) analyze the wealth distribution in economies that mix the features of life-cycle models and dynastic models.

II.A. Economy

Time is discrete and open ended. Labor income of household i in period t is denoted by y_{it}. Labor income is random and defined by an initial level y_{i0} and the law of motion

$$y_{i,t+1} = (1 + \eta_{i,t+1}) \, y_{it} \, , \tag{1}$$

where the random variable $\eta_{i,t+1}$ describe shocks to the labor income of worker i.

We make the following assumptions about the process of individual labor income defined by (1). We assume that $\eta_{it} = \eta(s_{it}, S_t)$, where s_{it} denotes a vector of idiosyncratic shocks to the income of worker i and S_t the aggregate state of the economy. We assume further that the joint process $\{s_{it}, S_t\}$ is Markov with transition probabilities $\pi(s_{i,t+1}, S_{t+1}|s_{it}, S_t)$ satisfying

$$\pi(s_{i,t+1}, S_{t+1}|s_{it}, S_t) = \pi(s_{i,t+1}, S_{t+1}|S_t) \, . \tag{2}$$

In short, conditional on the history of aggregate states, idiosyncratic labor income shocks are identically distributed across workers and independently distributed over time (idiosyncratic shocks are unpredictable).

Specification (1) in conjunction with our assumption that idiosyncratic shocks are unpredictable implies that, conditional on the aggregate state, the growth rate of individual labor income is unpredictable. Using the approximation $log(1 + x) \approx x$, equation (1) yields:

$$log y_{i,t+1} = log y_{it} + \eta_{i,t+1} \, . \tag{3}$$

That is, log-labor income approximately follows a random walk. Clearly, depending on the specification of the function $\eta(.)$, this random walk might exhibit state-dependent drift and a heteroscedastic error term. This shows that the idiosyncratic labor income shocks discussed in this paper correspond to permanent shocks to labor income, and explains why self-insurance does not work very well in the current model. Moreover, when calibrating

8

the model, one should focus on the permanent component of labor income risk as estimated by Meghir and Pistaferri (2004) and Storesletten et al. (2004) for total labor income and by Jacobson et al. (1993) for income losses due to job displacement. See the quantitative section IV.B for more details on this.

Each household begins life with no initial financial wealth. Households have the opportunity to buy and sell short-lived securities $j = 1, \ldots, J$ in zero net supply. The sequential budget constraint of worker i reads

$$c_{it} + Q_t \cdot a_{i,t+1} = y_{it} + D_t \cdot a_{it} \tag{4}$$

$$Q_t \cdot a_{i,t+1} \geq -M \ , \quad a_{i0} = 0 \ .$$

Here c_{it} denotes consumption of household i in period t , $a_{it} = (a_{i1t}, \ldots, a_{iJt})$ his vector of asset holdings at the beginning of period t, $D_t = (D_{1t}, \ldots, D_{Jt})$ the vector of asset payoffs, and $Q_t = (Q_{1t}, \ldots, Q_{Jt})$ the vector of asset prices. The budget constraint (4) simply says that total spending on consumption and assets is equal to total income from labor and asset holdings (capital). The real number M represents an explicit debt constraint that rules out Ponzi schemes (Krebs, 2004). We assume that asset payoffs may depend on the aggregate state, but not on idiosyncratic shocks: $D_t = D(S_t)$. Thus, asset markets are incomplete in the sense that they do not provide explicit insurance against idiosyncratic labor income risk. Notice that we allow one of the assets, say the first, to be risk-free, in which case the asset payoffs are state-independent: $D_1(S_t) = 1$. Thus, households have the opportunity to self-insure through saving, even though in equilibrium they will not do so.

Households have identical preferences that allow for a time-additive expected utility representation:

$$U(\{c_{it}\}) = E \left[\sum_{t=0}^{\infty} \beta^t u(c_{it}) \right] \ . \tag{5}$$

Moreover, we assume that the one-period utility function, u, is given by $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $\gamma \neq 1$, or $u(c) = \log c$, that is, preferences exhibit constant degree of relative risk aversion γ. Finally,

9

we assume that the following condition is satisfied:

$$\beta E\left[(1 + \eta(s_{i,t+1}, S_{t+1}))^{1-\gamma}|S_t\right] < 1 . \tag{6}$$

This inequality ensures that in equilibrium the expected lifetime utility of workers is finite for any budget-feasible plan, and that any solution to the Euler equation also solves a corresponding transversality condition (Krebs, 2004). Notice that this condition is automatically satisfied if $\gamma = 1$ (log-utility).

II.2. Equilibrium

For a given process of asset prices and asset payoffs, each household chooses a consumption-saving plan that maximizes expected lifetime utility (5) subject to the budget constraint (4). In equilibrium, asset markets must clear:[12]

$$\sum_i a_{it} = 0 . \tag{7}$$

Notice that (4) and (7) imply goods market clearing (Walras' law). We now show that there is an asset price process for which the plan $a_{it} = 0$ and $c_{it} = y_{it}$ is individually optimal. That is, there is an asset price process for which workers choose to consume all their income in every period. Since this plan also satisfies the market clearing condition (7), it is an equilibrium.

The Euler equations associated with the consumption-saving problem of household i read

$$c_{it}^{-\gamma}Q_{jt} = \beta E[c_{i,t+1}^{-\gamma}D_{j,t+1}|s_i^t, S^t] , \tag{8}$$

where $s_i^t = (s_{i0}, \ldots, s_{it})$ and $S^t = (S_0, \ldots, S_t)$ is the information that is available to household i in period t. The Euler equation (8) says that the marginal utility cost of buying (saving)

---

[12]The notation suggests that there are a finite number of households, but all propositions in this paper remain valid for the case of a continuum of households.

one more unit of asset j is equal to the expected marginal utility gain of doing so. In Krebs (2004) it is shown that any plan solving the Euler equation (8) and the budget constraint (4) also satisfies a corresponding transversality condition if (6) is satisfied. Thus, we can focus on Euler equations when discussing optimal consumption/saving plans.

Suppose asset prices are given by

$$Q_{jt} = \beta\, E\left[\left(\frac{y_{i,t+1}}{y_{it}}\right)^{-\gamma} D_{j,t+1}|S_t\right]\ ,\tag{9}$$

where the individual labor income growth rate $y_{i,t+1}/y_{it}$ is given by (1). Given these asset prices, the Euler equation (8) is satisfied. Since for $a_{it} = 0$ the market clearing condition (7) holds, we have found an equilibrium. Notice that we have used the fact that s_i^t does not predict future income growth. Without this assumption, the Euler equation (8) would not hold at $a_{it} = 0$ and $c_{it} = y_{it}$.

Proposition 1. The consumption-saving plan $\{a_{it}, c_{it}\}$, where $a_{it} = 0$ and $c_{it} = y_{it}$, in conjunction with asset prices (9) constitute an equilibrium.

III. General Welfare Analysis

In this section, we use the model laid out in section II to derive an explicit formula for the welfare cost of business cycles (proposition 2). To this end, we first discuss in part III.A how to eliminate business cycles in economies with idiosyncratic risk following the analysis in Krebs (2003a), Krusell and Smith (1999,2002), and Lucas (2003), and then apply the general approach in part III.B to the current model.

III.A. How to Eliminate Business Cycles

We are interested in the change in expected lifetime utility of workers when business cycles are eliminated. In our model, the elimination of business cycles amounts to moving from an

economy with fluctuations in the aggregate state S_t (the economy with business cycles) to an economy with constant S_t (the economy without business cycles). For simplicity, assume that the aggregate state process is serially uncorrelated so that $\pi(s_{i,t+1}, S_{t+1} | S_t) = \pi(s_{i,t+1}, S_{t+1})$. In this case, we move from an economy with individual labor income process described by probabilities $\pi(s_i, S)$ and individual growth-rate realizations $\eta = \eta(s_i, S)$ to an economy with an individual labor income process described probabilities $\pi'(s_i)$ and individual growth-rate realizations $\eta' = \eta'(s_i)$. The question that arises is what relationship holds between $(\pi, \eta(.))$ and the corresponding $(\pi', \eta'(.))$.

It seems natural to compute the probabilities in the economy without business cycles as the simple marginals of the probabilities in the economy with business cycles. Using an analogous principle for outcome functions, we have:[13]

$$
\begin{aligned}
\pi'(s_i) &= \sum_S \pi(s_i, S) \\
\eta'(s_i) &= \sum_S \eta(s_i, S)\pi(S | s_i) \ .
\end{aligned}
\tag{10}
$$

In short, we have the following relationship between random variables describing individual income risk in the two economies:

$$
\eta_i' = E[\eta_i | s_i] \ .
\tag{11}
$$

In words: conditional on the idiosyncratic shock, the expected value of all exogenous random variables is the same in both economies. The procedure of integrating out aggregate states in this way has been first put forward by Krusell and Smith (1999). They call it the "integration principle", and we follow their terminology in this paper. This principle is a natural extension of the methodology used by Lucas (1987) for representative-agent economies in the sense that in both cases random variables are replaced by their expected value, but in (11) we use the mean conditional on the idiosyncratic shock. Krebs (2003a), Lucas (2003), and Krusell and

---

[13]For random variables with uncountable support summation is replaced by integration.

