Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 55 (2008) 501–525
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Insurance and opportunities: A welfare analysis of labor market risk ☆

Jonathan Heathcote[a,b,c,f], Kjetil Storesletten[d,*], Giovanni L. Violante[e,f,g]

[a]Federal Reserve Bank of Minneapolis, USA
[b]Federal Reserve Board, USA
[c]Georgetown University, USA
[d]Department of Economics, University of Oslo, Box 1095 Blindern, N-0317 Oslo, Norway
[e]New York University, USA
[f]Centre for Economic Policy Research, UK
[g]National Bureau of Economic Research, USA

Received 5 July 2006; received in revised form 13 February 2008; accepted 14 February 2008
Available online 29 February 2008

Abstract

Using a model with constant relative risk-aversion preferences, endogenous labor supply and partial insurance against idiosyncratic wage risk, this paper provides an analytical characterization of three welfare effects: (a) the welfare effect of a rise in wage dispersion, (b) the welfare gain from completing markets, and (c) the welfare effect from eliminating risk. The analysis reveals an important trade-off for these welfare calculations. On the one hand, higher wage uncertainty increases the cost associated with missing insurance markets. On the other hand, greater wage dispersion presents opportunities to raise aggregate productivity by concentrating market work among more productive workers. Welfare effects can be expressed in terms of the underlying parameters defining preferences and wage risk or, alternatively, in terms of changes in observable second moments of the joint distribution over individual wages, consumption and hours.

JEL classification: D31; D58; D91; E21; J22; J31

Keywords: Insurance; Labor supply; Productivity; Wage dispersion; Welfare

☆A predecessor of this paper circulated as "Insurance and opportunities: The welfare implications of rising wage dispersion." We would particularly like to thank Randall Wright. Several ideas in this paper originated from a conversation we had with him. We also thank Orazio Attanasio, Marco Bassetto, Jeremy Greenwood, Guido Lorenzoni, Victor Ríos-Rull, one anonymous referee, and participants at many seminars and conferences for comments. Heathcote and Violante thank the Economics Program of the National Science Foundation (Grant SES 0418029) for financial support. Storesletten—a fellow of the CEPR—thanks the Norwegian Research Council for financial support. The opinions expressed here are those of the authors and not necessarily those of the Board of Governors of the Federal Reserve System or its staff.
*Corresponding author. Tel.: +47 2284 4009; fax: +47 2285 5035.
E-mail address: kjetil.storesletten@econ.uio.no (K. Storesletten).

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2008.02.001

1. Introduction

Cross-sectional wage dispersion and individual wage volatility over the life-cycle are large. For example, the variance of the growth rate of individual wages in the United States in the cross-section is over *100 times* larger than the variance of the growth rate of average wages over time.[1] Moreover, there has been a sharp increase in wage dispersion in the United States over the past 30 years.[2] An important task for macroeconomists is to study the welfare consequences of this phenomenon.

This paper develops a tractable class of dynamic heterogeneous-agent economies with partial insurance against idiosyncratic labor productivity (wage) risk and with endogenous labor supply. The process for idiosyncratic wages has two orthogonal components: an uninsurable piece, and a component that may be fully insured. This assumption implies that the equilibrium allocations can be solved for analytically which, in turn, permits a transparent welfare analysis.

Several authors have examined the welfare consequences of changes in earnings or income risk.[3] This paper focuses instead on *wage-rate risk* given flexible labor supply. Endogenizing labor supply is important because the ability to adjust hours can mitigate the welfare cost of rising wage dispersion via two alternative channels. First, agents may use hours worked to mitigate fluctuations in earnings by increasing (decreasing) labor supply when wages fall (rise). Alternatively, agents may choose to work more hours in periods when individual wages are high, thereby increasing average earnings per hour. A negative wage–hour correlation is more likely to be observed if agents cannot smooth income by other means, such as by purchasing explicit insurance against wage risk. Conversely, the wage–hour correlation will be positive if wage inequality can be insured directly within financial markets. Thus, the model highlights an interesting interaction between the asset market structure and the role of endogenous labor supply in absorbing idiosyncratic wage shocks.[4]

The paper considers two standard classes of time-additive preferences with constant relative risk aversion; one where agents have CRRA preferences over a Cobb–Douglas composite of consumption and leisure, and one where preferences are additively separable between consumption and hours worked. For each preference specification, explicit analytical solutions for equilibrium allocations are provided. Individual consumption and hours are shown to be log-linear in the insurable and uninsurable components of individual labor productivity. When these components are assumed to be log-normally distributed, expected lifetime utility can be expressed as a tractable function of preference parameters and the variances of the insurable and uninsurable components of the wage process. Using these expressions, three distinct questions about welfare and inequality are addressed.[5]

First, what are the *welfare effects of rising wage dispersion*, holding constant the asset market structure? Second, what are the *welfare costs of market incompleteness*, defined as the difference between expected lifetime utility in the baseline incomplete-markets economy versus a complete-markets economy, holding constant the wage-generating process? Third, what are the *welfare effects from eliminating individual wage risk*? This last welfare calculation is the cross-sectional equivalent of the calculation underlying the large literature on the welfare costs of business cycles fluctuations (for a survey, see Lucas, 2003). Note that these three inquiries reflect changes in different primitives of the model: technology in the first (e.g., skill-biased technical change), markets in the second (e.g., the emergence of new financial instruments), and policies in the third (e.g., redistributive taxation schemes that align ex post wages across all workers).

When labor supply is flexible, increased wage inequality impacts not just consumption inequality, but also leisure inequality and the average values for consumption and leisure. More precisely, welfare effects are driven by two offsetting forces: an increase in idiosyncratic wage risk increases the need for *insurance*, but also presents *opportunities* to increase the level of aggregate productivity, measured as output per hour worked, by

[1] This number is calculated from the PSID, 1967–1996. The variance of the mean wage growth over the period is 0.0012 and the cross-sectional variance of individual wage growth, averaged over the period, is 0.161. See Section 7 for details on the sample selection.

[2] For surveys on the causes of the changes in inequality, see Katz and Autor (1999), Acemoglu (2002), Aghion (2002), and Hornstein et al. (2005).

[3] See e.g. Attanasio and Davis (1996), Blundell and Preston (1998), Krueger and Perri (2006), and Krebs et al. (2005).

[4] Low (2005) explores the implications of this interaction for the life-cycle profiles of consumption, hours and asset holdings.

[5] The paper focuses on inequality in individual labor productivities within a competitive labor market. Thus, our analysis abstract from "frictional inequality"—pure wage dispersion arising between ex ante identical workers because of search frictions (e.g., Mortensen, 2003).

concentrating the work effort among more productive workers. To clarify the trade-off between risk and opportunities, the overall welfare effects are decomposed into the relative contributions of changes in aggregate consumption and leisure on the one hand, and changes in the cross-sectional dispersion of these variables on the other (see also Benabou, 2002; Flodén, 2001).

A related point is that there is an important difference between *insuring* risk and *eliminating* risk when labor supply is flexible. In fact, eliminating risk will always lead to smaller welfare gains than insuring risk, because removing risk also takes away opportunities to increase average labor productivity.

The results discussed so far pertain to welfare expressions incorporating structural model parameters defining preferences and the insurability of wage risk. As a complement to this approach, the paper provides an alternative set of expressions for the welfare effects. This alternative route expresses welfare effects as simple functions of various moments of the cross-sectional joint distribution over wages, hours and consumption. For example, in the separable-preferences case, the welfare effect associated with a change in the wage process can be expressed as the sum of the changes in (i) the covariance between log wages and log hours, (ii) the variance of log-consumption weighted by the coefficient of relative risk aversion, and (iii) the variance of log-hours weighted by the inverse of the labor supply elasticity.

This representation of welfare effects has two advantages relative to the first set of welfare expressions, which were based on a structural model. First, it is more general, since it does not depend on the particular market structure assumed. In particular, the expression applies to any economy where the standard intratemporal consumption–leisure first-order condition is satisfied, and where equilibrium allocations and wages are jointly log-normally distributed. The second advantage is that one does not have to take a stand on the fraction of wage risk that is insurable. Thus, welfare effects can be estimated simply by computing the relevant moments in repeated cross-sections and assigning values to preference parameters. However, a drawback with the cross-sectional-moment-based representation is that it requires high-quality data on consumption and hours, while the first approach only requires panel data on wages. Therefore, we view the two alternative approaches as complementary. In the rest of the paper, the first set of welfare expressions will be referred to as "model-based" and the second as "observables-based".

The quantitative part of the paper provides answers to the three welfare questions by calibrating the model to the U.S. economy. With Cobb–Douglas preferences and a coefficient of relative risk aversion equal to two, the welfare cost of the rise in labor market risk in the U.S. over the past 30 years in the partial-insurance economy is 2.5% of lifetime consumption. This number is the combination of a welfare loss of 7.5% due to larger uninsurable fluctuations in individual consumption and hours, and a welfare gain of 5% from an increase in aggregate labor productivity.

For the same preferences, households would be willing, ex ante, to give up almost 40% of their expected lifetime consumption in exchange for access to complete markets. One might suspect that this welfare gain stems from reducing inequality in the cross-sectional distributions for consumption and leisure. Instead, it turns out that two thirds of the welfare gains from completing markets stems from increasing average productivity. Thus, the analysis highlights an important cost of missing markets that has largely been overlooked to date, namely the loss in aggregate labor productivity that arises when low-productivity agents work too much (because lack of insurance makes them inefficiently poor), while high-productivity agents work too little (because lack of insurance makes them inefficiently rich).

Finally, the analysis suggests that eliminating all individual wage risk through distortionary taxation delivers a welfare gain which is only about half the size of the gain from completing markets, but nevertheless more than two orders of magnitude larger than Lucas' estimates of the potential welfare gains from stabilizing business cycles (i.e., less than 0.1% of average consumption).

The main contribution of the paper is to clarify what drives the welfare effects of changes in the wage process, emphasizing the role of labor supply. In addition, the simple framework can also shed light on the *quantitative* findings of richer incomplete-markets models with more complex interaction between wages and the wealth distribution. In particular, when properly calibrated, the model delivers quantitatively similar results to those of Krueger and Perri (2003) and Pijoan-Mas (2006). The advantage of the approach taken here is that welfare effects can be solved for in closed form (rather than via numerical solution and simulation) and, consequently, the roles of preference parameters, wage risk parameters and market structure are all transparent.

The rest of the paper is organized as follows. Section 2 describes the model economies and Section 3 defines our three welfare measures. Sections 4 and 5 characterize equilibrium allocations and the analytical model-based welfare expressions that are obtained under the two alternative preference specifications considered. Section 6 describes the alternative observables-based welfare representation. Section 7 contains the calibration and the quantitative results. Section 8 concludes.

2. The economy

The economy is populated by a unit mass of infinitely lived agents. Each agent has the same time-separable utility function over streams of consumption $\{c_t\}_{t=0}^{\infty}$ and hours worked $\{h_t\}_{t=0}^{\infty}$

$$\mathscr{W} = (1-\beta)\mathbb{E}_{-1}\sum_{t=0}^{\infty}\beta^t u(c_t, h_t),$$

where $\beta \in (0,1)$ is the agents' discount factor. Two alternative specifications for the period utility function will be considered. In the first, consumption and leisure $(1-h_t)$ enter in a Cobb–Douglas fashion. In the second, period utility is separable between consumption and hours worked.

Production and individual labor productivity: The aggregate production function exhibits constant returns to scale with labor as the only input. Output cannot be stored. The labor market and the goods market are perfectly competitive, so individual wages equal individual productivity. Since the analysis does not focus on growth or aggregate short-term fluctuations, the hourly rental rate per efficiency unit of labor is normalized to unity.

Individuals' wage rates vary stochastically over time and are independently and identically distributed across agents in the economy. The wage is assumed to comprise two orthogonal components: a fixed effect $\alpha \in A \subseteq \Re$, and a transitory *iid* shock $\varepsilon_t \in E \subseteq \Re$;

$$\log w_t = \alpha + \varepsilon_t. \tag{1}$$

The fixed effect α is drawn in an initial period prior to the start of period 0. Then, for every $t \geqslant 0$, each agent draws a value for ε_t.[6] Let Φ_v denote the normal cumulative distribution function with mean $-v/2$ and variance v. Then, $\varepsilon_t \sim \Phi_{v_\varepsilon}$ and $\alpha \sim \Phi_{v_\alpha}$. As a result, $\log w \sim \Phi_v$, where $v = v_\varepsilon + v_\alpha$, which implies that the population mean wage (in levels) is equal to one. This feature is convenient when studying comparative statics with respect to the variances v_ε and v_α.

