The dynamics of risk-sensitive allocations

Evan W. Anderson[*]

February 20, 2004

Abstract

This paper describes Pareto optimal allocations when agents have risk-sensitive preferences as formulated by Hansen and Sargent (1995). Necessary and sufficient conditions are given for the existence of a steady state at which Pareto weights are time-invariant. Sufficient conditions are given for steady states to be stable or unstable. When all agents are risk-sensitive with power reward functions there is a unique interior steady state. This steady state is stable when the power is positive and unstable when the power is negative. When there is at least one agent with time-additive preferences eventually all risk-sensitive agents have zero Pareto weights.

1 Introduction

In standard dynamic macroeconomic models agents are assumed to have time-additive preferences over consumption of the form

$$E_{i0} \sum_{t=0}^{\infty} \beta^t u_i(c_{it}) \tag{1}$$

where c_{it} is the consumption of agent i at time t, u_i is a reward function, β is a time-discount factor and E_{i0} is a time zero expectation operator. Preferences are time-additive since they are linear in current and future rewards. Agents are usually assumed to have the same discount factors and the same information.

When agents have time-additive preferences, Pareto optimal allocations are given by a time-invariant function of aggregate consumption. At time t, the consumption allocation for agent i can be written as

$$c_{it} = c_i(e_t, \theta)$$

where e_t is aggregate consumption at time t, c_i is a function, and θ is a time-invariant vector of Pareto weights. Unless the distribution of aggregate consumption varies over time, the distribution of individual consumption allocations will not vary over time.[1] For example, at any two in dates in which aggregate consumption is identical, individual consumption must also be identical.

Evidence suggests that consumption allocations are not a time-invariant function of aggregate consumption.[2] Cochrane (1991) and Hayashi, Altonji, and Kotlikoff (1996) show consumption allocations are affected by health and employment. Attanasio and Davis (1996) find changes in consumption parallel changes in income. To account for consumption the standard model (which assumes complete markets, full information and time-additive preferences) needs to relax some of its assumptions. Papers by Phelan and Townsend (1991), Atkeson and Lucas (1995) and others have looked at the implications of surrendering perfect information. Papers by Heaton and Lucas (1995), Alvarez and Jermann (2000), and others have looked at the implications of surrendering complete markets. This paper maintains the assumptions of complete markets and full information but surrenders the assumption of time-additive preferences.

There are good reasons for surrendering the assumption of time-additive preferences. Time-additivity imposes strong restrictions on preferences which require agents to be risk-neutral in future *utility*. Con-

---

[1] With time-additive preferences, if agents have different discount factors or different information then optimal allocations may vary over time even if aggregate consumption does not vary. If agent's preferences display habit persistence or if consumption goods are durable then consumption allocations may depend on additional state variables. See Hansen (1987) and chapter 13 of Hansen and Sargent (1997) for a description of how this can be done for quadratic preferences.

[2] See Browning, Hansen, and Heckman (2000) for a useful survey and arguments for including microeconomic factors and heterogeneity in macroeconomic models.

sider the following two gambles which have been discussed by Duffie and Epstein (1992), Schroder and Skiadas (1999), and others. In gamble A a coin is tossed every period in the future. In periods in which heads obtains the lottery pays ten consumption units. In periods in which tails obtains the lottery pays one consumption unit. In gamble B a coin is tossed only once. If heads obtains then the lottery pays ten consumption units every period in the future. If tails obtains then the lottery pays one consumption unit every period in the future. Assume the payoffs of the gambles represent the final consumption of agents. If agents have time-additive preferences then they must be indifferent between the two lotteries. Yet, it seems rational for agents to have preferences over the gambles. An agent might prefer gamble A because he views the distribution of utilities as more favorable under this gamble. A theory of preferences should not *require* agents to be indifferent between these gambles any more than it should require agents to be indifferent between consumption today and consumption tomorrow. The preferences described in this paper allow agents to have preferences over the gambles.

This paper studies allocations when agents have risk-sensitive preferences of the form

$$U_{it} = u_i\left(c_{it}\right) + \frac{\beta}{\sigma_i}\log E_{it}\exp\left(\sigma_i U_{it+1}\right) \tag{2}$$

where σ_i is a risk-sensitivity parameter and U_{it} is the lifetime utility from time t onward. Versions of risk-sensitive preferences have been discussed in the control theory literature dating back to Jacobson (1973). The form of risk-sensitive preferences used in this paper is due to Hansen and Sargent (1995) and incorporates discounting in a different way than the control theory literature. Risk-sensitive preferences have been previously studied in representative agent economies by Hansen, Sargent, and Tallarini (1999) and Tallarini (2000).

Risk-sensitive preferences are interesting to consider for several reasons. First, risk-sensitive preferences are not time-additive in future utility and allow agents to be risk averse in future utility in addition to being risk averse in future consumption. When $\sigma_i < 0$, risk-sensitive agents always prefer gamble A to gamble B. Second, risk-sensitive preferences have been argued by Hansen, Sargent, and Tallarini (1999) to be attractive since they can be used to model preferences for robustness. Third, as discussed by Tallarini (2000), risk-sensitive preferences can allow for a partial separation between risk aversion and the elasticity of intertemporal substitution.[3]

---

[3]See Epstein and Zin (1989) for more discussion on non time-additive preferences and the separation between risk aversion and the elasticity of intertemporal substitution.

This paper shows the implications for individual consumption when agents are risk-sensitive but possibly differ in their preferences and their initial wealth. Following Lucas and Stokey (1984) and Kan (1995) we show that for any date $t \geq 0$ Pareto optimal optimal allocations can be written as

$$c_{it} = c_i(e_t, \theta_t)$$

where e_t is aggregate consumption at time t, c_i is a function, and θ_t is a vector of time t Pareto weights. The time t Pareto weights are history dependent. Optimal allocations at any date are determined solely by the current period aggregate endowment and the current period Pareto weights. Typically Pareto weights evolve over time, so even when the distribution of aggregate consumption does not evolve, the distribution of individual consumption can evolve. For example, at any two dates in which aggregate consumption is identical, individual consumption will differ if the Pareto weights differ.

To study the dynamics of Pareto optimal allocations it is convenient to study the dynamics of the Pareto weights. This paper gives necessary and sufficient conditions for the existence of a steady state at which Pareto weights are time-invariant. Steady states are particularly interesting if they are stable. For deterministic economies with agents who have non time-additive preferences, Lucas and Stokey (1984) showed if preferences display *increasing marginal impatience* then interior steady states are stable. Increasing marginal impatience implies, when all agents have identical preferences, "wealthy" agents discount the future more than "poor" agents. In this case the consumption allocations of wealthy (respectively poor) agents decrease (increase) over time until all agents have identical consumption allocations.[4]

In stochastic economies increasing marginal impatience is neither necessary nor sufficient for the stability of interior steady states. Another mechanism, the extent to which agents distort probabilities, in conjunction with marginal impatience determines the stability of steady states. Agents have *increasing* distortions if at the Pareto optimal allocations deviations from time-additive preferences are increasing in wealth. When all agents have the same constant impatience then increasing distortions is sufficient to guarantee the stability of interior steady states. When agents have the same constant impatience then *decreasing* distortions is sufficient to guarantee interior steady states are not stable.

The major findings of this paper are that

1. Steady states often provide bounds on the evolution of Pareto weights. When technical conditions are satisfied the ratio of the Pareto weights of any two agents can never cross a steady state ratio of Pareto weights. If there are a multiple steady states, then none of them can be crossed. The

---

[4]In discrete time Benhabib, Jafarey, and Nishimura (1988) and in continuous time Epstein (1987) also discuss the dynamics of allocations in deterministic economies.

technical conditions required a very weak and are typically only violated in degenerate economies in which the exogenous aggregate state evolves in deterministic patterns.

2. If there is at lease one agent with time-additive preferences then the weights of all risk-sensitive agents tend toward zero over time with probability one. Though there are histories, which occur with probability zero, in which risk-sensitive agents see the sum of their Pareto weights converge to one.

3. When there are two agents who are both risk-sensitive with logarithmic rewards, the weight of the agent who has a more negative risk-sensitivity parameter converges to zero with probability one.

4. When there are two agents who are both risk-sensitive with the same power reward function, there is a unique interior steady state. The steady state is stable or not stable depending upon the magnitude of the power. If the power is positive then the steady state is stable. In this case if agents have the same risk-sensitivity parameter then the agents will eventually be equal with probability one even if one agent is initially richer. If the power is negative then it is possible with positive probability that one agent will end up consuming everything and have a Pareto weight of one.

This paper is organized as follows. Section two describes resource constraints, information and preferences. Section three discusses optimal allocations and describes a recursive social planning problem for computing allocations. Sections four, five, and six characterize the dynamics of optimal allocations by looking at steady states, off steady state dynamics, and stability. Section seven interprets the stability results in terms of distorted probabilities and decentralizes the allocations discussed in earlier sections. Section eight studies economies in which there are agents who have both risk-sensitive and non risk-sensitive preferences. Section nine studies economies in which agents have power reward functions. Section 10 concludes.

2 Economic environment

This section describes resource constraints, information and preferences.

2.1 Resource constraints

Let x_t be the exogenous state vector of the economy at time t. Let $e_t = c(x_t)$ be the aggregate endowment at time t where c is a time-invariant function of the exogenous state. We make the following assumptions on exogenous states and endowments:

Assumption 1 *The exogenous state vector x follows a first-order Markov process. There are a finite number of possible values for the exogenous state vector. Let $\mathcal{X}$ be the set of all of the possible values. Let m be the number of possible values.*

Assumption 2 *The aggregate endowment is bounded below by $e_{min} > 0$ and bounded above by $e_{max} < +\infty$. For all states x*

$$e_{min} \leq c(x) \leq e_{max}. \tag{3}$$

All of the results in this paper will use Assumption 1 and 2. These assumptions make the results of this paper easier to prove but are not essential.

Resource constraints require aggregate consumption to be less than or equal to the aggregate endowment. No individual is allowed to consume negative quantities. Individual consumption must satisfy

$$\sum_{i=1}^{n} c_{it} \leq c(x_t), \tag{4a}$$

$$c_{it} \geq 0, \qquad\qquad 1 \leq i \leq n, \tag{4b}$$

at all dates $t \geq 0$ and in all states. Here $c(x_t)$ is the aggregate endowment at time t, c_{it} is the consumption of agent i at time t, and n is the number of agents.[5] Storage is not permitted and there is no production. Since later it will be assumed that agents always prefer more to less, at any Pareto optimal allocation aggregate consumption will equal the aggregate endowment. Consequently we sometimes will write the constraint in equation 4a as $\sum_{i=1}^{n} c_{it} = c(x_t)$.

2.2 Information

All agents observe the current exogenous state and remember previous values. If this period's state is x then let the probability next period's state is y be $\pi(x, y)$. Let $\pi_i(x, y)$ be the probability that agent i believes y will be next period's state. The following assumptions on beliefs will be used in all of the results of this paper

Assumption 3 *For all agents i and for all states x and y, it is the case that $\pi(x, y) > 0$ if and only if $\pi_i(x, y) > 0$.*

This assumption guarantees beliefs are absolutely continuous with respect to the true probabilities and the true probabilities are absolutely continuous with respect to beliefs.

---

[5]Through out the paper we will let there be n agents. All of the results of the paper apply if there are equal numbers of n types of agents.

Even though all agents observe the current exogenous state and remember previous values, agents are not allowed to learn about the true probabilities $\pi(x, y)$. When agents have inaccurate beliefs they live in a permanent state of error.

2.3   Preferences

Agents have risk-sensitive preferences which can be written in several different ways. The first way is informal and was used in the introduction. This section shows how to write preferences in two additional ways. One method writes preferences in terms of current consumption and future anticipated utility. The other method writes preferences in terms of current consumption and future anticipated consumption. These additional ways of writing preferences will be useful in later sections.

In order to write preferences in terms of current consumption and future anticipated utility, define q_i to be a vector which gives the anticipated utility, from next period on, contingent on next period's state. Let q_{iy} be the element of q_i which gives the utility remaining from next period on when next period's state is y. Agent's preferences over current consumption and future anticipated utility can be written as

$$W_i(x, c_i, q_i) = u_i(c_i) + \frac{\beta}{\sigma_i} \log \sum_{y \in \mathcal{X}} \pi_i(x, y) \exp(\sigma_i q_{iy}). \tag{5}$$

The reward function u_i gives the contribution of current consumption to utility and the logarithmic term gives the contribution of future anticipated utility to current utility. This method of writing preferences is related to the method described in the introduction, in equation two, since

$$U_{it} = W_i(x_t, c_{it}, U_{it+1}) \tag{6}$$

when we interpret the U_{it+1} on the right hand side of equation six as a vector giving the future anticipated utility for each possible value of the time $t + 1$ state.