Smith (2002) have recently argued that this principle is intuitively convincing, and that it imposes a useful discipline upon the analysis. Notice that for random variables that exhibit a dependence on S, it means that the elimination of business cycles unambiguously reduces risk. In other words, the random variables μ and η are mean-preserving spreads of the random variables μ' and η'.

III.B. A Welfare Formula

The integration principle defines how the elimination of business cycles affects the labor income process of workers. We now apply the integration principle to the model discussed in part II, and derive an explicit formula for the welfare cost of business cycles.

Changes in the labor income process change the equilibrium consumption and welfare of workers. We define the welfare cost of business cycles as the number Δ that solves

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t u \left(c_{it}(1 + \Delta) \right) \right] = E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c'_{it}) \right] , \tag{12}$$

where c_{it} is consumption in the economy with business cycles and c'_{it} is consumption in the economy without business cycles.[14] That is, we define the welfare cost of business cycles as the percentage of lifetime consumption that workers have to receive in order to be fully compensated for the cyclical variations in labor income risk. Using the fact that in equilibrium individual consumption equals individual income, we find the following welfare expressions

$$
\begin{aligned}
E_0 \left[\sum_{t=0}^{\infty} \beta^t log c_{it} \right] &= \frac{1}{1-\beta} log c_{i0} + \frac{\beta}{(1-\beta)^2} \left(E[log(1+\eta)] \right) \\
E_0 \left[\sum_{t=0}^{\infty} \frac{c_{it}^{1-\gamma}}{1-\gamma} \right] &= \frac{c_{i0}^{1-\gamma}}{(1-\gamma)\left(1 - \beta E[(1+\eta)^{1-\gamma}]\right)} .
\end{aligned} \tag{13}
$$

---

[14]Note that (12) follows the bulk of the previous literature and defines the total cost of business cycle fluctuations. Alvarez and Jermann (2003) define a concept of marginal cost of business cycles by considering small reductions in consumption fluctuations.

Substituting the welfare expression (13) into (12) and using the integration principle (11) yields

$$
\begin{aligned}
\Delta &= \exp\left(\frac{\beta}{1-\beta}\left(E[log(1+\eta')] - E[log(1+\eta)]\right)\right) - 1 \\
\Delta &= \left(\frac{1 - \beta E[(1+\eta)^{1-\gamma}]}{1 - \beta E[(1+\eta')^{1-\gamma}]}\right)^{\frac{1}{1-\gamma}} - 1\,,
\end{aligned}
\tag{14}
$$

where the random variables η and η' have been defined above. Notice that Δ is the same for all worker, and that the expectations is taken over idiosyncratic and aggregate shocks. Note also the fact that η is a mean-preserving spread of η', which in conjunction with the concavity of the utility function implies that $\Delta \geq 0$. In other words, cyclical variations in idiosyncratic labor income risk never decrease the welfare cost of business cycles. Finally, note that the strict inequality holds if and only if the outcome function $\eta = \eta(s_i, S)$ depends in a non-trivial way on S. That is, we have $\Delta > 0$ if and only if the aggregate state affects the magnitude of income changes. We summarize the preceding discussion in the following proposition:

Proposition 2. The welfare cost of business cycles is given by (14). The welfare cost of business cycles is non-negative, $\Delta \geq 0$, and strictly positive if and only if the outcome function $\eta(.)$ depends in a non-trivial way on the aggregate state: $\eta(s_i, S) \neq \eta(s_i, \hat{S})$ for some $s_i, S, \hat{S}$. In other words, cyclical variations in idiosyncratic labor income risk never decrease the welfare cost of business cycles, and they increase the welfare cost of business cycles if and only if the magnitude of idiosyncratic labor income shocks depends on the aggregate state.

Proposition 2 makes clear why several papers in the previous literature have found small welfare cost of business cycles. For example, the papers by Atkeson and Phelan (1994), Imrohoroglu (1989), and Krusell and Smith (1999, 2002) all study models of worker unemployment in which the individual income process that has a two-state support that is

independent of the aggregate state. With the exception of Krusell and Smith (2002), all these papers also assume that the idiosyncratic variable s_i is the individual employment state itself. That is, $s_i = 0$ if worker i is unemployed and $s_i = 1$ if worker i is employed. Given that in this case the outcome function $\eta(.)$ does not depend on the aggregate state, proposition 2 implies $\Delta = 0$, and this is exactly the result reported by Atkeson and Phelan (1994). Krusell and Smith (1999) find positive welfare cost of business cycles due to indirect price effects, but these effects turn out to be quantitatively small as long as this approach is taken. Imrohoroglu (1989) finds small welfare cost of business cycles not due to indirect price effects, but this effect is due to the fact that her procedure of eliminating business cycles does not satisfy the integration principle.

Notice that if we identify s_i with the employment state of an individual worker, then the idiosyncratic state is necessarily correlated with the aggregate state (as long as unemployment probabilities depend on the aggregate state). In order to avoid such a correlation, Krusell and Smith (2002) introduce a new idiosyncratic variable that by definition is independent of the aggregate state, and then define the observed employment state as a non-linear function of this idiosyncratic variable and the aggregate state. In this case, the welfare cost of business becomes positive (see also the appendix for details). Thus, depending on the choice of the state space, the welfare cost might be either zero or strictly positive even if the (marginal) process of individual labor income in the economy with business cycles is known, a point that has also been made by Krebs (2003a) in the context of the normal-distribution setting discussed in the next section. However, if we impose the additional assumption that idiosyncratic and aggregate shocks are stochastically independent (see equation 15), then the welfare cost of business cycles is uniquely determined as long as idiosyncratic risk is one-dimensional.[15] In contrast, once we allow for multi-dimensional idiosyncratic risk, there is a

---

[15]See section IV.A and the appendix for details. The proof of uniqueness requires in general an additional monotonicity assumption.

wide range of welfare cost of business cycles that is consistent with the (marginal) process of individual labor income in the economy with business cycles. In the next section, we provide a general characterization of this range.

IV. Multi-Dimensional Idiosyncratic Risk

In this section, we discuss the welfare cost of business cycles when there are multiple dimensions of idiosyncratic risk and the independence assumption (15) is satisfied. In the first part (section IV.A), we parameterize idiosyncratic risk and use the parameterization to provide a general characterization of the welfare cost of business cycles when preferences and the marginal process of individual labor income in the economy with business cycles are given. This analysis also shows that the introduction of multiple sources of idiosyncratic risk never decreases the welfare cost of business cycles, and strictly increases it if there are cyclical fluctuations across the different dimensions of idiosyncratic risk. In the second part (section IV.B), we conduct a quantitative analysis that computes the welfare cost of business cycles for a version of the model calibrated to match U.S. labor market data. In this part, we calculate the welfare cost of business cycles as a function of one parameter, namely the correlation coefficient between idiosyncratic labor income shocks during booms and idiosyncratic labor income shocks during recessions. This parameter, in turn, is monotonically related to the degree to which the relative importance of the different dimensions of risk varies over the business cycle. In the final part (section IV.C), we provide some preliminary evidence that suggesting that the U.S. labor market is characterized by substantial cyclical variations across two particular dimensions of human capital risk.

IV.A. Qualitative Analysis

To render the analysis as transparent as possible, we assume that individual labor income shocks η_i are normally distributed, an approach also taken in Krebs (2003a) and Storesletten

16

et al. (2001).[16] Empirically, this corresponds to assuming that η_i represents the permanent component of labor income risk as estimated by Meghir and Pistaferri (2004) and Storesletten et al. (2004). When evaluating the welfare expression (14), we use a second-order Taylor approximation of the utility function, which combined with the normal-distribution assumption leads to a simple and intuitive welfare formula. To check the accuracy of this approximation for quantitative applications, in section IV.B we also compute welfare using the exact utility function and a simple two-state approximation of the normal distribution. Finally, the appendix contains a general discussion of the case in which η_i has a distribution with two-state support that does not use the Taylor approximation of the utility function.

Following the previous literature, we consider the case of two aggregate shocks, $S = L, H$, where $S = L$ corresponds to the event that the economy is contracting (low level of economic activity) and $S = H$ that the economy is expanding (high level of economic activity). These two aggregate states occur with probability π_L, respectively π_H. We also assume that the idiosyncratic shock variable, s_i, has two components: $s_i = (s_{i1}, s_{i2})$. Krebs (2003a), Krusell and Smith (2002), and Lucas (2003) have recently argued that idiosyncratic shocks should by definition be stochastically independent of the aggregate state (independence assumption), and we follow their lead in this paper. More precisely, we assume that:

$$\pi(s_{i1}, s_{i2}, S) = \pi(s_{i1})\pi(s_{i2})\pi(S) . \tag{15}$$

We call (15) the independence assumption. Notice that the independence assumption (15) ensures that the application of the integration principle (11) removes any correlation across individual income shocks (agents). That is, in general we have $corr(\eta_i, \eta_j) \neq 0$ in the economy with business cycles, but we always have $corr(\eta_i', \eta_j') = 0$ in the economy without business cycles. Thus, if (15) is satisfied, then the approach taken in Atkeson and Phelan (1994), which requires that the elimination of business cycles removes any correlation across

---

[16]Storesletten et al. (2001) assume log-normally distributed income shocks.

individual income shocks (agents), coincides with the approach proposed in Krusell and Smith (1999), namely to use the integration principle (11).[17]

As mentioned before, we assume that idiosyncratic shocks are normally distributed: $s_{1i} \sim N(0, \sigma_{1s}^2)$ and $s_{2i} \sim N(0, \sigma_{2s}^2)$. Further, we make the assumption that individual income shocks are a linear function of the two sources of idiosyncratic risk, where the coefficients of this linear function may depend on the aggregate state:

$$\eta(s_{1i}, s_{2i}, S) = \alpha_S \left(\omega_S s_{1i} + (1 - \omega_S) s_{2i} \right) + g \ . \tag{16}$$

In (16) α_S are parameters that measure cyclical fluctuations in total labor income risk, ω_S are parameters that measure cyclical fluctuations across the different dimensions (sources) of idiosyncratic labor income risk, and g is the (constant) growth rate of aggregate income. Notice that if we have $\omega_S = \omega$, then (16) is equivalent to the case of one-dimensional idiosyncratic risk in the sense that income shocks and welfare only depend on $\tilde{s}_i = \omega s_{1i} + (1 - \omega) s_{2i}$ and S. Finally, note that we have by construction

$$E[\eta(s_i, S)|S] = g \ . \tag{17}$$

Equation (17) says that there is no aggregate risk in the economy. Thus, for the economy considered here, the welfare cost of business cycles is nil if there is no uninsurable idiosyncratic risk.