Market structure: Households have access to perfect insurance against transitory ε-shocks and no insurance against permanent α-shocks. Since α is not insured, the equilibrium of this economy will only offer *partial insurance*. The extent to which insurance is incomplete depends on the size of v_α relative to v_ε. The competitive equilibria are defined sequentially: all traded assets are Arrow securities paying out next period and are in zero net supply.

Budget constraints: The period-t budget constraint is given by

$$c_t + \int_E p_t(\varepsilon')b_t(\varepsilon')\,\mathrm{d}\varepsilon' = b_{t-1} + w_t h_t, \quad t \geqslant 0, \tag{2}$$

where b_{t-1} is the realized gross income from assets purchased in $t-1$, and $p_t(\varepsilon')$ and $b_t(\varepsilon')$ are functions defining, respectively, the price and quantity purchased of securities that pay one unit of output contingent on the realization ε' in period $t+1$. An arbitrarily loose constraint on borrowing rules out Ponzi schemes.

In period $t=-1$, the timing is as follows. First α is drawn. Then, financial markets open offering state-contingent claims conditional on the realization of ε_0. Agents are born with zero financial wealth, so the initial

[6] These assumptions on the statistical representation of the shocks are made for ease of exposition. In Heathcote et al. (2007b), we demonstrate that the analysis can be extended to allow for a richer specification of the wage process, while still retaining analytical tractability. In particular, the process for α (the uninsurable component) can incorporate permanent shocks, and the process for ε (the insurable component) can follow virtually any ARIMA process.

portfolio purchased must satisfy the budget constraint

$$\int_E p_t(\varepsilon')b_t(\varepsilon')\,\mathrm{d}\varepsilon' = 0, \quad t = -1.$$

Discussion: Our model imposes a specific exogenous market structure leading to partial insurance. In this respect, it belongs to the set of models in which markets are exogenously incomplete, a set which also includes ‘‘Bewley models’’ in which asset trade is limited to a non-contingent bond. This approach to modelling partial insurance is designed to capture, in a tractable way, the consensus that actual economies feature some degree of risk-sharing, but not perfect risk-sharing.[7] To this end, risk is assumed to be of two types: either transitory and fully insurable, or permanent and uninsurable. This assumption can be motivated as follows.

First, one natural interpretation of our framework is as an approximation to models in the Bewley (1986) tradition where a single risk-free asset is traded. Even though these models do not have explicit insurance markets, allocations in the two environments are very similar because borrowing and saving through a risk-free asset allows for near-perfect smoothing of transitory shocks, but provides no insurance against permanent productivity differences.[8] However, while Bewley models must be solved numerically, equilibrium allocations in our economies can be characterized analytically, as shown below. Section 7 compares welfare effects in these two models.

Second, a literal interpretation of our model is that there exists explicit insurance against some risks (such as short spells of unemployment or illness) but not against others (such as being endowed with low-ability or being born to poor or uneducated parents). In environments where market incompleteness emerges endogenously as a result of informational or enforcement frictions, it is easier to provide insurance against transitory risks than against permanent risks.[9]

Solving for the equilibrium: It is instructive to sketch the approach for finding the competitive equilibrium allocations (see Appendix A for full details). Start by guessing that agents with a particular realization of α trade claims to insurable shocks with each other but do not trade with agents who have other realizations of α. If so, the economy is equivalent to a world in which agents are distributed across segregated ‘‘α-islands’’, where each island is a closed economy with complete insurance against the ε shocks. Allocations on the island can then be derived from a static planner's problem with equal weights (equal since all members of an island share the same uninsurable component α and have zero initial financial wealth), subject to a resource constraint that equates aggregate island consumption to aggregate island production. Given these allocations, the implied prices of Arrow securities in the corresponding within-island competitive equilibria can be computed. It turns out that these prices do not depend on α, which confirms the initial guess of no trade across α-islands.[10]

A convenient property of equilibrium allocations is that the pair (α, ε) is a sufficient statistic for equilibrium individual consumption, hours worked, and asset holdings. In particular, it is not necessary to include individual financial wealth as a separate state variable, because equilibrium net savings of agents with the same α are zero, while wealth dispersion across agents with the same α reflects different returns on state-contingent claims (conditional on the realization of ε). This feature of the economy simplifies the solution considerably, since instead of including the endogenous evolution of individual wealth as a state variable, it is sufficient to keep track of the exogenous evolution of individual labor productivity. In contrast, in the typical Bewley model, no shocks can be perfectly insured, and the welfare theorems do not apply. Therefore, individual wealth must be included as an endogenous state variable when solving for a competitive equilibrium.

[7]Cochrane (1991) finds evidence of full insurance against short-lived transitory income shocks (e.g. short spell of illness, absence from work due to strikes). Altonji et al. (1992) argue that some income shocks are fully absorbed within the family. Guiso et al. (2005) show that a sizeable fraction of firm-level productivity shocks are insured by the firm and do not transmit to workers. Livshits et al. (2007) demonstrate that bankruptcy laws act as an effective insurance against some states with low-income realizations. Finally, Blundell et al. (2008) and Heathcote et al. (2007b) find evidence of substantial but not full risk sharing in U.S. data.

[8]See Deaton (1991) and Carroll (1997) for discussions on the role of precautionary saving in smoothing income shocks which are not too persistent.

[9]See, for example, Huggett and Parra (2006) and Krueger and Perri (2006) for discussions of the link between persistence of shocks and ability to provide insurance in economies with private information and limited commitment frictions, respectively.

[10]In Heathcote et al. (2007c), we show that this approach for solving for the equilibrium allocations remains valid even when agents face recurrent permanent stochastic innovations to α.

The time-invariant functions defining equilibrium individual wages, consumption, hours and start-of-period asset holdings are labelled $w(\alpha,\varepsilon)$, $c(\alpha,\varepsilon)$, $h(\alpha,\varepsilon)$ and $b(\alpha,\varepsilon)$, respectively.

3. Three welfare questions

Allocations are ranked using the following utilitarian social welfare function:

$$\mathscr{W} = (1-\beta)\mathbb{E}_{-1}\sum_{t=0}^{\infty}\beta^t u(c_t,h_t) = \int_A\int_E u(c(\alpha,\varepsilon),h(\alpha,\varepsilon))\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)\,\mathrm{d}\Phi_{v_\alpha}(\alpha). \tag{3}$$

This expression for welfare has two interpretations. First, it is the value for a utilitarian planner who weights all agents equally. Second, it is the expected lifetime utility for an agent at time $t=-1$ "under the veil of ignorance", i.e., before uncertainty is realized. The welfare costs associated with labor market uncertainty are assessed from three different perspectives. First, given the insurance market structure, what is the welfare effect of a rise in labor market risk? Second, for a given level of risk, what are the welfare gains from completing markets? Third, what is the welfare gain from eliminating all labor market risk?

Welfare effect of rising labor market risk: Suppose the variances of permanent and transitory shocks rise from v_α and v_ε to $\widehat{v}_\alpha$ and $\widehat{v}_\varepsilon$, respectively. Let $\Delta v_\alpha = \widehat{v}_\alpha - v_\alpha$ and $\Delta v_\varepsilon = \widehat{v}_\varepsilon - v_\varepsilon$. Let ω denote the associated welfare gain, expressed in units of the equivalent variation in lifetime consumption under the baseline wage variance:

$$\int_A\int_E u((1+\omega)c(\alpha,\varepsilon),h(\alpha,\varepsilon))\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)\,\mathrm{d}\Phi_{v_\alpha}(\alpha) = \int_A\int_E u(\widehat{c}(\alpha,\varepsilon),\widehat{h}(\alpha,\varepsilon))\,\mathrm{d}\Phi_{\widehat{v_\varepsilon}}(\varepsilon)\,\mathrm{d}\Phi_{\widehat{v_\alpha}}(\alpha), \tag{4}$$

where $\widehat{c}(\alpha,\varepsilon)$ and $\widehat{h}(\alpha,\varepsilon)$ denote equilibrium allocations in the economy with $\widehat{v}_\alpha$ and $\widehat{v}_\varepsilon$.

A theme of our paper is that increases in wage dispersion can have an impact on aggregate productivity (by changing the covariance between hours worked and individual productivity) in addition to affecting the measures of inequality. Therefore, it is useful to decompose the overall welfare effect ω in two parts—a *level effect* and a *volatility effect*. The level effect captures the welfare effect associated with changes in the size of the aggregate pie. The volatility effect captures the welfare effect associated with changes in how evenly the pie is distributed.

Formally, our strategy for identifying these two components closely follows that outlined by Flodén (2001) who, in turn, builds on earlier work by Benabou (2002). Let capital letters denote population averages. The *level effect* associated with an increase in wage dispersion (in units of consumption) is defined as the value for ω^{lev} that solves the following equation:

$$u((1+\omega^{\text{lev}})C,H) = u(\widehat{C},\widehat{H}). \tag{5}$$

Next, for an agent behind the veil of ignorance, the *cost of uncertainty* (in terms of consumption) is defined as the value for p that solves the following equation:

$$u((1-p)C,H) = \int_A\int_E u(c(\alpha,\varepsilon),h(\alpha,\varepsilon))\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)\,\mathrm{d}\Phi_{v_\alpha}(\alpha). \tag{6}$$

Note that the cost of uncertainty is a measure of the utility difference between drawing a lottery over $c(\alpha,\varepsilon)$ and $h(\alpha,\varepsilon)$ versus receiving the expected values of consumption and leisure associated with this lottery. Analogously, the cost of uncertainty in an economy with variances $\widehat{v}_\alpha$ and $\widehat{v}_\varepsilon$ is denoted as $\widehat{p}$. The *volatility effect* of an increase in wage dispersion is defined as the value for ω^{vol} that solves the following equation:

$$(1+\omega^{\text{vol}})(1-p) = 1-\widehat{p}. \tag{7}$$

Thus, the volatility effect is the percentage change in the cost of uncertainty associated with the increase in wage dispersion.

Below it is shown that for both types of preferences, the two components approximately sum to the total welfare effect, i.e., $\omega \simeq \omega^{\text{lev}} + \omega^{\text{vol}}$.

Welfare gain from completing markets: The welfare gain from completing insurance markets (holding v_α and v_ε constant) is measured as the percentage increase in consumption in the partial-insurance economy required to achieve the same welfare as in the economy with complete markets. Thus, the welfare gain is defined as the value for χ that solves

$$\int_A \int_E u((1+\chi)c(\alpha,\varepsilon), h(\alpha,\varepsilon))\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)\,\mathrm{d}\Phi_{v_\alpha}(\alpha) = \int_A \int_E u(c(0,\alpha+\varepsilon), h(0,\alpha+\varepsilon))\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)\,\mathrm{d}\Phi_{v_\alpha}(\alpha), \tag{8}$$

where the expression on the right-hand side reflects welfare when markets are complete and fluctuations in both α and ε are insurable.

Completing markets amounts to reducing the variance of uninsurable risk, and simultaneously increasing the variance of insurable risk by the same amount v_α. Thus, the welfare effect can be read directly from the expression for ω in (4) by setting $\hat{v}_\varepsilon = v_\varepsilon + v_\alpha$ and $\hat{v}_\alpha = 0$.

Welfare effect of eliminating risk: In computing the welfare cost of business cycles, Lucas (1987) compared welfare associated with the actual U.S. time series for aggregate consumption to welfare associated with the trend of the actual path.[11] Thus, he calculated the hypothetical welfare gain of eliminating aggregate fluctuations. The welfare gain from eliminating idiosyncratic risk is computed by applying the same actual-versus-trend comparison as Lucas, but at the individual rather than the aggregate level. Thus, individuals' wages are set equal to unity—their unconditional expected value.

For Lucas, eliminating aggregate fluctuations was a hypothetical thought experiment. In contrast, our experiment of eliminating risk could—in the context of our model—be achieved via an appropriate policy of full wage compression. In particular, wage risk can be eliminated by a system of distortionary wage taxes and subsidies that guarantees each worker an after-tax hourly wage rate equal to average labor productivity, which in turn equals one. Thus, the tax (subsidy) rate paid by a worker with current pre-tax wage w is given by $\tau(w) = 1 - 1/w$.[12] The welfare gain from eliminating wage risk, κ, can then be read directly from the expression for ω in (4) by setting $\hat{v}_\varepsilon = 0$ and $\hat{v}_\alpha = 0$.

Finally, note that the solutions for ω, χ, and κ represent welfare comparisons across two steady states characterized by different variances for wages. However, this does not imply that the welfare expressions ignore transitional dynamics. This is due to the fact that the transition to a new steady state in response to a change in the wage process is immediate in this environment.[13]

4. Cobb–Douglas preferences

Consider first preferences that are Cobb–Douglas between consumption and leisure, i.e.,

$$u(c,h) = \frac{(c^\eta(1-h)^{1-\eta})^{1-\theta}}{1-\theta},$$

where $\eta \in (0,1)$ determines the relative taste for consumption versus leisure. Cobb–Douglas preferences are widely used in the macroliterature, since they are consistent with balanced growth, irrespective of the choice for θ. Moreover, this specification implies non-separability between consumption and leisure, which is consistent with some empirical evidence emphasized in labor economics (Heckman, 1974; Browning and Meghir, 1991).

