

Schulhofer-Wohl, Sam

Working Paper

Heterogeneous Risk Preferences and the Welfare Cost of Business Cycles

Discussion Papers in Economics, No. 235

Provided in Cooperation with:

Woodrow Wilson School of Public and International Affairs (WWS), Princeton University

Suggested Citation: Schulhofer-Wohl, Sam (2007) : Heterogeneous Risk Preferences and the Welfare Cost of Business Cycles, Discussion Papers in Economics, No. 235, Princeton University, Woodrow Wilson School of Public and International Affairs, Princeton, NJ

This Version is available at:

<https://hdl.handle.net/10419/59030>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Princeton University

**Heterogeneous Risk Preferences and the
Welfare Cost of Business Cycles**

by
Sam Schulhofer-Wohl
Princeton University

Discussion Paper #235

January 2008

Discussion Papers in Economics

Woodrow Wilson School
of Public and International Affairs



Heterogeneous Risk Preferences and the Welfare Cost of Business Cycles

Sam Schulhofer-Wohl*

January 25, 2008

Abstract

I study the welfare cost of business cycles in a complete-markets economy where some people are more risk averse than others. Relatively more risk-averse people buy insurance against aggregate risk, and relatively less risk-averse people sell insurance. These trades reduce the welfare cost of business cycles for everyone. Indeed, the least risk-averse people benefit from business cycles. Moreover, even infinitely risk-averse people suffer only finite and, in my empirical estimates, very small welfare losses. In other words, when there are complete insurance markets, aggregate fluctuations in consumption are essentially irrelevant not just for the average person – the surprising finding of Lucas (1987) – but for everyone in the economy, no matter how risk averse they are. If business cycles matter, it is because they affect productivity or interact with uninsured idiosyncratic risk, not because aggregate risk *per se* reduces welfare.

Keywords: business cycles; risk aversion; risk sharing; heterogeneity.

JEL classification: E32, E21.

*Department of Economics and Woodrow Wilson School of Public and International Affairs, Princeton University, 363 Wallace Hall, Princeton, NJ 08544. E-mail: sschulho@princeton.edu. Phone: (609) 258-7392. I thank Pierre-André Chiappori, James Heckman, Robert Shimer and especially Robert Townsend for invaluable discussions. The coordinating editor, associate editor and an anonymous referee also made many helpful suggestions, as did numerous seminar participants. The National Institute on Aging provided generous financial support while I wrote the first draft of this paper. Portions of this research circulated previously as part of a longer manuscript titled “Heterogeneity, Risk Sharing and the Welfare Costs of Risk.”

1 Introduction

Ever since Lucas (1987) demonstrated that business cycles have minuscule welfare costs in a representative-agent model, researchers have tried to find alternative contexts in which business cycles do have meaningful welfare costs. One of the most fruitful ideas has been to consider the possibility that people are heterogeneous and, therefore, that business cycles have larger welfare costs for some people than for others. This paper studies a competing and, to my knowledge, previously unexplored phenomenon: Heterogeneity creates more opportunities for trade, thereby reducing the welfare cost of business cycles for everyone. I use theory and data to show that when markets are complete and some people are more risk averse than others, aggregate shocks generate small and bounded welfare costs *even for consumers whose risk aversion approaches infinity*. Furthermore, the least risk-averse people can be better off with business cycles than without them, because business cycles create the opportunity to sell insurance against aggregate risk.

To gain intuition for the results, consider an economy that contains some risk-averse agents as well as some risk-neutral agents whose consumption is allowed to be negative in equilibrium. The risk-neutral agents will fully insure the risk-averse agents at a fair price; all agents are indifferent between a world with business cycles and one without. The welfare cost of business cycles thus is zero for everyone, even if the risk-averse agents are extremely risk averse and regardless of the numbers of risk-averse and risk-neutral agents in the economy. Next, and more realistically, suppose the risk-neutral agents' consumption must be non-negative. Full insurance may now be infeasible if the total endowment of the risk-neutral agents is sufficiently small. The risk-neutral agents hence may have to charge a risk premium so that the risk-averse agents demand a feasible amount of insurance. The risk premium makes the risk-neutral agents better off with business cycles than without; the risk-averse agents experience a welfare loss, but it is smaller than they would experience if no insurance

were available. In this paper, I consider an economy where all agents are risk averse but some can be arbitrarily close to risk neutral. The intuition from risk-neutral agents whose consumption must be non-negative carries through to my model because, in the limit as risk aversion goes to zero, an agent with constant relative risk aversion becomes one who is risk neutral but faces a non-negativity constraint.

My results offer two cautions for the literature on the welfare cost of business cycles. First is a non-aggregation result. One cannot calculate the welfare cost in an economy with heterogeneous preferences by averaging the welfare costs of representative agents with various levels of risk aversion, because the cost of a representative agent is higher than the cost of an agent who has the same preferences in an economy with heterogeneity. Second is a caveat to one of the responses to Lucas' result. Some researchers have argued that Lucas assumed too small a coefficient of relative risk aversion and that business cycles matter more if one assumes people are very risk averse. One might think, therefore, that some people could suffer greatly from business cycles just because they are very risk averse compared with the average person. But this is not necessarily true: I show that when markets are complete, very risk-averse people need not experience large welfare costs, because these people may be well insured as long as they can trade with others who are less risk averse.

My results also clarify the role of heterogeneity in understanding the costs of aggregate risk. Heterogeneity has two effects on welfare costs because it comes in two flavors: heterogeneity in experiences and heterogeneity in initial conditions. Heterogeneity in experiences means that people start out life identical, but they experience different uninsured shocks and, as a result, their consumption paths diverge over time. Heterogeneity in experiences is intimately tied to market structure. When all shocks are fully insured, heterogeneity vanishes, while in general the amount of heterogeneity depends on which shocks are insured and to what extent. Heterogeneity in initial conditions, by contrast, means that people are not identical even at the beginning of life; some start with different preferences, endow-

ments or technology than others. In general, these initial differences translate into persistent differences in consumption patterns throughout life, regardless of market structure.

The separate roles of these two kinds of heterogeneity have not been entirely clear in the literature because influential papers on heterogeneous agent models, in particular the work of Krusell and Smith (1998, 1999), have considered heterogeneity in preferences and experiences at the same time. I use a model with heterogeneity only in preferences to highlight the differing effects of the two kinds of heterogeneity. Agents in the model own shares of the aggregate endowment and trade a complete set of contingent claims. In equilibrium, more risk-averse agents have smoother consumption and bear less aggregate risk, in return for lower average consumption. The result that more risk-averse agents bear less aggregate risk has been known since at least Wilson (1968), but its implications for business cycles appear not to have been studied previously. My model shows that because heterogeneity in initial conditions creates the opportunity to reallocate aggregate risk, it reduces welfare costs. The heterogeneity that can increase welfare costs is heterogeneity in experiences or, put another way, market incompleteness.

The paper proceeds as follows. Section 2 sets out the model and derives a formula for the welfare gain from eliminating aggregate fluctuations. Section 3 performs computational experiments to show how welfare gains depend on an individual's preferences and on the distribution of preferences in the economy. Section 4 discusses econometric methods for estimating the model using microdata on household consumption. Section 5 presents estimates of the welfare costs of business cycles, and section 6 concludes.

2 Model

Lucas (1987) calculated the expected utility of a representative agent who receives a random consumption stream and computed the amount of certain consumption that would give

the agent the same utility. The difference between this certainty-equivalent consumption and the mean of the risky consumption stream is the welfare cost of risk. If the random consumption stream is (filtered or detrended) aggregate consumption, the calculation describes the welfare cost of business cycles.

My goal is simply to show the role of preference heterogeneity in the calculation, so my model deviates from that of Lucas (1987) in only two ways. First, instead of a representative agent, there are many agents, and they do not all have the same risk preferences. Second, Lucas obtained analytic results by assuming a log-normal distribution of aggregate shocks, but distributional assumptions do not simplify the calculations in my model and so I allow the shocks to come from any distribution with a finite number of states.¹

2.1 Preferences, endowments, markets and equilibrium

When agents have time-separable expected utility preferences, risk aversion is the inverse of the elasticity of intertemporal substitution, and heterogeneous preferences will motivate agents to make intertemporal consumption trades even in the absence of aggregate risk. Indeed, with constant relative risk aversion preferences, the least risk-averse agent in a growing economy will ultimately consume all of the aggregate endowment. To abstract from these intertemporal issues and focus on risk aversion, while recognizing that I will need time series data to estimate the model, I consider a sequence of one-period economies indexed by dates t . Each economy can be in one of several possible states, $s = 1, \dots, S$, each with probability π_s . The possible states and their probabilities are the same for all economies. Before the state of the economy is known, agents in the economy trade contingent claims; then the state is realized, the claims pay off, and the agents consume. There is one good, denoted c .

¹Assuming a finite number of states allows me to avoid technical concerns about infinite-dimensional commodity spaces and the existence of expectations over states but is not otherwise crucial to the results.

Each economy contains a continuum of agents. An agent i is characterized by an endowment share w_i and a coefficient of relative risk tolerance $\theta_i \in (0, \infty)$ (the inverse of the coefficient of relative risk aversion). An agent with endowment share w_i receives a fraction w_i of the aggregate income in each state, and $\int w_i dF(w_i, \theta_i) = 1$. This is an endowment economy; there is no production or (since the economy lasts only one period) investment. Agents maximize expected utility and have constant relative risk aversion preferences:

$$u_i(c) = \frac{c^{1-1/\theta_i}}{1-1/\theta_i}.$$

The economy may be larger at some dates t than at other dates, but I assume that aggregate risk is the same in all economies: Aggregate income in economy t in state s is $g_t m_s$, where g_t is a non-random sequence. I normalize $\sum_{s=1}^S \pi_s m_s = 1$, so that the expected value of aggregate income in economy t is g_t . I also assume that the joint distribution of preferences and endowments $F(w_i, \theta_i)$ is the same at all dates t .²

Markets are complete. Let $\pi_s p_{st}$ be the price of a claim to one unit of consumption in economy t in state s . Agent i in economy t solves

$$\max_{\{c_{ist}\}_{s=1}^S} \sum_{s=1}^S \pi_s \frac{c_{ist}^{1-1/\theta_i}}{1-1/\theta_i} \quad \text{subject to} \quad \sum_{s=1}^S \pi_s p_{st} c_{ist} = \sum_{s=1}^S \pi_s p_{st} w_i g_t m_s. \quad (1)$$

An equilibrium in economy t is a set of prices $\{p_{st}\}$ and a consumption allocation $\{c_{ist}\}$ such that the consumption allocation solves each agent's problem (1) given the prices and such that markets clear:

$$\forall s \quad \int c_{ist} dF(w_i, \theta_i) = g_t m_s. \quad (2)$$

²In a model with agents living multiple periods, one could obtain a stationary joint distribution of preferences and endowments by having agents die each period with some probability and be replaced by a draw from some (not necessarily stationary) distribution, but one does not need so much machinery to obtain the results in this paper.