Several interpretations of (16) are possible. One interpretation assumes that labor income of worker i in period t is given by

$$y_{it} = exp\left(r_{1it} h_{1i} + r_{2it} h_{2i} \right) \ , \tag{18}$$

---

[17]Notice, however, that the examples considered in Atkeson and Phelan (1994) and Krusell and Smith (1999) do not satisfy the independence assumption (15) since in both papers the idiosyncratic state, s_i, is identified with the employment state of an individual worker.

where h_{1i} and h_{2i} are the endowments of worker i of each of the two skill-types, and r_{1it} and r_{2it} are the random returns to the two skills. Equation (18) assumes that the skill levels are constant over time and that returns to individual skills evolve stochastically, but an alternative specification in which skill levels are random is also possible. Notice that the definition of labor income (18) is the typical specification used in any Mincer-type wage regression (Mincer, 1974), which assumes that the log-wage of an individual worker is the sum of several components reflecting the return to the different types of human capital (education, experience, tenure). For more on this, see section IV.C.

To ensure that all workers are ex-ante identical, suppose also that their endowments of the different skills are identical: $h_{1i} = h_1$ and $h_{2i} = h_2$. Given the specification (18), we have

$$
\begin{aligned}
\eta_{i,t+1} &\approx log(1 + \eta_{i,t+1}) \\
&= log\left(\frac{y_{i,t+1}}{y_{it}}\right) \\
&= (r_{1i,t+1} - r_{1it})h_1 + (r_{2i,t+1} - r_{2it})h_2 \ .
\end{aligned}
\tag{19}
$$

Thus, if we set $(r_{1i,t+1} - r_{1it})h_1 = \alpha_{t+1}\omega_{t+1}s_{1i,t+1} + \nu g$ and $(r_{2i,t+1} - r_{2it})h_2 = \alpha_{t+1}(1 - \omega_{t+1})s_{2i,t+1} + (1 - \nu)g$, where $\nu\epsilon[0,1]$, then we obtain (16). Notice that this identification implies $\sigma_{r1,t+1}^2 h_1^2 = \alpha_{t+1}^2\omega_{t+1}^2\sigma_{s1}^2$ and $\sigma_{r2,t+1}^2 h_2^2 = \alpha_{t+1}^2(1 - \omega_{t+1})^2\sigma_{s2}^2$, where we defined $\sigma_{r1,t+1}^2 = var[(r_{1i,t+1} - r_{1it})|S_{t+1}]$ and $\sigma_{r2,t+1}^2 = var[(r_{2i,t+1} - r_{2it})|S_{t+1}]$. In section IV.C we use these equalities to recover the parameter ω_{t+1} from knowledge of $\sigma_{1s}^2/\sigma_{2s}^2$, $\sigma_{r1,t+1}^2 h_1$, and $\sigma_{r2,t+1}^2 h_2$.

Let η_{iL} and η_{iH} stand for individual income shocks when $S = L$, respectively $S = H$. Given our specification of the outcome function (16) and the assumption that idiosyncratic shocks are normally distributed, $s_{1i} \sim N(0, \sigma_{1s}^2)$ and $s_{2i} \sim N(0, \sigma_{2s}^2)$, the two random variables η_{iL} and η_{iH} are jointly normally distributed:

$$
\begin{aligned}
(\eta_{iL}, \eta_{iH}) &\sim (g, g, \sigma_L^2, \sigma_H^2, \rho) \\
\sigma_L^2 &= \alpha_L^2\left(\omega_L^2\sigma_{1s}^2 + (1 - \omega_L)^2\sigma_{2s}^2\right)
\end{aligned}
\tag{20}
$$

$$\sigma_H^2 = \alpha_H^2 \left(\omega_H^2 \sigma_{1s}^2 + (1 - \omega_H)^2 \sigma_{2s}^2 \right)$$

$$\rho = \frac{\omega_L \omega_H \frac{\sigma_{1s}^2}{\sigma_{2s}^2} + (1 - \omega_L)(1 - \omega_H)}{\sqrt{\omega_L^2 \frac{\sigma_{1s}^2}{\sigma_{2s}^2} + (1 - \omega_L)^2} \sqrt{\omega_H^2 \frac{\sigma_{1s}^2}{\sigma_{2s}^2} + (1 - \omega_H)^2}} .$$

Note that to derive (20) we also used the assumption (15) that idiosyncratic shocks are independently distributed. Empirical work on labor income risk (Meghir and Pistaferri, 2004, Storesletten, et al., 2004) provides estimates of the variances σ_L^2 and σ_H^2, but we have no direct estimates of the correlation coefficient ρ. Thus, we can think of σ_L^2 and σ_H^2 as being determined by the labor income data, but the correlation coefficient ρ depends on the particular nature of idiosyncratic risk. If $\omega_L = \omega_H$, then there are no variations in the relative importance of the individual sources of idiosyncratic risk, and we are back to the case of one-dimensional idiosyncratic risk in the sense that idiosyncratic labor income shocks depend on the one-dimensional variable $\tilde{s}_i = \omega s_{1i} + (1 - \omega)s_{2i}$. This one-dimensional case is the one considered by the previous literature, and in our general framework it corresponds to a correlation coefficient of $\rho = 1$. In contrast, as long as we have cyclical variations across dimensions of risk, $\omega_L \neq \omega_H$, the correlation coefficient is strictly less than one: $\rho < 1$. Finally, notice that for any income parameters $(\sigma_L^2, \sigma_H^2, \rho)$, there are underlying parameters $(\alpha_L, \alpha_H, \omega_L, \omega_H, \sigma_{1s}^2, \sigma_{2s}^2)$ so that (20) is satisfied.

Applying the integration principle (10) to the random variable (η_{iL}, η_{iH}), we find $\eta_i' = \pi_L \eta_{iL} + \pi_H \eta_{iH}$. Thus, individual labor income risk in the economy without business cycles is again normally distributed, but now its variance is constant:

$$\eta_i' \sim N(g, \sigma'^2) \tag{21}$$
$$\sigma'^2 = \pi_L^2 \sigma_L^2 + \pi_H^2 \sigma_H^2 + 2\pi_L \sigma_L \pi_H \sigma_H \rho .$$

Notice that for the one-dimensional case, $\rho = 1$, we have $\sigma' = \pi_L \sigma_L + \pi_H \sigma_H$. That is, if idiosyncratic risk is one-dimensional, then the elimination of business cycles amounts to replacing the two normal distributions with state-dependent standard deviations by one nor-

mal distributions with a standard deviation that is equal to the mean of the state-dependent standard deviations. Note further that a decrease in ρ decreases the variance of individual labor income shocks in the economy without business cycles: $\frac{d\sigma'^2}{d\rho} > 0$. In other words, for given labor income risk in the economy with business cycles, defined by σ_L^2 and σ_H^2, business cycle variations across the individual sources of idiosyncratic risk reduce labor income risk in the economy without business cycles.

To provide some intuition for this risk-reduction effect, consider the special case in which $\pi_L = \pi_H = 1/2$, $\alpha_S = \alpha$, $\sigma_L^2 = \sigma_H^2 = \sigma^2$, $\omega_L = 1$, and $\omega_H = 0$. That is, during recessions $(S = L)$ only skill-type 1 is useful, and during booms $(S = H)$ only skill-type 2 is useful. Moreover, both skill types have the same amount of idiosyncratic income risk, and total labor income risk in the economy with business cycles is therefore constant and equal to σ^2. Clearly, any one-dimensional view of labor income risk would lead to the conclusion that the welfare cost of business cycles is zero. However, the two-dimensional view described above suggests that the welfare cost of business cycles is strictly positive, since the elimination of business cycles reduces the variance of labor income risk from σ^2 to $\sigma'^2 = 1/4\sigma^2 + 1/4\sigma^2 = 1/2\sigma^2$. Of course, this is simply a standard asset diversification result: in the economy with business cycles, the labor income of a worker at each point in time is determined by the random return to one asset only, whereas in the economy without business cycles, the labor income of the same worker is determined by the simple mean return of two assets.