In order to write preferences in terms of current and future anticipated consumption, define a consumption plan for agent i for all current and future dates as

$$\mathbf{C}_i = (\mathbf{c}_i, \mathbf{c}_{i1}, \mathbf{c}_{i2}, \ldots). \tag{7}$$

Here $\mathbf{c}_i$ is a scalar representing current period consumption and $\mathbf{c}_{it}$ is a m^t dimensional vector which gives a plan for consumption t periods ahead, contingent on the states between periods one and t. To write preferences in terms of consumption it is convenient to be able to shift forward the consumption plans in $\mathbf{C}_i$ one period, after the state next period is realized. When time is shifted forward there is

no longer any reason to keep track of consumption plans in future states which no longer can occur. Denote a one period ahead version of $\mathbf{C}_i$ as

$$\tau_y \mathbf{C}_i = (\mathbf{c}_{i1}(y), \mathbf{c}_{i2}(y), \mathbf{c}_{i3}(y), \ldots)$$

where y is next period's state. Here $\mathbf{c}_{it}(y)$ is a m^{t-1} dimensional vector which gives the plan specified in $\mathbf{c}_{it}$ for consumption t periods ahead when y is next period's state.

We can write preferences over current and future anticipated consumption as $U_i(x, \mathbf{C}_i)$ where U_i is a fixed point of the functional equation

$$U_i(x, \mathbf{C}_i) = u_i(c_i) + \frac{\beta}{\sigma_i} \log \sum_{y \in \mathcal{X}} \pi_i(x, y) \exp\left[\sigma_i U_i(y, \tau_y \mathbf{C}_i)\right]. \tag{8}$$

When the fixed point is not unique, U_i can be uniquely defined using the arguments in Ozaki and Streufert (1996). This way of writing preferences is identical to writing preferences in terms of current consumption and future utility when the anticipated future utility is equal to the utility received from the consumption stream $\mathbf{C}_{iy}$. Thus

$$U_i(x, \mathbf{C}_i) = W_i(x, c_i, q_i)$$

where for all y, $q_{iy} = U_i(y, \tau_y \mathbf{C}_i)$.

The following assumptions on preferences will be made:

Assumption 4 (Reward function) *The reward functions $u_i(c_i)$ for all agents are strictly concave, strictly increasing, continuously differentiable (when $c_i > 0$) and such that*

$$u_i(e_{max}) < +\infty,$$
$$u_i(c_i) > -\infty, \;\; \text{when } c_i > 0,$$
$$\lim_{c_i \to 0^+} u_i'(c_i) = +\infty,$$

where e_{max} is the maximum possible value of the aggregate endowment.

Assumption 5 (Discount factor and risk-sensitivity parameter) *The time-discount factor β is a constant which is greater than zero and less than one. All agents have the same time-discount factor. The risk-sensitivity parameter σ_i for all agents is constant and less than or equal to zero.*

A reward function is a function which maps a non-negative real number into a real number. It gives the contribution of current consumption to current utility. Assumption four allows for the two reward

functions that will be studied in this paper: the logarithmic reward function of the form $u_i(c_i) = \log c_i$ and the power reward function of the form $u_i(c_i) = c_i^{1-\gamma_i}/(1 - \gamma_i)$ for $\gamma_i > 0$ and $\gamma_i \neq 1$. We require all agents have the same discount factor to abstract from the effects of heterogeneous discounting.

The parameter σ_i affects agents' attitudes towards risk in future *utility*. Standard time-additive preferences allow agents to be risk-averse in future consumption, but not future utility. As σ_i approaches zero from the left preferences collapse to the usual time-additive specification given in equation one. In this case agents are risk-neutral towards future utility. When $\sigma_i < 0$ agents are risk-averse towards future utility. In this case, agents will prefer gamble A (discussed in the introduction) to gamble B. When an agent's risk-sensitivity parameter is equal to zero we will say they have time-additive preferences or degenerate risk-sensitive preferences. When an agent's risk-sensitivity parameter is negative we will say they are strictly risk-sensitive or have non-degenerate risk-sensitive preferences. For a fixed reward function, as an agents risk-sensitivity parameter becomes smaller they become *more* risk-sensitive.

The extent to which agents are risk-sensitive is jointly determined by their reward function and their risk-sensitivity factor. To see this, we discuss an equivalence result between preference orderings. Let z be a positive constant. The preference ordering induced by the reward function u_i, the risk-sensitivity parameter σ_i, and the discount factor β is equivalent to the preference ordering induced by the reward function

$$\left(\frac{1}{z}\right) u_i, \tag{9}$$

the risk-sensitivity parameter $z\sigma_i$, and the discount factor β. So, if the risk-sensitivity factor of an agent is scaled upward by a factor of z then preference orderings can be maintained if their reward function is scaled downward by a factor of z. Unlike with time-additive preferences, scaling reward functions will change preference orderings unless σ_i is altered appropriately.[6] We also see one does not want to only use the magnitude of σ_i to measure the degree an agent is risk-sensitive since both the reward function and the risk-sensitivity parameter matter. If the reward function is held fixed then as σ_i decreases agents are more risk-sensitive. But as the reward function is increased in magnitude, holding everything else fixed, the agent also becomes more risk-sensitive.

---

[6]To see this divide both sides of equation eight by z and multiply the terms inside the exponential by z/z to obtain

$$\frac{U_i(x, \mathbf{C}_i)}{z} = \frac{u_i(c_i)}{z} + \frac{\beta}{z\sigma_i} \log \sum_{y \in \mathcal{X}} \pi_i(x,y) \exp\left[z\sigma_i \frac{U_i(y, \tau_y \mathbf{C}_i)}{z}\right]. \tag{10}$$

If U_i is a fixed point of equation eight then U_i/z is a fixed point of equation 10. The preference ordering given by U_i/z is equivalent to the preference ordering given by U_i.

3 Optimal allocations

Following Lucas and Stokey (1984) and Kan (1995), this section describes a recursive social planning problem for computing Pareto optimal allocations.[7] The social planning problem is a non standard dynamic programming problem in which Pareto weights evolve over time and are endogenous state variables. Current period Pareto weights summarize the effect of the history of aggregate consumption and the initial Pareto weights on current and future consumption allocations.

After some preliminary definitions in Section 3.1 and discussion in Section 3.2, Section 3.3 describes a dynamic programming problem for computing optimal allocations. Section 3.4 describes solutions to the dynamic programming problem.

3.1 Definitions

This section defines feasible allocations, Pareto optimal allocations, and Pareto weights. Consider the environment described in Section two with n agents. Let $\mathbf{C} = \{\mathbf{C}_i\}_{i=1}^n$ and $\mathbf{D} = \{\mathbf{D}_i\}_{i=1}^n$ be consumption plans for all the agents which prescribe values for the consumption in all current and future possible states of the world.

Definition 1 *A* feasible allocation *is an allocation which satisfies equations* (4a) *and* (4b) *for all dates and states.*

Definition 2 *A feasible allocation* $\mathbf{C}$ *is a* Pareto optimal allocation *if there is no other feasible allocation* $\mathbf{D}$ *such that* $U_i(x, \mathbf{D}_i) \geq U_i(x, \mathbf{C}_i)$ *for all i and* $U_i(x, \mathbf{D}_i) > U_i(x, \mathbf{C}_i)$ *for some i.*

Definition 3 *A* Pareto weight vector *is an n dimensional vector whose elements are non-negative and sum to one. An* interior Pareto weights vector *is a Pareto weight vector whose elements are all strictly positive.*

3.2 Discussion

A Pareto optimal allocation can be computed by maximizing a weighted sum of utilities by choice of feasible consumption allocations. Let θ be a Pareto weight vector and x_0 be the initial (time zero) state. A Pareto optimum can be computed by maximizing

$$\sum_{i=1}^n \theta_i U_i(x_0, \mathbf{C}_i) \tag{11}$$

---

[7]In continuous time, Duffie, Geoffard, and Skiadas (1994) and Dumas, Uppal, and Wang (2000) discuss methods for computing optimal allocations with recursive utility.

by choice of a feasible consumption plan for all agents, $\mathbf{C}$. Here θ_i is the Pareto weight on agent i. Call this problem the social planning problem. By varying the Pareto weight vector all Pareto optima can be computed.

Given the assumptions on preferences made in Section 2.3, a Pareto optimal allocation can be computed from the first-order conditions of the social planning problem. The first-order conditions dictate equating the ratios of the time zero marginal values of current and future consumption for any two agents across states. Let the history of states between time zero and time t be

$$x^t = \{x_0, x_1 \ldots x_t\}. \tag{12}$$

Let

$$U_{it}\left(x^t\right) = U_i\left(x_t, \left[\tau_{x_t}\tau_{x_{t-1}} \cdots \tau_{x_1}\right]\mathbf{C}_i\right) \tag{13}$$

be the utility remaining from time t, state x^t onward when consumption from time zero onward is specified in $\mathbf{C}_i$. The time zero marginal value of time t, state x^t consumption is

$$\frac{\partial U_{i0}\left(x^0\right)}{\partial c_{it}\left(x^t\right)} = \frac{\partial U_{i0}\left(x^0\right)}{\partial U_{it}\left(x^t\right)}\frac{\partial U_{it}\left(x^t\right)}{\partial c_{it}\left(x^t\right)} \tag{14}$$

$$= \mathcal{M}_{it}\left(x^t\right)u_i'\left[c_{it}\left(x^t\right)\right] \tag{15}$$

where

$$\mathcal{M}_{it}\left(x^t\right) = \frac{\partial U_{i0}\left(x^0\right)}{\partial U_{it}\left(x^t\right)} \tag{16}$$

is the marginal contribution of time t, history x^t utility to time zero utility. First-order conditions for the social planning problem are

$$\frac{u_i'\left[c_{it}\left(x^t\right)\right]\mathcal{M}_{it}\left(x^t\right)}{u_j'\left[c_{jt}\left(x^t\right)\right]\mathcal{M}_{jt}\left(x^t\right)} = \frac{\theta_j}{\theta_i} \tag{17}$$

for all dates, for all states and for all agents i and j.

When agents have time-additive preferences (with the same beliefs and the same discount factors) the marginal contributions of future utility are constant across agents. So $\mathcal{M}_{it}/\mathcal{M}_{jt}$ is equal to one. Thus the first-order conditions of the social planning problem dictate setting

$$\frac{u_i'\left[c_{it}\left(x^t\right)\right]}{u_j'\left[c_{jt}\left(x^t\right)\right]} = \frac{\theta_j}{\theta_i} \tag{18}$$

at all dates and states. In this case, optimal allocations are easy to compute since the left hand side of equation 18 depends *only* upon time t consumption.

When agents have risk-sensitive preferences optimal allocations are typically difficult to compute since the marginal value of time t utility at time zero depends non-trivially on future anticipated utility. The marginal contribution of time zero utility to time zero utility is one. The marginal contribution of time $t \geq 1$, history x^t utility to time zero utility is

$$\mathcal{M}_{it}\left(x^t\right) = \Pi_{s=1}^t \frac{\beta \pi_i(x_{s-1}, x_s) \exp\left[\sigma_i U_{is}\left(x^s\right)\right]}{\sum_{y \in \mathcal{X}} \pi_i(x_{s-1}, y) \exp\left[\sigma_i U_{is}\left(\{x^{s-1}; y\}\right)\right]} \tag{19}$$

where $\{x^{s-1}; y\}$ denotes a hypothetical history of states at time t which includes the history in x^{t-1} for dates zero to $t-1$ and the state y at time t. $U_{is}\left(x^s\right)$ is the anticipated utility from time s, history x^s onward. Computing $\mathcal{M}_{it}$ is not easy since it depends upon agent's anticipated utilities which in turn depend upon all future anticipated consumption. The next section describes a dynamic programming problem which makes the computation of Pareto optima tractable.

The dynamic programming problem makes the planning problem recursive by defining time t, history x^t, Pareto weights for agent i as

$$\theta_{it}\left(x^t\right) = \frac{\theta_i \mathcal{M}_{it}\left(x^t\right)}{\sum_{k=1}^n \theta_k \mathcal{M}_{kt}\left(x^t\right)} \tag{20}$$

where θ_i is the initial, time zero, Pareto weight used in equation 11. With this definition the first-order conditions for time t consumption can be written as

$$\frac{u_i'\left[c_{it}\left(x^t\right)\right]}{u_j'\left[c_{jt}\left(x^t\right)\right]} = \frac{\theta_{jt}\left(x^t\right)}{\theta_{it}\left(x^t\right)}. \tag{21}$$

If we know the time t Pareto weights we can compute the time t consumption allocations by solving a one period social planning problem with weights θ_{it}. The Pareto weights index inequality at time t and summarize the effects of the history of the exogenous state on consumption allocations. The time t Pareto weights are also the weights that a social planner could use at date t to recompute allocations from time t onward. Viewed from time zero, there is in a sense only one vector of Pareto weights, θ, that applies at all dates. In future dates its effect on consumption allocations depends upon the marginal value of future utility thru $\mathcal{M}_{it}$. However, by interpreting the Pareto weights as evolving over time, we can conveniently index time-varying consumption allocation rules by a vector of parameters and, as the next section discusses, time-varying Pareto weights allow the Pareto optimal problem to be written recursively.