The parameter η can be identified by the share of disposable time agents devote to market work. This implies that the parameter θ governs both the intertemporal elasticity of substitution for consumption and the corresponding elasticity for hours worked. In particular, the intertemporal elasticity of substitution for

[11]More recently, Krusell and Smith (1999), Storesletten et al. (2001), and Krebs (2003) have made similar calculations in models with heterogeneous agents.

[12]To verify that this system of wage taxes and subsidies is feasible, it is necessary to verify that it is revenue neutral. Since every agent faces the same after-tax wage, each agent works the same number of hours per period and enjoys the same level of consumption. Per-capita consumption will equal per-capita after-tax income which, in turn, is equal to (constant) hours times the after-tax wage, which is equal to one given the tax function $\tau(w)$. Since average labor productivity is also equal to one, output per capita will equal consumption per capita. Thus, the tax-subsidy scheme is revenue neutral.

[13]More precisely, the transition due to a one-off change in the variances of either or both components of the wage process is immediate in the sense that our expected welfare measure (3) takes the same value in the period the wage process changes as in all subsequent periods. The reason is that assets are in zero net supply.

consumption is given by $1/\theta$. The coefficient of relative risk aversion is

$$\bar{\gamma} \equiv \gamma(\theta, \eta) \equiv -\frac{cu_{cc}}{u_c} = 1 - \eta + \eta\theta. \tag{9}$$

The Frisch elasticity of labor supply depends on hours worked, and is given by $\phi(\theta, \eta, h) = \lambda(1-h)/h$, where $\lambda \equiv (1 - \eta + \eta\theta)/\theta$ defines the Frisch elasticity for leisure.[14] It is useful to define a "non-stochastic Frisch" elasticity of labor supply corresponding to a non-stochastic version of the model, in which case $h = H = \eta$, where H denotes average hours worked:

$$\bar{\phi} \equiv \phi(\theta, \eta, H) \equiv \left.\frac{u_h}{u_{hh}h - u_{ch}^2 h/u_{cc}}\right|_{h=H} = \frac{\lambda(1-\eta)}{\eta}. \tag{10}$$

4.1. Equilibrium allocations with Cobb–Douglas preferences

The equilibrium consumption and leisure allocations in our partial-insurance economy are

$$\begin{aligned} \log c(\alpha, \varepsilon) &= \log(\eta) + \alpha + (1-\lambda)\varepsilon + \lambda(1-\lambda)\frac{v_\varepsilon}{2}, \\ \log(1 - h(\alpha, \varepsilon)) &= \log(1-\eta) - \lambda\varepsilon + \lambda(1-\lambda)\frac{v_\varepsilon}{2}. \end{aligned} \tag{11}$$

The insurable transitory shock ε reduces leisure proportionately to the Frisch elasticity for leisure λ. Moreover, leisure is independent of the permanent uninsurable component α since the income and substitution effects associated with an uninsurable change in the wage exactly offset each other when preferences are Cobb–Douglas.

Because α has no impact on hours worked, consumption is directly proportional to α. Given non-separability between consumption and leisure, current consumption depends on the insurable shock ε as long as $\lambda \neq 1$. For $\lambda < 1$ (which is equivalent to $\theta > 1$), consumption and leisure are substitutes, in the sense that the marginal utility of consumption is decreasing in leisure. In this case, in order to equate the marginal utility of consumption intertemporally, individuals who draw a high-value for ε and who therefore enjoy relatively little leisure must be compensated with relatively high consumption. When $\theta = 1$ (in which case $u(c, h) = \eta \log c + (1-\eta)\log(1-h)$), consumption is constant and equal to $\eta \exp(\alpha)$.

Note that individual consumption and leisure also depend on the variance of the insurable component of the log wage, v_ε. For $\lambda \in (0, 1)$, both consumption and leisure are increasing in insurable wage dispersion. We will return to this point when examining the welfare effects of a rise in wage dispersion.

Appendix A contains the derivations of the above expressions for $c(\alpha, \varepsilon)$ and $h(\alpha, \varepsilon)$. The cost of an individual's portfolio of Arrow securities turns out to be zero in every state (see the Appendix). Hence, the budget constraint (2) implies that for each possible realization ε', the payoff from the corresponding Arrow security can be expressed as the equilibrium value for current consumption net of current labor income, or

$$b(\varepsilon'; (\alpha, \varepsilon)) = c(\alpha, \varepsilon') - w(\alpha, \varepsilon')h(\alpha, \varepsilon') = \exp\left(\alpha + (1-\lambda)\varepsilon' + \lambda(1-\lambda)\frac{v_\varepsilon}{2}\right) - \exp(\alpha + \varepsilon').$$

Consequently, the dynamics of individual asset income inherit the process for the insurable component of wages, which in this paper is i.i.d. over time.

4.2. Welfare analysis with Cobb–Douglas preferences

The following proposition addresses the three welfare questions posed above:

[14] The Frisch elasticity of labor supply (leisure) measures the elasticity of hours worked (leisure) to changes in wages, keeping the marginal utility of consumption constant.

Proposision 1. *With Cobb– Douglas preferences the (approximate) welfare effect of a change in labor market risk* $(\Delta v_\varepsilon, \Delta v_\alpha)$ *is*

$$\omega(\Delta v_\alpha, \Delta v_\varepsilon) \simeq -\bar{\gamma}\frac{\Delta v_\alpha}{2} + \bar{\phi}\frac{\Delta v_\varepsilon}{2} = \underbrace{\bar{\phi}\Delta v_\varepsilon}_{\omega^{\text{lev}}} \underbrace{-\bar{\phi}\frac{\Delta v_\varepsilon}{2} - \bar{\gamma}\frac{\Delta v_\alpha}{2}}_{\omega^{\text{vol}}}.$$

Proof. See Appendix B.

Corollary 1. *Let* $\chi(v_\alpha)$ *denote the (approximate) welfare gain from completing markets in an economy with uninsurable risk variance equal to* v_α. *Let* $\kappa(v_\alpha, v_\varepsilon)$ *denote the (approximate) welfare gain from eliminating risk in an economy with variances* $(v_\alpha, v_\varepsilon)$. *Then*

$$\chi(v_\alpha) = \omega(-v_\alpha, v_\alpha) \simeq \bar{\gamma}\frac{v_\alpha}{2} + \bar{\phi}\frac{v_\alpha}{2},$$
$$\kappa(v_\alpha, v_\varepsilon) = \omega(-v_\alpha, -v_\varepsilon) \simeq \bar{\gamma}\frac{v_\alpha}{2} - \bar{\phi}\frac{v_\varepsilon}{2}.$$

In the Proof of Proposition 1, exact closed-form solutions for the welfare effects are derived. However, these expressions are cumbersome and not particularly transparent. Through a set of log-approximations of the type $\ln(1+x) \simeq x$ and $e^x \simeq 1+x$, the simple and useful solutions stated in Proposition 1 are obtained. The linearity of the welfare effects in Δv_α and Δv_ε is a feature of the approximation. Section 7.2 documents the accuracy of these approximations.

Welfare effect of rising labor market risk (ω): The first term in the expression for ω captures the welfare loss associated with a rise in the dispersion of the uninsurable component of wages. This loss is equal to Lucas' welfare cost of aggregate consumption fluctuations in an economy with inelastic labor supply. Note that the welfare loss is proportional to the risk-aversion parameter $\bar{\gamma}$.

The second term shows that increasing insurable productivity dispersion *increases* welfare in proportion to the Frisch elasticity of labor supply $\bar{\phi}$. The intuition is that given flexible labor supply, an unconstrained planner can achieve better allocative efficiency with larger productivity dispersion, without any loss in terms of consumption smoothing, by commanding longer hours from high-productivity workers and higher leisure from less productive workers. This result is closely related to the classical consumer-theory result that the indirect utility function of a static consumer is quasi-convex in prices, so a mean-preserving spread of the price distribution raises welfare (see, for example, Mas-Colell et al., 1995, p. 59). In the decomposition of welfare effects into level and volatility components, the level effect, ω^{lev}, captures the welfare gain associated with the increase in aggregate labor productivity. Why is there a negative volatility effect related to Δv_ε, notwithstanding full insurance against this source of risk? The reason is that to exploit greater dispersion in productivity across workers, the planner must increase dispersion in hours. Since utility is concave in leisure, this is welfare reducing. At the margin, the welfare gain for the planner from additional specialization in terms of increased average labor productivity is exactly offset by the loss associated with greater dispersion in leisure.

As noted above, hours worked do not respond to uninsurable wage changes under Cobb–Douglas preferences. Therefore, the overall welfare impact from additional uninsurable dispersion (i.e., Lucas' expression) is equal to the negative volatility effect.

Fig. 1 provides a picture of how our welfare effects vary with θ, the parameter defining agents' willingness to substitute intertemporally (note that η is held constant). Panel (A) plots ω for different values for the inverse of the Frisch elasticity, $1/\bar{\phi}$, which is increasing in θ.[15] Raising the Frisch elasticity (reducing $1/\bar{\phi}$) reduces the welfare cost of higher dispersion. If labor supply is sufficiently elastic, rising dispersion is actually welfare-improving. In part, this finding reflects the fact that the productivity gain associated with larger insurable risk is increasing in the Frisch elasticity, as discussed above. A second effect, working in the same direction, is that

[15] The values for $(v_\alpha, v_\varepsilon)$ and $(\Delta v_\alpha, \Delta v_\varepsilon)$ used to produce the plots are from the calibration described in Section 6.1. Plotting welfare effects against $1/\bar{\phi}$ rather than against θ facilitates a comparison with the separable-preferences specification. Low values for $\bar{\phi}$ cannot be considered in the context of the Cobb–Douglas utility specification, since $\lim_{\theta\to\infty}\bar{\phi}(\theta,\eta) = 1-\eta$.

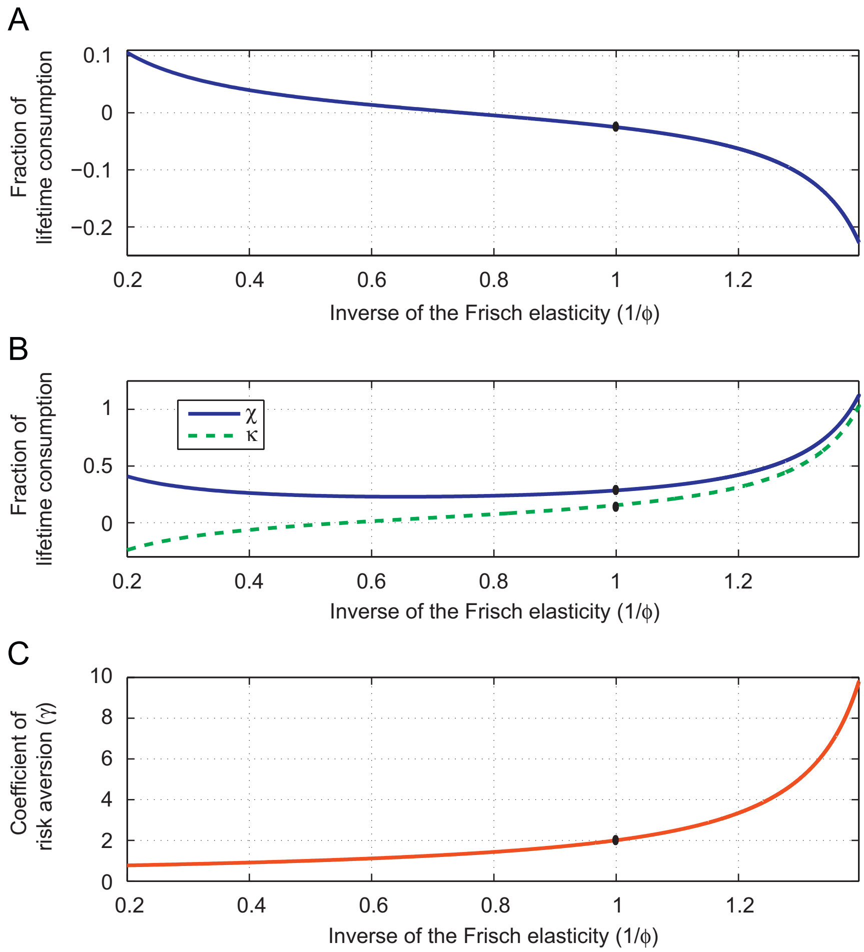

Fig. 1. Cobb–Douglas preferences. (A) Welfare effect from rising labor market risk (ω) ($\Delta v_\alpha = \Delta v_\varepsilon = 0.05$); (B) welfare gain from completing markets (χ) and from eliminating risk (κ) ($v_\alpha = 0.22,\ v_\varepsilon = 0.13$); and (C) coefficient of risk aversion as a function of $1/\phi$.

with the Cobb–Douglas utility function, a higher Frisch elasticity means a lower coefficient of risk aversion (see panel (C)) which, in turn, implies a lower cost of rising uninsurable risk.