Returning to the general case defined by (20) and (21), let us now evaluate the welfare formula (14) for the log-utility case using a second-order Taylor approximation for the functions $log(1 + \eta_i)$ and $log(1 + \eta_i')$ around $\eta_i = \eta_i' = 0$. Since $E[\eta_i] = E[\eta_i'] = g$, we find:

$$
\begin{aligned}
\Delta &= \exp\left(\frac{\beta}{1-\beta}\left(E[log(1+\eta'_i)] - E[log(1+\eta_i)]\right)\right) - 1 \qquad (22) \\
&\approx \exp\left(\frac{\beta}{1-\beta}\left(E[-\tfrac{1}{2}\eta_i'^2] - E[-\tfrac{1}{2}\eta_i^2]\right)\right) - 1
\end{aligned}
$$

$$= \exp\left(\frac{\beta}{2(1-\beta)}\left(\pi_L\sigma_L^2 + \pi_H\sigma_H^2 - \sigma'^2\right)\right) - 1,$$

where σ'^2 is given in (21). For the general case $\gamma \neq 1$, a similar reasoning shows:

$$\Delta = \left(\frac{1 - \beta\left(1 - \frac{\gamma(1-\gamma)}{2}(\pi_L\sigma_L^2 + \pi_H\sigma_H^2)\right)}{1 - \beta\left(1 - \frac{\gamma(1-\gamma)}{2}\sigma'^2\right)}\right)^{\frac{1}{1-\gamma}}. \tag{23}$$

As discussed before, a decrease in ρ decreases σ'^2. Thus, it follows from (22) and (23) that for given σ_L^2 and σ_H^2, a decrease in ρ increases the welfare cost of business cycles. Clearly, the welfare cost of business cycles (22) and (23) achieves its minimum if $\rho = 1$, which corresponds to the case of one-dimensional idiosyncratic risk. Moreover, the welfare cost of business cycles (22) and (23) achieves its maximum if $\rho = -1$, in which case the elimination of business cycles leads to the elimination of all income risk: $\sigma'^2 = 0$.

We summarize the preceding discussion in the following proposition:

Proposition 3. Suppose that the independence assumption (15) holds, that income shocks are normally distributed, and that welfare is calculated using a second-order Taylor approximation of utility functions. Then the welfare cost of business cycles, Δ, is given by (22) and (23). In particular, for a given marginal process of individual labor income in the economy with business cycles (defined by σ_L^2 and σ_H^2), the welfare cost of business cycles is a strictly decreasing function of ρ, where the parameter ρ measures the correlation between idiosyncratic income shocks during booms and idiosyncratic income shocks during recessions. If $\rho = 1$, then there are no cyclical fluctuations across the different dimensions of idiosyncratic risk ($\omega_L = \omega_H = \omega$), and idiosyncratic labor income risk is basically one-dimensional. If $\rho < 1$, then there are cyclical fluctuations across the different dimensions of idiosyncratic risk ($\omega_L \neq \omega_H$), and the welfare cost of business cycles is strictly larger than the minimum that is achieved in the one-dimensional case. Finally, if $\rho = -1$, eliminating business cycles amounts to eliminating all idiosyncratic risk, and the welfare cost of business cycles is maximal.

IV.B. Quantitative Analysis

In this section, we analyze the quantitative importance of the effects discussed in the previous section. For the sake of simplicity, we assume that both aggregate states are equally likely: $\pi_L = \pi_H = .5$. We choose an annual discount factors of $\beta = .95$, which is in line with previous work. For example, Imrohoroglu (1989) chooses $\beta = .96$, Storesletten et al. (2001) have $\beta = .95$, and Krusell and Smith (1999,2002) assume a stochastic discount factor with a mean of .95. At this stage, we will not try to pin down the correlation parameter ρ. Thus, it remains to specify the parameters σ_L and σ_H defining the marginal process of individual labor income in the economy with business cycles.

Equation (3) shows that η_i represents permanent income shocks since labor income follows (approximately) a logarithmic random walk with innovation term equal to η_i. The empirical literature on labor income risk has often used this random walk specification for the permanent component of labor income (Carroll and Samwick, 1997, Gourinchas and Parker, 2002, and Meghir and Pistaferri, 2004), or has used an AR(1) specification and then estimated a serial correlation coefficient close to one (Hubbard, Skinner, and Zeldes, 1994, and Storesletten et al., 2004). All these studies use annual income data drawn from the PSID data set, and estimate an average standard deviation of permanent income shocks between .15 and .18 for the pooled household/worker sample. This estimate of the average standard deviation corresponds to the expression $.5\sigma_L + .5\sigma_H$ in the theoretical model, and we choose σ_L and σ_H so that the value of this average is equal to .15 (the lower bound of the range of estimates).

Meghir and Pistaferri (2004) and Storesletten et al. (2004) are the only studies estimating cyclical variations in idiosyncratic labor income risk. More specifically, Meghir and Pistaferri (2004) estimate σ_t for each year t (see their table A4). After applying the Hodrick-Prescott filter to their time-series estimates of σ_t, we find that the standard deviation of the cyclical

component of σ_t is equal to .035. Storesletten et al. (2004) restrict σ_t to take on two values, and estimate $\sigma_L = .21$ and $\sigma_H = .12$, which corresponds to a standard deviation of .045 when both states are equally likely: $\pi_L = \pi_H = .5$. We choose values of σ_L and σ_H so that the implied standard deviation of $\sigma(S)$ is equal to .03, which is clearly a conservative value for the degree of cyclical variations in income risk. In short, we require:

$$.5\sigma_L + .5\sigma_H = .15 \tag{24}$$

$$\sqrt{.5(\sigma_L - .15)^2 + .5(\sigma_H - .15)^2} = .03$$

The restrictions (24) translate into the values $\sigma_L = .18$ and $\sigma_H = .12$.

Choosing $\sigma_{\eta L} = .18$ and $\sigma_{\eta H} = .12$ for labor income risk and using the formula (22) and (23) to calculate the welfare cost of business cycles is likely to overestimate these welfare costs for the following reason. In the model, households do not hold capital (financial wealth) and labor income shocks are permanent, which implies that the volatility of individual consumption growth, which is the main determinant of welfare, is equal to the volatility of individual labor income growth, σ_η. Krebs (2003a,b) considers a model with permanent labor income shocks in which households hold a positive amount of capital. He shows that (with log-utility) the volatility of individual consumption growth is equal to $\beta^2(\alpha^2\sigma_k^2 + (1-\alpha)^2\sigma_h^2)$, where α is the fraction of income that is capital income, σ_k^2 is the variance of changes in capital income, and σ_h^2 is the variance of the change in labor income. If we assume a capital share of $\alpha = .67$, a discount factor of $\beta = .95$, and risk-free capital, $\sigma_k^2 = 0$, then this suggests that we multiply our estimates of labor income risk by .63 when using the formula (22) and (23). This suggests to use as a benchmark model $\sigma_L = .63 * .18 = .1134$ and $\sigma_H = .63 * .12 = .0756$ resulting in a mean of $\sigma(S)$ of .0945 and a standard deviation of $\sigma(S)$ of .019.[18]

---

[18]For log-utility preferences, the implied price of a risk-free, one period bond is $Q_f = \beta E[(1 + \eta_i)^{-1}]$. Using an aggregate growth rate of $g = .02$ and a second-order Taylor approximation of $(1 + \eta_i)^{-1}$ around $\eta_i = 0$, we find an implied annual interest rate of 6.87%. In comparison, the corresponding complete-market economy would imply an annual interest rate of 7%. Notice also that the implied average consumption

The solid line in figure 1 shows the welfare cost of business cycles for different values of ρ using the formula (22) (quadratic approximation of log-utility preferences) and the benchmark choice of $\sigma_L = .1134$ and $\sigma_H = .0756$. The dotted line in figure 1 shows the results when the normal distribution is approximated by a two-state distribution using the exact welfare formula (14). Notice first that the two lines are very close together, which suggests that the second-order Taylor approximation of the welfare formula (14) is quite accurate. More importantly, we notice that the welfare cost of business cycles is very sensitive to relatively small changes in ρ. For example, if $\rho = 1$ (the one-dimensional case), then Δ is equal to .34 percent of lifetime consumption, but if $\rho = .66$ the welfare cost of business cycles increases to 1.78 percent of lifetime consumption. A value of $\rho = .66$ could result if $\omega_L = .67$, $\omega_H = .21$, and $\sigma_{1s} = \sigma_{2s}$ (see equation 20), and in section IV.C we will argue that this degree of cyclical variation in ω is quite realistic using estimates of U.S.labor market risk and a particular interpretation of the two dimensions of human capital risk. However, even if we choose a more moderate degree of variation in ω, for example $\omega = .59$ and $\omega = .30$, we still find a correlation coefficient that results in a welfare cost of business cycles that is substantial, namely one percent of lifetime consumption. Finally, we note that if $\rho = -1$, then eliminating business cycles amounts to eliminating all idiosyncratic labor income risk, and the welfare cost of business cycles becomes 8.86 percent of lifetime consumption. To summarize, realistic variations across the individual sources of idiosyncratic risk have substantial effects on the welfare cost of business cycles, and in principle such variations across dimensions of risk could increase the welfare cost of business cycles tremendously.[19]

---

volatility is equal to the income volatility, which has an annual standard deviation of .0945. Using the CEX data set, Brav et al. (2002) estimate an average individual consumption volatility that is somewhat larger. Attanasio (1999) argues that the consumption volatility is significantly smaller, but his conclusion is based on synthetic panel data sets ruling out by assumption truly idiosyncratic shocks to consumption.