3.3 A dynamic programming problem

Define the social planner's value function to be

$$Q\left(x,\theta\right) = \sup_{\mathbf{C}\in\mathcal{C}(x)} \sum_{i=1}^{n} \theta_i U_i(x,\mathbf{C}_i) \tag{22}$$

where $\mathcal{C}(x)$ is the set of all feasible consumption allocations. The social planner's value function is the maximum value that can be achieved for the weighted sum of individual utilities. Under the assumptions on preferences given in Section 2.3 the supremum is obtained and the sup in equation 22 can be replaced with a max.

This section shows how to compute the social planner's value function recursively. To do this we need to know the minimum amount of utility that an agent could receive. Let

$$\underline{q}_i = (1-\beta)^{-1} u_i(0) \tag{23}$$

be the extended real number which gives the lifetime utility obtained from a consumption plan which entails consuming zero at all dates. Depending upon the agent i's reward function, $\underline{q}_i$ may be minus infinity or finite. This is the lowest utility agent i could receive at any Pareto optimum since agents are not allowed to consume negative quantities.

By virtue of the properties of risk-sensitive preferences one can show that the social planner's value function Q, defined in equation (22), satisfies the following functional equation proposed by Lucas and Stokey (1984) and Kan (1995):

$$Q\left(x,\theta\right) = \max_{\{c_i\geq 0, q_i\geq \underline{q}_i\}_{i=1}^{n}} \sum_{i=1}^{n} \theta_i W_i(x,c_i,q_i) \tag{24a}$$

where

$$\sum_{i=1}^{n} c_i = c(x) \tag{24b}$$

and for each possible value, y, of next period's state

$$0 \leq \min_{\phi_y\in\Gamma^n} \left[Q\left(y,\phi_y\right) - \sum_{i=1}^{n} \phi_{iy} q_{iy} \right]. \tag{24c}$$

Recall W_i was defined in Section 2.3. The set Γ^n is the set of all Pareto weight vectors which sum to one. In equation (24c) q_{iy} is the continuation value for agent i when next period's state is y. $\phi_y = \{\phi_{iy}\}_{i=1}^{n}$ is

the Pareto weight vector next period when next period's state is y. The weight on agent i next period is ϕ_{iy}.

This is a non standard dynamic programming in which the social planner chooses current consumption allocations, continuation values for next period and Pareto weights for next period. The choice of consumption allocations is subject to the feasibility constraint (24b). The choice of continuation values is also subject to the m feasibility constraints in equation (24c). Kan (1995) shows that the equations in (24c) are equivalent to the requirement that the continuation values, q_i for all i, are feasible. The Pareto weights that obtain the minimum in constraint (24c) for $y \in \mathcal{X}$ will be the optimal choice of next period's Pareto weights when next period's state is y.

This formulation has been previously justified in a general deterministic setting by Lucas and Stokey (1984) and in a stochastic setting by Kan (1995). In virtue of the properties of risk-sensitive preferences all of the assumptions that Kan (1995) uses to justify the dynamic program, except for one, are met. The assumption that the reward functions are bounded from below is not met and some of our examples will have reward functions that are unbounded from below. However since aggregate consumption is bounded below, the social planner's value function is also bounded from below. Using this fact it is possible to extend the results of Lucas and Stokey and Kan to the setting in this paper.

The dynamic problem can be used for computation. One can iterate on an initial guess of the social planner's value function Q to compute the the social planner's value function. See the appendix for details. We will exploit this in later sections to compute numerical approximations to the solution of the social planning problem.

3.4   Solution to the dynamic programming problem

This section gives the first-order conditions for the dynamic programming problem. These conditions are used in later section to analyze dynamics. The first-order conditions will be written in terms of the value functions for individual agents and the marginal values of future utilities.

Let $Q_i(x, \theta)$ give the value remaining in the optimal allocation for agent i from this period on if current state is x and the Pareto weights are θ. Q_i will be referred to as the value function for agent i. The value function for an individual agent can be written in terms of the social planner's value function. To see this, begin by writing the social planner's value function in terms of the individual value functions

$$Q(x, \theta) = \sum_{j=1}^{n} \theta_j Q_j(x, \theta).$$

Differentiating the social planner's value function with respect to the Pareto weight for agent i yields

$$\frac{\partial Q\left(x,\theta\right)}{\partial\theta_i} = Q_i\left(x,\theta\right) + \sum_{j=1}^{n}\theta_j\frac{dQ_j\left(x,\theta\right)}{d\theta_i} \tag{25}$$

$$= Q_i\left(x,\theta\right). \tag{26}$$

The summation in equation 25 is zero because of optimality.[8] These derivatives must exist when θ_i is interior. When θ_i is zero, the solution to the social planner's problem entails allocating zero consumption to agent i at all dates and states. So in this case the value remaining for agent i is $\underline{q}_i$. When $\theta_i = 1$ the solution to the social planner's problem entails allocating all aggregate consumption to agent i. So, in this case the value remaining for agent i is identical to the social planner's value function. In summary the individual value functions are

$$Q_i\left(x,\theta\right) = \begin{cases} Q\left(x,\theta\right) & \theta_i = 1 \\ \frac{\partial Q\left(x,\theta\right)}{\partial\theta_i} & 1 > \theta_i > 0 \\ \underline{q}_i & \theta_i = 0. \end{cases} \tag{27}$$

We now write the first-order conditions of the social planning problem in terms of the marginal value of future utility. Let

$$\mathcal{M}_i\left(x,y,q_i\right) = \frac{\beta\pi_i\left(x,y\right)\exp\left(\sigma_i q_{iy}\right)}{\sum_{z\in\mathcal{X}}\pi\left(x,z\right)\exp\left(\sigma_i q_{iz}\right)}. \tag{28}$$

be the marginal value of utility next period if this period's state is x, next period's state is y, and the vector q_i gives the anticipated utility for all possible values of next period's state. This marginal value depends upon the anticipated utility in all states that could have occurred next period, not only the anticipated utility when y is the state.

When θ is an interior Pareto weight vector first-order conditions for the dynamic programming problem include for all agents i and j

$$\theta_i u_i'\left(c_i\right) = \theta_j u_j'\left(c_j\right) \tag{29a}$$

---

[8] If it wasn't zero then $Q\left(x,\theta\right)$ could not be the social planner's value function. For example, if the sum were positive then there would be some Pareto weight vector ψ such that $\sum_{j=1}^{n}\theta_j Q_j(x,\psi) > \sum_{j=1}^{n}\theta_j Q_j\left(x,\theta\right).$

for current period consumption, and for each possible value, y, of next period's state

$$\left(\frac{\phi_{iy}}{\phi_{jy}}\right)\mathcal{M}_j\left(x,y,q_j\right) = \left(\frac{\theta_i}{\theta_j}\right)\mathcal{M}_i\left(x,y,q_i\right),\tag{29b}$$

$$q_{iy} = Q_i\left(y,\phi_y\right).\tag{29c}$$

For each y, $\phi_y = \{\phi_i\}_{i=1}^n$ is normalized so that its elements sum to one. ϕ_{iy} is next period's Pareto weight on agent i when next period's state is y. The social planner's value function is embedded in equation 29b since Q_i is defined in terms of the social planner's value function.

The first-order conditions require that the ratio of the marginal values of consumption this period (in equation 29a) and next period (in equation 29b) be equal across agents. To find the optimal allocations for consumption this period one only has to compute the solution to the first-order equation 29a and the resource constraint 24b. This equivalent to solving the one-period social planning problem:

$$\max_{\{c_i\geq 0\}_{i=1}^n}\sum_{i=1}^n\theta_i u_i(c_i)$$

subject to

$$\sum_{i=1}^n c_i = c(x).$$

The solution to the one-period problem entails that the optimal consumption allocations are given by a function of x and θ. Denote this function as $c_i\left(e,\theta\right)$ where $e = c(x)$. For example if all agents have logarithmic functions of the form $u_i(c) = \log(c)$ then the one-period consumption allocation rules are

$$c_i\left(e,\theta\right) = \theta_i e.\tag{30}$$

The one-period consumption allocation rules are often called risk-sharing rules.

Finally, we note that it is possible to extend the proofs in Lucas and Stokey, given our assumptions in Section two, to guarantee that the social planner's value function is strictly convex in the Pareto weight vector. The optimal choice of next period's Pareto vector, for a given value of this period's exogenous state, will be continuous in this period's Pareto weight vector. It also can be shown that there is a unique solution to the social planner's first-order conditions, equations 29a thru 29c.

4   Steady states

The next three sections characterize the dynamics of allocations by studying the dynamics of Pareto weights. Pareto weights are convenient for studying dynamics since they summarize the effects of

the history of states on current allocations since current allocations are completely determined by the current Pareto weights and the current exogenous state. We will show that in many economies Pareto weights converge over time to steady states in which they are constant over time.

Let x be the current state, θ the current Pareto weight vector, and y next period's state. The solution to the social planner's dynamic programming problem implies that the optimal choice of next period's Pareto weight vector can be written as $\phi_y(x, \theta)$. Given a realization for the state $\{x_t\}_{t=0}^{\infty}$ for all time periods and an initial Pareto weight vector, θ_0, the solution to the social planner's problem implies that the Pareto weight vector evolves as

$$\theta_t = \phi_{x_t}(x_{t-1}, \theta_{t-1}).$$

The next three sections study how they evolve. Given the Pareto weights at time t, current consumption allocations can be computed by solving the one-period social planning problem described in Section 3.4. Current period consumption allocations can be written as $c_{it} = c_i\left[c(x_t), \theta_t\right].$

In order to characterize dynamics it is convenient to make additional assumptions:

Assumption 6 *For all $x, y \in \mathcal{X}$; and all agents i and j, $\pi_i(x, y) = \pi_j(x, y) = \pi(x, y)$.*

Assumption 7 *$\mathcal{X}$ is an ergodic set.*

Assumption 8 *For all $x, y \in \mathcal{X}$, $\pi(x, y) > 0$.*

Assumption 9 *There exists states $x \in \mathcal{X}$ and $y \in \mathcal{X}$ such that $c(x) \neq c(y)$.*

Assumption six guarantees that all agents have the same, correct, beliefs and implies Assumption three. Assumption seven entails that there is no proper subset of $\mathcal{X}$ which is ergodic. It also entails that all states in $\mathcal{X}$ will be realized with strictly positive probability sometime in the future. Assumption eight implies Assumption seven and guarantees that all values of the exogenous state vector can occur at every date $t \geq 1$ with strictly positive probability regardless of the history of the exogenous state vector. Assumption nine guarantees that there is uncertainty. For any value of the state this period, at least two distinct values of the aggregate endowment can occur next period with strictly positive probability. Many of our results will not require Assumptions seven thru nine.

4.1 Definition and characterization of steady states

This section will provide necessary and sufficient conditions for the existence of steady states. Several examples are provided in Section 4.4.

Although ultimately we are interested in how Pareto weights evolve, it is convenient to first characterize steady states in which Pareto weights are constant.

Definition 4 *A steady state for all agents is a Pareto weight vector $\bar{\theta}$ such that if $\theta_t = \bar{\theta}$ at some date t then it is Pareto optimal to set $\theta_s = \bar{\theta}$ for all dates $s > t$, in all future states of the world.*

Even though the world is uncertain and the exogenous state may be always varying over time, at a steady state the Pareto weights are constant. If the state space $\mathcal{X}$ was not an ergodic set then whether or not $\bar{\theta}$ is a steady state can depend upon the current value of the exogenous state.