Agent i 's first-order condition in economy t is

$$\log c_{ist}^* = \log A_{it}^* + \log g_t - \theta_i \log p_{st}^*, \quad (3)$$

where $(g_t A_{it}^*)^{-1/\theta_i}$ is the Lagrange multiplier on i 's budget constraint in economy t . Substituting (3) into the aggregate resource constraint (2) gives

$$\forall s \quad \int A_{it}^* (p_{st}^*)^{-\theta_i} dF(w_i, \theta_i) = m_s. \quad (4)$$

Substituting (3) into the agent's budget constraint gives

$$\forall i \quad A_{it}^* \sum_{s=1}^S \pi_s (p_{st}^*)^{1-\theta_i} = w_i \sum_{s=1}^S \pi_s p_{st}^* m_s. \quad (5)$$

Equations (4) and (5) together determine the Lagrange multipliers A_{it}^* and the prices p_{st}^* in economy t . The size of the economy g_t does not enter these equations, and nothing else in the model depends on t . Therefore, neither the prices nor the Lagrange multipliers depend on t , and we can henceforth write A_i^* instead of A_{it}^* and p_s^* instead of p_{st}^* . It will also be helpful in interpreting the results to normalize prices such that the value of the endowment equals 1, i.e., $\sum_{s=1}^S \pi_s p_s^* m_s = 1$. With this normalization, and with the realization that prices and Lagrange multipliers do not depend on t , we can write the allocation as

$$\log c_{ist}^* = \log A_i^* + \log g_t - \theta_i \log p_s^*, \quad (6)$$

and – rearranging (5) – the Lagrange multipliers as

$$\forall i \quad A_i^* = w_i H(p_1^*, \dots, p_S^*, \theta_i) \quad (7)$$

where $H(p_1^*, \dots, p_S^*, \theta) = \left(\sum_{s=1}^S \pi_s (p_s^*)^{1-\theta} \right)^{-1}$. The market-clearing condition (4) then becomes

$$\forall s \quad \int H(p_1^*, \dots, p_S^*, \theta) (p_s^*)^{-\theta} dW(\theta) = m_s \quad (8)$$

where $dW(\theta) = \int w_i dF(w_i, \theta_i = \theta)$ is the total endowment of agents with risk tolerance θ .

Equation (8) does not have an analytic solution for general joint distributions of risk tolerance and endowment shares, but we can derive some simple facts about the prices. First, (8) implies that p_s^* is strictly decreasing in m_s . Second, the normalization on prices requires either that $p_s^* = 1$ for all s or $p_s^* < 1$ for some s and $p_{s'}^* > 1$ for some $s' \neq s$, since otherwise the normalizations $\sum_{s=1}^S \pi_s m_s = 1$ and $\sum_{s=1}^S \pi_s p_s^* m_s = 1$ cannot both hold. Third, because p_s^* is strictly decreasing in m_s , we must then have $p_s^* < 1$ for some s and $p_{s'}^* > 1$ for some $s' \neq s$ in any economy where there is aggregate risk. Fourth, (8) shows that, with respect to prices, there is aggregation within groups of agents classified by risk tolerance; all that matters is the total endowment share $dW(\theta)$ of agents with each possible value of risk tolerance, not the division of endowment shares among agents with the same risk tolerance.

To interpret the Lagrange multipliers A_i^* , first consider the case where all agents have identical preferences $\theta_i = \bar{\theta}$. Since $H(p_1^*, \dots, p_S^*, \theta_i)$ is then the same constant for all i , (7) then says that A_i^* is proportional to w_i : When everyone has the same preferences, the Lagrange multipliers are (up to a normalization) the endowment shares. Now consider the case where preferences vary. Equation (7) shows that the Lagrange multipliers A_i^* are increasing in the endowment shares w_i , but the Lagrange multipliers are adjusted away from the endowment shares by a factor $H(p_1^*, \dots, p_S^*, \theta_i)$ that depends on risk tolerance; the adjustment implies that agent i 's average consumption is adjusted away from i 's endowment share, since according to (6) each agent's consumption is increasing in A_i^* in every state.

We can see the precise way in which consumption is adjusted away from endowment

shares by considering the ratio of i 's expected consumption to i 's expected endowment:

$$\frac{\mathbb{E}[c_{ist}^*]}{\mathbb{E}[w_i g_t m_s]} = \frac{\sum_{s=1}^S \pi_s c_{ist}^*}{w_i g_t} = \frac{A_i^*}{w_i} \sum_{s=1}^S \pi_s (p_s^*)^{-\theta_i} = \frac{\sum_{s=1}^S \pi_s (p_s^*)^{-\theta_i}}{\sum_{s=1}^S \pi_s (p_s^*)^{1-\theta_i}}. \quad (9)$$

If this ratio is less than 1, the expected value of i 's consumption is less than the expected value of i 's endowment, which means i pays a risk premium to receive insurance. If the ratio is greater than 1, the expected value of i 's consumption exceeds the expected value of i 's endowment: i receives a risk premium for providing insurance. The ratio is strictly increasing in θ_i as long as the prices p_s^* are not all the same.³ Feasibility requires either that the ratio be exactly 1 for all agents or that it be greater than 1 for some agents and less than 1 for others. Thus, since the ratio is strictly increasing in θ_i , agents with low θ_i have a ratio less than 1 and pay a risk premium, while agents with sufficiently high θ_i have a ratio greater than 1 and receive a risk premium. The more risk tolerant an agent is, the larger the risk premium the agent receives. The only cases where no one pays or receives a risk premium are those where there is no aggregate risk or all the agents have the same preferences.

2.2 Removing aggregate risk

The Lucas experiment removes aggregate risk while leaving the trend of aggregate consumption unchanged. The calculation thus abstracts from dynamic effects on production

³To prove that the ratio is strictly increasing in θ_i , observe that its derivative with respect to θ_i is

$$\frac{d}{d\theta_i} \frac{\sum_s \pi_s (p_s^*)^{-\theta_i}}{\sum_s \pi_s (p_s^*)^{1-\theta_i}} = \frac{(\sum_s \pi_s (p_s^*)^{1-\theta_i} \log p_s^*) (\sum_s \pi_s (p_s^*)^{-\theta_i}) - (\sum_s \pi_s (p_s^*)^{-\theta_i} \log p_s^*) (\sum_s \pi_s (p_s^*)^{1-\theta_i})}{(\sum_s \pi_s (p_s^*)^{1-\theta_i})^2}.$$

Define $Q_i = \sum_s \pi_s (p_s^*)^{-\theta_i}$ and $q_{is} = \pi_s (p_s^*)^{-\theta_i} / Q_i$. Then the numerator of the derivative equals

$$Q_i^2 \left[\sum_s q_{is} p_s^* \log p_s^* - \left(\sum_s q_{is} \log p_s^* \right) \left(\sum_s q_{is} p_s^* \right) \right] = Q_i^2 \left[\tilde{\mathbb{E}}[p_s^* \log p_s^*] - \tilde{\mathbb{E}}[p_s^*] \tilde{\mathbb{E}}[\log p_s^*] \right] = Q_i^2 \widetilde{\text{Cov}}[p_s^*, \log p_s^*],$$

where $\tilde{\mathbb{E}}[\cdot]$ and $\widetilde{\text{Cov}}[\cdot]$ denote an expectation and a covariance under the probability measure q_{is} . The logarithm is strictly increasing, so the covariance is positive. Further, $Q_i^2 > 0$, and the denominator is positive. Thus the derivative is positive.

and investment from removing aggregate risk. Dynamic effects may be especially significant in an economy with heterogeneous preferences. For example, in related work (Schulhofer-Wohl, 2007) I show that workers in a risky economy choose jobs in part on the basis of risk preferences rather than comparative advantage in productivity, so that removing risk could change the assignment of workers to jobs and make the entire economy more productive. This could change the growth rate or the level of consumption; either way, there would be a welfare gain. However, my goal here is to study only the allocational consequences of heterogeneous risk preferences, so I follow Lucas in assuming that removing risk leaves the level and trend of consumption unchanged.

Recall that the expected value of consumption in economy t is g_t . Suppose agent i in economy t gives up a fraction k of the expected value but also eliminates all aggregate risk. Then the agent's endowment will be $w_i(1 - k)g_t$ in every state of the world. There is no scope for trade in a risk-free one-period economy with one good, so the agent will consume his endowment, and his utility will be

$$\hat{U}_{it}(k) = \frac{(w_i g_t)^{1-1/\theta_i}}{1 - 1/\theta_i} (1 - k)^{1-1/\theta_i}. \quad (10)$$

Meanwhile, in the original risky economy, agent i 's expected utility given the allocation (6) and (7) is

$$U_{it}^* = \sum_{s=1}^S \pi_s u_i(c_{is}^*) = \frac{(w_i g_t)^{1-1/\theta_i}}{1 - 1/\theta_i} \left(\sum_{s=1}^S \pi_s (p_s^*)^{1-\theta_i} \right)^{1/\theta_i}. \quad (11)$$

To measure i 's willingness to pay to remove aggregate risk in economy t , we calculate the fraction k_{it} that makes $\hat{U}_{it}(k_{it})$ equal to the expected utility in the risky economy U_{it}^* . Setting (10) equal to (11) and solving for k_{it} yields

$$k_{it} = 1 - \left(\sum_{s=1}^S \pi_s (p_s^*)^{1-\theta_i} \right)^{-1/(1-\theta_i)}.$$

The willingness to pay k_{it} depends only on risk tolerance θ_i , not on the size of the economy g_t , the date t or the endowment share w_i , so from now on I investigate the willingness to pay as a function of risk tolerance, which I denote $k(\theta)$:

$$k(\theta) = 1 - \left(\sum_{s=1}^S \pi_s (p_s^*)^{1-\theta} \right)^{-1/(1-\theta)}. \quad (12)$$

The willingness to pay is decreasing in risk tolerance θ . The first reason for this result is obvious: More risk-tolerant agents suffer less disutility from a given amount of variance in consumption. The second reason is less obvious but is central to the results of this paper: As shown in (9), the presence of aggregate risk creates the opportunity for more risk-tolerant agents to sell insurance to less risk-tolerant agents. Indeed, there exists a finite value of risk tolerance – call it θ^* – such that all agents with risk tolerance $\theta > \theta^*$ have $k(\theta) < 0$ and experience a welfare *gain* from aggregate risk, even though they are risk averse.⁴ Aggregate risk allows these less risk-averse agents to sell so much insurance that the risk premium they receive more than offsets the disutility from the consumption fluctuations they experience.