[19]Figure 1 also shows that the welfare cost of business cycles is almost linear in ρ, which is not surprising given that formula (22) implies this linearity as long as $ln(1 + x) \approx x$ is a good approximation.

To compare our results with the case of complete markets (representative-agent economy), notice that in the current economy the welfare cost of business cycles is zero if markets are complete since there are no cyclical variations in aggregate income. If we introduce cyclical variations in aggregate income, then for the same preferences the welfare cost of business cycles increases to .05% of lifetime consumption (Lucas, 1987). Thus, introducing idiosyncratic risk into the analysis increases the welfare cost of business cycles significantly even in the one-dimensional case (from .05% to .34%), but introducing idiosyncratic risk with multiple dimensions has the potential to increase the welfare cost by a much larger amount (welfare costs around 1.5% are likely and 8.86% is the upper bound).

Figure 2 shows the welfare cost of business cycles for different degrees of risk aversion using the second-order Taylor approximation implicit in formula (23). More specifically, we choose $\gamma = .5$, $\gamma = 1$ (log-utility), and $\gamma = 1.5$. Notice first that an increase in the degree of risk aversion from $\gamma = .5$ to $\gamma = 1$ roughly doubles the welfare cost of business cycles. For example, if $\rho = 1$ then $\Delta = .17\%$ for risk aversion of $\gamma = .5$ and $\Delta = .34$ for risk aversion of $\gamma = 1$. If $\rho = .66$, then $\Delta = .87\%$ for a risk aversion parameter of $\gamma = .5$ and $\Delta = 1.78\%$ for a risk aversion parameter of $\gamma = 1$. Second, increasing the degree of risk aversion from $\gamma = 1$ to $\gamma = 1.5$ roughly increases the welfare cost of business cycles by 50 percent. For example, if $\rho = 1$ then $\Delta = .55$ for a risk aversion parameter of $\gamma = 1$ and if $\rho = .66$ the welfare cost becomes $\Delta = 2.85\%$ (as opposed to $\Delta = .34\%$, respectively $\Delta = 1.78\%$ if $\gamma = 1$).

Figure 3 reports the welfare cost of business cycles when there are no cyclical variations in total labor income risk: $\sigma_{\eta L} = \sigma_{\eta H} = .0935$. In this case, the welfare cost of business cycles is nil, $\Delta = 0$, if $\rho_\eta = 1$ (the one-dimensional case). However, even though there are no cyclical variations in total labor income risk, there can still be cyclical variations across dimensions of risk, and this will lead to positive welfare cost of business cycles. For example, if $\rho = .66$, the welfare cost is 1.45% percent of lifetime consumption with log-utility

preferences. Thus, even if the one-dimensional analysis suggests no welfare cost of business cycles, a multi-dimensional view of labor income risk implies a substantial welfare cost of business cycles.

Finally, figure 4 shows the welfare cost of business cycles as a function of ρ when $\sigma_L = .18$ and $\sigma_H = .12$, that is, when we do not make the adjustment that accounts for the fact that, contrary to the model, workers do owns some capital. In this case, the general points made before still hold, but overall the welfare cost of business cycles is substantially larger. For example, with log-utility preferences, the one-dimensional case ($\rho = 1$) leads to $\Delta = .86\%$, and for a correlation coefficient of $\rho = .66$ we find $\Delta = 4.44\%$.[20]

IV.C. Some Preliminary Evidence on ρ

In this section, we provide some preliminary evidence regarding the degree to which the U.S. labor market exhibits cyclical variations across two particular dimensions of idiosyncratic labor income risk yielding an estimate of the parameter ρ. It is worth pointing out that the evidence presented in this section is highly tentative, and that the resulting estimate of ρ is subject to a wide margin of error.

As in equation (18), suppose that total human capital of a worker is composed of two types of human capital, and that the income of each worker is the sum of the returns to the two different types of human capital. In contrast to the previous section, however, assume now that there are two groups of workers, one group that is only endowed with the first type of human capital (group 1) and one group that is only endowed with the second type of human capital (group 2). Denote by $\sigma_{1\eta}^2$, respectively $\sigma_{2\eta}^2$, the variance of the (permanent) income changes of group 1, respectively group 2. Using equation (16) and the assumptions

---

[20]For log-utility preferences, the implied annual interest rate becomes now 6.12% (using again $g = .02$ and a second-order Taylor approximation).

made so far, we find

$$\sigma_{1\eta t} = \alpha_t \omega_t \sigma_{1s} \tag{25}$$

$$\sigma_{2\eta t} = \alpha_t (1 - \omega_t) \sigma_{2s} \ .$$

Suppose we assume that $\sigma_{1s} = \sigma_{2s}$. Combining the two equations in (25) then yields

$$\omega_t = \frac{1}{1 + \frac{\sigma_{1\eta t}}{\sigma_{2\eta t}}} \ . \tag{26}$$

Thus, given estimates of $\sigma_{1\eta t}$ and $\sigma_{2\eta t}$, we can use (26) to back out estimates of ω_t, the parameter measuring the extent to which the relative importance of the two dimensions of idiosyncratic human capital risk varies of the cycle. Estimates of ω_t, in turn, allow us to estimate the correlation coefficient ρ, which is the main parameter of interest. We now present one application of this approach.

Suppose we identify the first type of human capital with specific human capital and the second type of human capital with general human capital. In order to use (26), we would need to estimate $\sigma_{\eta t}$ for two groups of workers, one group of workers with no specific human capital but some general human capital, and another group of workers with some specific human capital but no general human capital. Suppose, for example, that we proxy specific human capital by tenure at a firm and general human capital by education, and that we make the additional assumption that workers with (almost) no tenure have no specific human capital and workers with the lowest level of education have no general human capital. In this case, estimates of $\sigma_{\eta t}^2$ for the two groups of workers (group 1 with no education but some tenure and group 2 with no tenure but some education) allow us to use equation (26). Unfortunately, such estimates do no yet exists in the empirical labor literature, but there is something that at least comes close to it, namely the estimates reported in Meghir and Pistaferri (2004).

Meghir and Pistaferri (2004) estimate the variance of permanent changes in individual

28

income, $\sigma_{\eta t}^2$, for three different groups of workers: low level of education (high-school drop-outs), medium level of education (high school graduates), and high level of education (college graduates). If, as before, we identify education with general human capital, and if we also assume that high-school drop-outs have no endowment of general human capital, then we can identify the group of high-school drop-outs with group 1. Thus, the estimates of the variance of the permanent component of labor income risk obtained for the group of high-school drop-outs by Meghir and Pistaferri (2004) can be identified with estimates of $\sigma_{1\eta t}^2$ in (26). As a second reference group, we pick workers with at least a college degree (highest level of education), and assume that these workers have no specific human capital.[21] In this case, we can identify the group of college graduates with group 2, and the estimates of the variance of the permanent component of labor income risk obtained for this group can be identified with estimates of $\sigma_{2\eta t}^2$ in (26). In short, this particular interpretation provide us with estimates of $\sigma_{1\eta t}^2$ and $\sigma_{2\eta t}^2$ that can be used to back out ω_t in (26).

We use the estimates of ω_t thus obtained to find values for ω_L and ω_H as follows. We first calculate the sample mean of the time series estimates of ω_t yielding $\frac{1}{T}\sum_t \omega_t = .44$. We then apply the Hodrick-Prescott filter to the time series of ω_t and calculate the standard deviation of the cyclical component of ω_t, which yields .23. Finally, we choose ω_L and ω_H so that the random variable ω has a mean equal to .44 and a standard deviation equal to .23. That is, we choose values ω_L and ω_H solving

$$\frac{1}{2}\omega_L + \frac{1}{2}\omega_H = .44 \tag{27}$$

$$\sqrt{\frac{1}{2}(\omega_L - .43)^2 + \frac{1}{2}(\omega_H - .43)^2} = .23 .$$

Equation (27) leads to $\omega_L = .21$ and $\omega_H = .67$. These estimates are then substituted into (20) in order to find ρ (using again $\sigma_{1s} = \sigma_{2s}$). This yields $\rho = .66$, which is the value of ρ

---

[21]Of course, this assumption is an extreme version of the somewhat more moderate assumption that most of the human capital of college-educated workers is general human capital.

considered "realistic" in the discussion in section IV.B.[22]

At this stage, it seems worth pointing out two implications of the specific approach taken in this section that are supported by the data, which provides indirect evidence in favor of the assumptions made so far. First, given that high-school drop-outs have no general human capital and college graduates have no specific human capital, it must be true that workers in the "middle-group", namely those workers with only a high-school education, are endowed with both general and specific human capital. Thus, income for this group of workers is the sum of the return to two types of human capital, whereas income of workers belonging to the other two groups is the return to only one type of human capital. This, in turn, suggests that total labor income risk for workers with high-shool education is smaller than for workers with either no high-school education or college education since the high-school educated workers have a more diversified portfolio. Indeed, using the estimates of Meghir and Pistaferri (2004), we find that the average variance of permanent labor income shocks, $\frac{1}{T}\sum_t \sigma^2_{\eta t}$, is equal to .0297 for high-school-educated workers, but equal to .0389, respectively .0423, for high-school drop-outs, respectively college-educated workers.