Welfare gain from completing markets (χ): Recall that completing markets means (i) a reduction $\Delta v_\alpha = -v_\alpha$ in the variance of uninsurable risk and (ii) a corresponding increase $\Delta v_\varepsilon = v_\alpha$ in the variance of insurable risk. The first term in the expression for χ—proportional to the coefficient of relative risk aversion $\bar{\gamma}$—captures the value of the additional insurance provided by increased risk sharing. The second term captures the gains from specialization, whereby more productive households work relatively harder and less productive households enjoy more leisure.

Panel (B) of Fig. 1 shows how the welfare gain from completing markets varies with the elasticity of labor supply. Interestingly, the welfare gain is non-monotone. Initially, as the Frisch elasticity falls ($1/\bar{\phi}$ rises), the welfare gain gets smaller, since it becomes harder to reallocate hours in favor of more productive workers. However, as $\bar{\phi}$ is reduced, $\bar{\gamma}$ rises (panel (C)), and eventually a point is reached where the value of additional insurance to shelter consumption fluctuations comes to dominate the welfare calculus.

Welfare effect from eliminating risk (κ): In a model with exogenous labor supply, there would be no difference between insuring and eliminating idiosyncratic labor income risk. Both changes would lead to income and consumption being equalized across individuals, with no changes in aggregate quantities. With endogenous labor supply, however, increasing risk sharing is not the same thing as reducing risk at the source. The reason is that additional insurable risk is welfare-improving when labor supply is endogenous, as discussed above.

Comparing κ with χ, it is clear that the welfare gains from eliminating risk are always smaller than those from insuring risk. Eliminating the uninsured part of wage dispersion is welfare-improving, since this reduces consumption dispersion. However, eliminating dispersion in the insurable component of wages is detrimental, since it eliminates the positive covariance between the insurable component of individual productivity and individual hours that boosts aggregate labor productivity. The cost associated with eliminating insurable dispersion is increasing in the Frisch elasticity, which explains why the gap between χ and κ is decreasing in $1/\bar{\phi}$ in panel (B) of Fig. 1.

Our finding that there is a down-side to reducing risk in the presence of flexible labor supply is mirrored in some work on the welfare costs of business cycles. Cho and Cooley (2001) noted that if aggregate hours are pro-cyclical, then eliminating aggregate business cycle risk may reduce average labor productivity. Gomes et al. (2001) provide an example where aggregate fluctuations may be welfare improving in an equilibrium search model when the agent can choose to allocate time between work and search.

Finally, throughout the analysis, we have emphasized that wage inequality generates both benefits and costs. A natural question then arises: Is there an optimal level of inequality? Other authors have formally addressed this question within models where there is a trade-off between inequality and growth (Cordoba and Verdier, 2007) or where inequality has benefits associated with incentive provision (Phelan, 2006). In the model studied here, the answer depends on whether dispersion is insurable. If it is, inequality is unambiguously good, otherwise it is unambiguously bad.

5. Separable preferences

Separability is a common assumption in the microliterature on consumption and labor supply (for a survey, see Browning et al., 1999). We focus on the following class of preferences:

$$u(c,h) = \frac{c^{1-\gamma}}{1-\gamma} - \psi \frac{h^{1+\sigma}}{1+\sigma}, \tag{12}$$

where $\gamma,\ \sigma \in [0,+\infty)$. The coefficient of relative risk aversion is γ, while the intertemporal elasticity of substitution for consumption is $1/\gamma$. The Frisch elasticity for labor supply is simply $1/\sigma$. When preferences are separable one can distinguish between agents' willingness to substitute consumption and hours intertemporally. Cobb–Douglas preferences do not admit such flexibility.

Without loss of generality, the parameter measuring the distaste for work relative to the taste for consumption—ψ—is normalized to unity. It is easy to verify that such a normalization has no impact on the welfare expressions. An important implication of this result is that, even if we were to allow for heterogeneity with respect to ψ, the final welfare expressions would remain unchanged. Thus, our analysis is robust to an important class of preference heterogeneity.

The derivations for equilibrium allocations and welfare effects with separable preferences follow those described in the Appendix for the Cobb–Douglas case very closely and are therefore omitted.[16]

5.1. Equilibrium allocations with separable preferences

When preferences are separable between consumption and hours worked, equilibrium allocations in the partial-insurance economy are given by

$$\begin{aligned} \log c(\alpha,\varepsilon) &= \left(\frac{1+\sigma}{\gamma+\sigma}\right)\alpha + \left(\frac{1+\sigma}{\gamma+\sigma}\right)\frac{1}{\sigma}\frac{v_\varepsilon}{2}, \\ \log h(\alpha,\varepsilon) &= \left(\frac{1-\gamma}{\gamma+\sigma}\right)\alpha + \frac{1}{\sigma}\varepsilon - \left(\frac{1+\sigma}{\gamma+\sigma}\right)\frac{\gamma}{\sigma^2}\frac{v_\varepsilon}{2}. \end{aligned} \tag{13}$$

The response of hours to uninsurable shocks is governed by the Marshallian (uncompensated) elasticity of labor supply $(1-\gamma)/(\gamma+\sigma)$. Whether hours increase or decrease with α depends on the relative strength of

[16] Heathcote et al. (2007c), the working paper version of this paper, contains detailed proofs for the economy with separable utility.

substitution versus income effects. With separable preferences, the income effect dominates the substitution effect if $\gamma > 1$. The Frisch elasticity $1/\sigma$ determines the responsiveness of individual hours to insurable shocks to individual wages.

Individual consumption is independent of the realization of the transitory shock ε, reflecting full insurance against this component of the wage process coupled with separability between consumption and hours in the utility function. The response of consumption to the uninsurable component of wages is equal to the response of earnings. Since log earnings are equal to log wages plus log hours, the pass-through coefficient from the uninsurable component of wages to earnings is given by $1 + (1-\gamma)/(\gamma+\sigma) = (1+\sigma)/(\gamma+\sigma)$.

As under the Cobb–Douglas specification, individual consumption and leisure also depend on the variance of the insurable component of the log wage, v_ε. For any individual state (α, ε), both consumption and leisure are increasing in v_ε.

5.2. Welfare analysis with separable preferences

We now state a pair of propositions analogous to Propositions 1 and 2.

Proposition 1a. *With separable preferences, the (approximate) welfare effect of a change in labor market risk* $(\Delta v_\varepsilon, \Delta v_\alpha)$ *is*

$$\begin{aligned}\omega(\Delta v_\alpha, \Delta v_\varepsilon) &\simeq -\left[\frac{\gamma-1}{\gamma+\sigma} + \gamma\left(\frac{1+\sigma}{\gamma+\sigma}\right)\right]\frac{\Delta v_\alpha}{2} + \frac{1}{\sigma}\frac{\Delta v_\varepsilon}{2} \\ &= \underbrace{-\frac{\gamma-1}{\gamma+\sigma}\Delta v_\alpha + \frac{1}{\sigma}\Delta v_\varepsilon}_{\omega^{\text{lev}}} + \underbrace{\left[\frac{\gamma-1}{\gamma+\sigma} - \gamma\left(\frac{1+\sigma}{\gamma+\sigma}\right)\right]\frac{\Delta v_\alpha}{2} - \frac{1}{\sigma}\frac{\Delta v_\varepsilon}{2}}_{\omega^{\text{vol}}}.\end{aligned}$$

Corollary 1a. *With separable preferences, the (approximate) welfare gains from completing markets and eliminating risk in an economy with variances* $(v_\alpha, v_\varepsilon)$ *are given, respectively, by*

$$\begin{aligned}\chi(v_\alpha) = \omega(-v_\alpha, v_\alpha) &\simeq \left[\frac{\gamma-1}{\gamma+\sigma} + \gamma\left(\frac{1+\sigma}{\gamma+\sigma}\right)\right]\frac{v_\alpha}{2} + \frac{1}{\sigma}\frac{v_\alpha}{2}, \\ \kappa(v_\alpha, v_\varepsilon) = \omega(-v_\alpha, -v_\varepsilon) &\simeq \left[\frac{\gamma-1}{\gamma+\sigma} + \gamma\left(\frac{1+\sigma}{\gamma+\sigma}\right)\right]\frac{v_\alpha}{2} - \frac{1}{\sigma}\frac{v_\varepsilon}{2}.\end{aligned}$$

Welfare effect from rising labor market risk (ω): As with Cobb–Douglas preferences, increasing insurable productivity dispersion strictly increases welfare in proportion to the Frisch elasticity. Once again, the intuition is simply that an unconstrained planner can achieve better allocative efficiency with larger dispersion by having the more productive agents specialize in market work. An interesting difference between the two preference specifications is that when preferences are separable, the productivity gain associated with greater wage dispersion translates into higher average consumption and leisure, whereas in the Cobb–Douglas case, productivity gains have no impact on average consumption, but do increase average leisure.[17]

The welfare effects of a rise in uninsurable uncertainty are more complex. When $\gamma > 1$, the income effect from a positive wage shock dominates the substitution effect, so that agents increase their work effort in bad times. In this case, flexible labor supply is used to improve consumption smoothing at the expense of productivity (the level effect is negative). When $\gamma < 1$, the substitution effect dominates the income effect, and agents increase work effort in good times. In this case, flexible labor supply actually increases consumption volatility, but it is still beneficial because agents are relatively unconcerned about fluctuations in consumption, and concentrating their work effort in high-wage periods raises average output per hour (the level effect is positive). This discussion implies that for flexible labor supply to mitigate the welfare cost of increases in *uninsurable* wage risk, it must be the case that $\gamma \neq 1$, implying that preferences are inconsistent with balanced growth.

[17]This can easily be seen by computing $\mathbb{E}[c(\alpha,\varepsilon)]$ and $\mathbb{E}[1-h(\alpha,\varepsilon)]$ under both preference specifications. See Appendix B for the Cobb–Douglas case and Appendix D in Heathcote et al. (2007c) for the separable case.

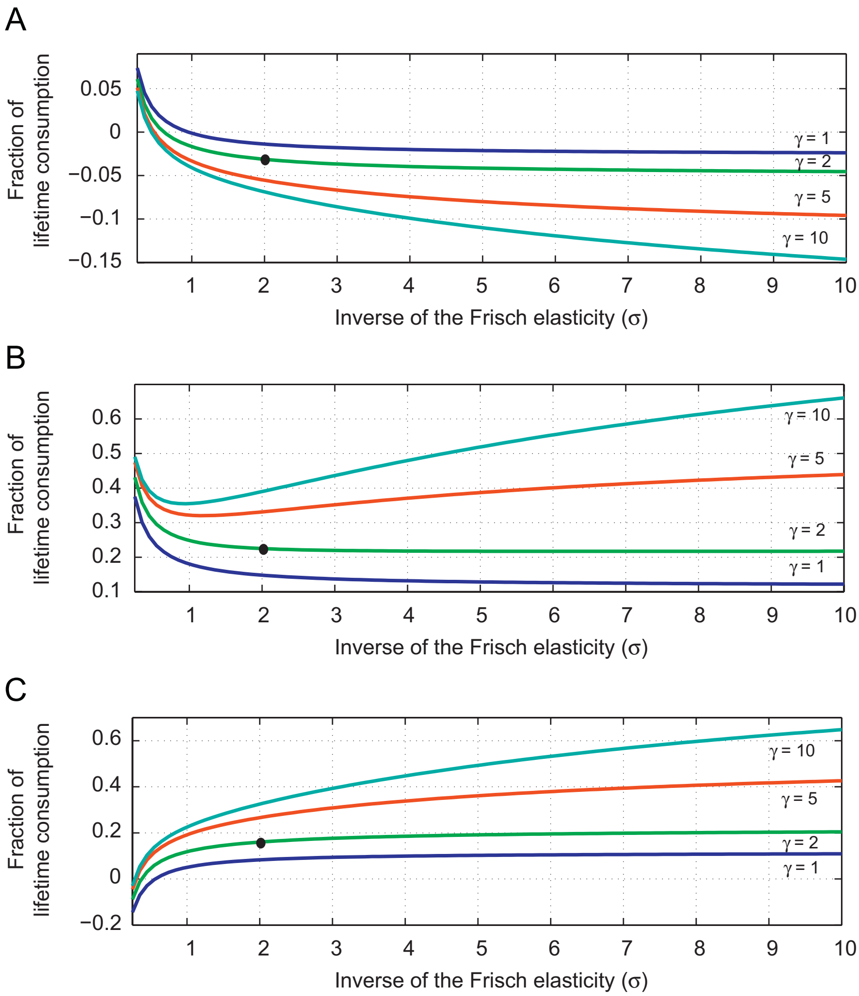

Fig. 2. Separable preferences. (A) Welfare effect from rising labor market risk (ω) ($\Delta v_{\alpha} = \Delta v_{\varepsilon} = 0.05$); (B) welfare change from completing markets (χ) ($v_{\alpha} = 0.22,\ v_{\varepsilon} = 0.13$); and (C) welfare change from eliminating risk (κ) ($v_{\alpha} = 0.22,\ v_{\varepsilon} = 0.13$).