We can also consider agents at the opposite extreme of preferences, those who are extremely risk averse and have risk tolerance near zero. Consider reducing one agent’s risk tolerance while holding the prices fixed:⁵

$$\lim_{\theta \rightarrow 0} k(\theta) = 1 - \left(\sum_{s=1}^S \pi_s p_s^* \right)^{-1}. \quad (13)$$

This limit provides an upper bound on welfare costs. Even if we do not know the distribution of risk aversion, the limit tells us the most that anyone in the economy would be willing to

⁴To prove that $k(\theta) < 0$ for θ sufficiently large, note that when there is aggregate risk, and under the chosen normalization on prices, $\min_s p_s^* < 1$. Therefore, $\lim_{\theta \rightarrow \infty} (\sum_{s=1}^S \pi_s (p_s^*)^{1-\theta})^{-1/(1-\theta)} = [\min_s \{p_s^*\}]^{-1} > 1$ and thus $\lim_{\theta \rightarrow \infty} k(\theta) < 0$. Because $k(\theta)$ is continuous in θ and because welfare costs cannot depend on the normalization of prices, this completes the proof.

⁵With a continuum of agents, any one agent’s preferences do not affect the prices.

pay to eliminate business cycles.

In addition, we can compare the welfare costs of aggregate risk in an economy with heterogeneous preferences to the costs we would calculate in an economy where all agents have the same preferences. Let $k^{rep}(\theta)$ be the willingness to pay to eliminate aggregate risk of a representative agent with risk tolerance θ . Since the expected utility of such an agent in the presence of aggregate risk is $U_t^{rep}(\theta) = \sum_{s=1}^S \pi_s (w_i g_t m_s)^{1-1/\theta} / (1 - 1/\theta)$, we have

$$k^{rep}(\theta) = 1 - \left(\sum_{s=1}^S \pi_s m_s^{1-1/\theta} \right)^{1/(1-1/\theta)}. \quad (14)$$

We can show that, for essentially all agents, the willingness to pay of a representative agent with risk tolerance θ in an economy containing only agents with the same risk tolerance is strictly larger than the agent's willingness to pay in an economy where risk preferences vary.

Proposition. *Suppose that there is aggregate risk and that not all agents have the same risk tolerance. Then $k^{rep}(\theta) \geq k(\theta)$ for all θ in the support of $F(w_i, \theta_i)$, and there is at most one value of θ in the support for which the inequality is not strict.*

Proof. See appendix A. □

The proposition follows from the fact that the competitive equilibrium is in the core. Since participants in the heterogeneous-agent economy can do as well as a representative agent if they remain in autarky, the statement is equivalent to the claim that in the economy with heterogeneity, all agents but one are strictly better off in the competitive equilibrium than in autarky. If two agents with different risk tolerances were no better off in the equilibrium than in autarky, they could form a blocking coalition, trade and be better off, contradicting the fact that the equilibrium is in the core.⁶

⁶To see why one agent can be indifferent between autarky and equilibrium, consider an economy with two states. For any i , individual maximization requires the first-order condition $p_2^*/p_1^* = u'_i(c_{i2})/u'_i(c_{i1})$. Since the endowment satisfies the budget constraint, i will consume the endowment if it also satisfies the

Consider computing an “average willingness to pay” of all agents in an economy by some weighted average of the individual agents’ willingness to pay. Many weights are possible. For example, agents could all have equal weight, or the agents’ weights could be proportional to their endowment shares. However, the proposition implies that, regardless of the weights used, any computation based on the representative agent willingness to pay will overstate the true welfare cost of business cycles. Specifically, let $dV(\theta)$ be any weights that are non-negative for all θ in the support of $F(w_i, \theta_i)$ and strictly positive for at least two values of θ in the support. Then $\int k(\theta) dV(\theta) < \int k^{rep}(\theta) dV(\theta)$. To put it another way, in an economy with heterogeneous risk preferences, the average willingness to pay – however the average is defined – is less than the average of representative agents’ willingness to pay.

Finally, it is interesting to compare the welfare cost of business cycles for an infinitely risk-averse representative agent to that for an infinitely risk-averse agent in the economy with heterogeneity. Equation (13) gives the upper bound on welfare costs in the heterogeneous-agent economy. The bound depends on the probability of each state and resources in that state. By contrast, the welfare cost of an infinitely risk-averse representative agent is

$$\lim_{\theta \rightarrow 0} k^{rep}(\theta) = 1 - \min_s m_s,$$

which depends only on resources in the worst state and not on the probability of this state or what happens in any state other than the worst. Disasters that occur with arbitrarily low first-order condition, i.e.,

$$p_2^*/p_1^* = (m_2/m_1)^{-1/\theta_i} \quad \Rightarrow \quad \theta_i = -[\log(p_2^*/p_1^*)]/[\log(m_2/m_1)]. \quad (*)$$

Note that $-[\log(p_2^*/p_1^*)]/[\log(m_2/m_1)] > 0$ since $p_2^* > p_1^*$ if and only if $m_1 < m_2$. If the distribution of risk tolerance is continuous and unbounded, then, regardless of the equilibrium prices, (*) holds for some i . This i consumes his endowment in equilibrium and is indifferent between equilibrium and autarky. An unbounded or discrete distribution may still contain the θ_i given by (*). With $S > 2$ states, (*) becomes a system of $S - 1 > 1$ nonlinear equations in one unknown, θ_i . Such a system generically has no solution, but I have not found a proof that no distribution of states, endowments and preferences can generate equilibrium prices for which the system does have a solution.

probability thus completely determine the infinitely risk-averse representative agent's cost, but not so when we introduce heterogeneity.

3 Numerical examples

In this section, I numerically compute the welfare costs of aggregate risk in some simple but quantitatively reasonable example economies. The computations may help to provide intuition for the theoretical results. The computations also show that plausible parameter values can generate quantitatively important departures from representative agent welfare costs, including cases where some agents benefit from aggregate risk.

To simplify the calculations, all of the economies I consider have two equally probable states of nature, with aggregate resources $m_1 = 0.98$ and $m_2 = 1.02$. These values correspond to a 2% standard deviation of aggregate shocks, which is similar to the behavior of detrended postwar U.S. GDP. All of the economies have two agents. One agent always has log utility. I experiment with varying the other agent's risk preferences and the endowment share of each agent. The economies can equivalently be interpreted as containing two types of agent, with the endowment shares representing the total endowment of each type.

3.1 An unconstrained risk-neutral agent

Suppose that one agent is risk neutral and that there are no non-negativity constraints on consumption. Let w_1 be the endowment share of the risk-averse agent, who has log utility. In the competitive equilibrium, the risk-averse agent must be fully insured at a fair price and consumes w_1 in each state. Meanwhile, the risk-neutral agent consumes $m_s - w_1$ in each state. Both agents are indifferent between this economy and one without business cycles, regardless of the endowment shares. Further, the risk-averse agent is strictly better off in this economy than he would be in a representative-agent economy where all agents had log

utility, while the risk-neutral agent is indifferent between his outcome in this economy and his outcome as a representative agent.

3.2 A risk-neutral agent with non-negativity constraints

The case with unconstrained risk-neutral agents is not directly comparable to the case where all agents are strictly risk averse because, when all agents are strictly risk averse with CRRA preferences, an Inada condition will force all agents' consumption to be strictly positive. For comparability, I examine the case where one agent is risk neutral but each agent's consumption is constrained to be non-negative.⁷

In this case, the risk-averse agent remains fully insured at a fair price if his endowment share w_1 does not exceed 0.98. However, full insurance at a fair price (i.e., $c_{1s} = w_1$ for all s) is infeasible if $w_1 > 0.98$. I have solved the agents' first-order conditions and budget constraints to find the allocation as a function of w_1 . It is

$$[c_{1,1}^*, c_{2,1}^*, c_{1,2}^*, c_{2,2}^*] = \begin{cases} [w_1, 0.98 - w_1, w_1, 1.02 - w_1] & w_1 \leq 0.98, \\ [0.98, 0, 1.02w_1/(2 - w_1), 1.02(2 - 2w_1)/(2 - w_1)] & w_1 > 0.98. \end{cases}$$

⁷To justify the comparison formally, we can show that in the limit as risk tolerance goes to infinity, CRRA preferences become the preferences of a risk-neutral agent facing a non-negativity constraint. Observe that $u(c; \theta) = c^{1-1/\theta}/(1-1/\theta)$ is not defined for $c < 0$. However, since the Inada condition guarantees the agent will choose $c > 0$, defining $u(c; \theta) = -\infty$ for $c < 0$ will not change the agent's choices. That is, CRRA preferences can be represented by

$$u(c; \theta) = \begin{cases} \frac{c^{1-1/\theta}}{1-1/\theta} & c \geq 0 \\ -\infty & c < 0. \end{cases}$$

Once we define preferences this way, we have

$$\lim_{\theta \rightarrow \infty} u(c; \theta) = \begin{cases} c & c \geq 0 \\ -\infty & c < 0, \end{cases}$$

which is the utility function of a risk-neutral agent whose consumption is constrained to be non-negative.

The risk-neutral agent's expected utility is

$$U_2^* = 0.5c_{2,1}^* + 0.5c_{2,2}^* = \begin{cases} 1 - w_1 & w_1 \leq 0.98, \\ (1 - w_1)^{\frac{1.02}{2-w_1}} & w_1 > 0.98. \end{cases}$$

In a risk-free economy, the risk-neutral agent's expected utility would be $1 - w_1$. Since $1.02/(2 - w_1) > 1$ for all $w_1 > 0.98$, the risk-neutral agent is better off in the risky economy than the risk-free economy whenever $w_1 > 0.98$. The intuition is that when full insurance is infeasible, the risk-neutral agent collects a positive risk premium to reduce the risk-averse agent's demand for insurance. The risk premium makes the risk-neutral agent better off than if there were no risk and no scope for insurance.

Figure 1 plots each agent's willingness to pay to remove aggregate risk as a function of the risk-averse agent's endowment share w_1 . When $w_1 \leq 0.98$, there is full insurance and both agents are indifferent between the risky and risk-free economies. For $w_1 > 0.98$, the risk-averse agent is willing to pay to remove risk, while the risk-neutral agent is better off with risk. As the risk-averse agent's endowment share approaches 1, his willingness to pay approaches 0.02%, which is what he would be willing to pay if he were a representative agent.

Figure 2 illustrates the sense in which the limit as risk tolerance goes to infinity is a risk-neutral agent subject to non-negativity constraints. The figure plots the welfare costs of aggregate risk for each agent when $\theta_2 = 50$, when $\theta_2 = 200$, when $\theta_2 = 500$ and when agent 2 is risk neutral but faces a non-negativity constraint. The curves for $\theta_2 = 50$, $\theta_2 = 200$ and $\theta_2 = 500$ converge to the curve for the risk-neutral constrained case.

3.3 Two risk-averse agents, varying endowment shares

Suppose that agent 1 has risk tolerance $\theta_1 = 0.25$, corresponding to a coefficient of relative risk aversion of 4, while agent 2 has log utility ($\theta_2 = 1$). Since agent 1 is more risk averse,

agent 2 will insure agent 1 against aggregate risk. This creates the possibility that, for some parameters, agent 2 collects a sufficiently large risk premium to benefit from aggregate risk.