A second implication of our approach is that labor income risk of high-school drop-out and labor income risk of college-educated workers should behave differently over the business cycle. Indeed, it turns out that the correlation between the cyclical component of $\sigma^2_{1\eta t}$ (labor income risk of high-school drop-outs) and $\sigma^2_{2\eta t}$ (labor income risk of college-educated workers) is negative (correlation coefficient of $-.24$). Notice that the one-dimensional approach to labor income and labor income risk would predict a correlation coefficient of $+1$. Moreover, the correlation coefficient between the cyclical components of $\sigma^2_{1\eta t}$ and real (annual) GDP

---

[22]Clearly, to the extent that $\sigma^2_{\eta t}$ is estimated with error, this procedure overestimates the degree of variations in ω_t, and therefore underestimates ρ. Note, however, that a similar criticism applies to Krebs (2003a) and Storesletten et al. (2001) when they discuss the welfare cost of business cycles due to cyclical variations in (one-dimensional) labor income risk since neither papers makes an adjustment for the fact that the values of σ^2_η are estimated with error.

is $-.10$, whereas the correlation coefficient between the cyclical components of $\sigma^2_{2\eta t}$ and real (annual) GDP is $+.31$. That is, for high-school drop outs the variance of permanent income changes goes up when economic growth is sluggish, whereas for college-educated workers this variance increases during economic booms.[23] Thus, the labor income risk of these two groups of workers, which by assumption are equipped with very different types of human capital, exhibit very different patterns over the business cycle, which according to the current interpretation is the result of cyclical variations across the two dimensions of human capital risk.

Finally, let us note that the welfare cost of business cycles calculated in section IV.B using $\rho = .66$ (as opposed to $\rho = 1$) does not apply to workers with no high-school education (group 1) or college education (group 2) since we have assumed that workers belonging to either one of these two groups are only endowed with one type of human capital. In other words, these workers will not reap the additional benefits that come from the elimination of cyclical variations across different dimensions of human capital risk since we have assumed that their human capital endowment, and therefore their human capital risk, is always one-dimensional. However, the middle-group of workers with only high-school education is truly two-dimensional, and will therefore gain from the elimination of cyclical variations across the two dimensions of human capital risk. Thus, using the particular interpretation of multi-dimensional risk put forward in this section, the welfare cost of business cycles calculated in section IV.B using $\rho = .66$ would only apply to workers belonging to the middle-group (high-school education only).

---

[23]Although the positive correlation between income variance and economic growth for college-educated workers seems somewhat counter-intuitive, it could be explained by the fact that positive income "shocks" become more likely and larger during booms, and that this effect is more pronounced for college-educated workers.

V. Conclusion

This paper analyzed the welfare costs of business cycles when workers face uninsurable idiosyncratic labor income risk. In contrast to the previous literature, this paper introduced multiple sources of idiosyncratic labor market risk, and used this approach to provide a general characterization of the welfare cost of business cycles when preferences and the marginal process of individual labor income in the economy with business cycles are given. The general analysis showed that the introduction of multiple sources of idiosyncratic risk never decreases the welfare cost of business cycles, and strictly increases it if there are cyclical variations across the different sources of risk. A quantitative analysis based on a calibrated version of the model showed that the welfare cost of business cycles due multi-dimensional idiosyncratic risk is likely to be substantial. More specifically, depending on the magnitude of cyclical variations across different dimensions of idiosyncratic labor income risk, the welfare cost of business cycles could be anything between 0.3 (the one-dimensional case) and 9 percent of lifetime consumption (for log-utility preferences). Moreover, for "realistic" cyclical variations across dimensions of labor market risk, the welfare cost of business cycles is 1.5% of lifetime consumption.

Clearly, the last result of welfare cost of business cycles of 1.5% of lifetime consumption is highly tentative in the sense that it relies on very preliminary evidence on the magnitude of cyclical variations across different dimensions of idiosyncratic labor income risk. In other words, in order to arrive at firm conclusions, much more empirical work is needed that identifies particular dimensions of human capital risk and then estimates the magnitude of variations across the individual dimensions over the business cycle. However, even though this paper does not provide final answers, it does point towards an important new channel through which idiosyncratic risk affects the welfare cost of business cycles, and it provides a simple framework that can be used to translate (future) empirical results about cyclical

variations across dimensions of labor market risk into precise statements about the welfare cost of business cycles.

Appendix: Job Displacement

Atkeson and Phelan (1994), Imrohoroglu (1989), and Krusell and Smith (1999, 2002) study a model with an individual income process that is Markov and has a two-state support. Their interpretation of the low-income realization is that the worker is unemployed. In this appendix, we consider a model that is similar in the sense that there are only two possible realizations of η, namely 0 and $-\bar{\eta}$. We interpret the event $\eta_i = -\bar{\eta}$ as job displacement (low income level) and the event $\eta_i = 0$ as no job displacement (high income level). Clearly, an important difference between the previous literature and the current model is that the previous literature has emphasized the (temporary) income losses during the unemployment period, whereas the current paper focuses on the (permanent) loss of income that persists even after the displaced worker has found a new job. However, for the main issue analyzed in this paper, namely the effect of multiple sources of idiosyncratic risk on the welfare cost of business cycles, this difference does not seem to be of first-order importance.

As in section IV, we assume that there are two aggregate states, $S = L, H$, that occur with probability π_L, respectively π_H. We also assume that the idiosyncratic shock variable, s_i, has two components: $s_i = (s_{i1}, s_{i2})$, and that the independence assumption holds: $\pi(s_{i1}, s_{i2}, S) = \pi(s_{i1})\pi(s_{i2})\pi(S)$. Finally, we also make the assumption that idiosyncratic shocks are normally distributed: $s_{1i} \sim N(0, \sigma_{1s}^2)$ and $s_{2i} \sim N(0, \sigma_{2s}^2)$.

Define the following function:

$$f(s_{1i}, s_{2i}, S) = \alpha_S \left(\omega_S s_{1i} + (1 - \omega_S)s_{2i} \right) . \tag{28}$$

Given the assumptions made so far, the random variables f_{iL} and f_{iH} are jointly normally distributed, $(f_{iL}, f_{iH}) \sim N(0, 0, \sigma_L^2, \sigma_H^2, \rho)$, where the parameters σ_L^2, σ_H^2, and ρ are defined as in (20), and the choice $\rho = 1$ again corresponds to the one-dimensional case. In contrast

to section IV, however, we do not identify the outcome function $\eta(.)$ with $f(.)$, but assume

$$\eta(s_{1i}, s_{2i}, S) = \begin{cases} -\bar{\eta} & if \ f(s_{1i}, s_{2i}, S) \leq c_S \\ 0 & if \ f(s_{1i}, s_{2i}, S) > c_S \ , \end{cases} \qquad (29)$$

where c_L and c_H are two real numbers. Notice that even though $f(.)$ is a linear function in s_{1i} and s_{2i}, the outcome function $\eta(.)$ is not. One interpretation of this outcome function is that worker i becomes displaced once the productivity drops below the threshold value c_S. Notice that for any S, the outcome function is monotone in s_{1i} and s_{2i}, something we would expect to hold for most optimal decision rules.

Denote the probability of job displacement in state S by p_S. Given our specification of the outcome function, we have:

$$prob(f_{iL} \leq c_H) = p_H \qquad (30)$$
$$prob(f_{iH} \leq c_L) = p_L \ .$$

Clearly, we can always find values of σ_L^2, σ_H^2, c_L, and c_H so that (30) holds. Notice also that the correlation coefficient ρ can be varied without changing the displacement probabilities p_L and p_H. We will now show how a decrease in ρ increases the welfare cost of business cycles.

The application of the integration principle yields a new labor income process defined by probabilities $\pi'(s_i) = \pi(s_i)$ and an outcome function:

$$\eta'(s_{1i}, s_{2i}) = \begin{cases} \eta_1' & if \ f(s_{1i}, s_{2i}, L) \leq c_L \ and \ f(s_{1i}, s_{2i}, H) \leq c_H \\ \eta_2' & if \ f(s_{1i}, s_{2i}, L) > c_L \ and \ f(s_{1i}, s_{2i}, H) \leq c_H \\ \eta_3' & if \ f(s_{1i}, s_{2i}, L) \leq c_L \ and \ f(s_{1i}, s_{2i}, H) > c_H \\ \eta_4' & if \ f(s_{1i}, s_{2i}, L) > c_L \ and \ f(s_{1i}, s_{2i}, H) > c_H \end{cases}$$

where the realizations of the new random variable η' are

$$\eta_1' = \pi_L * (-\bar{\eta}) + \pi_H * (-\bar{\eta}) = -\bar{\eta}$$
$$\eta_2' = \pi_L * 0 + \pi_H(-\bar{\eta}) = -\pi_H \bar{\eta}$$

$$\eta'_3 = \pi_L * (-\bar{\eta}) + \pi_H * 0 = -\pi_L \bar{\eta}$$

$$\eta'_4 = \pi_L * 0 + \pi_H * 0 = 0 .$$

Notice that the four-state support, $\{\eta'_1, \eta'_2, \eta'_3, \eta'_4\}$, of the random variable η' does not depend on the structure of the state space. More specifically, changes in σ^2_{1s}, σ^2_{2s}, and α_S, which correspond to changes in σ^2_L, σ^2_H, and ρ, do not affect the support $\{\eta'_1, \eta'_2, \eta'_3, \eta'_4\}$ of the random variable η'_i. However, these variations do change the probabilities with which the individual realizations of η'_i occur, which in turn changes the welfare cost of business cycles.