Recall that the first-order conditions for the Pareto optimal problem described in Section three require the ratio of the marginal contributions of future consumption to current utility for two agents to be constant across time:

$$\frac{\mathcal{M}_{it}\left(x^t\right)u_i'\left[c_{it}\left(x^t\right)\right]}{\mathcal{M}_{jt}\left(x^t\right)u_j'\left[c_{jt}\left(x^t\right)\right]} = \frac{\theta_j}{\theta_i}.$$

The left hand side can be broken into two components. The first component is the ratio of the marginal contributions of future utility to current utility $\mathcal{M}_{it}/\mathcal{M}_{jt}$. The second component is the ratio of the marginal contributions of future consumption to future utility, $u'(c_{it})/u'(c_{jt})$. At an optimal allocation the ratio of the products of these components must be constant across time, but the ratio of the individual components can differ across time. The difference in the ratios of the components is what drives the evolution of allocations. When the ratio of the components, in addition to the ratio of the product of the components, are equal across time (for all pairs of agents in all possible states of the world) the economy is in a steady state. In particular, if the ratio $\mathcal{M}_{it}/\mathcal{M}_{jt}$ equals one in all future dates and states then i and j agents will be in a steady state in which

$$\frac{u_i'\left[c_{it}\left(x^t\right)\right]}{u_j'\left[c_{jt}\left(x^t\right)\right]} = \frac{\theta_j}{\theta_i}$$

is satisfied at all dates and states.[9] In other words at a steady state agents have the same marginal value of future utility.

To formalize this notion and provide conditions which guarantee that a Pareto weight vector is a steady state, we define the following two weighted differences

$$\Delta_{ij}\left(y,\theta\right) = \sigma_i Q_i\left(y,\theta\right) - \sigma_j Q_j\left(y,\theta\right), \tag{31}$$

$$\delta_{ij}\left(y,\theta\right) = \sigma_i u_i\left(c_i\left[c(y),\theta\right]\right) - \sigma_j u_j\left(c_j\left[c(y),\theta\right]\right) \tag{32}$$

---

[9]The ratio can not equal a constant other than one in a steady state since at time zero $\mathcal{M}_{i0} = \mathcal{M}_{j0} = 1$.

of utility allocations. The weights are the individual's risk-sensitivity parameters. Δ_{ij} is the weighted difference of agents i and j's lifetime utility allocation. δ_{ij} is the weighted difference of agents i and j's one-period utility allocation.

Also define

$$\Omega_{ij}(x,\theta) = \{d \in \mathcal{R} : d = \Delta_{ij}(y,\theta) \text{ for } y \in \mathcal{X} \text{ and } \pi(x,y) > 0\}, \tag{33}$$

$$\omega_{ij}(x,\theta) = \{d \in \mathcal{R} : d = \delta_{ij}(y,\theta) \text{ for } y \in \mathcal{X} \text{ and } \pi(x,y) > 0\}. \tag{34}$$

The set $\Omega_{ij}(x,\theta)$ is the smallest set which contains all possible values of the weighted differences of lifetime utility allocations for agents i and j from next period on, when this period's states is x and the Pareto weight vector *next* period is θ. Likewise $\omega_{ij}(x,\theta)$ is the smallest set which contains all possible values of the weighted differences of next period utility allocations for agents i and j when this period's states is x and next period's Pareto weight vector is θ. In defining Ω and ω we are fixing the value of next period's Pareto weight vector.

We now provide conditions for steady states. The conditions are stated in terms of weighted differences of individual value functions in Theorem one and in terms of weighted differences of reward functions in Corollary one.

Theorem 1 *Let $\bar{\theta}$ be a Pareto weight vector.*

1. *Let Assumptions one thru six hold. If for each pair of agents i and j who have strictly positive Pareto weights, $\Omega_{ij}(x,\bar{\theta})$ is a singleton for all states x then $\bar{\theta}$ is a steady state.*

2. *Let Assumptions one thru seven hold. If $\bar{\theta}$ is a steady state then for each pair of agents i and j who have strictly positive Pareto weights, $\Omega_{ij}(x,\bar{\theta})$ is a singleton for all states x.*

Corollary 1 *Let $\bar{\theta}$ be a Pareto weight vector.*

1. *Let Assumptions one thru six hold. If for each pair of agents i and j who have strictly positive Pareto weights, $\omega_{ij}(x,\bar{\theta})$ is a singleton for all states x then $\bar{\theta}$ is a steady state.*

2. *Let Assumptions one thru eight hold. If $\bar{\theta}$ is a steady state then for each pair of agents i and j who have strictly positive Pareto weights, $\omega_{ij}(x,\bar{\theta})$ is a singleton for all states x.*

See the appendix for proofs. The proofs show the equivalence between these conditions and a situation in which the ratio of the marginal contributions of future utility to current utility is constant across agents. If it was assumed that the state space was not ergodic then Ω and ω would only have to be singletons for states that can occur at future dates.

20

Theorem one and Corollary one are easier to think about if Assumption eight happens to hold. If this assumption holds then $\bar{\theta}$ is a steady state if and only if the weighted individual value functions are equal up to a constant for all values of the exogenous state. In other words, for each pair of agents i and j who have a strictly positive weight

$$\sigma_i Q_i\left(x,\bar{\theta}\right) - \sigma_j Q_j\left(x,\bar{\theta}\right)$$

does not vary with the state x, though it may vary with i and j. In terms of reward functions, $\bar{\theta}$ is a steady state if and only if the weighted individual reward functions are equal to a constant for all states x. In other words for each pair of agents i and j who have a strictly positive weight

$$\sigma_i u_i\left(c_i\left[c\left(x\right),\bar{\theta}\right]\right) - \sigma_j u_j\left(c_j\left[c\left(x\right),\bar{\theta}\right]\right) \tag{35}$$

does not vary with the state x, though it may vary with i and j.

When Assumption eight does not hold then the theorem tells us that it is possible to have steady states in which the weighted differences equal different values at different dates. However, this will only happen in degenerate economies in which the state moves around over time in deterministic patterns and there is very little uncertainty. See Section 4.4.5 for an example.

The corollary is easy to apply in practice and can be conveniently used to classify steady states even when an analytical solution to the social planner's dynamic programming problem is not available since the corollary requires only that the one-period optimal allocations are known. These are easy to compute (see Section 3.4 and the examples in Section 4.4.)

There also are steady states in which some agents have zero Pareto weights. The conditions in Theorem one and Corollary one only have to hold for agents who have strictly positive Pareto weights. Agents who have zero Pareto weights today will always have zero Pareto weights in every future date and state.

4.2   Regular steady states

Regular steady states are a refinement of steady states. Regular steady states are a stronger notion of invariance that is useful for analyzing dynamics when there are more than two agents. Regular steady states are defined independently of the state space whereas steady states depend upon the exact specification of the exogenous environment. All regular steady states are steady states, but there can exist steady states which are not regular.

Definition 5 *A regular steady state for all agents is a Pareto weight vector, $\bar{\theta}$, which is a steady state regardless of the specification of the exogenous state space $\mathcal{X}$, aggregate consumption $c(x)$, and transition probabilities $\pi(x, y)$.*

Regular steady states have the appealing property that the set of Pareto weight vectors which are regular steady states is invariant to the specification of the state space, endowments, and probabilities. A steady state which is not regular may cease to be a steady state when the state space is expanded or when additional uncertainty is introduced. When the state space is contracted or when uncertainty is removed more Pareto weights can be steady states.

The following theorem shows that regular steady states depend only on the reward functions and the risk-sensitivity parameters.

Theorem 2 *Let Assumptions one thru six hold. Let $\bar{\theta}$ be a Pareto weight vector. $\bar{\theta}$ is a regular steady state for all agents if and only if for each pair of agents i and j who have strictly positive weights*

$$\sigma_i u_i \left(c_i \left(e, \bar{\theta} \right) \right) - \sigma_j u_j \left(c_j \left(e, \bar{\theta} \right) \right) \tag{36}$$

does not vary with e for all positive real numbers e, though it may vary with i and j.

We have seen that at a steady state it always is the case that equation 36 does not vary with the *possible* values of aggregate consumption at a steady state for all agents. At a regular steady state it will not vary with aggregate consumption even if we consider values of aggregate consumption which can not occur at a steady state for all agents. Theorem two also tells us that all regular steady states are steady states since if equation 36 does not vary with e, $\omega_{ij}(x, \bar{\theta})$ must be a singleton for all x.

The notion of a regular steady state will be convenient for analyzing dynamics when there are more than two (types of) agents and the one-period consumption allocation rules are nonlinear. Section 5.1 will describe why regular states are useful when there are more than two (types) of agents. We conclude this section by showing that if there is uncertainty (so that Assumption nine is satisfied) and the one-period consumption allocation rules are linear then all steady states are regular. The one-period consumption allocation rules are linear if they can be written as

$$c_i(e, \theta) = s_i(\theta)e \tag{37}$$

where s_i is a function which can depend upon the Pareto weights, but not the aggregate endowment.

Corollary 2 *Let Assumptions one thru six and Assumptions eight thru nine hold. Then if the one-period consumption allocations rules are linear in aggregate consumption, all steady states are regular.*

22

See Section 4.4.5 for an example of a steady state which is not regular.

4.3 Steady states for a subset of agents

There can exists situations in which the *ratio* of Pareto weights for two agents is time-invariant even though the Pareto weights for these agents are time-varying. These agents can be viewed as being in their own private steady state. The marginal values of future utility for these two agents will be equal over time, as will the ratios of the marginal contributions of future consumption to future utility for these agents.

Definition 6 *A steady state Pareto ratio for agents i and j is a non-negative number $\bar{\lambda}_{ij}$ such that if ratio of Pareto weights for agents i and j is equal to $\bar{\lambda}_{ij}$*

$$\frac{\theta_{it}}{\theta_{jt}} = \bar{\lambda}_{ij} \tag{38}$$

at some date t then it is Pareto optimal to set the ratio of Pareto weights for these agents equal to $\bar{\lambda}_{ij}$

$$\frac{\theta_{is}}{\theta_{js}} = \bar{\lambda}_{ij} \tag{39}$$

for all dates $s > t$, in all future states of the world.

At a private steady state for agents i and j the ratio of the Pareto weights for i and j agents is constant over time even though their Pareto weights may be varying over time. The weights may be varying over time since i and j agents may not be in a steady state with other agents. If they are in a steady state with all agents then the Pareto weights for agents i and j will also be constant over time.

We provide two results on steady states for a subset of agents. Theorem three provides sufficient conditions which guarantee that a Pareto ratio is not a steady state Pareto ratio. Theorem four provides sufficient conditions which guarantee that a Pareto ratio is a steady state.

Theorem 3 *Let Assumptions one thru eight hold. Let the one-period consumption allocation rules be linear in the aggregate endowment. If for each Pareto weight vector θ such that*

$$\frac{\theta_i}{\theta_j} = \bar{\lambda}_{ij} \tag{40}$$

there exists $x \in \mathcal{X}$ and $y \in \mathcal{X}$ such that

$$\sigma_i u_i \left(c_i \left[c(x), \theta \right] \right) - \sigma_j u_j \left(c_j \left[c(x), \theta \right] \right) \neq \sigma_i u_i \left(c_i \left[c(y), \theta \right] \right) - \sigma_j u_j \left(c_j \left[c(y), \theta \right] \right) \tag{41}$$

then $\bar{\lambda}_{ij}$ is not a steady state Pareto ratio for agents i and j.

Theorem 4 *Let Assumptions one thru six hold. If $\bar{\theta}$ is a regular steady state then for each pair of agents i and j*

$$\bar{\lambda}_{ij} = \frac{\bar{\theta}_i}{\bar{\theta}_j} \tag{42}$$

is a steady state Pareto ratio for agents i and j.

See the appendix for proofs.

Theorem four tells that a regular steady state, $\bar{\theta}$, has the appealing property that any two agents whose have the Pareto ratio, $\bar{\theta}_i/\bar{\theta}_j$ will be in a private steady state regardless of the Pareto weights of other agents.

4.4 Examples

This section concludes by looking at steady states in several example economies. The examples will require, unless stated otherwise, Assumptions one thru nine. In examples that discuss specific reward functions, the properties listed in Assumption four will be satisfied. Additional examples are discussed in Sections eight and nine.

4.4.1 Time-additive preferences

Let all agents have time-additive preferences. For any Pareto weight vector θ, $\delta_{ij}(y, \theta)$ is equal to zero for all states x since $\sigma_i = \sigma_j = 0$. Hence by Corollary one all Pareto weight vectors are steady states. Since equation 36 is always satisfied all of the steady states are regular. We already knew with time-additive preferences all Pareto weight vectors are steady states from much simpler arguments, but it is interesting to see how the machinery developed in this paper can show the same result.

4.4.2 Identical preferences

If all agents have identical preferences ($u_i = u_j$ and $\sigma_i = \sigma_j$ for all i and j) then the Pareto weight vector $\bar{\theta}$ which assigns equal Pareto weights to all individuals is a regular steady state. To see this, note that if $\bar{\theta}_i = 1/n$ for all i then agents have identical consumption and

$$\sigma_i u_i \left(c_i \left(e, \bar{\theta} \right) \right) = \sigma_j u_j \left(c_j \left(e, \bar{\theta} \right) \right) \tag{43}$$

for all $e > 0$. Hence $\bar{\theta}$ is a regular steady state. So if all agents have identical preferences there always exists at least one interior steady state in which agents have identical Pareto weights, identical consumption, and identical utility.