I have numerically solved the agents' first-order conditions and budget constraints to find the competitive equilibrium prices and allocation for varying levels of agent 1's endowment share. Figure 3 plots each agent's corresponding willingness to pay to remove aggregate risk. As agent 1's endowment share rises, his willingness to pay also rises, because agent 2 has less resources and can provide less insurance. Increases in agent 1's endowment share reduce agent 2's willingness to pay, because when agent 2 has a small endowment share, he can collect a large risk premium from agent 1. Indeed, when agent 1's endowment share is larger than about 0.66, agent 2 has a negative willingness to pay. That is, in an economy where less than one-third of the people have log utility and more than two-thirds have a coefficient of relative risk aversion of 4, and where aggregate shocks have a standard deviation of 2%, the agents with log utility benefit from business cycles.

3.4 Two risk-averse agents, varying preferences

Finally, I consider holding the agents' endowment shares fixed and varying one agent's preferences. Agent 2 has log utility and an endowment share of 0.333. Agent 1 has an endowment share of 0.667; I vary agent 1's risk tolerance and then compute the competitive equilibrium allocation and each agent's willingness to pay.

Figure 4 plots the results. Agent 1 is always willing to pay to remove aggregate risk, and his willingness to pay rises with his risk aversion. However, it rises less rapidly than the willingness to pay that we would compute by treating him as a representative agent.

Agent 2's willingness to pay has an inverted U shape. When agent 1 has log utility, agent 2's willingness to pay in the competitive equilibrium exactly matches the representative agent calculation because both agents have the same preferences and there is, in fact, a representative agent. Otherwise, agent 2's willingness to pay is less than the representative

agent calculation would suggest. When agent 1 has very low risk aversion, agent 2's willingness to pay is low because he can obtain insurance from agent 1. As agent 1's risk aversion rises, agent 2's willingness to pay also rises because the agents are becoming more alike and the opportunities for insurance are reduced. Once agent 1 becomes more risk averse than agent 2, further increases in agent 1's risk aversion decrease agent 2's willingness to pay, because agent 2 can now sell insurance to agent 1 and collect a risk premium. If agent 1 has a coefficient of relative risk aversion greater than about 4, agent 2 benefits from cycles.

4 Empirical analysis: data and econometric methods

Equation (12) expresses the willingness to pay to eliminate aggregate risk as a function of risk tolerance θ , prices p_s^* , aggregate shocks m_s , and probabilities π_s . (Aggregate shocks do not appear directly in the equation but are implicit in the normalization $\sum_{s=1}^S \pi_s p_s^* m_s = 1$ that the equation assumes.) Thus, for a real economy, we can calculate the willingness to pay for any given value of θ if we know the aggregate shocks, the prices and the probabilities. (My goal will be to estimate the willingness to pay as a function of θ , not to estimate any individual agent's risk tolerance, the distribution of risk tolerance in the population or the average willingness to pay of all agents. Nonetheless, given estimates of the function $k(\theta)$, a reader who has in mind a distribution of θ can integrate $k(\theta)$ against that distribution to find the average willingness to pay.)

The probabilities are straightforward to handle: Since the model is stationary, averages over possible states in one economy are the same as averages over time in a sequence of economies where different states are realized. Suppose we collect data on a sequence of dates indexed by $\tau = 1, \dots, T$. Let m_τ be the realized value of the aggregate shock at τ , and let p_τ^* be the price corresponding to the state realized at τ . Then, if we observed the aggregate shocks and prices, we could normalize the prices to satisfy $T^{-1} \sum_{\tau=1}^T p_\tau^* m_\tau = 1$ and then

replace the sums over states in (12) by sums over time and estimate the willingness to pay of an agent with risk tolerance θ by

$$\tilde{k}(\theta; T) = 1 - \left(\frac{1}{T} \sum_{\tau=1}^T (p_{\tau}^*)^{1-\theta} \right)^{-1/(1-\theta)} .$$

In the limit as T goes to infinity, $\tilde{k}(\theta; T)$ would converge in probability to $k(\theta)$ by a law of large numbers.

In principle, we could measure the aggregate shocks m_{τ} from the National Income and Product Accounts. The prices are more difficult. Although it might seem natural to obtain prices from financial market data, we cannot do so because the prices in question are those of a non-traded asset: a one-period-ahead claim to a share of the aggregate endowment in various states.⁸ Thus we must turn elsewhere. My strategy is to exploit the relationship between prices and individual consumption to recover the prices from cross-sectional consumption data. For consistency, I then estimate aggregate shocks from the same cross-sectional data. Estimating the aggregate shocks from NIPA would run the risk of inconsistency with the estimated prices from cross-sectional data, especially since the aggregate data (which are derived from surveys of firms) and microdata (from surveys of consumers) have been diverging over time for reasons that remain unclear (Garner et al., 2006). Below, when I report my results, I check whether my estimates of aggregate consumption from microdata are in line with NIPA data. The rest of this section concerns how to estimate prices and aggregate shocks from cross-sectional data and how to use the estimated prices and shocks to compute the willingness to pay as a function of θ .

⁸Methods are known for bounding the expected return on non-traded assets using no-arbitrage conditions (e.g., Alvarez and Jermann, 2004), but the expected return is not sufficient for calculating $k(\theta)$: We need to know the actual price for a claim on each separate state so that we can calculate the expectation of a nonlinear function of the prices. Alvarez and Jermann (2004) also estimate prices date-by-date under the assumption of common preferences, but this assumption would be inappropriate in the present context.

4.1 Data

I analyze data from the U.S. Consumer Expenditure Survey. Implicit in the results, therefore, will be the assumption that a complete markets allocation reasonably approximates the U.S. economy. Although this assumption is undoubtedly controversial, a growing body of evidence (e.g., Schulhofer-Wohl, 2007; ?) suggests that insurance in the U.S. economy is quite good. U.S. households also make substantial transactions that insure more risk-averse people against aggregate risk. More risk-averse investors put more money in bonds and less in stocks (Barsky et al., 1997), and we can interpret the risk premium on stocks as an insurance premium that bondholders pay to avoid aggregate risk.

The Consumer Expenditure Survey is a rotating panel with four quarterly observations per household. Different households begin the survey in different months of a quarter, so although consumption is measured at a quarterly frequency, I can construct data on aggregate fluctuations at a monthly frequency. I use data on consumption of nondurable goods and services from 1982 to 2002. As these data are well known, I do not dwell on them here. Table 1 gives summary statistics, and appendix B describes the sample selection in detail.

4.2 Notation and assumptions

I assume that we observe a random sample of agents at each date τ . Let $\bar{E}_\tau[\zeta]$ denote the population mean of a random variable ζ in the cross-section of agents at date τ , and let $\hat{E}_\tau[\zeta]$ denote the sample mean in the cross-section of observed agents.

I assume that agents' consumption is measured with error: We observe

$$\tilde{c}_{i\tau} = c_{i\tau}^* e_{i\tau}, \tag{15}$$

where $e_{i\tau}$ is a strictly positive random variable that is independent of $c_{i\tau}^*$, i.i.d. over individ-

uals i at each date τ , independent over dates τ , and satisfies

$$\bar{E}_\tau[e_{i\tau}] = \exp(\gamma_0 + \gamma_1\tau), \quad (16a)$$

$$\bar{E}[\log e_{i\tau}] = \delta_0 + \delta_1\tau, \quad (16b)$$

for some unknown constants $\gamma_0, \gamma_1, \delta_0, \delta_1$. This formulation allows the mean and variance of measurement error to change over time, in case data quality changes over time. Measurement error can also have a stationary distribution if $\gamma_1 = \delta_1 = 0$.

Finally, I assume that the economy grows exponentially and, since I will analyze monthly data, that it experiences predictable monthly fluctuations:

$$g_\tau = \exp(\alpha_0 + \alpha_1\tau + \boldsymbol{\alpha}'\mathbf{x}_\tau), \quad (17)$$

where \mathbf{x}_τ is a vector of indicator variables for months. My results are potentially sensitive to this assumption because I will measure the variability of aggregate shocks by examining deviations of aggregate consumption from the exponential trend and monthly seasonals. If the overall growth path is not exponential, or if seasonal effects are not the same in all years, then the assumption in (17) will fail and I will mismeasure the economy's aggregate risk.

4.3 Aggregate shocks

According to (15) and (16a), for any τ we have

$$N\bar{E}_\tau[\tilde{c}_{i\tau}] = \exp(\gamma_0 + \gamma_1\tau)N\bar{E}_\tau[c_{i\tau}^*] = \exp(\gamma_0 + \gamma_1\tau)g_\tau m_\tau, \quad (18)$$

where I have let N denote the measure of agents in the economy and used the fact that N times average consumption equals total consumption, which is $g_\tau m_\tau$. Substituting (17) into

(18) and rearranging gives

$$\log(\bar{E}_\tau[\tilde{c}_{i\tau}]) = \alpha_0 + \gamma_0 - \log N + (\alpha_1 + \gamma_1)\tau + \boldsymbol{\alpha}'\mathbf{x}_\tau + \log m_\tau. \quad (19)$$

Equation (19) shows that the aggregate shock $\log m_\tau$ is the error term in a time-series regression of the logarithm of mean consumption $[\log(\bar{E}_\tau[\tilde{c}_{i\tau}])]$ on an intercept, a time trend and month dummies. The population mean of consumption is not known, but we can estimate it by a sample average. I therefore let $\widehat{\log m_\tau}$ be the residual from a regression of the log of the sample average of consumption $[\log(\hat{E}_\tau[\tilde{c}_{i\tau}])]$ on an intercept, trend and month dummies. I then estimate the aggregate shock m_τ by $\hat{m}_\tau = \exp(\widehat{\log m_\tau})$. The law of large numbers implies that \hat{m}_τ converges in probability to m_τ for each τ in the limit as the number of individuals sampled at each date and the number of dates T both go to infinity. To see this, notice that 1) the regression residuals converge to the regression errors as T goes to infinity, and 2) the sample average of consumption converges to the population average of consumption as the number of individuals sampled goes to infinity.

It is worth noting that, in principle, we need not estimate (19) with cross-sectional data. We could instead obtain $\bar{E}_\tau[\tilde{c}_{i\tau}]$ from per capita consumption in aggregate data, and then run the same time-series regression to obtain $\widehat{\log m_\tau}$. The only reason to estimate (19) with microdata is for consistency with the estimates of prices, which I discuss next.