Let us introduce the following notation:

$$prob(\eta'_i = \eta'_1) = prob\,(f_{iL} \leq c_L \text{ and } f_{iH} \leq c_H) = q_1 \tag{31}$$

$$prob(\eta'_i = \eta'_2) = prob\,(f_{iL} > c_L \text{ and } f_{iH} \leq c_H) = q_2$$

$$prob(\eta'_i = \eta'_3) = prob\,(f_{iL} \leq c_L \text{ and } f_{iH} > c_H) = q_3$$

$$prob(\eta'_i = \eta'_4) = prob\,(f_{iL} > c_L \text{ and } f_{iH} > c_H) = q_4 = 1 - q_1 - q_2 - q_3 .$$

Notice that the above construction partitions the idiosyncratic state space into four regions. The parameter q_1 is the probability that the worker is displaced regardless of the aggregate state, q_2 is the probability that the worker is displaced during a boom but not during a recession, q_3 is the probability that the worker is displaced during a recession but not during a boom, and q_4 is the probability that the worker is not displaced regardless of the aggregate state.

Given the above outcome function $\eta(.)$, we have the following relationship between the probabilities q_1, q_2, and q_3 and the observed displacement probabilities p_L and p_H:

$$p_H = q_1 + q_2 \tag{32}$$

$$p_L = q_1 + q_3 ,$$

For given p_L and p_H, this imposes two equality restrictions on the three probabilities q_1, q_2, and q_3. We also have the equality restrictions $1 - p(H) = q_2 + q_4$ and $1 - p(L) = q_3 + q_4$, but

these restrictions automatically hold if $p(H) = q_1 + q_2$ and $p(L) = q_1 + q_3$. Put differently, (32) is an exhaustive representation of all linearly independent equality restrictions. Thus, we can freely choose q_1, and then pick the remaining probabilities q_2 and q_3 so as to match the observed displacement probabilities $p(L)$ and $p(H)$ in the sense that (32) is satisfied.

Clearly, when varying q_1, we have to satisfy additional inequality restrictions since probabilities have to lie in $[0, 1]$. Thus, we need $0 \leq q_1 \leq 1$. Moreover, the conditions $q_2 \geq 0$ and $p_H \geq 0$ imply $q_1 \leq p_H$, and the conditions $q_3 \geq 0$ and $p_L \geq 0$ imply $q_1 \leq p_L$. Finally, $q_4 = 1 - q_1 - q_2 - q_3 \geq 0$ implies $q_1 \geq p_L + p_H - 1$. If we assume $p_L > p_H$ (displacement probabilities increase during economic contractions) and $p_L + p_H \leq 1$ (average displacement probabilities are less than $1/2$), then all the inequalities taken together result in the restriction $q_1 \epsilon [0, p_H]$.

Notice also that for fixed σ_L^2, σ_H^2, c_L, and c_H, any choice of q_1 correspond to a particular value of the correlation coefficient ρ, so that we can write $q_1 = q_1(\rho)$. If $\rho = 1$, we have $q_1 = p_H$, the maximum value the free parameter q_1 can take on and still be consistent with the data. If $\rho = -1$, we have $q = 0$, the minimum value the free parameter q_1 can take and still be consistent with the data. Finally, it follows from the properties of joint-normally distributed random variables that q_1 is monotonically increasing in ρ: $\frac{dq_1}{d\rho} > 0$.

To save space, assume now that the utility function is of the logarithmic type. If we apply the general welfare formula (14) to this case taking into account the previous discussion, then we find the following welfare expression:

$$
\begin{aligned}
\Delta &= \exp\left(\frac{\beta}{1-\beta}x\right) - 1 \\
x &= q_1(\rho)log(1 - \bar{\eta}) + (p_H - q_1(\rho))log\left(1 - \pi_H \bar{\eta}\right) + (p_L - q_1(\rho))log(1 - \pi_L \bar{\eta}) \\
&\quad - ((\pi_L p_L + \pi_H p_H)log(1 - \bar{\eta}))) \ .
\end{aligned}
\tag{33}
$$

Equation (33) defines a linear function $\Delta = \Delta(\rho)$. It immediately follows from (33) that

$\Delta(\rho) > 0$ if and only if $p_L \neq p_H$. Since $log(.)$ is a strictly concave function, differentiation of (33) yields

$$\frac{d\Delta}{d\rho} = (log(1 - \bar{\eta}) - log(1 - \pi_H \bar{\eta}) - log(1 - \pi_L \bar{\eta})) \frac{dq_1}{d\rho} < 0 \ . \tag{34}$$

Equation (34) in conjunction with $\rho \epsilon [-1, 1]$ implies that the minimum cost of business cycles is achieved at $\rho = +1$. Since $\rho = +1$ is the one-dimensional case, we have shown that the introduction of additional sources of idiosyncratic risk increases the welfare costs of business cycles. Thus, from a qualitative point of view, the results obtained in this appendix are identical to the results obtained in section IV, with the only difference that the maximum welfare cost of business cycles, which is again achieved when $\rho = -1$, is in general not equal to the welfare cost of all idiosyncratic risk (see equation 33). In summary, we have the following version of proposition 3:

Proposition 3'. Suppose that idiosyncratic risk is defined by (28) and (29). Then the welfare cost of business cycles, Δ, is given by (33). In particular, for a given marginal process of individual labor income in the economy with business cycles defined by σ_L^2 and σ_H^2, the welfare cost of business cycles is a strictly decreasing function of ρ, where the parameter ρ measures the correlation between idiosyncratic productivity shocks during booms and idiosyncratic productivity shocks during recessions. If $\rho = 1$, then there are no cyclical variations across the different dimensions of idiosyncratic risk ($\omega_L = \omega_H = \omega$), and idiosyncratic labor income risk is basically one-dimensional. If $\rho < 1$, then there are cyclical variations across the different dimensions of idiosyncratic risk ($\omega_L \neq \omega_H$), and the welfare cost of business cycles is strictly larger than the minimum that is achieved in the one-dimensional case.

References

Attanasio, O. (1999) "Consumption,", In: Taylor, J., and M. Woodford (Eds.), *Handbook of Macroeconomics.* North-Holland.

Alvarez, F., and U. Jermann (2003) "Using Asset Prices to measure the Cost of Business Cycles," Forthcoming, *Journal of Political Economy.*

Atkeson, A., and C. Phelan (1994) "Reconsidering the Cost of Business Cycles with Incomplete Markets," *NBER Macroeconomics Annual*, MIT Press, Cambridge, MA, pp. 187-207.

Beaudry, P., and C. Pages (2001) "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance," *European Economic Review*, 45: 1545-1572.

Brav, A., Constantinides, G., and C. Geczy (2002) "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy* 110: 793-825.

Carroll, C., and A. Samwick (1997) "The Nature of Precautionary Wealth," *Journal of Monetary Economics* 40: 41-72.

Castanedas, A., Diaz-Gimenez, J., and V. Rios-Rull (2003) "Accounting for the Earnings and Wealth Distribution," *Journal of Political Economy* 111: 818-857.

Constantinides, G. and D. Duffie (1996) "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104: 219-240.

Gomez, J., Greenwood, J., and S. Rebelo (2001) "Equilibrium Unemployment," *Journal of Monetary Economics* 48: 109-152.

Gourinchas, P., and J. Parker (2002) "Consumption over the Life Cycle," *Econometrica*

Hall, R. (1995) "Lost Jobs," *Brookings Papers on Economic Activity*, 1: 221-273.

Hubbard, G., Skinner, J., and S. Zeldes (1994) "The Importance of Precautionary Motives in Explaining Individual and Aggregate Savings," *Carnegie-Rochester Conference Series on Public Policy* 40: 59-126

Huggett, M., (1996) "Wealth Distribution in Life-Cycle Economies," *Journal of Monetary Economics* 38: 469-494.

Imrohoroglu, A. (1989) "Costs of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 1364-83.

Jacobson, L., LaLonde, R., and D. Sullivan (1993) "Earnings Losses of Displaced Workers," *American Economic Review* 83: 685-709.

Krebs, T. (2003a) "Growth and Welfare Effects of Business Cycles in Economies with

Idiosyncratic Human Capital Risk," *The Review of Economic Dynamics* 6: 846-868.

Krebs, T. (2003b) "Human Capital Risk and Economic Growth," *The Quarterly Journal of Economics* 118: 709-744.

Krebs, T. (2004) "Testable Implications of Consumption-Based Asset Pricing Models with Incomplete Markets," *Journal of Mathematical Economics*, 40: 191-206.