For $\gamma \geqslant 1$, the expression for ω indicates that additional uninsurable risk is unambiguously welfare-reducing. However, a surprising finding is that when $\gamma < 1/(2+\sigma)$, a rise in v_{α} has a positive welfare effect.[18] The intuition is that when risk aversion is sufficiently small and labor supply elasticity sufficiently large, agents willingly substitute labor supply intertemporally to raise average productivity, and are relatively unconcerned about the resulting fluctuations in consumption. One interesting benchmark is risk neutrality ($\gamma = 0$), in which case $\omega = (1/2)(1/\sigma)(\Delta v_{\alpha} + \Delta v_{\varepsilon})$. Thus, when insurance is not valued, the distinction between insurable and uninsurable shocks becomes irrelevant.

Panel (A) of Fig. 2 provides a graphical summary of how the overall cost of rising dispersion, ω, varies with risk aversion, γ, and the inverse of the Frisch elasticity, σ.

Welfare gain from completing markets (χ): As in the Cobb–Douglas case, there are two sources of welfare gains from insuring risk. The first is the gain from the additional insurance provided by increased risk sharing. The second is the allocative efficiency gain associated with elastic labor supply: under complete markets, more productive households work relatively longer hours and less productive households enjoy more leisure.

The welfare gain from completing markets is strictly increasing in relative risk aversion (γ). Consider three special cases. First, under risk neutrality ($\gamma = 0$), $\chi = 0$ since consumption fluctuations are not costly.

[18]Recall that in the Cobb–Douglas case, increases in uninsurable wage dispersion v_{α} always reduce welfare.

Second, in the absence of flexible labor supply ($\sigma \to \infty$), the welfare gain is $\chi \simeq \gamma v_\alpha/2$, the Lucas expression for the welfare cost of consumption fluctuations. Third, if $\gamma = \sigma = 1$, then $\chi \simeq v_\alpha$.

Panel (B) of Fig. 2 shows that χ is non-monotone in σ. For $\sigma < 1$, χ is always increasing in the Frisch elasticity. However, for $\sigma \geqslant 1$, whether or not χ is increasing in the Frisch elasticity depends on whether $\gamma \leqslant 2\sigma/(\sigma - 1)$. The intuition is that, given high aversion to consumption fluctuations, an increase in the willingness to substitute hours intertemporally can have a larger positive impact on welfare under autarky (by effectively improving self-insurance) than under complete markets (by increasing average productivity).

Welfare effect from eliminating risk (κ): As in the Cobb–Douglas specification, eliminating labor market risk amounts to reducing to zero the variances of both components of the wage process, which is welfare reducing for insurable risk, and likely to be welfare-improving for uninsurable risks. Comparing panels (B) and (C) of Fig. 2, it is clear that the welfare gains from eliminating risk are similar to those from completing markets when the Frisch elasticity is low (σ is high), but are much smaller—and in some cases negative—when the Frisch elasticity is high.

6. Observables-based welfare analysis

It is possible to derive alternative representations for the welfare effects of rising inequality (and for the level and volatility components) as functions only of preference parameters and second moments of the joint cross-sectional distribution for wages, hours and consumption.

The key advantage of the observables-based expressions, relative to the parametric expressions described above, is that they are more general. They apply to any economy in which (i) the standard intratemporal optimality condition between consumption and leisure/hours worked is satisfied, and (ii) wages, consumption and leisure/hours are jointly log-normal. Moreover, in order implement this approach, one does not have to take a stand on how the variances of uninsurable versus insurable wages risks have changed over time (Δv_α,Δv_ε)—it is sufficient to compute the relevant moments in repeated cross-sections and to choose preference parameters. This requires high-quality data on consumption and hours, however. We therefore view the two alternative approaches as complementary.

The following assumptions are convenient.

Assumption A1. Preferences are Cobb–Douglas and wages w, consumption c, and leisure $1 - h$ are log-normally distributed in the cross-section.

Assumption A1′. Preferences are separable and wages w, consumption c and hours worked h are log-normally distributed in the cross-section.

Assumption A2. Wages and allocations satisfy individual intratemporal optimality, aggregate consumption equals aggregate labor income, and the average wage equals one (so $\mathbb{E}[\log w] = -var(\log w)/2$).

We are now ready to state the main result of this section.

Proposision 2. *Under Assumptions* A1–A2 *or* A1′–A2, *the* (*approximate*) *welfare effect* ω *of a rise in wage dispersion can be expressed as*

$$\omega \simeq \underbrace{\Delta\, cov(\log w, \log h)}_{\omega^{\text{lev}}} \underbrace{-\frac{1}{2}\left(-\frac{u_{CC}C}{u_C}\right)\Delta\, var(\log c) - \frac{1}{2}\left(\frac{u_{HH}H}{u_H}\right)\Delta\, var(\log h) + \frac{u_{CH}H}{u_C}\Delta\, cov(\log c, \log h)}_{\omega^{\text{vol}}},$$

where $C = \mathbb{E}(c)$ *and* $H = \mathbb{E}(h)$. *In the Cobb–Douglas case*

$$-\frac{u_{CC}C}{u_C} = \bar{\gamma}, \quad \frac{u_{HH}H}{u_H} \simeq \bar{\gamma} - 1 + \frac{\eta}{1-\eta} \quad \text{and} \quad \frac{u_{CH}H}{u_C} \simeq \bar{\gamma} - 1$$

and in the separable case, the expression becomes

$$-\frac{u_{CC}C}{u_C} = \gamma, \quad \frac{u_{HH}H}{u_H} = \sigma \quad \text{and} \quad \frac{u_{CH}H}{u_C} = 0.$$

Moreover, the level effect ω^{lev} *(approximately) equals the percentage change in aggregate labor productivity* $\Delta \log(C/H)$.

Proof. See Appendix C.

The expression for ω in Proposition 2 comprises four terms. The first term is the change in the covariance between hours and wages: a higher positive correlation between individual hours and individual productivities increases average welfare. The change in the covariance is equal to the level effect ω^{lev} as defined in (5) and can also be shown to equal the change in aggregate labor productivity in the economy.

The second and third terms capture the volatility cost of a rise in wage dispersion: an increase in the variance of log consumption translates into a welfare cost proportional to the risk-aversion coefficient, and an increase in the variance of log hours translates into a welfare cost that is proportional to the coefficient $u_{HH}H/u_H$, which measures aversion to fluctuations in hours worked. In the separable case, this term is exactly the inverse of the Frisch elasticity, σ. In the Cobb–Douglas case, when $\eta = \frac{1}{2}$ (consumption and leisure receive equal weight in utility), it is equal to the risk-aversion coefficient $\bar{\gamma}$.

The fourth term, involving the change in the covariance between consumption and hours worked, is only present when utility is non-separable in consumption and leisure. It is zero in the separable case. In the Cobb–Douglas case, when $\theta > 1$ (which implies $\bar{\gamma} > 1$), consumption and leisure (hours worked) are substitutes (complements); thus, households gain from a rise in the comovement between consumption and hours worked.

Clearly, our model economy satisfies Assumptions A1–A2. It can be shown that the two different representations for the welfare effects are equivalent. More precisely, given preference parameters and variances of uninsurable and insurable shocks $(v_\alpha, v_\varepsilon)$, the expressions in Proposition 2 and Propositions 1 and 1a are identical.[19] Note also that given log-normality of the allocations, the Cobb–Douglas case delivers an exact expression in terms of observables (i.e., no approximations). However, that expression is in terms of moments of *leisure*. In order to obtain the *common* representation for welfare change in Proposition 2, the Cobb–Douglas case requires an approximation to translate cross-sectional moments involving leisure into moments involving hours worked.

Assumptions A1 and A1′ can be relaxed. In fact, the observables-based expression for the welfare effect in Proposition 2 can alternatively be obtained from a second-order Taylor approximation of any continuously differentiable concave utility function (where the Taylor approximation is taken over $\log c$ and $\log h$ around average consumption and hours worked).[20] Log-normality of the allocations is required to show that the welfare gain from changes in aggregate consumption and leisure—the level effect—is approximately equal to the change in the covariance between log hours and log wages. Details are available upon request.

7. Quantitative welfare analysis

This section describes the calibration and the implied measured welfare effects. The analytical results are then compared to the welfare effects implied by a standard incomplete-markets model.

7.1. Calibration and measurement

The section begins with a discussion of baseline choices for preference parameters and the measurement of cross-sectional moments of the joint wage, hours and consumption distribution. These are used to implement the alternative observables-based approach to quantifying the welfare effects of rising wage dispersion (from

[19] To show this, compute the cross-sectional moments using the equilibrium allocations and the distributions of shocks. Substituting these expressions into the welfare representation of Proposition 2, and rearranging terms, yields the model-based welfare effects of Propositions 1 and 1a. See Heathcote et al. (2007c) for details.

[20] We thank an anonymous referee for making this point.

Proposition 2). Finally, estimates of the variances of insurable and uninsurable wage risk before (v_α, v_ε) and after ($\widehat{v}_\alpha, \widehat{v}_\varepsilon$) the recent well-documented surge in wage dispersion are reported (see Katz and Autor, 1999; Eckstein and Nagypal, 2004, for empirical surveys). These variances are a key input to the model-based welfare expressions of Propositions 1 and 1a.

Preference parameters: Consider first the separable case. Estimates for the risk-aversion coefficient between one and three are typical in the empirical consumption literature (see Attanasio, 1999, for a survey), γ is set to two. Domeij and Flodén (2006) sample the empirical literature on male labor supply and conclude that the typical estimates of Frisch elasticities for male labor supply range between 0.1 and 0.3. However, they argue that these estimates are downward-biased because the standard estimation methods ignore the possibility that borrowing constraints may bind. By simulation, they show that the unbiased estimates can be up to twice as large. Moreover, estimates of this elasticity for females are estimated to be 3–4 times larger than those for men (see Blundell and MaCurdy, 1999, Table 2). Therefore, the Frisch elasticity is set to 0.5, corresponding to $\sigma = 2$.

With Cobb–Douglas utility, the Frisch labor supply elasticity and the coefficient of risk aversion are not independent since they are both functions of the pair of parameters (θ, η), as discussed in Section 4.2. Moreover, the parameter η has a natural counterpart in the fraction of the time endowment devoted to work activities. Following the macroeconomic literature on business cycles, η is set to $\frac{1}{3}$ (see e.g. Cooley, 1995).[21] Moreover, θ is set to four so that the implied coefficient of risk aversion $\bar{\gamma}$ equals two, as in the separable case. These parameters imply a Frisch elasticity $\bar{\phi}$ equal to one—a higher number than in the separable case.[22]

We recognize that there is disagreement regarding appropriate values for preference parameters, and that some may object to our particular choices. One advantage of our closed-form expressions for welfare is that alternative values can easily be plugged in. Figs. 1 and 2 present results for a large set of alternative parameterizations.

Measurement of wage, hours and consumption dispersion: We use the 1968–1997 waves of the Panel Study of Income Dynamics (PSID). The sample consists of roughly 2,400 observations/year and includes every head of household aged between 20 and 59 with positive earnings (not top-coded and not below half of the current minimum wage) and with annual hours worked between 520 and 5824.[23] Hourly wages are computed as annual pre-tax earnings divided by annual hours worked, and both wages and hours are regressed on race dummies and a quartic in age in order to filter out predictable life-cycle variation.[24] Variances and covariances are constructed from the (log) residuals of these regressions. The variance of log wages rose by 0.10 (from 0.25 to 0.35) over this time period, the variance of log-hours worked rose by 0.01 (from 0.082 to 0.092), and the covariance between hours and wages rose by 0.017 (from −0.023 to −0.006).