4.4 Prices

According to (6), (15), (16b) and (17), for any τ we have

$$\bar{E}_\tau[\log \tilde{c}_{i\tau}] = \bar{E}_\tau[\log A_i^*] + \alpha_0 + \delta_0 + (\alpha_1 + \delta_1)\tau + \boldsymbol{\alpha}'\mathbf{x}_\tau - \bar{\theta} \log p_\tau^*, \quad (20)$$

where $\bar{\theta} = \bar{E}_\tau[\theta_i]$ denotes the population mean risk tolerance, which the model assumes is constant over time. Equation (20) shows that the price $\log p_\tau^*$ is the error term in a regression of the population mean of log consumption on an intercept, a time trend and month dummies. Again, the population mean is not known, but we can estimate it by the sample average of log consumption. For a given value of $\bar{\theta}$, I let $\widehat{\log p_\tau^*(\bar{\theta})}$ be $(-1/\bar{\theta})$ times the residual from a regression of the sample average of log consumption, $\hat{E}_\tau[\log \tilde{c}_{i\tau}]$, on an intercept, trend and month dummies. I then estimate the price p_τ^* for this value of $\bar{\theta}$ by $\hat{p}_\tau^*(\bar{\theta}) = \exp[\widehat{\log p_\tau^*(\bar{\theta})}]$. The estimated prices will not automatically satisfy the normalization $\sum_{s=1}^S \pi_s p_s^* m_s = 1$; I impose it by scaling prices such that $T^{-1} \sum_{\tau=1}^T \hat{p}_\tau^*(\bar{\theta}) \hat{m}_\tau = 1$. As with the aggregate shocks, \hat{p}_τ^* converges in probability to p_τ^* as the number of individuals and dates both go to infinity.

The foregoing analysis identifies prices up to a factor of $\bar{\theta}$, the population average risk tolerance. We cannot determine the average risk tolerance from consumption data alone, since multiplying all agents' risk tolerance by the same constant would not change the set of Pareto-optimal consumption allocations and, unless we know the initial endowments, we cannot learn anything from observing which competitive equilibrium arises from among all the Pareto optima.⁹ Therefore, in section 5 I report estimates of welfare costs for several

⁹Formally, for any $x > 0$, if we replace θ_i by $x\theta_i$ for all i ; p_s^* by $\tilde{p}_s^* = (p_s^*)^{1/x} (\sum_{s=1}^S \pi_s (p_s^*)^{1/x} m_s)^{-1}$ for all s ; and w_i by

$$\tilde{w}_i \equiv w_i \frac{\sum_{s=1}^S \pi_s (p_s^*)^{1/x - \theta_i}}{\left(\sum_{s=1}^S \pi_s (p_s^*)^{1/x} m_s\right) \left(\sum_{s=1}^S \pi_s (p_s^*)^{1 - \theta_i}\right)}$$

for all i , the equation governing the consumption allocation (6) still holds and (since the consumption allocation has not changed) the aggregate resource constraint is still satisfied. Further,

$$\begin{aligned} \int \tilde{w}_i dF(w_i, \theta_i) &= \int w_i \frac{\sum_{s=1}^S \pi_s (p_s^*)^{1/x - \theta_i}}{\left(\sum_{s=1}^S \pi_s (p_s^*)^{1/x} m_s\right) \left(\sum_{s=1}^S \pi_s (p_s^*)^{1 - \theta_i}\right)} dF(w_i, \theta_i) \\ &= \left(\sum_{s=1}^S \pi_s (p_s^*)^{1/x} m_s\right)^{-1} \sum_{s=1}^S \pi_s (p_s^*)^{1/x} \left(\int w_i \frac{(p_s^*)^{-\theta_i}}{\left(\sum_{s=1}^S \pi_s (p_s^*)^{1 - \theta_i}\right)} dF(w_i, \theta_i)\right) \\ &= \left(\sum_{s=1}^S \pi_s (p_s^*)^{1/x} m_s\right)^{-1} \sum_{s=1}^S \pi_s (p_s^*)^{1/x} m_s = 1, \end{aligned}$$

where the next-to-last equality follows from (8). Since $\tilde{w}_i > 0$ and $\int \tilde{w}_i dF(w_i, \theta_i) = 1$, we can interpret \tilde{w}_i

possible values of $\bar{\theta}$.

4.5 Accounting for random sampling

Given estimates of aggregate shocks and prices, we can estimate the welfare cost of business cycles for an agent with risk tolerance θ , as a function of θ and of the population mean risk tolerance $\bar{\theta}$, by

$$\hat{k}(\theta, \bar{\theta}; T) = 1 - \left(\frac{1}{T} \sum_{\tau=1}^T [\hat{p}_{\tau}^*(\bar{\theta})]^{1-\theta} \right)^{-1/(1-\theta)}, \quad (21)$$

where the prices are normalized such that $T^{-1} \sum_{\tau=1}^T \hat{p}_{\tau}^*(\bar{\theta}) \hat{m}_{\tau} = 1$. Since \hat{p}_{τ}^* and \hat{m}_{τ} converge in probability to p_{τ}^* and m_{τ} , and since \hat{k} is a continuous function of \hat{p}_{τ}^* and \hat{m}_{τ} , we have that $\hat{k}(\theta, \bar{\theta}; T)$ converges in probability to $k(\theta, \bar{\theta})$ in the limit as the number of individuals and dates sampled both go to infinity.

Despite this consistency result, the finite-sample behavior of \hat{k} is important. The estimated aggregate shocks \hat{m}_{τ} and estimated prices \hat{p}_{τ}^* will vary across dates τ in the data both because the economy experienced shocks and because different random samples of agents are observed at different dates. In other words, at any date τ , the estimated price \hat{p}_{τ}^* will be imprecise: $\hat{p}_{\tau}^* = p_{\tau}^* + \epsilon_{\tau}$, where ϵ_{τ} is an estimation error, and similarly for \hat{m}_{τ} . The estimated shocks and estimated prices will therefore be more variable than the true shocks and the true prices. This added variance will make the economy appear more risky than it truly is, and a riskier economy will appear to have larger welfare costs of risk. In consequence, the estimator \hat{k} will be biased away from zero relative to the true welfare costs.

The structure of the survey creates a second possible source of bias. The use of over-

as an endowment share. Thus if the allocation $\{c_{ist}^*\}$ is a competitive equilibrium in an economy with prices p_s^* , preferences θ_i and endowment shares w_i , the same allocation is also an equilibrium in an economy with prices \tilde{p}_s^* , preferences $x\theta_i$ and endowment shares \tilde{w}_i . We have now described two economies with identical consumption allocations but different mean risk tolerance, $\bar{\theta}$ vs. $x\bar{\theta}$; we cannot distinguish between them on the basis of consumption data alone.

lapping quarters of data to create a monthly time series implies that the observed monthly shocks will be a moving average of the true monthly shocks and, therefore, less variable than the true shocks. This reduced variance could bias my estimated welfare costs toward zero.

The total bias in my estimates from these two sources of bias could be positive or negative. I correct for the bias using the bootstrap. Let \hat{k} be the estimated willingness to pay calculated from (21) for some θ and $\bar{\theta}$ using the original data. Let k_1, \dots, k_Q be estimates of the willingness to pay calculated using Q different samples, of the same size as the original sample, drawn from the original data with replacement.¹⁰ Horowitz (2001) considers bootstrap bias correction for estimators that are smooth functions of sample moments, a class that includes the estimator \hat{k} considered here. He shows that an estimate of the bias of \hat{k} is the difference between the average of the bootstrap estimates and the original estimate, $\hat{B} = \sum_{j=1}^Q k_j/Q - \hat{k}$, and that a bias-corrected estimate of k is $\hat{k}^* = \hat{k} - \hat{B}$. According to equations 3.4 and 3.6 of Horowitz (2001), the correction removes bias up to order $O(N^{-1})$, where N is the sample size; higher-order bias can remain.

5 Empirical analysis: results

Figure 5 shows the time series of estimated prices and aggregate shocks. There is substantially uncertainty in the point estimates of the price and aggregate shock at each date. The uncertainty emphasizes the importance of accounting for sampling error in the calculations: Since much of the variability in the estimated aggregate shocks and prices is due to sampling error rather than true aggregate risk, we are at risk of overestimating the amount of aggregate risk and thus overestimating the welfare costs of risk. The bias correction discussed above is designed precisely to fix this problem. It allows us to obtain accurate estimates of

¹⁰In the data I employ, each household is observed at multiple dates. I therefore construct the bootstrap samples by resampling households and then, to account for serial correlation in aggregate shocks, resampling 18-month blocks of estimated prices and aggregate shocks.

welfare costs despite the noisy estimates of shocks and prices, in essence by estimating how much of the volatility in the estimated prices and shocks comes from sampling error rather than true risk.

Despite the sampling error, the estimated prices and aggregate shocks are consistent with what we expect from the model and from macroeconomic data. First, the prices and aggregate shocks are strongly negatively related, with a correlation coefficient of -0.94 , matching the prediction that prices are decreasing in aggregate resources.¹¹ Second, the aggregate shocks estimated from the Consumer Expenditure Survey match up well to the National Income and Product Accounts. The fourth panel of figure 5 shows the estimated aggregate shocks as well as detrended, seasonally adjusted quarterly per capita real personal consumption expenditures on nondurable goods and services. The NIPA series looks like a smoothed version of the CEX series, which is unsurprising since the CEX series contains sampling error. The correlation between the two series – using only every third observation from the CEX, since the NIPA data are quarterly rather than monthly – is 0.52. The correlation between the NIPA series and a three-month moving average of the CEX series is 0.61.

Figure 6 graphs the estimated willingness to pay to eliminate aggregate fluctuations as a function of the individual risk tolerance θ , for economies with different values of the average risk tolerance $\bar{\theta}$. (Given the distribution of risk tolerance, which I cannot estimate from the consumption data, one could convert these graphs into a distribution of the welfare costs of aggregate risk.) The willingness to pay is small. If the average person has a risk tolerance of 0.25, equivalent to a coefficient of relative risk aversion of 4, then the willingness to pay is less than three-tenths of a percent of consumption for even the most risk-averse

¹¹My estimation technique does not mechanically produce a negative relationship between prices and aggregate resources since prices are estimated from the mean of log consumption and aggregate shocks from the mean of the level of consumption. The cross-sectional mean of the log and the cross-sectional mean of the level can move in opposite directions if the variance increases when the level increases.

people, and anyone with a coefficient of relative risk aversion less than about 2 benefits from business cycles. If the average person has a risk tolerance of 1, or log utility, then the willingness to pay is less than one-tenth of a percent of consumption even for the most risk-averse person, and anyone with a coefficient of relative risk aversion less than about 0.5 has a welfare gain. Figure 6 also shows the willingness to pay of a representative agent with various levels of risk tolerance. As the theoretical analysis showed, a representative agent is always willing to pay more than an agent with the same risk tolerance in an economy where risk preferences vary. The representative agent's willingness to pay also diverges sharply as risk tolerance approaches zero, in contrast to the willingness to pay of agents in an economy where preferences vary. Figure 7 focuses on the areas in figure 6 where the representative agent's cost and the cost in a heterogeneous-agent economy are close; although the curves approach each other, the representative agent's cost remains strictly higher.

Table 2 lists the estimated willingness to pay of an infinitely risk averse household, as a function of the average risk aversion. These calculations show that even making people extremely risk averse does not produce enormous welfare costs from business cycles. Consider an economy where the average person has risk tolerance of 0.1, corresponding to a coefficient of relative risk aversion of 10. Then take an infinitely risk-averse household in this already quite risk-averse economy. The infinitely risk-averse household would be willing to pay less than three-quarters of a percent of consumption to eliminate aggregate risk.