4.4.3 Log risk-sensitivity

Consider an economy in which all agents have logarithmic reward functions of the form $u_i(c) = \log(c)$. Preferences of this form have been previously studied by Tallarini (2000) in representative agent models. The optimal allocation rules for current period consumption are

$$c_i(e, \theta) = \theta_i e$$

when the current period aggregate consumption is e and the current period Pareto weight vector is θ. In this example

$$\sigma_i u_i \left(c_i \left[e, \theta \right] \right) - \sigma_j u_j \left(c_j \left[e, \theta \right] \right) = \sigma_i \log(\theta_i e) - \sigma_j \log(\theta_j e)$$
$$= \sigma_i \log(\theta_i) - \sigma_j \log(\theta_j) + (\sigma_i - \sigma_j) \log(e).$$

This varies with e if and only if $(\sigma_i - \sigma_j) \log(e)$ varies with e. Hence by Theorem two when all agents have the same risk-sensitivity parameter every Pareto weight vector is a regular steady state since equation 36 can not vary with e.

When agents have heterogeneous risk-sensitivity parameters there are no interior regular steady states since $(\sigma_i - \sigma_j) \log(e)$ varies with e for any two agents with different risk-sensitivity parameters. In virtue of Assumptions eight and nine there also are no interior steady states. Any two agents with the same risk-sensitivity parameter will always be in a private steady state in which their Pareto ratios are constant.

4.4.4 Steady states where a single agent has a weight of one

Regardless of agents' reward functions and risk-sensitivity parameters, there always exists a steady state in which one agent has a Pareto weight of one and all other agents have a Pareto weight of zero. ω is a singleton for all pairs of agents who have strictly positive weights since there are not two distinct agents with strictly positive weights and ω_{ii} is always a singleton.

4.4.5　Steady states which are not regular

Examples of non-regular steady states are abundant in a deterministic economy which violates Assumptions eight or nine. Consider an economy in which there is only one possible value of the aggregate endowment at each date. Let there be two agents with logarithmic rewards. Let the risk-sensitivity parameter for the first agent be -1 and for the second agent be -2. For all x in $\mathcal{X}$, let there exist a y such that $\pi(x, y) = 1$. Then since ω will always be a singleton, all Pareto weight vectors are steady states. Most of these steady states will not be regular since if uncertainty is introduced there will not be any interior steady states. (See Section 4.4.3.)

4.5　A representative agent

At a steady state consumption is shared just as in a time-additive economy since the Pareto weights are constant over time. Moreover at a regular state equation 36 tell us there is a tight link between the preferences of agents. This opens up the possibility of the existence of a useful representative agent. In this section we explore this idea and show that at an interior regular steady state there always exists a representative agent. The representative agent is unique but can be described using the preferences of any of the agents. For any agent, a representative agent can be constructed with a similar reward function, the same risk-sensitivity parameter, and the same discount factor.

For there to exist an interior regular steady state either all agents have time-additive preferences or all agents have risk-sensitive preferences with strictly negative risk-sensitivity parameters. If they all have time-additive preferences then standard results are available for when a representative agent exists. In this section we focus on economies in which all agents have risk-sensitive preferences.

The following theorem is the main result of this section and shows that there exists a representative agent with preferences over aggregate consumption $e = c(x)$.

Theorem 5 *Let Assumptions one thru six hold. Assume all agents have risk-sensitive preferences with strictly negative risk-sensitivity parameters. Let $\bar{\theta}$ be an interior regular steady state. For any agent i a representative agent can be defined with the reward function*

$$f_i(e) = u_i\left(c_i\left(e, \bar{\theta}\right)\right),$$

the risk-sensitivity parameter σ_i, and the discount factor β. The representative is unique but can be described using the preferences and risk-sensitivity parameters of any of the agents.

See the appendix for a proof. It is possible to use the preferences of any agent to describe the representative agent because at a regular steady state we know from Theorem two that the agent's

reward functions are equal up to a constant. Hence all agents can be viewed as having the same preferences over the aggregate endowment at an interior regular steady state. The requirement that agents be strictly risk-sensitive is important. If all agents have time-additive preferences then reward functions typically will not be equal up unto a constant at a steady state. Away from steady states it is unlikely that there will exist a representative agent with preferences similar to any particular agent.

5   Off steady state dynamics

What do optimal allocations look like away from a steady state? Unfortunately for most specifications of preferences analytic solutions are not available away from steady states. We instead have to rely on qualitative theorems and numerical solutions. This sections presents two qualitative theorems and a numerical illustration of risk-sharing in the short run. Section 5.1 shows that Pareto weights can never cross a steady state Pareto ratio. Section 5.2 qualitatively and numerically studies dynamic risk-sharing in the short run. Subsequent sections will build on the qualitative theorems in this section and provide additional illustrative numerical solutions.

5.1   The no-crossing property

Steady states are especially interesting if they are stable. Section six will study stability. This section shows that steady states are interesting even when they are not stable because they provide bounds on the evolution of ratios of Pareto weights. This result is interesting in itself and will be a crucial ingredient of our stability proofs in Section six.

Theorem 6 (No-crossing property) *Let Assumptions one thru six hold. Let θ_t be the Pareto weight vector at time t and let $\bar{\theta}$ be a regular steady state. If*

$$\frac{\theta_{it}}{\theta_{jt}} > \frac{\bar{\theta}_i}{\bar{\theta}_j} \tag{44}$$

then

$$\frac{\theta_{is}}{\theta_{js}} > \frac{\bar{\theta}_i}{\bar{\theta}_j} \tag{45}$$

for all $s > t$.

See the appendix for a proof.

This theorem tells us that the Pareto ratios can never cross a regular steady state Pareto ratio. If the Pareto ratio, for two agents, is initially greater than a regular steady state ratio then it will always be greater than that regular steady state ratio in the future. By inversion, it follows that if a

Pareto ratio, for two agents, is initially less than a regular steady state then it will always be less than that regular state in the future. Thus regular steady states provide bounds on the evolution of Pareto weights. If the economy consists of only two agents (or two types of agents) then this result can be strengthened to show that it is never possible to cross a steady state Pareto ratio even if the Pareto ratio does not correspond to a regular steady state.[10]

Consider an example in which all agents have the same preferences. We know from previous sections there is a regular steady state in which all agents have identical weights. Assume agent one is initially richer than the subset ξ_p of other agents, initially poorer than the subset ξ_r of other agents, and initially equal to the subset ξ_e of other agents. Here "richer" means has a larger Pareto weight, "poorer" means has a smaller Pareto weight, and "equal" means has the same Pareto weight. The no-crossing property tells us that agents in set ξ_p will never be richer than agent one and agents in set ξ_r will never be poorer than agent one. Since in this example equality is a steady state Pareto ratio between any two agents, the agents in ξ_e will always be equal to agent one. So, in this sense the rich stay rich and the poor stay poor. There is no possibility for an agent who is initially poor to become rich. It is possible that agents in ξ_p will become essentially equal to the agents in ξ_r and equal to agent one. (Exact equality can only happen in the limit as time goes to infinity. At any finite date in the future agent one will be strictly richer than all agents in set ξ_p and strictly poorer than all agents in set ξ_r)

When there are many regular steady states the no-crossing property tells us that all of the regular steady states provide bounds. When there are no interior steady states then it is possible for an agent who is initially very poor to become very rich, and vice-versa. Consider the example in Section 4.4.3 in which agents have logarithmic rewards but *different* risk-sensitivity parameters. In that example there are no interior steady states. It is possible that eventually any agent, can be very poor or very rich regardless of their initial wealth (as long as their initial Pareto weight is not zero or one.)

5.2   Dynamic risk-sharing in the short run

How is risk shared between risk-sensitive agents? At any date agents can be viewed as sharing two things: current consumption and future utility. Given the current Pareto weights current consumption is shared subject to a linear resource constraint, equation 24b. Future utility is shared subject to the nonlinear feasibility constraint, equation 24c. When some agents have risk-sensitive preferences optimal sharing of future utility often entails that Pareto weights move over time to provide some agents with smoother allocations than other agents.

---

[10]The proof of this is very similar to the proof of Theorem six. See Anderson (1998) for the details.

In many cases risk-sharing is determined by how the weighted differences of rewards [as measured by $\delta_{ij}(y, \theta)$ which was defined in equation 32] varies with the aggregate state and the current Pareto weights. To show this we let the state be equal to aggregate consumption each period, assume that the state is i.i.d. over time, and limit the analysis to two agents. Our main result is stated in Theorem seven.

Assumption 10 *Assume the exogenous state x_t at any date t is a scalar and is equal to aggregate consumption: $x_t = c(x_t)$. For any $a, b, y \in \mathcal{X}$ it is the case that*

$$\pi(a, y) = \pi(b, y). \tag{46}$$

Theorem 7 *Let there be two agents who are denoted by $i = 1, 2$. Let Assumptions one thru eight and Assumption 10 hold. Let Φ be set of the possible values for next period's Pareto weight vector.*

1. *If for any given $\phi \in \Phi$, $\delta_{21}(y, \phi)$ is a strictly increasing function of y then the optimal choice of next period's Pareto weight for agent two is a strictly increasing function of next period's state y.*

2. *If for any given $\phi \in \Phi$, $\delta_{21}(y, \phi)$ is a strictly decreasing function of y then the optimal choice of next period's Pareto weight for agent two is a strictly decreasing function of next period's state y.*

See the appendix for a proof. The result also applies if there are equal numbers of two types of agents. Because of the no-crossing property, the set Φ in Theorem seven need not be the set of all Pareto weight vectors since, in some economies, many Pareto weight vectors for next period never will be selected.

Theorem seven can readily be applied to an example in which all agents have logarithmic utility. Consider the environment described in Section 4.4.3 when there are only two agents and the assumptions required for Theorem seven hold. For this environment

$$\delta_{21}(y, \phi) = \sigma_2 \log(\phi_2) - \sigma_1 \log(\phi_1) + (\sigma_2 - \sigma_1) \log(y)$$

is strictly increasing in y when $\sigma_2 > \sigma_1$ and strictly decreasing in y when $\sigma_2 < \sigma_1$. Hence by Theorem seven the optimal choice of next period's Pareto weight for agent two is a strictly increasing function of next period's state y when $\sigma_2 > \sigma_1$ and a strictly decreasing function of y when $\sigma_2 < \sigma_1$. When $\sigma_2 = \sigma_1$, δ_{21} is constant and, as we have previously shown, agent one and agent two will be in a steady state. If agent two's risk-sensitivity parameter is lower than agent one's, $\sigma_2 < \sigma_1$, then when the aggregate endowment is low, agent two's Pareto weight will be relatively large, and when the aggregate endowment is high, agent two's Pareto weight will be relatively small.

State	π^A	π^B
0.10	0	1×10^{-11}
0.30	0	1×10^{-9}
0.50	0	1×10^{-7}
0.70	0	1×10^{-5}
0.90	0	1×10^{-3}
0.95	0.01	0.01
0.96	0.02	0.02
0.97	0.04	0.04
0.98	0.06	0.06
0.99	0.20	0.20
1.00	0.34	0.33797979798
1.01	0.20	0.20
1.02	0.06	0.06
1.03	0.04	0.04
1.04	0.02	0.02
1.05	0.01	0.01
1.10	0	1×10^{-3}
1.30	0	1×10^{-5}
1.50	0	1×10^{-7}
1.70	0	1×10^{-9}
1.90	0	1×10^{-11}

Table 1: This table lists the values of the state and their corresponding probabilities in the numerical examples. Two different specifications for the probabilities are considered which we label specification A and specification B.

We now describe some numerical examples. In the examples the state is i.i.d. over time and equal to the aggregate endowment. We will consider two different specifications for the probabilities of states, labeled A and B, which are listed in Table 1. In setting A there are 11 possible values, between 0.95 and 1.05, of the state. In setting B there are 21 possible values of the state. Setting B assigns small probabilities to extreme values of the state. In setting B there is about a 0.101% probability that the state will be less than 0.95 and about a 0.101% probability that the state will be greater than 1.05. The probabilities of states near one were designed to *roughly* be in the same ballpark as the probabilities for detrended aggregate real annual consumption per-person in the United States between 1950 and 2002 (discretized and scaled so that its mean is one), assuming that aggregate consumption is i.i.d. Of course aggregate consumption is not i.i.d. and this specification may overstate the one-period ahead uncertainty in aggregate consumption and understate the many-period ahead uncertainty in aggregate consumption. The probabilities for extreme events were arbitrarily chosen. In all of our numerical examples β is set at 0.95 so that the natural interpretation of the length of a period is one year.