For consumption dispersion, we rely on existing studies based on the Consumer Expenditure Survey (CEX). For consistency with individual wage and hours data, the consumption data are expressed in adult-equivalent units. According to Slesnick (2001), the increase in the variance of log-consumption between 1980 and 1995 was small, around 0.01 (0.20 in 1980, 0.21 in 1995). Krueger and Perri (2006) and Attanasio et al. (2007) argue that consumption inequality rose by about 0.05 over the same period. Since there are important measurement issues that are not yet settled in this literature, we simply adopt a mid-point estimate of 0.03 for our calculations. Finally, Krueger and Perri (2003) report that the covariance between hours and consumption declined by 0.007 (from 0.037 to 0.030).[25]

[21] More precisely, the first-order condition for hours worked in a non-stochastic version of the model implies $h = \eta$.

[22] We chose to equate the coefficient of risk aversion across alternative preference specifications, rather than the Frisch elasticity for labor supply, because the lower bound on the Frisch elasticity under the Cobb–Douglas specification is $\overline{\phi} = 1 - \eta = \frac{2}{3}$.

[23] This latter restriction serves the purpose of reducing the extent to which measurement error in hours, which is well known to be pervasive, can affect our statistics. In general, the levels of variances and covariances are potentially affected by measurement error. However, as long as the measurement error (1) is multiplicative in levels, (2) is orthogonal to the true value, and (3) exhibits constant variance over the period, the *changes* in these measured cross-sectional moments, which are the inputs to our cross-sectional calculations, will not be affected.

[24] This first-stage regression ensures consistency with the consumption data, since Krueger and Perri (2006) report cross-sectional variances for log consumption using residuals from a similar regression.

[25] Note that our PSID sample has an earlier starting date than the CEX, which is only available on a consistent basis since 1980. Fortunately, almost all the observed rise in wage inequality occurred after this date.

Table 1
Welfare effects (% of lifetime consumption)

Welfare effect of rise in wage dispersion				Welfare gain from completing markets		Welfare effect from eliminating risk	
Model-based		Observables-based					
Cobb–Douglas preferences							
ω		ω		χ		κ	
−2.47% (−2.50%)		−2.75%		+39.1% (+33.0%)		+16.9% (+15.5%)	
Volat.	Level	Volat.	Level	Volat.	Level	Volat.	Level
−7.50%	+5.00%	−4.45%	+1.70%	+11.0%	+22.0%	+28.5%	−13.0%
Separable preferences							
ω		ω		χ		κ	
−3.06% (−3.13%)		−2.30%		+29.2% (+24.8%)		+17.8% (+16.0%)	
Volat.	Level	Volat.	Level	Volat.	Level	Volat.	Level
−4.38%	+1.25%	−4.00%	+1.70%	+8.3%	+16.5%	+17.0%	−1.0%

Measurement of insurable/uninsurable wage components: A simple permanent/transitory model for the variance of log wages is estimated, exactly the process specified in the description of the model economy. The estimated variance of the transitory component v_ε starts around 0.08 in the late 1960s and levels off 30 years later at around 0.13. The variance of the permanent component v_α starts at a value around 0.17 and rises to 0.22 in the mid 1990s. In light of these results, the changes in variances are set to $\Delta v_\varepsilon = \Delta v_\alpha = 0.05$. Moreover, focusing on the levels of labor market uncertainty for the 1990s, the levels of variances are set to $v_\alpha = 0.22$ and $v_\varepsilon = 0.13$.[26]

An alternative approach to estimating $(\Delta v_\varepsilon, \Delta v_\alpha)$ would involve using expressions for the variances and covariances of wages, hours and consumption, which can be derived in closed form given the equilibrium decision rules in Sections 4 and 5. The idea is that observed changes in second moments involving endogenous variables are informative about changes in the variances of underlying insurable and uninsurable shocks. We pursue this strategy in a companion paper (Heathcote et al., 2007b).

7.2. Results

The results are summarized in Table 1. To gauge the quality of our approximations relative to the exact welfare expressions contained in the Appendix, the values implied by the approximated welfare expressions described in Propositions 1 and 1a are reported in parentheses. Below the total welfare changes, the approximated welfare effects are decomposed into level and volatility components.

Welfare effects of rising dispersion, model-based approach: The welfare losses associated with the observed rise in wage dispersion are quite similar across the two alternative preference specifications, between 2.5% and 3% of lifetime consumption. With Cobb–Douglas preferences, the welfare loss due to the volatility component is 7.5% of lifetime consumption, while the partially offsetting welfare gain due to improved aggregate labor productivity is 5%. With separable preferences, both components are smaller in absolute value. One reason is that the Frisch elasticity is lower under the separable specification, which implies that additional insurable risk translates into a smaller increase in dispersion in hours worked, and a smaller increase in aggregate productivity.[27]

Welfare effects of rising dispersion, observables-based approach: For the separable preferences case (assuming $\gamma = \sigma = 2$), the observed changes in the variances of hours, consumption, and in the covariance between hours

[26] Our findings can be summarized as follows: (i) the transitory component accounts for roughly $\frac{1}{3}$ of the total dispersion; (ii) the two components each account for about half the rise in wage dispersion. These results are broadly in line with the findings by Gottschalk and Moffitt (1994).

[27] In addition, the fact that $\gamma > 1$ means that additional uninsurable risk reduces average labor productivity when preferences are separable since strong wealth effects induce agents who are permanently more productive to increase leisure.

and wages can be plugged into the expression for ω in Proposition 2, which gives $\omega = -2.3\%$. A similar computation for the Cobb–Douglas case, which involves the covariance between hours and consumption, yields $\omega = -2.75\%$.

These estimates are very close to those from the model-based approach, which is encouraging given that the two sets of calculations rely on very different inputs. The reason these two approaches give similar answers is that the observed changes in empirical cross-sectional variances and covariances for wages, hours and consumption are quantitatively close to the changes in these moments that are implied by our partial-insurance model, given the calibration of preference parameters and $(\Delta v_{\varepsilon,} \Delta v_{\alpha})$. In Heathcote et al. (2007b) we investigate this point further.

Krueger and Perri (2003) propose evaluating welfare effects using individual consumption data. Using the panel-dimension of CEX, they estimate Markov transition matrices for consumption and hours worked in two sub-samples (before and after the rise in wage inequality) and feed these stochastic processes directly into preferences to compute welfare effects. However, in constructing their data, they abstract from the level effect by demeaning all observations, so their calculations should be compared to our volatility effect. They assume Cobb–Douglas preferences and set $\bar{\phi} = \bar{\gamma} = 1.33$. Given the observed changes in $cov(\log h, \log w)$, $var(\log c)$ and $var(\log h)$, this parameterization maps into a volatility effect of $\omega^{\text{vol}} = -2.5\%$, which is quite close to their estimated welfare loss of -2.1%.[28]

From Proposition 2, it follows that the degree to which a society is able to allocate labor efficiently—labor productivity—has the simple empirical representation $cov(\log h, \log w)$, irrespective of preferences. In our PSID sample, labor productivity, measured as the ratio of aggregate earnings to aggregate hours, increased by 13% from 1975 to 1995. Thus, the increase in the wage–hours covariance (1.7%) can alone account for more than a tenth of the increase in aggregate labor productivity over this period.

Welfare gains from completing markets: With Cobb–Douglas preferences, a household in the partial-insurance economy values the availability of a complete set of insurance markets against the permanent component of wages at 39% of her lifetime consumption. With separable preferences, this estimate is smaller, around 29%. The striking feature of these results is that, in both cases, the gains associated with better productive opportunities in complete markets are twice as large as the gains from reduced dispersion. Recall that in the separable case, since $\gamma > 1$, households with low-permanent (uninsurable) wage components work longer hours than those with high-permanent components. However, efficiency dictates a positive correlation between wages and hours. Our calculations indicate that the aggregate productivity loss due to this inefficient assignment is huge, accounting for two thirds of the welfare cost of market incompleteness.

Attanasio and Davis (1996, Table 6) calculated the gains from insuring all consumption risk between age/educational groups to be around 2.67% for a risk aversion value of 2. This number is of an order of magnitude lower than ours for two reasons. First, the empirical data show a large amount of consumption dispersion even *within* groups, which is captured in our calculations. Second, by abstracting from labor supply, they miss the level effect that we find to be the largest source of welfare gains from completing markets.

Welfare gains from eliminating risk: The welfare gains from eliminating labor market risk (i.e., from full wage compression) are $\kappa = 16.9\%$ in the Cobb–Douglas case. It is instructive to compare this number to the value of insuring risk, i.e., $\chi = 39.1\%$. Thus, eliminating risk implies a welfare gain less than half that from completing markets. The corresponding numbers for the separable preferences case are $\kappa = 17.8\%$ and $\chi = 29.2\%$.

The fact that eliminating the *insurable* component of wage risk is welfare-reducing leaves open the theoretical possibility that the welfare effect from eliminating all idiosyncratic wage risk through some redistributive policy might be negative. In our calibration to the United States, however, most wage dispersion is uninsurable, and given plausible choices for preference parameters, the welfare gains from eliminating uninsurable risk exceed the costs of eliminating insurable risk.

[28] This number is obtained as follows. Krueger and Perri (2003) report that when using consumption data only, welfare losses are about -1.6%. Incorporating leisure into their analysis subtracts another 0.5% from their benchmark estimate.

7.3. Relation to numerically-solved Bewley models

How do the results from our analytical model compare to standard incomplete-market models relying on self-insurance through hours worked and borrowing and lending? To provide a natural and comparable benchmark, we compute the equilibrium of an economy identical to our partial-insurance model, except that instead of having access to a complete set of state-contingent claims providing perfect insurance against transitory wage shocks, agents trade only a non-contingent bond (e.g., Bewley, 1986; Imrohoruglu, 1989; Huggett, 1993; Aiyagari, 1994; Ríos-Rull, 1994). At the aggregate level, bonds are in zero net supply.

In the Bewley economy, the welfare effect associated with an increase in wage dispersion will depend on two additional parameters that could be left unspecified in the partial-insurance model: the borrowing limit, and the discount factor β.[29] The borrowing constraint is set to the "natural" limit (see, for example, Aiyagari, 1994) which ensures that interest payments never exceed earnings, given maximum labor effort. Preferences are Cobb–Douglas and the discount factor is $\beta = 0.97$, which implies a final steady state interest rate of 3.05%. The expected welfare effects are computed for individuals born with zero wealth who draw lifetime wage profiles at random from the unconditional wage distribution.

The expected welfare effect in the Bewley economy associated with the measured rise in wage dispersion is a 2.77% loss. This number should be compared to the 2.37% loss in our partial-insurance economy.[30] Increases in wage risk are slightly more costly in the model with a single bond, because with a positive interest rate a transitory wage shock has some effect on lifetime income. Nevertheless, the two models deliver surprisingly similar answers to our main welfare question.

For comparisons with other papers, consider e.g. Pijoan-Mas (2006). He calculates the welfare gain from completing markets to be about 16% of lifetime consumption in an infinitely lived-agent, production economy with separable preferences and flexible labor supply. Since fixed effects implicitly remain uninsured under his interpretation of what it means for markets to be complete, this number should be compared to the welfare gain of moving from autarky to an environment where the transitory component of wages is fully insured, while the permanent component remains uninsured (i.e. the partial-insurance economy). This is given by $\omega(-v_\varepsilon, v_\varepsilon)$. Setting $\gamma = 0.98$ and $\sigma = 0.61$ (the parametrization used by Pijoan-Mas) gives 18.6%.

We conclude that our simple and transparent framework can shed light on the economics underlying numerical findings in richer models in the Bewley tradition.

However, our framework cannot be expected to match the *quantitative* welfare effects when additional channels of opportunities and insurance are introduced, over and above savings and hours worked. For example, in Heathcote et al. (2007a), we explore two additional choices that allow households to increase average labor productivity and mitigate the welfare loss in response to changes in the wage structure: the flexibility to adjust enrollment decisions in response to a widening college premium, and the flexibility to reallocate market work within the household in response to a narrowing gender wage gap.

8. Concluding remarks

The main contributions of the paper are (1) the analytical characterization of the welfare effects from an increase in the dispersion of labor productivity, and (2) the focus on the role of endogenous labor supply. In addition, welfare effects are shown to have a common representation in terms of observable second moments (variances and covariances) of the joint equilibrium distribution of wages, hours worked and consumption. This is true for both Cobb–Douglas and separable preferences.

Analytical insights, together with a simple calibration exercise, show that eliminating idiosyncratic wage risk implies a welfare gain that is at least two orders of magnitude larger than most estimates of the welfare gains

[29] Levine and Zame (2002) study an endowment economy with infinitely lived agents, CRRA preferences, and a non-contingent bond as the only traded asset. They show that as the discount rate goes to zero, agents achieve arbitrarily good insurance against non-permanent shocks.