Krusell, P., and A. Smith (1999) "On the Welfare Effects of Business Cycles," *Review of Economic Dynamics*, 2: 247-272.

Krusell, P., and A. Smith (2002) "Revisiting the Welfare Effects of Eliminating Business Cycles," Working Paper, Carnegie-Mellon University.

Lucas, R. *Models of Business Cycles.* New York: Blackwell, 1987.

Lucas, R. (2003) "Macroeconomic Priorities," *American Economic Review* 93: 1-14.

Meghir, C. and L. Pistaferri (2004) "Income Variance Dynamics and Heterogeneity," *Econometrica*

Mincer, J., *Schooling, Experience, and Earnings.* New York: Columbia University Press, 1974.

Storesletten, K., Telmer, C., and A. Yaron (2001) "The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variations in Idiosyncratic Risk," *European Economic Review.*

Storesletten, K., Telmer, C., and A. Yaron (2004) "Cyclical Dynamics of Labor Market Risk," Working Paper, *Journal of Political Economy*

Figure 1. Cost of business cycles for log-utility and normally distributed income shocks with (solid line) and without (dashed line) Taylor approximation.
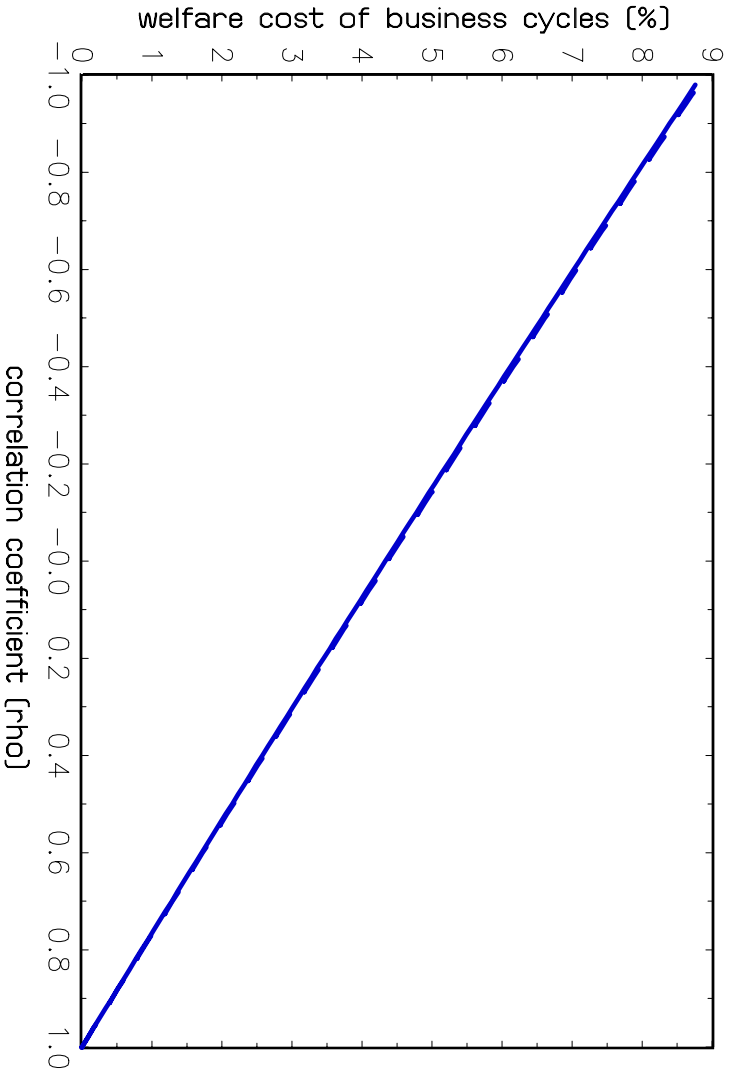
Figure 2. Cost of business cycles for benchmark labor income risk and risk aversion of 1 (solid line), 1.5 (dashed line), and .5 (dotted line).
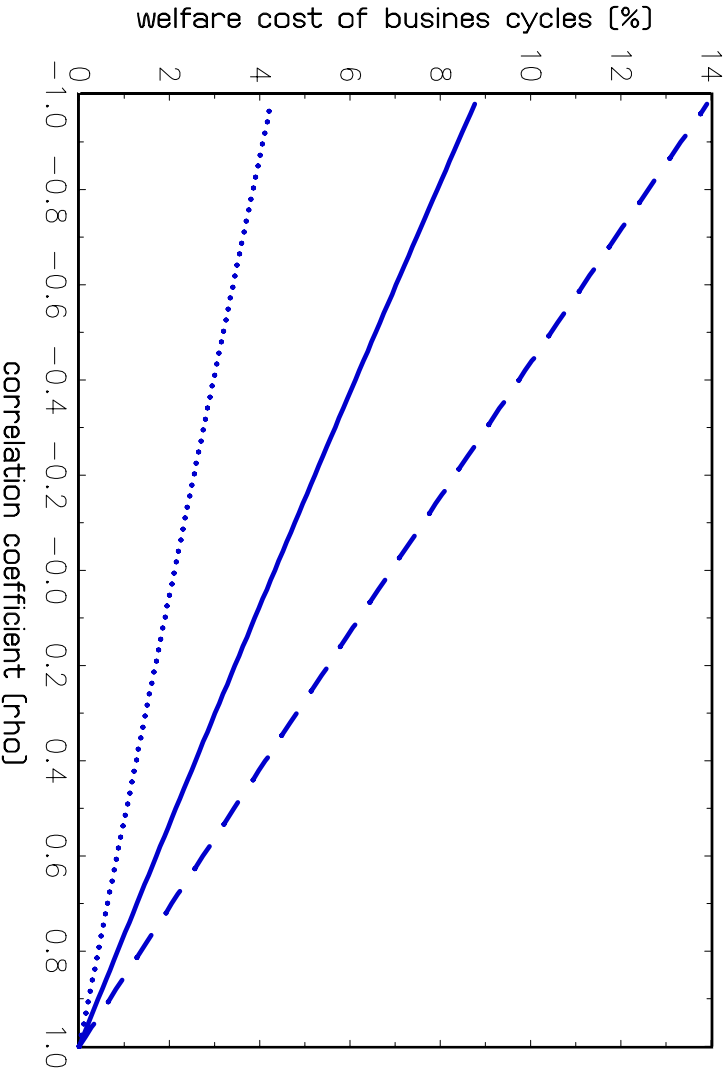
Figure 3. Cost of business cycles for constant labor income risk and risk aversion of 1 (solid line), 1.5 (dashed line), and .5 (dotted line).
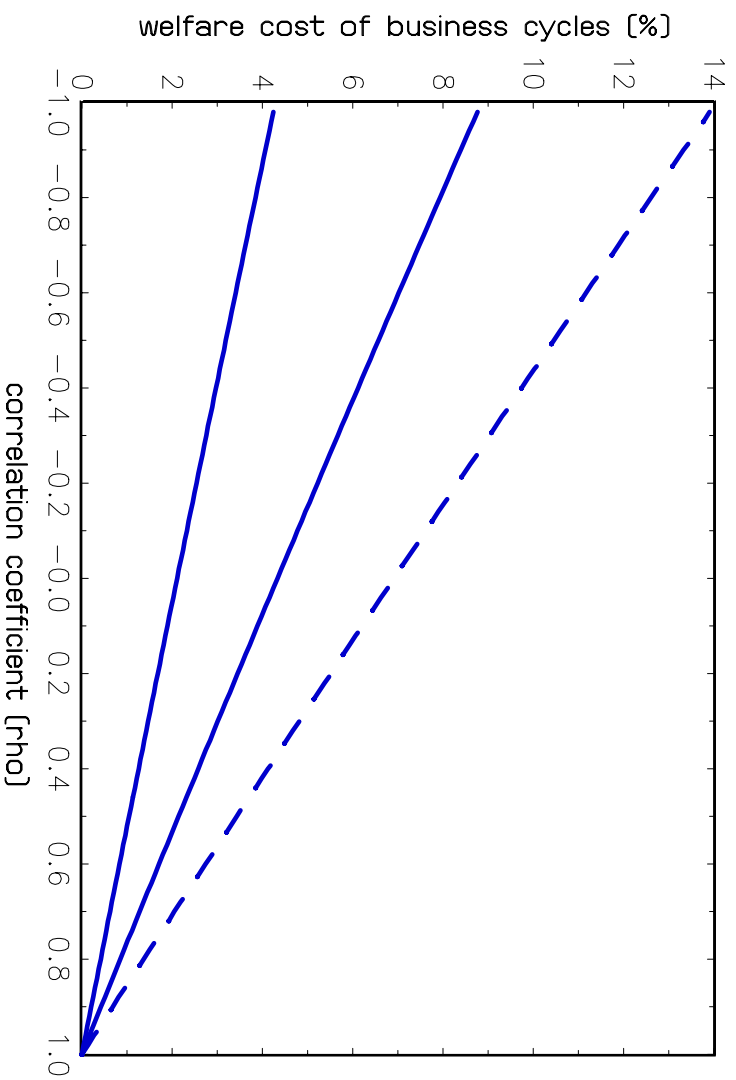
Figure 4. Cost of business cycles for high degree of labor income risk and risk aversion of 1 (solid line), 1.5 (dashed line), and .5 (dotted line).