The results are, of course, conditional on the aggregate shocks observed in my data. If the 1982-2002 period is not representative of the true aggregate risk facing the U.S. economy, for example due to the possibility of rarely observed disasters as in Barro (2007), my results could underestimate the true welfare cost of business cycles. However, even if the 1982-2002 period does not provide a representative sample of the distribution of aggregate shocks, the theoretical analysis shows that the central result of this paper – allowing heterogeneous preferences reduces welfare costs – would obtain under *any* trend-stationary i.i.d. distribution

of aggregate shocks, even a much riskier one. An aggregate shock process with rare disasters could raise the overall level of welfare costs but would not change the effect of heterogeneity.

A separate issue is whether aggregate income is i.i.d. and trend stationary. If aggregate income were a random walk, my estimates would be incorrect because the time series averages I use to compute welfare costs would not converge to the true costs. If aggregate income were trend stationary but not i.i.d., my results would understate the true welfare costs because persistent shocks with a small variance can have large effects on lifetime consumption. One of the objections to the results of Lucas (1987) has indeed been that persistent shocks would generate larger welfare costs (Obstfeld, 1994). To estimate welfare costs in a world with persistent shocks, however, we would need to find the competitive equilibrium in a dynamic model with agents who live more than one period. Such an equilibrium would be difficult to compute. Suppose agents live more than one period but die with some probability, leaving their assets to offspring who may have different preferences. Such a model has a non-degenerate long-run joint distribution of preferences and wealth, since new agents can be born with any combination of wealth and risk aversion. But the joint distribution of wealth and preferences may be non-stationary, since shocks will change the distribution of assets, leading to differences in the distribution of bequests depending on the history of shocks. If the joint distribution of wealth and assets is non-stationary, prices may also be non-stationary. However, the Krusell and Smith (1998) approximation method for non-stationary prices would be difficult to implement because we would need to keep track of prices for many different contingent claims, rather than a single risk-free rate. Further, arguably we should consider non-time-separable preferences to separate the roles of risk aversion and intertemporal substitution. Due to these technical challenges, I leave the analysis of a dynamic model for future research. Still, two points are worth noting:

- The results in this paper hinge on gains from trade that are present regardless of the distribution of shocks.

- The main reason to consider a dynamic model is that small persistent shocks have a large lifetime impact; in this sense, persistent shocks are similar to transitory shocks with a large variance.

Thus, while the small estimated variance of shocks in my data leads to small welfare costs, a larger variance would not change the results that heterogeneity in risk aversion generates gains from trade and reduces welfare costs. A reader who is concerned about the persistence of shocks may wish to conclude that the magnitude of welfare effects I estimate is too small, but the qualitative results should not be at issue.

6 Conclusion

One might think that, even if the average person does not suffer much from business cycles, a very risk-averse person could suffer greatly. My results show that this is not necessarily the case. In a complete-markets endowment economy where some people are very risk averse and others are not, the very risk-averse agents will buy insurance from less risk-averse agents and will not experience substantial consumption fluctuations; welfare losses are reduced for everyone. In other words, we cannot undo Lucas' (1987) result simply by appealing to the possibility that some people strongly dislike risk.

Business cycles may have welfare consequences for many reasons other than the variability of aggregate consumption *per se*. Aggregate shocks may increase the welfare losses associated with uninsured idiosyncratic risk (e.g., Krusell and Smith, 1999). Alternatively, reducing aggregate risk might increase aggregate income, for example because firms would not make *ex post* inefficient investments (Ramey and Ramey, 1991), because government policies could raise output in recessions without lowering it in booms (DeLong and Summers, 1988), or because removing fluctuations would raise the economy's growth rate (Barlevy, 2004). The point of this paper is simply that if business cycles matter, it is primarily for these other

reasons – not because anyone, even a hypothetical infinitely risk-averse person, suffers much disutility directly from fluctuations in the aggregate endowment.

It would be valuable for future work to investigate empirically the extent to which people with different preferences actually share aggregate risk, and the mechanisms they use to do so. The central implication of the model in this paper is that more risk-averse people's consumption moves less with aggregate shocks. Some evidence already exists on this issue. As noted earlier, Barsky et al. (1997) show that people who express greater risk aversion in a survey also report holding more bonds and fewer stocks; thus, their asset income is less correlated with aggregate shocks. Since financial markets do not offer pure state-contingent claims on the aggregate resources of the economy, people who share aggregate risk potentially use mechanisms other than financial markets to do so. Human capital is one possibility: In related work (Schulhofer-Wohl, 2007), I show that more risk-averse people have labor income that is less correlated with aggregate shocks. However, it is not known whether consumption responds to asset and labor income in a way that makes more risk-averse people's *consumption* move less with aggregate shocks.

Besides the implication about the relationship between consumption and preferences, the model in this paper can generate other implications given sufficient assumptions about market structure. For example, suppose that people trade a complete set of one-period Arrow securities and that they report as assets the market value of the Arrow-security portfolios. When there is a bad shock, relatively less risk-averse people consume less than their endowment, while relatively more risk-averse people consume more than their endowment. If someone consumes less than his endowment, he must save the rest; his assets increase. Hence, when there is a bad shock, the value of relatively less risk-averse people's portfolios should rise relative to the value of relatively more risk-averse people's portfolios. The opposite is true when there is a good shock. One could test this implication by combining panel data on people's portfolios and preferences with data on aggregate shocks.

A Proof of proposition

Since agents in the heterogeneous-agent economy can attain the utility (and hence the welfare cost) of the representative agent if they remain in autarky, it suffices to show that:

1. $U_t^*(\theta) \geq U_t^{rep}(\theta)$ for all θ in the support of $F(w_i, \theta_i)$, where U_t^* and U_t^{rep} are expected utility in the heterogeneous-agent competitive equilibrium and in the representative-agent economy at t , respectively, and both utilities are normalized by w_i^{1-1/θ_i} .
2. The inequality $U_t^*(\theta) \geq U_t^{rep}(\theta)$ is weak for at most one θ in the support.

Agents in the heterogeneous-agent economy attain utility U_t^{rep} if they remain in autarky. Since all agents must weakly prefer the competitive equilibrium to autarky, we thus have $U_t^*(\theta) \geq U_t^{rep}(\theta)$ for all θ in the support of $F(w_i, \theta_i)$. We will show by contradiction that the inequality is weak for at most one θ in the support. Suppose to the contrary that there are two agents i and j , $\theta_i \neq \theta_j$, such that the inequality is weak for both agents. The competitive equilibrium is in the core; therefore, i and j cannot do better than U_t^* by leaving the competitive equilibrium and trading with each other. Since i and j can attain U_t^{rep} by consuming their endowments, and since by hypothesis $U_t^{rep} = U_t^*$ for both i and j , it must be that i and j cannot improve on their endowments by trading with each other. However, if there is aggregate risk and $\theta_i \neq \theta_j$, it is not Pareto optimal for i and j to consume their endowments, since $u'_i(w_i g_t m_s) / u'_i(w_i g_t m_{s'}) \neq u'_j(w_j g_t m_s) / u'_j(w_j g_t m_{s'})$ when risk tolerances differ and $m_s \neq m_{s'}$. Hence i and j can do better by trading, a contradiction. \square

B Consumer Expenditure Survey sample selection

I use the Consumer Expenditure Survey for 1982 to 2002. As is common in work with this dataset, I drop all of the following due to concerns about data quality: incomplete income responders, non-urban households, individuals in student housing, households where the age

of the reference person or spouse changes by other than zero or one year between interviews, and all 1980 and 1981 data. I also drop households where the reference person or spouse is younger than 21 or older than 85 and where the marital status of the reference person changes during the survey period. I drop six observations where the reported consumption data do not span a three-month period and 18 observations that were not separated by three months from either the previous or subsequent observation. I then drop all households with fewer than four quarters of data. This leaves 124,348 observations on 31,087 households. Because changes in the survey prevent matching households across 1985-1986 and 1995-1996, I have two months without data: December 1985 and December 1995. In addition, there are relatively few observations on October and November 1985 and October and November 1995. From 92 to 733 households contribute to my estimates for each month.

I use data on nondurable goods and services. I will provide a list of the consumption categories I include upon request. I sum all expenditures by a household in a given month, deflate by the nondurable goods GDP deflator for that month, then sum the three months covered by an interview to create a quarterly consumption observation for the household. I divide consumption by effective household size, defined as 1 for a one-person household, 2 for a two-person household, and 0.4 additional units for each person after the second. I weight all results by the survey weights.

References

- Alvarez, Fernando and Urban J. Jermann, 2004, "Using Asset Prices to Measure the Cost of Business Cycles," *Journal of Political Economy* 112(6), 1223–1256.
- Barlevy, Gadi, 2004, "The Cost of Business Cycles Under Endogenous Growth," *American Economic Review* 94(4), 964–989.
- Barro, Robert J., 2007, "On the Welfare Costs of Consumption Uncertainty," Manuscript, Harvard University.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball and Matthew D. Shapiro, 1997,

- “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics* 112(2), 537–579.
- Blundell, Richard, Luigi Pistaferri and Ian Preston, 2005, “Consumption Inequality and Partial Insurance,” Manuscript, University College London and Stanford University.
- DeLong, J. Bradford and Lawrence H. Summers, 1988, “How Does Macroeconomic Policy Affect Output?” *Brookings Papers on Economic Activity* 1988(2), 433–480.
- Garner, Thesia I., George Janini, Laura Paszkiewicz and Mark Vendemia, 2006, “The CE and the PCE: a comparison,” *Monthly Labor Review* 129(9), 20–46.
- Horowitz, Joel L., 2001, “The Bootstrap,” in Heckman, James J. and Edward Leamer, editors, *Handbook of Econometrics*, volume V, pp. 3159–3228, Amsterdam: Elsevier.
- Krusell, Per and Anthony A. Smith, Jr., 1998, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy* 106(5), 867–896.
- Krusell, Per and Anthony A. Smith, Jr., 1999, “On the Welfare Effects of Eliminating Business Cycles,” *Review of Economic Dynamics* 2(1), 245–272.
- Lucas, Robert E., Jr., 1987, *Models of Business Cycles*, New York: Basil Blackwell.
- Obstfeld, Maurice, 1994, “Evaluating risky consumption paths: The role of intertemporal substitutability,” *European Economic Review* 38(7), 1471–1486.
- Ramey, Garey and Valerie A. Ramey, 1991, “Technology Commitment and the Cost of Economic Fluctuations,” Working paper No. 3755, National Bureau of Economic Research.
- Schulhofer-Wohl, Sam, 2007, “Heterogeneity, Risk Sharing and the Welfare Costs of Idiosyncratic Risk,” Manuscript, Princeton University.
- Wilson, Robert, 1968, “The Theory of Syndicates,” *Econometrica* 36(1), 119–132.

Table 1: Summary statistics for Consumer Expenditure Survey consumption data, 1982-2002.