[30] To solve this model numerically, both the permanent component of the wage α and the transitory component ε are drawn from symmetric two-point distributions. Given this two-point distribution, the welfare effect from increased wage dispersion in our benchmark partial-insurance economy is -2.37% as compared to the -2.47% loss reported in Table 1 for the continuous normal distribution. More details on the numerical implementation are available upon request.

from eliminating business cycle risk. Thus, according to our model, the potential gain for a society from applying progressive taxes and wage compression is much larger than the potential gain from aggregate stabilization.

However, we also emphasized that the welfare gains from *eliminating* wage risk (through policies that compress after-tax wages) are only around half as large as the gains that would accrue from perfectly *insuring* wage risk. From a policy perspective, an important implication is that the government should develop the legal and institutional frameworks that will allow new insurance markets to develop. Sargent (2001) and Shiller (2003) discuss a range of proposals along these lines.[31]

Throughout the analysis, income shocks have been assumed to be verifiable and contracts have been assumed to be perfectly enforceable. Informational asymmetries and imperfect commitment may limit the amount of insurance that one could ever hope to see provided. In this sense, our estimates of the welfare costs of market incompleteness are upper bounds. At the same time, to the extent that risk sharing is limited by fundamental frictions, these frictions can interact with changes in labor market risk in interesting ways. For example, Krueger and Perri (2006) study a calibrated endowment economy in which debt contracts can only be imperfectly enforced. They show that a rise in income dispersion might *increase* welfare by making default more painful, thereby increasing the amount of credit that can be supported in equilibrium. The implications of increased labor market risk in a private information environment have not yet been addressed. More broadly, an important challenge in introducing these sorts of frictions is to do this in a way that maintains tractability, thereby allowing for a transparent characterization of the various mechanisms at work.

Finally, the trade-off between insurance and opportunities emphasized in relation to the labor supply decision could also apply to other margins of adjustment. For example, the widening gap between the wages of college and high-school graduates offers opportunities to increase average earnings if agents can respond by extending their education. As discussed above, Heathcote et al. (2007a) incorporate an explicit education choice, and explore how introducing this margin of adjustment (as well as joint labor supply decisions in two-member households) mediates the effect of changes in the wage structure on labor productivity and welfare.

Appendix A. Derivation of equilibrium allocations (Cobb–Douglas case)

Consider the "α-island representation" of the economy outlined at the end of Section 2. Start by guessing that agents on different α-islands do not trade with each other. This implies that allocations within islands can be obtained from an island-planner problem. Finally, the initial guess is verified.

The static equal-weight planner problem for an island indexed by a specific value of α can be written as

$$\max_{\{c(\alpha,\varepsilon)h(\alpha,\varepsilon)\}} \int_E u(c(\alpha,\varepsilon)h(\alpha,\varepsilon))\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)$$

subject to a static resource constraint that reflects the absence of interisland trade and the lack of a storage technology

$$\int_E w(\alpha,\varepsilon)h(\alpha,\varepsilon) - c(\alpha,\varepsilon)\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon) = 0. \tag{14}$$

With Cobb–Douglas utility, the planner's first-order condition for hours is

$$w(\alpha,\varepsilon)h(\alpha,\varepsilon) = w(\alpha,\varepsilon) - c(\alpha,\varepsilon)\frac{1-\eta}{\eta}. \tag{15}$$

Substituting the right-hand side of this latter equation into Eq. (14) and collecting terms determines total consumption on the island

$$\int_E c(\alpha,\varepsilon)\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon) = \eta\int_E w(\alpha,\varepsilon)\,\mathrm{d}\Phi_{v_\varepsilon}(\varepsilon) = \eta\exp(\alpha). \tag{16}$$

[31]For example, Shiller proposed six types of insurance that should be further developed, namely "livelihood insurance", "home equity insurance", "macromarkets", "income-linked loans", "inequality insurance" and "intergenerational social security".

The first-order condition for consumption is

$$\mu = \eta c(\alpha, \varepsilon)^{\eta(1-\theta)-1}(1 - h(\alpha, \varepsilon))^{(1-\eta)(1-\theta)}, \tag{17}$$

where μ is the Lagrange multiplier on the resource constraint (14). Using (15) to substitute out for leisure in (17) and rearranging gives

$$c(\alpha, \varepsilon) = \left(\frac{\eta}{\mu}\right)^{1/\theta} \left(\frac{1-\eta}{\eta}\right)^{(1-\eta)(1-\theta)/\theta} w(\alpha, \varepsilon)^{-(1-\eta)(1-\theta)/\theta}. \tag{18}$$

Integrating (18) across the population yields an alternative expression for total consumption

$$\begin{aligned}\int_E c(\alpha, \varepsilon)\, \mathrm{d}\Phi_{v_\varepsilon}(\varepsilon) &= \left(\frac{\eta}{\mu}\right)^{1/\theta} \left(\frac{1-\eta}{\eta}\right)^{(1-\eta)(1-\theta)/\theta} \int_E \exp\left(-\frac{(1-\eta)(1-\theta)}{\theta}(\alpha+\varepsilon)\right) \mathrm{d}\Phi_{v_\varepsilon}(\varepsilon) \\ &= \left(\frac{\eta}{\mu}\right)^{1/\theta} \left(\frac{1-\eta}{\eta}\right)^{(1-\eta)(1-\theta)/\theta} \exp\left(\frac{(1-\eta)(\theta-1)}{\theta}\left(\alpha - \left(\frac{1-\eta+\eta\theta}{\theta}\right)\frac{v_\varepsilon}{2}\right)\right),\end{aligned}$$

where the last step exploits the fact that ε is log-normal. Combining this last equation with (16) yields an expression for μ. Substituting this expression into (18) to solve for consumption, and then using (15) to solve for hours yields the candidate equilibrium allocations, as functions of primitive parameters, reported in Section 4.1 in the main text.

The last step of the proof requires verifying the no-trade guess. At the candidate allocations $c(\alpha, \varepsilon)$ and $h(\alpha, \varepsilon)$, the agent's Euler equation

$$u_c(c(\alpha, \varepsilon), h(\alpha, \varepsilon)) = \beta R \int_E u_c(c(\alpha, \varepsilon'), h(\alpha, \varepsilon'))\, \mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)$$

yields an interest rate of $R = 1/\beta$ which supports the equilibrium without trade across α-islands, since it is independent of α.

Suppose that net savings are zero for every agent. The budget constraint then implies

$$b(\varepsilon'; (\alpha, \varepsilon)) = c(\alpha, \varepsilon') - w(\alpha, \varepsilon')h(\alpha, \varepsilon') = \exp\left(\alpha + (1-\lambda)\varepsilon' + \lambda(1-\lambda)\frac{v_\varepsilon}{2}\right) - \exp(\alpha + \varepsilon').$$

The first-order condition for the purchase of Arrow securities paying one unit of consumption in the event that an individual with state (α, ε) receives shock $\varepsilon' \in \mathscr{E}$ next period is

$$u_c(c(\alpha, \varepsilon), h(\alpha, \varepsilon))p(\mathscr{E}) = \beta \int_{\mathscr{E}} u_c(c(\alpha, \varepsilon'), h(\alpha, \varepsilon'))\, \mathrm{d}\Phi_{v_\varepsilon}(\varepsilon),$$

which yields $p(\mathscr{E}) = \beta \int_{\mathscr{E}} \mathrm{d}\Phi_{v_\varepsilon}(\varepsilon)$, i.e., asset prices are discounted probabilities. It is then straightforward to verify the guess that the net savings (i.e., the cost of the entire portfolio of Arrow securities) is zero:

$$\int_E b(\varepsilon'; (\alpha, \varepsilon))p(\varepsilon')\, \mathrm{d}\varepsilon' = \beta \int_E b(\varepsilon'; (\alpha, \varepsilon))\, \mathrm{d}\Phi_{v_\varepsilon}(\varepsilon') = 0.$$

Appendix B. Proof of Proposition 1

Start by computing unconditional expected utility

$$\mathscr{W} = \mathbb{E}\left[\frac{(c(\alpha, \varepsilon)^\eta (1 - h(\alpha, \varepsilon))^{1-\eta})^{1-\theta}}{1-\theta}\right] = \frac{((1-\eta/\eta))^{(1-\eta)(1-\theta)}}{1-\theta} \mathbb{E}[\exp(-(1-\eta)(1-\theta)(\alpha+\varepsilon))c(\alpha, \varepsilon)^{(1-\theta)}],$$

where the second equality follows from the intratemporal first-order condition (15). Substituting for the equilibrium expression for $c(\alpha, \varepsilon)$ in (11), expected utility becomes

$$\begin{aligned}\mathscr{W} &= \kappa \mathbb{E} \exp\left(\eta(1-\theta)\alpha - (1-\eta)(1-\theta)\varepsilon + \frac{(1-\theta)(1-\eta)(\theta-1)}{\theta}\left(\varepsilon + \frac{1-\eta+\eta\theta}{\theta}\frac{v_\varepsilon}{2}\right)\right) \\ &= \kappa \exp\left(\frac{(1-\eta+\eta\theta)(1-\eta)(1-\theta)}{\theta}\frac{v_\varepsilon}{2} - (1-\eta+\eta\theta)\eta(1-\theta)\frac{v_\alpha}{2}\right), \end{aligned} \tag{19}$$

where $\kappa \equiv ((1-\eta)/\eta)^{(1-\eta)(1-\theta)}\eta^{1-\theta}/(1-\theta)$, and where the second equation follows from α and ε being log-normal. Recall that ω is defined by Eq. (4). Substituting (19) into (4) and collecting terms yields an exact expression for ω

$$1+\omega = \exp\left(\frac{1-\eta}{\eta}\left(\frac{1-\eta+\eta\theta}{\theta}\right)\frac{\Delta v_\varepsilon}{2} - (1-\eta+\eta\theta)\frac{\Delta v_\alpha}{2}\right) = \exp\left(\bar{\phi}\frac{\Delta v_\varepsilon}{2} - \bar{\gamma}\frac{\Delta v_\alpha}{2}\right).$$

Taking logarithms on both sides and using a log-approximation of the type $\ln(1+\omega) \simeq \omega$ for the left-hand side yields the expression stated in Proposition 1.

We now show how to decompose ω into a level effect and a volatility effect. From Eq. (5), the level effect of changing variances from $(v_\alpha, v_\varepsilon)$ to $(\hat{v}_\alpha, \hat{v}_\varepsilon)$ is given by

$$\frac{((1+\omega^{\text{lev}})^\eta C^\eta (1-H)^{1-\eta})^{1-\theta}}{1-\theta} = \frac{(\widehat{C}^\eta (1-\widehat{H})^{1-\eta})^{1-\theta}}{1-\theta}, \tag{20}$$

where aggregate consumption and leisure are given by

$$\begin{aligned} C &= \mathbb{E}[c(\alpha,\varepsilon)] = \eta, \\ 1-H &= \mathbb{E}[1-h(\alpha,\varepsilon)] = (1-\eta)\exp(\lambda v_\varepsilon). \end{aligned}$$

Since C is invariant to wage dispersion, it follows that $(1+\omega^{\text{lev}})^\eta(1-H)^{1-\eta} = (1-\widehat{H})^{1-\eta}$, which yields the level effect of Proposition 1. Flodén (2001) shows that if $u(\cdot)$ is such that $u(xc,h) = g(x)u(c,h)$, then

$$1+\omega = (1+\omega^{\text{lev}})(1+\omega^{\text{vol}}) \Rightarrow \omega \simeq \omega^{\text{lev}} + \omega^{\text{vol}}, \tag{21}$$

up to second-order terms. Since Cobb–Douglas preferences satisfy this homogeneity property, Eq. (21) defines ω^{vol} residually, given ω^{lev}.