Figure one graphs the decision rules for next period's Pareto weight on agent two when there are two agents who have logarithmic reward functions and heterogeneous risk-sensitivity parameters. In
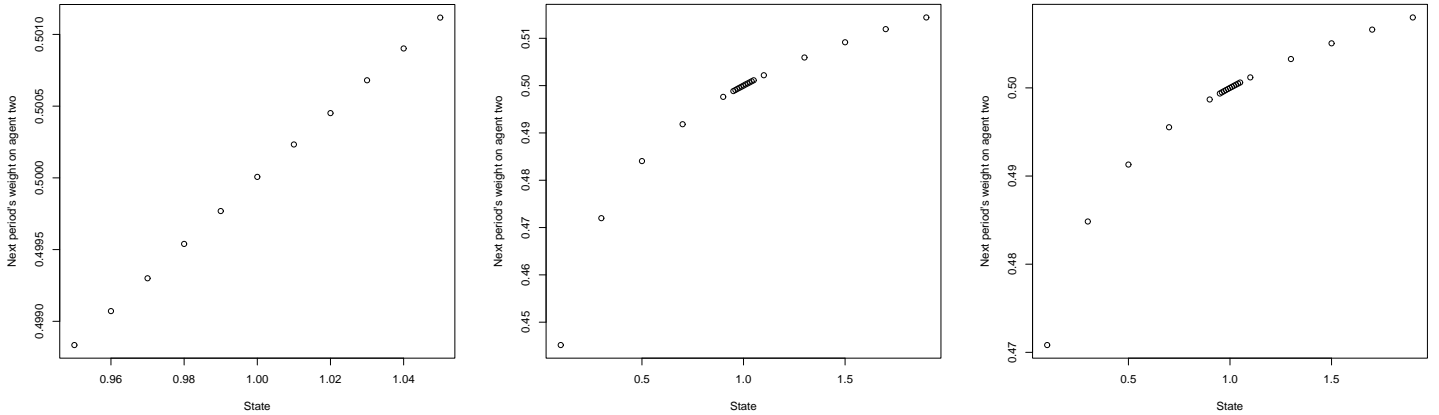
Figure 1: The decision rules for next period's Pareto weight on agent two when there are two agents who have logarithmic reward functions and heterogeneous risk-sensitivity parameters. In all of the figures agent two has time-additive preferences so that $\sigma_2 = 0$. The discount factor for both agents is $\beta = 0.95$. This period's Pareto weight on agent two is 0.50. In the two left most figures agent one have a risk-sensitivity parameter of $\sigma_1 = -1$. In the right most figure agent one has a risk-sensitivity parameter of -0.1. In the left most figure the probabilities are given by specification A and in the two right most figures the probabilities are given by specification B.

all of the graphs agent two has time-additive preferences and agent one is strictly risk-sensitive. The graphs vary the magnitude of agent one's risk-sensitivity parameter and use probability specifications A and B. The qualitative properties of the graphs are known from Theorem seven: agent two's Pareto weight is increasing in next period's state. We also see that the Pareto weights only move substantially when extreme values of the aggregate endowment occur. For moderate values of the endowment the Pareto weights do not move very much. In particular for specification A of the probabilities in the left most graph we see the maximum fall/rise in the Pareto weights over one period is about 0.0011. In the center graph, we see that for probability specification B the Pareto weights can move substantially but only when a catastrophic event or miraculous windfall occurs. The presence of the extreme possibilities does change the decision rules for highly probable realizations of the endowment. The movements, for highly probable realizations of the state, in the Pareto weights are still small but not as small as in the previous economy. This suggests that in order to get the decision rule for the Pareto weights correct for moderate values of the endowment it is important to have the correct specification of states and probabilities for extreme possibilities. In the right most figure we see that the weights move even less when agent one's risk-sensitivity parameter is closer to zero.

In Section nine we apply Theorem seven to an example in which all agents have power reward functions. In Section eight when risk-sharing between risk-sensitive and non-risk-sensitive agents is discussed, we will be able to show a stronger result which does not require Assumption 10 and allows the economy to have many agents.

6 Stability

This section describes the long-run behavior of Pareto weights and characterizes the stability of steady states. The proofs of stability presented in this section rely on Doob's Martingale Convergence Theorem. The probabilities under which we consider convergence are the true probabilities π. We begin be defining a stable Pareto weight vector, a martingale and a supermartingale. To make the paper self-contained we then state Doob's Theorem, and a related result.

Definition 7 *A Pareto weight vector, θ^*, is stable if $\{\theta_t\}_{t=0}^{\infty}$ converges with probability one to θ^* starting from any initial interior Pareto weight vector θ_0 at time zero.*

Definition 8 *The stochastic process $\{\lambda_t\}_{t=0}^{\infty}$ is a* supermartingale *(resp.* martingale*) if $E_t\lambda_{t+1} \leq \lambda_t$ (resp. $E_t\lambda_{t+1} = \lambda_t$) for all t.*

Theorem 8 (Martingale Convergence Theorem (Doob (1953))) *If the stochastic process $\{\lambda_t\}_{t=0}^{\infty}$ is a supermartingale and $\lambda_t \geq b$ for all t where b is finite then λ_t converges with probability one to a finite random variable.*

Theorem 9 *If for a stochastic process $\{\lambda_t\}_{t=0}^{\infty}$*

1. *There is a probability $\underline{\pi}$ such that for any $\lambda_t > 0$ there exists a $d > 0$ such that*

$$prob_t\left(|\lambda_{t+1} - \lambda_t| \geq d\right) \geq \underline{\pi} > 0. \tag{47}$$

   For any constants $\underline{\lambda}$ and $\overline{\lambda}$, if λ_t satisfies $+\infty > \overline{\lambda} \geq \lambda_t \geq \underline{\lambda} > 0$ the constant d can be chosen independent of λ_t.

2. *$E_t\lambda_{t+1} \leq \lambda_t$ and $\lambda_t \geq 0$.*

then $\{\lambda_t\}_{t=0}^{\infty}$ converges with probability one to zero.

See the appendix for a proof of Theorem nine. Doob's Martingale Convergence Theorem states that supermartingales (and martingales) bounded from below converge to a random variable. Theorem nine strengthens the Martingale Convergence Theorem and provides sufficient conditions for a martingale to converge to zero with probability one. Our proofs of stability will apply the Martingale Convergence Theorem and Theorem nine to the first-order conditions of the social planner's dynamic programming problem.

The following result allows us to apply the theorems.

Lemma 1 *The expected value of next period's Pareto ratio is*

$$E_t \left[\frac{\theta_{it+1}}{\theta_{jt+1}} \right] = \frac{\theta_{it}}{\theta_{jt}} - \frac{\text{cov}_t \left(e^{\sigma_j Q_j (x_{t+1}, \theta_{t+1})}, \frac{\theta_{it+1}}{\theta_{jt+1}} \right)}{E_t \left[e^{\sigma_j Q_j (x_{t+1}, \theta_{t+1})} \right]}. \tag{48}$$

The proof of Lemma one is in the appendix and involves taking expected values of the first-order conditions of the social planning problem and rearranging terms. The proofs of stability which follow will work by generating conditions under which the covariance term in equation (48) can be shown to be greater than or equal to zero. When it can the Pareto ratios are a supermartingale. Sometimes it is very easy to sign the covariance term. For example if $\sigma_j = 0$ then the covariance term is always zero. Section eight uses this insight to analyze stability when there are economies with some agents who have risk-sensitive preferences and some who have time-additive preferences. We now show that even when all agents are strictly risk-sensitive the covariance term can sometimes be signed.

The following theorem gives sufficient conditions for a steady state Pareto weight vector to be stable or not stable in terms of δ_{21} which was defined in equation 34. The proof of the theorem signs the covariance term in Lemma one and verifies that the hypotheses of Theorem nine are satisfied. The conditions are primarily useful when all agents have risk-sensitive preferences. They also apply when some agents have time-additive preferences but in this case Section eight will show that stronger results are available. The conditions of the theorem require the state to be i.i.d. over time and limit the analysis to two agents.

Theorem 10 *Let there be two agents who are denoted by $i = 1, 2$. Let Assumptions one thru 10 hold. Let $\bar{\theta}$ be a regular steady state Pareto weight vector.*

1. *If*

   (a) *for any fixed Pareto weight vector ϕ such that $1 > \phi_2 > \bar{\theta}_2$ it is the case that $\delta_{21}(y, \phi)$ is a strictly decreasing function of y and*

   (b) *for any fixed Pareto weight vector ϕ such that $0 < \phi_2 < \bar{\theta}_2$ it is the case that $\delta_{21}(y, \phi)$ is a strictly increasing function of y*

   then $\bar{\theta}$ is stable.

2. *If*

   (a) *for any fixed Pareto weight vector ϕ such that $1 > \phi_2 > \bar{\theta}_2$ it is the case that $\delta_{21}(y, \phi)$ is a strictly increasing function of y and*

   (b) *for any fixed Pareto weight vector ϕ such that $0 < \phi_2 < \bar{\theta}_2$ it is the case that $\delta_{21}(y, \phi)$ is a strictly decreasing function of y*

then $\bar{\theta}$ is not stable.

See the appendix for a proof. The result also applies if there are equal numbers of two types of agents. Since there are only two agents (or two types of agents) it is possible to generalize the result so that it applies when $\bar{\theta}$ is any steady state Pareto weight vector, not necessarily a regular steady state (Anderson 1998). Theorem 10 can readily be applied to an example in which all agents have logarithmic utility. Consider the example described in Section 4.4.3 and Section 5.2 in which agents have logarithmic reward functions. Let there be two agents and assume agent two has a lower risk-sensitivity parameter. Recall that $\delta_{21}(y, \phi)$ is strictly decreasing in y when $\sigma_2 < \sigma_1$. Hence by Theorem 10 the regular state Pareto weight vector which sets $\bar{\theta}_1 = 1$ and $\bar{\theta}_2 = 0$ is stable since the condition in (1a) is satisfied. The condition in (1b) does not apply since ϕ_2 cannot be less than zero. The weight of the agent who has a lower risk-sensitivity parameter converges to zero with probability one.

There is a very close connection between risk-sharing in the short run and long run. The conditions in Theorem seven on δ_{21} to guarantee that the optimal choice of next period's Pareto weight is a strictly increasing (respectively decreasing) function of next period's state are identical to the conditions in Theorem 10 to guarantee that a Pareto weight vector is not stable (resp. stable). This suggests that agents face a long run penalty for smoothing future utility when the smoothness is achieved by varying Pareto weights.

Figure two graphs the evolution of the Pareto weight on agent two when there are two agents who have logarithmic reward functions. Agent two has time-additive preferences and agent one has risk-sensitive preferences with $\sigma_1 = -1$. The probabilities of the state are given by specification B. The graph is for one draw of values for the exogenous state.

In Section nine we apply Theorem 10 to an example in which all agents have power reward functions. In Section eight when risk-sharing between risk-sensitive and non risk-sensitive agents is discussed, we will be able to show a stronger result which does not require Assumption 10 and allows the economy to have many agents.

7 Distorted beliefs and equilibria

Prior to considering more examples we discuss two notions that will help us understand the mechanisms behind dynamic risk-sharing. Section 7.1 discusses an equivalent formulation of the social planning problem where agents have distorted beliefs and Section 7.2 decentralizes Pareto optimal allocations.
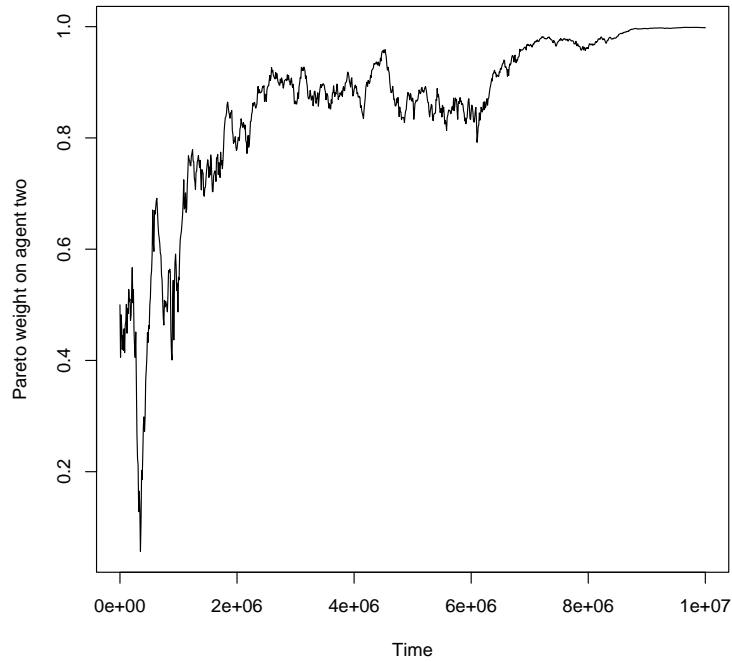
Figure 2: Evolution with logarithmic rewards

An example of the evolution of the Pareto weight on agent two when there are two agents who have logarithmic reward functions. The discount factor for both agents is $\beta = 0.95$. Agent two has time-additive preferences so that $\sigma_2 = 0$ and agent one has risk-sensitive preferences with $\sigma_1 = -1$. The probabilities of the state are given by specification B. The initial Pareto weight on agent two is 0.50. The Pareto weight is graphed for $1\mathrm{x}10^7$ time periods. (On the graph one point is plotted for every 10,000 periods so that 1000 points are plotted. This masks the local variability in the Pareto weights.)