Variable	mean	s.d.
quarterly consumption ^a	4716	2991
log(quarterly consumption) ^a	8.28	0.60
adjusted per capita consumption ^{a,b}	2431	1415
log(adjusted per capita consumption) ^{a,b}	7.66	0.51
Observations	124,348	
Households	31,087	
Quarters	250	
Households observed per quarter:		
mean	497	
minimum	92	
25th percentile	476	
median	537	
75th percentile	575	
maximum	733	

Data are on nondurable goods and services consumption. See appendix B for sample restrictions. Observations are one quarter's consumption for one household. Households can enter the survey in any month, so there are observations for 12 different quarterly consumption periods each year. ^aDeflated by GDP deflator for personal consumption expenditures on nondurable goods; 2000 dollars. ^bAdjusted per capita consumption is total consumption divided by effective household size, defined as 1 for a one-person household, 2 for a two-person household, and 0.4 additional units for each person after the second.

Table 2: Willingness to pay to eliminate aggregate fluctuations.

Mean risk tolerance	Willingness to pay of			
	most risk-averse household	average household	representative agent	representative agent
0.1	0.659% [1.467%]	0.339% [0.706%]	0.356% [0.791%]	0.356% [0.791%]
0.2	0.296 [0.660]	0.159 [0.334]	0.167 [0.375]	0.167 [0.375]
0.3	0.193 [0.428]	0.106 [0.219]	0.111 [0.246]	0.111 [0.246]
0.5	0.116 [0.251]	0.066 [0.130]	0.069 [0.146]	0.069 [0.146]
0.7	0.085 [0.178]	0.049 [0.092]	0.051 [0.104]	0.051 [0.104]
1.0	0.062 [0.124]	0.037 [0.064]	0.038 [0.072]	0.038 [0.072]
1.5	0.044 [0.082]	0.028 [0.043]	0.029 [0.048]	0.029 [0.048]
2.0	0.035 [0.062]	0.023 [0.032]	0.024 [0.036]	0.024 [0.036]

Willingness to pay is percentage of mean consumption. Most risk-averse household: a household whose risk tolerance (inverse of coefficient of relative risk aversion) approaches zero in an economy where the mean risk tolerance is that shown. Average household: a household whose risk tolerance is the mean risk tolerance shown, in an economy where risk preferences vary. Representative agent: a household in an economy where all households have the risk tolerance shown. Estimated using Consumer Expenditure Survey, 1982-2002. The first number in each column is a bootstrap bias-corrected estimate using 10,000 bootstrap replications, drawing households from the original sample with replacement and then drawing blocks of 18 consecutive months from the estimated aggregate shocks with replacement. The second number in each column, in brackets, is the estimate before bias correction.




Figure 1: Willingness to pay to eliminate aggregate fluctuations in example economies with a risk-neutral agent.

Each economy contains one agent who is risk neutral and one with log utility. The experiment varies the endowment share of the agent with log utility. Each agent's consumption is constrained to be non-negative. The representative agent line in each panel shows the willingness to pay in an economy where all agents have the same preferences as the agent illustrated in that panel. Distribution of aggregate shocks: $m_1 = 0.98$, $m_2 = 1.02$, $\pi_1 = \pi_2 = 0.5$.

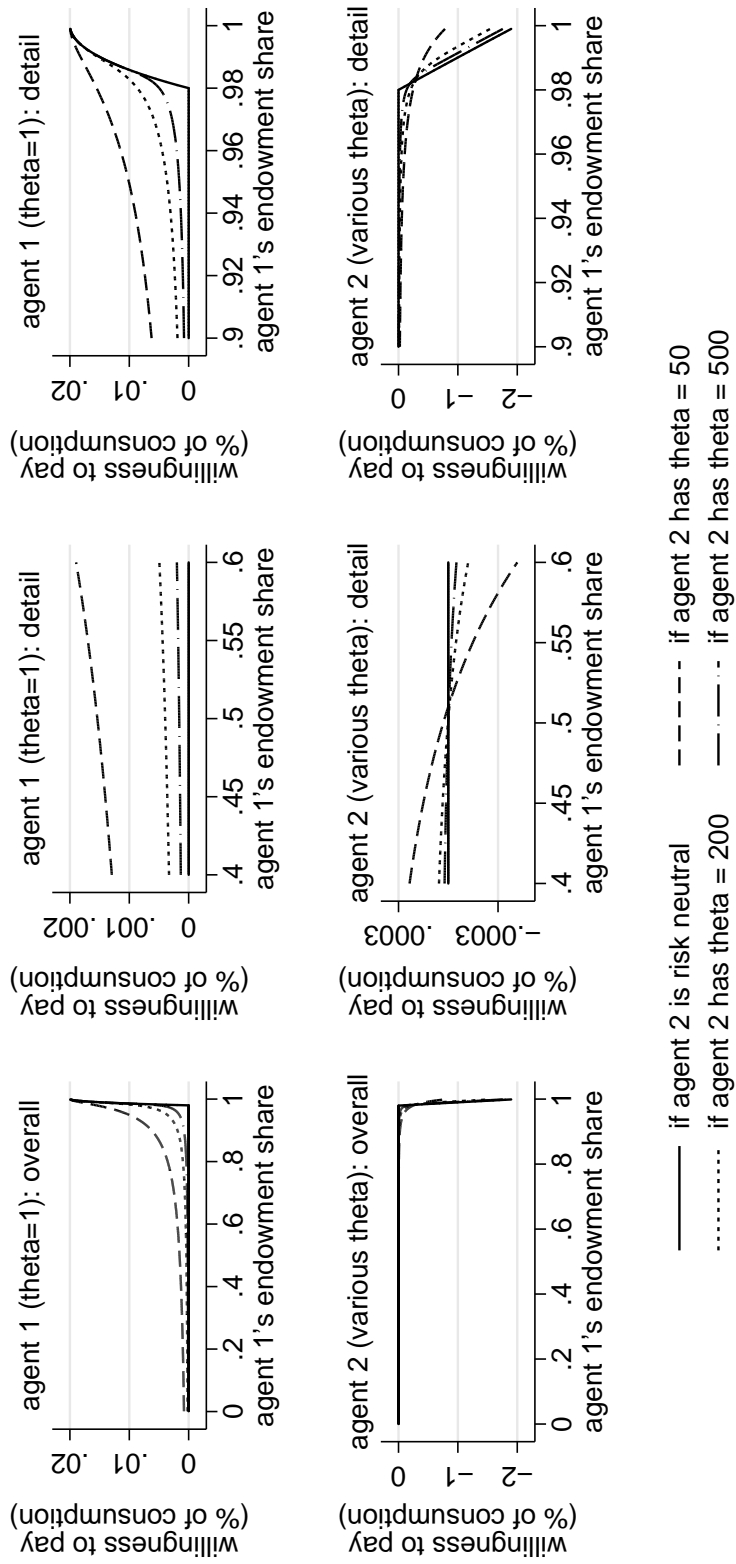


Figure 2: Willingness to pay to eliminate aggregate fluctuations in example economies as one agent's risk tolerance approaches infinity.

Each economy contains one agent with log utility. The experiment varies the endowment share of the agent with log utility and the risk tolerance of the other agent. Each agent's consumption is constrained to be non-negative. Distribution of aggregate shocks: $m_1 = 0.98$, $m_2 = 1.02$, $\pi_1 = \pi_2 = 0.5$.

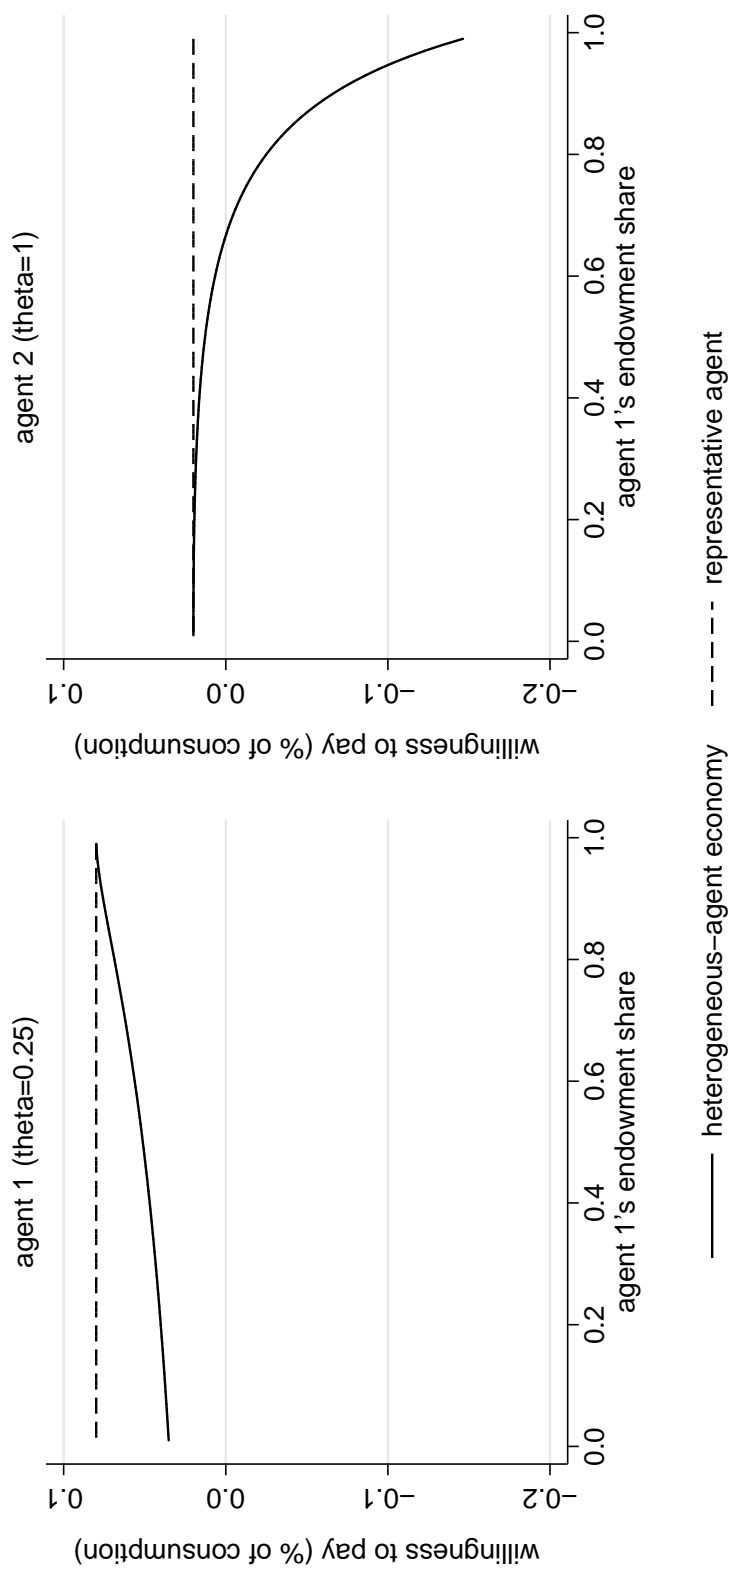


Figure 3: Willingness to pay to eliminate aggregate fluctuations in example economies with two risk-averse agents and varying endowment shares.