Appendix C. Proof of Proposition 2 (Cobb–Douglas case)

From the log-normality of c and $1-h$ (Assumption A1), expected utility is given by

$$\begin{aligned}\mathbb{E}\left[\frac{1}{1-\theta}(c^\eta(1-h)^{1-\eta})^{1-\theta}\right] &= \frac{1}{1-\theta}\mathbb{E}[\exp((1-\theta)\eta\log c + (1-\theta)(1-\eta)\log(1-h))] \\ &= \frac{1}{1-\theta}\exp\Big((1-\theta)\eta\mu_c + (1-\theta)(1-\eta)\mu_l \\ &\quad + \frac{(1-\theta)^2[\eta^2 v_c + (1-\eta)^2 v_l + 2\eta(1-\eta)v_{cl}]}{2}\Big),\end{aligned} \tag{22}$$

where the notation $\mathbb{E}[\log x] \equiv \mu_x$, $var(\log x) \equiv v_x$, and $cov(\log x, \log y) \equiv v_{xy}$ have used for any variables x and y. The welfare effect ω of changing to a new distribution of allocations (denoted with hats) is defined as

$$\mathbb{E}[u(\hat{c},\hat{h})] = \mathbb{E}[u((1+\omega)c,h)]. \tag{23}$$

Substituting (22) into (23), implies

$$\begin{aligned}&\eta(1-\theta)\log(1+\omega) + (1-\theta)\eta\mu_c + (1-\theta)(1-\eta)\mu_l + \frac{(1-\theta)^2[\eta^2 v_c + (1-\eta)^2 v_l + 2\eta(1-\eta)v_{cl}]}{2} \\ &\quad = (1-\theta)\eta\hat{\mu}_c + (1-\theta)(1-\eta)\hat{\mu}_l + \frac{(1-\theta)^2[\eta^2 \hat{v}_c + (1-\eta)^2 \hat{v}_l + 2\eta(1-\eta)\hat{v}_{cl}]}{2}.\end{aligned}$$

Rearranging terms, using the fact that $\log(1+\omega) \simeq \omega$ for ω small, and noting that $C \equiv \mathbb{E}(c) = \exp(\mu_c + v_c/2)$, and $1 - H \equiv \mathbb{E}(1-h) = \exp(\mu_l + v_l/2)$ yields

$$\begin{aligned}\omega &\simeq \Delta \log C + \frac{1-\eta}{\eta}\Delta\log(1-H) - \frac{1-(1-\theta)\eta}{2}\Delta v_c - \frac{1-\eta}{2\eta}[1-(1-\theta)(1-\eta)]\Delta v_l + (1-\theta)(1-\eta)\Delta v_{cl} \\ &= \Delta\log C + \frac{1-\eta}{\eta}\Delta\log(1-H) - \frac{1}{2}\bar{\gamma}\Delta v_c - \frac{1}{2}\left[\bar{\gamma} - 1 + \frac{\eta}{1-\eta}\right]\left(\frac{1-\eta}{\eta}\right)^2 \Delta v_l + (1-\bar{\gamma})\left(\frac{1-\eta}{\eta}\right)\Delta v_{cl}. \quad (24)\end{aligned}$$

By Assumption A2, the individual intratemporal first-order condition is satisfied:

$$(1-\eta)c = \eta w(1-h). \quad (25)$$

Taking expectations of (25), and using $W \equiv \mathbb{E}(w) = 1$ by Assumption A1, as well as $\mathbb{E}(wh) = C$ by Assumption A2, yields $C = \eta$ (and therefore, $\Delta\log C = 0$). At the same time, Assumption A2 and the log-normality of the allocations implies that

$$C = \mathbb{E}(wh) = W - \mathbb{E}[w(1-h)] = 1 - \exp\left(\mu_l + \frac{v_l}{2} + v_{wl}\right) = 1 - (1-H)\exp(v_{wl}), \quad (26)$$

where the fact that $W = 1$ has been used. Setting $C = \eta$,

$$\Delta v_{wl} = -\Delta\log(1-H). \quad (27)$$

Now, note that, for small deviations of h from its mean

$$\begin{aligned}\log(1-h) &= \log(1-H) + \log\frac{1-h}{1-H} \simeq \log(1-H) + \frac{1-h}{1-H} - 1 \\ &= \log(1-H) - \frac{H}{1-H}\left(\frac{h}{H} - 1\right) \simeq \log(1-H) - \frac{H}{1-H}\log\left(\frac{h}{H}\right) \\ &\simeq \log(1-H) + \frac{H}{1-H}\log H - \frac{\eta}{1-\eta}\log h,\end{aligned}$$

where the last approximation uses $H \simeq \eta$, which is true for v_{wl} small (see Eq. (26)). Exploiting this relationship between $\log(1-h)$ and $\log h$, it is easily seen that

$$v_h \simeq \left(\frac{1-\eta}{\eta}\right)^2 v_l, \quad v_{wh} \simeq -\frac{1-\eta}{\eta}v_{wl} \quad \text{and} \quad v_{ch} \simeq -\frac{1-\eta}{\eta}v_{cl}. \quad (28)$$

Substituting (27) and (28) into Eq. (24) immediately yields the representation for ω in Proposition 2 for the Cobb–Douglas case.

It remains to be shown how to decompose ω into a level effect and a volatility effect. After rearranging terms, the definition of the level effect in (20) for the Cobb–Douglas case implies

$$\log(1+\omega^{\text{lev}}) \simeq \omega^{\text{lev}} = \Delta\log C + \frac{1-\eta}{\eta}\Delta\log(1-H). \quad (29)$$

Since $\Delta\log C = 0$, Eqs. (27) and (28) imply that $\omega^{\text{lev}} \simeq v_{wh}$. Note that since $H \simeq \eta$, the following approximation applies

$$\frac{1-\eta}{\eta}\Delta\log(1-H) \simeq \left(\frac{1-\eta}{\eta}\right)\frac{\Delta(1-H)}{1-H} \simeq \frac{-\Delta H}{H} \simeq -\Delta\log H.$$

This approximation together with Eq. (29) imply that $\omega^{\text{lev}} \simeq \Delta\log(C/H)$. As argued in the proof of Proposition 1, the volatility effect is defined as the residual $\omega - \omega^{\text{lev}}$.

References

Acemoglu, D., 2002. Technical change, inequality and the labor market. Journal of Economic Literature 40, 7–72.
Aghion, P., 2002. Schumpeterian growth theory and the dynamics of income inequality. Econometrica 70, 855–882.
Aiyagari, R., 1994. Uninsured idiosyncratic risk and aggregate savings. Quarterly Journal of Economics 109, 659–684.

Altonji, J., Hayashi, F., Kotlikoff, L., 1992. Is the extended family altruistically linked: direct tests using micro data. American Economic Review 82, 1177–1198.

Attanasio, O., 1999. Consumption. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 1B. Elsevier, Amsterdam, pp. 741–812.

Attanasio, O., Davis, S., 1996. Relative wage movements and the distribution of consumption. Journal of Political Economy 104, 1227–1262.

Attanasio, O., Battistin, E., Ichimura, H., 2007. What really happened to consumption inequality in the U.S.? In: Berndt, E., Hulten, C. (Eds.), Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches. Chicago University Press, Chicago, pp. 515–543.

Benabou, R., 2002. Tax and education policy in a heterogeneous agent economy: what levels of redistribution maximizes growth and efficiency? Econometrica 70, 481–517.

Bewley, T., 1986. Stationary monetary equilibrium with a continuum of independently fluctuating consumers. In: Hildenbrand, W., Mas-Colell, A. (Eds.), Contributions to Mathematical Economics: In Honor of Gerard Debreu. North-Holland, New York, pp. 79–102.

Blundell, R., MaCurdy, T., 1999. Labor supply: a review of alternative approaches. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3. Elsevier, Amsterdam, pp. 1559–1695.

Blundell, R., Preston, I., 1998. Consumption inequality and income uncertainty. Quarterly Journal of Economics 113, 603–640.

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. American Economic Review, forthcoming.

Browning, M., Meghir, C., 1991. The effects of male and female labor supply on commodity demands. Econometrica 59, 925–951.

Browning, M., Hansen, L.P., Heckman, J., 1999. Micro data and general equilibrium models. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 1A. Elsevier, Amsterdam, pp. 543–633.

Carroll, C., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. Quarterly Journal of Economics 112, 1–56.

Cho, J.O., Cooley, T.F., 2001. Business cycle uncertainty and economic welfare. Mimeo, Stern School of Business, New York University.

Cochrane, J., 1991. A simple test of consumption insurance. Journal of Political Economy 99, 957–976.

Cooley, T., 1995. Frontiers of Business Cycle Research. Princeton University Press, Princeton.

Cordoba, J., Verdier, G., 2007. Lucas vs. Lucas: on inequality and growth. IMF Working Paper No. 07/17.

Deaton, A., 1991. Saving and liquidity constraints. Econometrica 59, 1221–1248.

Domeij, D., Flodén, M., 2006. The labor-supply elasticity and borrowing constraints: why estimates are biased. Review of Economic Dynamics 9, 242–262.

Eckstein, Z., Nagypal, E., 2004. The evolution of U.S. earnings inequality: 1961–2002. Federal Reserve Bank of Minneapolis Quarterly Review 28, 10–29.

Flodén, M., 2001. The effectiveness of government debt and transfers as insurance. Journal of Monetary Economics 48, 81–108.

Gomes, J., Greenwood, J., Rebelo, S., 2001. Equilibrium unemployment. Journal of Monetary Economics 48, 109–152.

Gottschalk, P., Moffitt, R., 1994. The growth of earnings instability in the U.S. labor market. Brookings Papers of Economic Activity 2, 217–272.

Guiso, L., Pistaferri, L., Schivardi, F., 2005. Insurance within the firm. Journal of Political Economy 113, 1054–1087.

Heathcote, J., Storesletten, K., Violante, G.L., 2007a. The macroeconomic implications of rising wage inequality in the United States. Mimeo, NYU.

Heathcote, J., Storesletten, K., Violante, G.L., 2007b. Consumption and labor supply with partial insurance: an analytical framework. Mimeo, NYU.

Heathcote, J., Storesletten, K., Violante, G.L., 2007c. Insurance and opportunities: the welfare implications of rising wage dispersion. Mimeo, NYU ⟨http://www.econ.nyu.edu/user/violante⟩.

Heckman, J., 1974. Life-cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life-cycle. American Economic Review 64, 188–194.

Hornstein, A., Krusell, P., Violante, G.L., 2005. The effects of technical change on labor market inequalities. In: Aghion, P., Durlauf, S. (Eds.), Handbook of Economic Growth 1(B). Elsevier, Amsterdam, pp. 1275–1370.

Huggett, M., 1993. The risk free rate in heterogeneous-agents incomplete-insurance economies. Journal of Economic Dynamics and Control 17, 953–969.

Huggett, M., Parra, J.C., 2006. How well does the U.S. social insurance system provide social insurance? Mimeo, Georgetown University.

Imrohoruglu, A., 1989. Cost of business cycles with indivisibilities and liquidity constraints. Journal of Political Economy 97, 1364–1383.

Katz, L.M., Autor, D., 1999. Changes in the wage structure and earnings inequality. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3A. Elsevier, Amsterdam, pp. 1463–1555.

Krebs, T., 2003. Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk. Review of Economic Dynamics 6, 846–868.

Krebs, T., Krishna, P., Maloney, W., 2005. Trade policy, income risk, and welfare. NBER Working Paper No. 11255.

Krueger, D., Perri, F., 2003. On the welfare consequences of the increase in inequality in the United States. In: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2003. MIT Press, Cambridge, pp. 83–121.

Krueger, D., Perri, F., 2006. Does income inequality lead to consumption inequality: evidence and theory. Review of Economic Studies 73, 163–193.

Krusell, P., Smith Jr., A., 1999. On the welfare effects of eliminating business cycles. Review of Economic Dynamics 2, 245–272.

Levine, D.K., Zame, W.R., 2002. Does market incompleteness matter? Econometrica 70, 1805–1839.

Livshits, I., MacGee, J., Tertilt, M., 2007. Consumer bankruptcy: a fresh start. American Economic Review 97, 402–418.

Low, H., 2005. Self-insurance in a life-cycle model of labor supply and savings. Review of Economic Dynamics 8, 945–975.

Lucas Jr., R.E., 1987. Models of Business Cycle. Basil Blackwell, Oxford.

Lucas Jr., R.E., 2003. Macroeconomic priorities. American Economic Review 93, 1–14.
Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. Microeconomic Theory. Oxford University Press, Oxford.
Mortensen, D., 2003. Wage Dispersion: Why are Similar Workers Paid Differently? MIT Press, Cambridge.
Phelan, C., 2006. Opportunity and social mobility. Review of Economic Studies 73, 487–505.
Pijoan-Mas, J., 2006. Precautionary savings or working longer hours? Review of Economic Dynamics 9, 326–352.
Ríos-Rull, V., 1994. On the quantitative importance of market incompleteness. Journal of Monetary Economics 34, 463–496.
Sargent, T., 2001. Discussion of "can market and voting institutions generate optimal intergenerational risk sharing," by A. Rangel and R. Zeckhauser. In: Campbell, J.Y., Feldstein, M. (Eds.), Risk Aspects of Investment-Based Social Security Reform. University of Chicago Press, Chicago, pp. 113–152.
Shiller, R.J., 2003. The New Financial Order: Risk in the 21st Century. Princeton University Press, Princeton.
Slesnick, R., 2001. Consumption and Social Welfare: Living Standards and Their Distribution in the United States. Cambridge University Press, Cambridge.
Storesletten, K., Telmer, C., Yaron, A., 2001. The welfare cost of business cycles revisited: finite lives and cyclical variation in idiosyncratic risk. European Economic Review 45, 1311–1339.