7.1 Distorted beliefs

This section shows that the social planning problem described in previous sections is equivalent to a social planning problem in which agents have time-additive preferences with distorted beliefs.

Assume we have a risk-sensitive heterogeneous agent economy and its solution, as described in previous sections. We have individual value functions for all agents and the lifetime utility for agent i from time t onward is

$$U_{it}\left(x^t\right) = Q_i\left[x_t, \theta_t\left(x^t\right)\right].\tag{49}$$

Here $\theta_t\left(x^t\right)$ is the Pareto weight vector at time t and the history of states is x^t at time t.

Consider an economy in which agents have distorted beliefs and time-additive preferences. Let agent i believe that the time zero probability that the history x^t will be realized at time t is

$$\frac{\mathcal{M}_{it}\left(x^t\right)}{\beta^t}.\tag{50}$$

These are well formed beliefs since $\mathcal{M}_{it} \geq 0$ and for any t

$$\sum_{x^t}\left(\frac{\mathcal{M}_{it}\left(x^t\right)}{\beta^t}\right) = 1.\tag{51}$$

The summation over x^t indicates summation over all histories that can be realized at time t. We assume the agent takes these probabilities as being exogenous, even though they will depend upon the Pareto optimal allocation. Let agent i's lifetime utility function be

$$\sum_{t=0}^{\infty}\sum_{x^t}\beta^t\left(\frac{\mathcal{M}_{it}\left(x^t\right)}{\beta^t}\right)u_i\left[c_{it}\left(x^t\right)\right].\tag{52}$$

where $\mathcal{M}_{it}$ is treated as exogenous.

A Pareto optimal allocation can be computed by maximizing

$$\sum_{i=1}^{n}\sum_{t=0}^{\infty}\sum_{x^t}\theta_i\beta^t\left(\frac{\mathcal{M}_{it}\left(x^t\right)}{\beta^t}\right)u_i\left[c_{it}\left(x^t\right)\right]\tag{53}$$

by choice of a feasible consumption allocation where θ_i is the Pareto weight of agent i. Since $\mathcal{M}_{it}$ for all i are exogenous and taken as given the first-order conditions are

$$\theta_i\mathcal{M}_{it}\left(x^t\right)u_i'\left[c_{it}\left(x^t\right)\right] = \theta_j\mathcal{M}_{jt}\left(x^t\right)u_j'\left[c_{jt}\left(x^t\right)\right]\tag{54}$$

for all dates, all histories x^t, and all pairs of agents i and j. These are identical to the first-order conditions for the social planner's problem given in equation 17. So, a risk-sensitive heterogeneous agent problem is identical to a problem in which agents have heterogeneous beliefs.

Our results about steady states have a nice interpretation in this equivalent economy with distorted beliefs. When agents are at an interior steady state in a risk-sensitive economy then in the equivalent economy agents have the same beliefs. Later sections will interpret our stability results in the context of the equivalent economy.

7.2 Equilibrium

This section decentralizes a Pareto optimal allocation.[11]

Let $w_i \geq 0$ be the initial wealth for agent i. Let $p_t\left(x^t\right)$ be the time zero price of time t, history x^t consumption. The agent's problem is to maximize their lifetime utility function (which is given either by U_{i0} or equation 52) subject to the lifetime budget constraint:

$$\sum_{t=0}^{\infty} \sum_{x^t} p_t\left(x^t\right) c_{it}\left(x^t\right) = w_i \tag{55}$$

where the prices p_t are taken as given. First-order conditions for the agent's problem are

$$u_i'\left[c_{it}\left(x^t\right)\right] \mathcal{M}_{it}\left(x^t\right) = \mu_i p_t\left(x^t\right) \tag{56}$$

for all dates and states where μ_i is a Lagrange multiplier on constraint (55). If we define

$$\mu_{it}\left(x^t\right) = \frac{\mu_i}{\mathcal{M}_{it}\left(x^t\right)} \tag{57}$$

then the first-order conditions for the agent's problem can be written as

$$u_i'\left[c_{it}\left(x^t\right)\right] = \mu_{it}\left(x^t\right) p_t\left(x^t\right). \tag{58}$$

μ_{it} can be thought of as a Lagrange multiplier for a problem in which at date t agent i maximizes U_{it} by choice of consumption allocations from date t on, using information available at time t, such that his lifetime budget constraint from date t on is satisfied.

An equilibrium is prices for all dates and states, initial wealth for all agents and consumption allocations for all agents such that

---

[11]Ma (1993) has studied the existence of equilibria for a wide range of recursive preferences.

1. Given prices, the Pareto optimal allocation for each agent is the solution to their problem.

2. The aggregate resource constraints are satisfied at all dates and states.

Given a Pareto optimal allocation it is always possible to decentralize the allocation by finding prices that satisfy our definition of an equilibrium. For some fixed agent i, prices can be taken as

$$p_t\left(x^t\right) = \frac{u_i'\left[c_{it}\left(x^t\right)\right]\mathcal{M}_{it}\left(x^t\right)}{u_i'\left[c_{i0}\left(x^0\right)\right]} \tag{59}$$

where c_{it} is the Pareto optimal consumption allocation for agent i. The initial wealth of any agent can be defined as the value of their Pareto optimal consumption allocations using the prices p_t and can be computed to satisfy equation 55. The ratio of the Lagrange multipliers for any two agents is equal to the inverse of the ratio of their Pareto weights

$$\frac{\mu_{it}\left(x^t\right)}{\mu_{jt}\left(x^t\right)} = \frac{\theta_{jt}\left(x^t\right)}{\theta_{it}\left(x^t\right)} \tag{60}$$

for all histories.

In this paper we often talk about the wealth distribution and often say some agents are richer than others even though we primarily deal with Pareto weights. This is justified (at least when agents have the same preferences) because of the connection between Pareto weights and the Lagrange multipliers on the agent's budget constraints described in this section. When reading this paper, it is also important to remember that at time zero all that matters to agents is their lifetime utility. Agent can be viewed as choosing their current and future consumption. When the consumption for some agents tends toward zero over time, that is optimal for the agent at the equilibrium prices. Agents do not directly care if their Pareto weights are converging to zero or not.

Although it is always possible to decentralize a Pareto optimal allocation, the decentralization may have some unappealing features for reasons outside of the model. For example, let the aggregate state be i.i.d. over time and equal to aggregate consumption. Consider a decentralization in which individuals are endowed with units of the consumption good that are i.i.d. over time. If an agent sees their Pareto weight converging to zero it must be that eventually a date is reached in which this agent has sold all of his future endowments to other agents. This is troubling for reasons outside of the model. First, it is not very plausible that someone is allowed to sell all their future endowments. Second, if a large fraction of the population falls into this situation the equilibrium may not be sustainable in the sense that agents would likely start a revolution and cancel past contracts.

8 Risk-sensitive and time-additive agents

In this section we consider economies in which there are both risk-sensitive and time-additive agents. We begin by showing that the mixture of time-additive and risk-sensitive agents precludes the existence of interior steady states. We then show how risk is shared in the short run and in the long run. We show that an exterior steady state in which all risk-sensitive agents receive a weight of zero is stable.

Consider an economy in which there is at least one agent who has time-additive preferences and at least one agent who has non-degenerate risk-sensitive preferences. Let agent j be any agent who has time-additive preferences so that $\sigma_j = 0$. Let agent i be any agent who has non-degenerate risk-sensitive preferences so that $\sigma_i < 0$. There may be other agents in the economy that have either time-additive or risk-sensitive preferences. Let Assumptions one thru six hold.

In this economy there can be no interior regular steady states since

$$\sigma_i u_i \left(c_i \left(e, \bar{\theta} \right) \right) - \sigma_j u_j \left(c_j \left(e, \bar{\theta} \right) \right) = \sigma_i u_i \left(c_i \left(e, \bar{\theta} \right) \right)$$

varies with e. This follow from Theorem two.[12] As long as Assumptions eight and nine hold there also are no interior steady states, since ω_{ij} will not be a singleton if $\bar{\theta}_i > 0$. Any two time-additive agents will always be in a private steady state, in which their Pareto ratios are constant over time. A risk-sensitive agent and a time-additive agent can never be in a private steady state in which they both have positive Pareto weights.

How risk is shared away from the steady state? Recall the first-order condition from Section 3.4 when this period's exogenous state is x and this period's Pareto weight vector is θ :

$$\frac{\phi_{iy}}{\phi_{jy}} = \left(\frac{\theta_i}{\theta_j} \right) \frac{\exp \left(\sigma_i q_{iy} \right)}{\sum_{z \in \mathcal{X}} \pi \left(x, z \right) \exp \left(\sigma_i q_{iz} \right)} \tag{61}$$

where $q_{iy} = Q_i \left(y, \phi_y \right).$ Given particular values for this period's state x and this period's Pareto weight vector the term $\theta_i / \left[\theta_j \sum_{z \in \mathcal{X}} \pi \left(x, z \right) \exp \left(\sigma_i q_{iz} \right) \right]$ on the right hand side of equation 61 is fixed. The only term that can vary with next period's state is $\exp \left(\sigma_i q_{iy} \right).$ Since $\sigma_i < 0$, it must be that the Pareto ratio ϕ_{iy} / ϕ_{jy} increases as q_{iy} decreases. Since the expected value of the right hand side of equation 61 is θ_i / θ_j it must be that when agent i has sufficiently large values of lifetime utility his Pareto weight decreases and when he has sufficiently low values his Pareto weight increases.[13]

---

[12] $\sigma_i u_i \left(c_i \left(e, \bar{\theta} \right) \right)$ must vary with e by Assumption four. Strict concavity of the reward functions guarantees that the consumption of every agent varies with the aggregate endowment for a fixed interior Pareto weight vector.

[13] q_{iy} can not be constant across states y because there are no steady state Pareto ratios for agents i and j.

First-order condition 61 tell us that the Pareto weights move in order to provide the risk-sensitive agent with a smoother lifetime utility allocation. In bad aggregate states the risk-sensitive agent will see his Pareto weight increase over last period's Pareto weight. In good aggregate states the risk-sensitive agent will see his weight decrease.

In the equivalent distorted beliefs economy described in Section 7.1, a risk-sensitive agent is pessimistic. He places a high probability (higher than the true probability) on those states in which he does poorly. In a decentralized economy he chooses to spend a large fraction of his wealth on states in which he does poorly. For example a pessimistic agent might worry about the prospects of a world-wide nuclear war. Because of this he spends a large fraction of his wealth on a bomb-shelter, canned foods, generators, batteries, etc., so that in the event of a nuclear war his utility is relatively close to his utility if there had been no war. If there is a war, he likely will have a high Pareto weight because most other agents have chosen not to prepare for a war. However, if there is no war he likely will have a low Pareto weight. A social planner who allocates consumption would want to give risk-sensitive relatively more consumption in states in which they place high probability, since the social planner takes an agent's objective, equation 52, as given. Even though a risk-sensitive agent may have a large Pareto weight when there is a nuclear war, he still has more utility when there is no war. Relative to time-additive agents he does better when there is a nuclear way but in absolute terms he does worse.

We now turn toward the main result of this section which describes risk-sharing in the long run. The result shows that a steady state in which all risk-sensitive agents receive zero weight is stable.