Each economy contains one agent with log utility and one with $\theta = 0.25$ (coefficient of relative risk aversion of 4). The experiment varies the endowment share of the agent with $\theta = 0.25$. The representative agent line in each panel shows the willingness to pay in an economy where all agents have the same preferences as the agent illustrated in that panel. Distribution of aggregate shocks: $m_1 = 0.98$, $m_2 = 1.02$, $\pi_1 = \pi_2 = 0.5$.




Figure 4: Willingness to pay to eliminate aggregate fluctuations in example economies with two agents and varying levels of risk aversion.

Each economy contains two agents. Agent 2 always has log utility and a total endowment share of 0.333. Agent 1 has an endowment share of 0.667, and the experiment varies this agent's risk aversion. The representative agent line in each panel shows the willingness to pay in an economy where all agents have the same preferences as the agent illustrated in that panel. Distribution of aggregate shocks: $m_1 = 0.98$, $m_2 = 1.02$, $\pi_1 = \pi_2 = 0.5$.

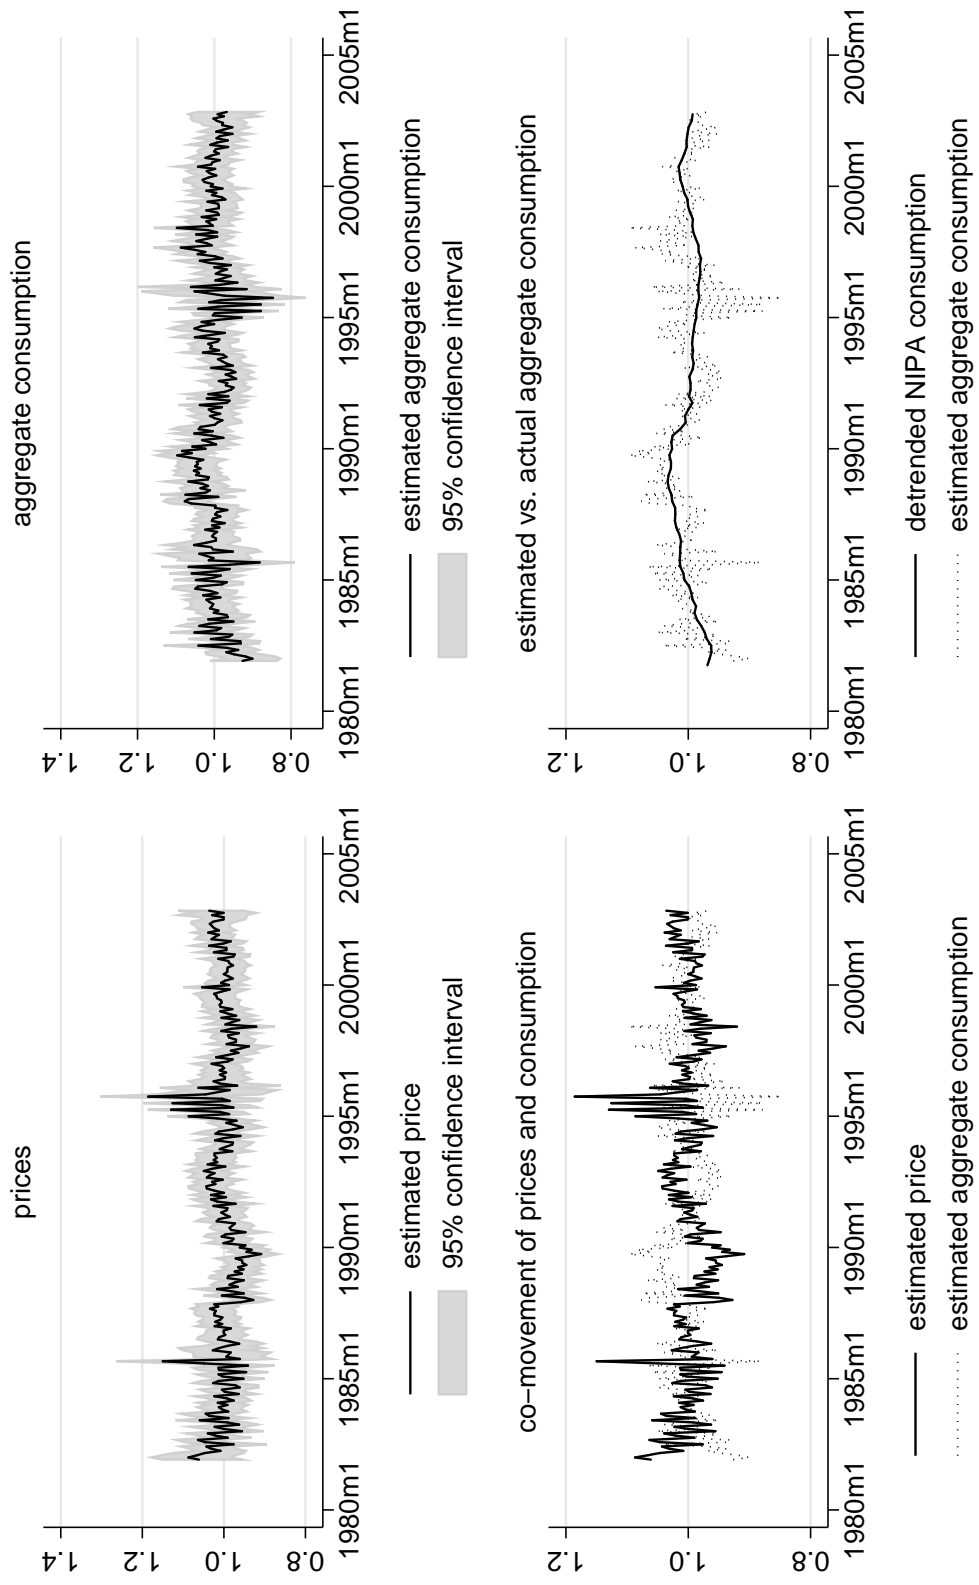


Figure 5: Estimated prices and aggregate consumption.

Estimated using Consumer Expenditure Survey, 1982-2002. Detrended NIPA consumption, for comparison, is exponentially detrended residual from regression on a trend of log per capita real, seasonally adjusted personal consumption expenditures on nondurables and services.




Figure 6: Willingness to pay to eliminate aggregate fluctuations, as a function of preferences.

Graph shows the percentage of mean consumption that a household would be willing to give up to eliminate aggregate fluctuations. Risk tolerance is inverse of coefficient of relative risk aversion. Negative numbers indicate a welfare loss from eliminating fluctuations. Estimated using Consumer Expenditure Survey, 1982-2002. Bootstrap bias correction and equal-tailed confidence interval use 10,000 bootstrap replications, drawing households from original sample with replacement and then drawing blocks of 18 consecutive months from estimated aggregate shocks with replacement.

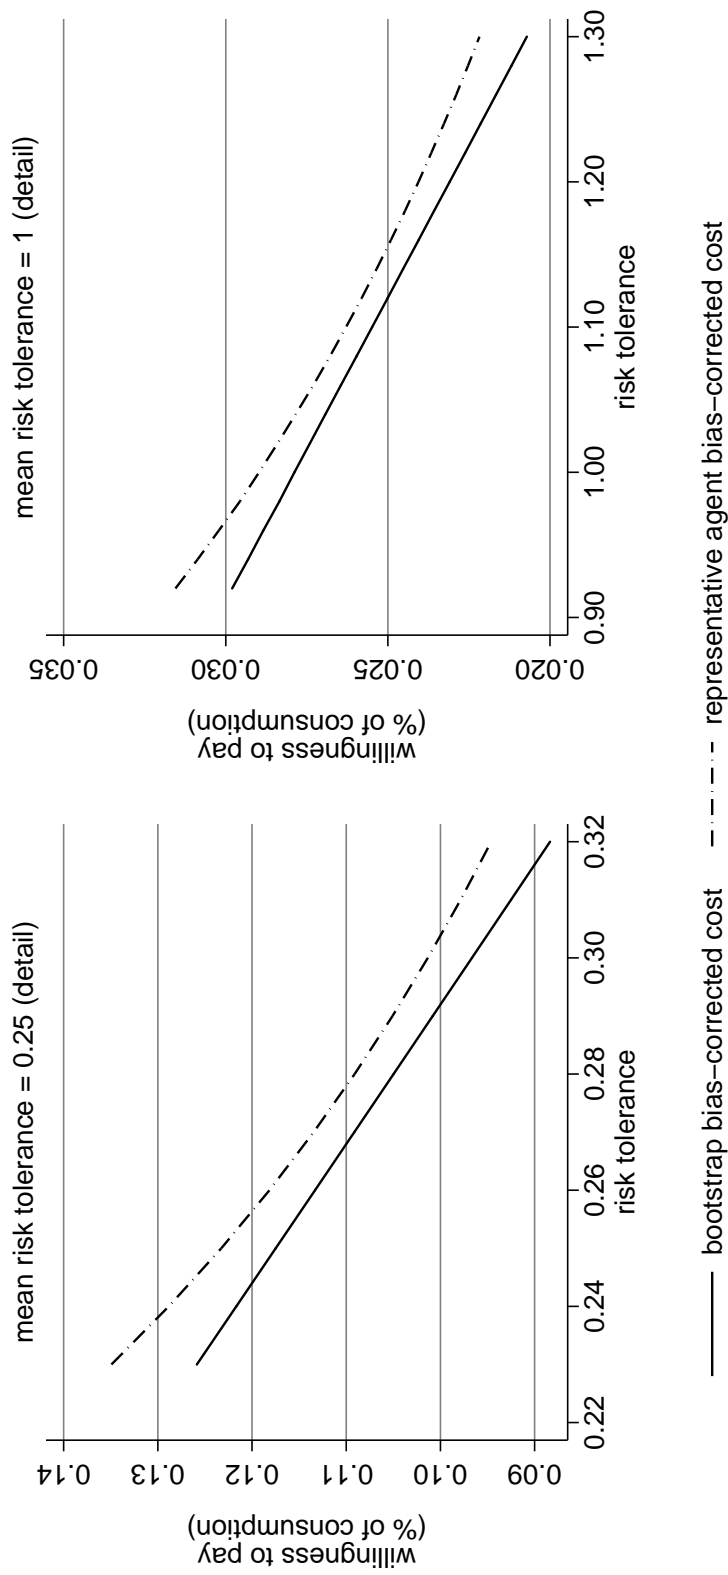


Figure 7: Willingness to pay to eliminate aggregate fluctuations, as a function of preferences: detail.

Graph shows the percentage of mean consumption that a household would be willing to give up to eliminate aggregate fluctuations. Risk tolerance is inverse of coefficient of relative risk aversion. Negative numbers indicate a welfare loss from eliminating fluctuations. Estimated using Consumer Expenditure Survey, 1982-2002. Bootstrap bias correction uses 10,000 bootstrap replications, drawing households from original sample with replacement and then drawing blocks of 18 consecutive months from estimated aggregate shocks with replacement.