Theorem 11 *Let Assumptions one thru nine hold. Consider an economy in which there is at least one agent who has time-additive preferences and at least one agent who has risk-sensitive preferences with a strictly negative risk-sensitivity parameter. Assume at least one of the time-additive agents has an initial Pareto weight which is strictly positive. Consider the steady state Pareto weight vector $\bar{\theta}$, which is not interior, that assigns weights of zeros to the risk-sensitive agents and weights such that*

$$\frac{\bar{\theta}_i}{\bar{\theta}_j} = \frac{\theta_{i0}}{\theta_{j0}} \tag{62}$$

for any two time-additive agents i and j such that $\theta_{j0} > 0$. This steady state is stable.

Theorem 11 tells us that if there is at least one-agent who has time-additive preferences then eventually the weights of all risk-sensitive agents converge to zero with probability one. In the end with probability one, time-additive agents consume all of the aggregate endowment and risk-sensitive agents consume nothing.

In the theorem θ_{i0}/θ_{j0} is the initial time zero ratio of Pareto weights for agents i and j. See the appendix for a proof. The proof works by applying Lemma one and letting any time-additive agent play the role of agent j. Assume agent one has time-additive preferences with $\theta_{10} > 0$ and set $j = 1$. For any time-additive agent i, θ_{it}/θ_{1t} is constant over time. For any risk-sensitive agent i, the Pareto ratios are a non-negative martingale:

$$E_t \left[\frac{\bar{\theta}_{it+1}}{\bar{\theta}_{1t+1}} \right] = \frac{\theta_{it}}{\theta_{1t}}.$$

This follows because the covariance term in equation 48 is zero. By the martingale convergence theorem θ_{it}/θ_{1t} converges to a constant (with probability one), when starting from a fixed initial Pareto weight vector. The proof in the appendix shows that zero is the only constant that it can converge to with positive probability.

Even though the Pareto weight of agent i converges to zero with probability one and $\bar{\theta}_{1t}$ converges to a positive number, the time zero *expected* value of the Pareto ratio $\bar{\theta}_{it}/\bar{\theta}_{1t}$ does not converge to zero as t tends toward infinity. For any t, the time zero expected value of the time t Pareto ratio

$$E_0 \left[\frac{\bar{\theta}_{it}}{\bar{\theta}_{1t}} \right] = \frac{\theta_{i0}}{\theta_{10}} \tag{63}$$

is just the initial value of the Pareto ratio. Since with probability one the Pareto ratio is converging to zero, it must be that there are states far into the future that occur with probability zero, such that the risk-sensitive agent has a very large Pareto weight and the time-additive agent has small Pareto weight. From the point of view of time zero, if we measure the expected future wealth distribution between a risk-sensitive and a time-additive agent with their Pareto ratio $\bar{\theta}_{it}/\bar{\theta}_{1t}$ it is the case that the expected future wealth distribution is equal to this period's wealth distribution even though with probability one agent i is eventually going to consume nothing and agent one is going to consume a positive amount.

The large variability in future Pareto weights may incorrectly be seen as introducing extra uncertainty into the economy. Agents don't know what their future Pareto weights will be: their weights could be very large or very low. The reason it is optimal for the Pareto weights to be variable is to ensure that risk-sensitive agents receive a relatively smooth (relative to an economy in which the Pareto weights are constant) allocation of future utility across states. The time-additive agents receive a relatively variable allocation across states.

Another interpretation of these results is that under the beliefs of risk-sensitive agents, in the equivalent distorted beliefs economy their Pareto weights are not converging to zero. If agents have distorted beliefs, as in the equivalent economy described in Section 7.1, they will, in a decentralized

economy, purchase consumption in states they believe are likely to occur. For dates far into the future agents with risk-sensitive preferences purchase most of their consumption in a group of states that they think will occur with probability one. In fact this group of states may occur with probability zero. In states that occur with high probability agents with distorted beliefs may purchase very little consumption as a consequence with probability one their consumption tends toward zero since time-additive agents have the correct beliefs and purchase most of their consumption in states that do occur with probability one.

Consider again the example in which a pessimistic agent worries about the prospects of a world-wide nuclear war. Because of this he spends a large fraction of his wealth on building a bomb-shelter, buying canned foods, generators, batteries, etc. If there is no war he likely will have a very low Pareto weight since he has spent so much of his wealth on items that are only useful when a war occurs. Since there is only a small probability that there is a war, it is very likely in the long run this agent will have a low Pareto weight.

The results of this section are related to results in Sandroni (2000). Sandroni (2000) shows in the context of heterogeneous belief economies, when all agents have the same discount factor, agents who eventually have accurate beliefs survive. If there is some agent who has accurate beliefs then all agents who have inaccurate beliefs eventually vanish from the economy. As we know from Section 7.1, a social planning problem with risk-sensitive agents is equivalent to a social planning problem problem with agents who have heterogeneous beliefs. In this equivalent economy time-additive agents always have accurate beliefs and strictly risk-sensitive agents never will have accurate beliefs.[14] Hence time-additive agents survive and risk-sensitive agents vanish.

9   Power reward functions

In this section we discuss examples in which all agents are risk-sensitive and have power reward functions of the form

$$u_i(c) = \frac{c^{1-\gamma_i}}{1 - \gamma_i} \tag{64}$$

where $\gamma > 0$ and $\gamma \neq 1$. Since power reward functions are often studied in macroeconomics when agents have time-additive preferences, it is interesting to see what happens to risk-sharing when power reward functions are combined with risk-sensitive preferences. Unlike examples in previous sections, it will not always be the case that agents who have a lower risk-sensitivity factor necessarily see their Pareto weights going to zero. Instead the initial Pareto weights (wealth), in addition to the risk-sensitivity

---

[14]Risk-sensitive agents would have accurate beliefs if there was no uncertainty in their future consumption. However, as long there is uncertainty in the aggregate endowment it is not optimal to allocate risk-sensitive agents a deterministic consumption sequence.

parameters, determine long run consumption. Section 9.1 shows that when all agents have identical power reward functions there exists a unique interior regular state. This steady state is stable when $\gamma_i < 1$ and unstable when $\gamma_i > 1$. In Section 9.2 we show that when agents have heterogeneous power reward functions there are no interior regular states.

9.1 Identical power rewards

We begin by characterizing the unique interior regular steady state.

Theorem 12 *Let all agents have identical power reward functions of the form $u_i(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$ and $\gamma \neq 1$. Let all agents be strictly risk-sensitive with $\sigma_i < 0$. Let Assumptions one thru three and Assumptions five thru six hold. In this economy there is a unique interior regular steady state Pareto weight vector, $\bar{\theta}$, which sets*

$$\bar{\theta}_i = \frac{\left(-\frac{1}{\sigma_i}\right)^{\frac{\gamma}{1-\gamma}}}{\sum_{j=1}^{n} \left(-\frac{1}{\sigma_j}\right)^{\frac{\gamma}{1-\gamma}}}$$

for all i.

See the appendix for a proof. The result also applies if there are equal numbers of two types of agents. As long as Assumptions eight and nine hold this is also the unique steady state since the one-period risk-sharing rules are linear in aggregate consumption. The steady state Pareto ratios can be written as

$$\frac{\bar{\theta}_i}{\bar{\theta}_j} = \left(\frac{\sigma_j}{\sigma_i}\right)^{\frac{\gamma}{1-\gamma}}$$

where $\sum_{j=1}^{n} \bar{\theta}_j = 1$. This will be a steady state Pareto ratio for agents i and j regardless of the preferences of other agents since $\bar{\theta}$ is a regular steady state. Notice that there exists a unique interior steady state even if the agents have different risk-sensitivity parameters. The unique interior steady state Pareto weight vector will be such that $\bar{\theta}_i = 1/n$ for all i only if σ_i is constant across agents.

It is interesting to compare the steady state consumption of two agents with different risk-sensitivity parameters. Let the agents be denoted by one and two and let agent two's risk-sensitivity parameter be lower (more negative) than agent one's. In a steady state, agent one receives more consumption (in every date-state) when $0 < \gamma < 1$ and less when $\gamma > 1$. As γ approaches one from the left the steady state ratio of Pareto weights, $\bar{\theta}_2/\bar{\theta}_1$, approaches 0. As γ approaches one from the right the steady state ratio of Pareto weights tends towards $+\infty$. Figure three graphs the steady state Pareto weight of agent two as function of γ when $\sigma_1 = -1$ and $\sigma_2 = -2$.

We now characterize how risk is shared away from the steady state and the stability of steady states when there are two agents and the aggregate endowment is i.i.d.
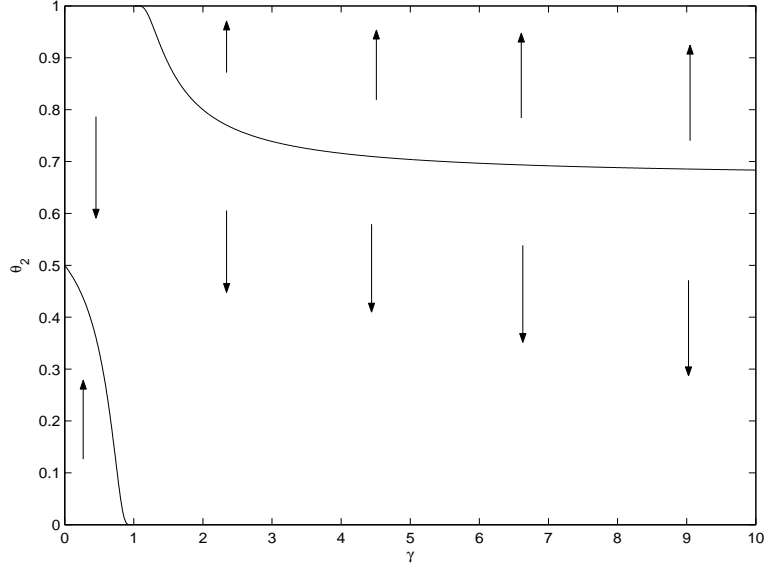
Figure 3: The interior steady state and stability when agents have the same reward functions
Let there be two agents. The sold line graphs the unique interior steady state Pareto weight for agent two as a function of γ when $\sigma_2 = -2$ and $\sigma_1 = -1$. When $\gamma = 1$ (which I take to correspond to log rewards) there are no interior steady states. The arrows indicate the stability of the interior steady state. For example, when $\gamma = 2$ the interior steady state is unstable and when $\gamma = 0.5$ the interior steady state is stable.

Theorem 13 *Let there be two agents who are denoted by $i = 1, 2$. Let Assumptions one thru three and Assumptions five thru 10 hold. Let both agents have power reward functions with the same γ. Let $\bar{\theta}$ be the unique interior regular steady state Pareto weight vector.*

1. **Positive powers.** *If $0 < \gamma < 1$ then*

   (a) *For any fixed Pareto weight vector θ such that $1 > \theta_2 > \bar{\theta}_2$ (respectively $0 < \theta_2 < \bar{\theta}_2$) it is the case that the optimal choice of next period's Pareto weight for agent two is a strictly decreasing (increasing) function of next period's state y.*

   (b) *The Pareto weight vector $\bar{\theta}$ is stable.*

2. **Negative powers.** *If $\gamma > 1$ then*

   (a) *For any fixed Pareto weight vector θ such that $1 > \theta_2 > \bar{\theta}_2$ (respectively $0 < \theta_2 < \bar{\theta}_2$) it is the case that the optimal choice of next period's Pareto weight for agent two is a strictly increasing (decreasing) function of next period's state y.*

   (b) *The Pareto weight vector $\bar{\theta}$ is not stable.*

For positive powers we have a very strong result. With probability one the Pareto weight vector is going to end up at an interior steady state. For negative powers our theoretical result is weaker. All that we have proved is that the unique interior steady state is unstable which means there is positive probability that when starting from an interior Pareto weight vector (that is not equal to the interior

steady state) that we end up away from the interior steady state. Since it can be shown the Pareto weights are a supermartingale in this case they still must be converge to some constant with probability one, for a fixed initial Pareto weight vector. (See the proof of Theorem 13 for a discussion.) The only constants that it could converge to are an interior steady state or an exterior steady state in which one agent has a Pareto weight of one. By the no-crossing property there is only one agent who could receive a Pareto weight of one in the long run. Substantial numerical explorations suggest that with probability one the Pareto weights converge to an exterior steady state with one agent receiving a weight of one. However, since only a finite number of simulations were run this result can not be unquestionably verified.

The dynamics depend on the relationship between the current period Pareto weight vector, the unique interior regular state Pareto weight vector $\bar{\theta}$, and the power γ in the agent's reward functions. Following Theorem seven and Theorem 10 the effect of these variables is summarized by whether or not $\delta_{21}(y, \phi)$ is increasing, decreasing or constant in y. For example let there be two agents, as in Figure three, with $\sigma_2 = -2$ and $\sigma_1 = -1$. If $\gamma = 2$ for both agents then if the initial weight on the more risk-sensitive agent, agent two, is 0.9 then eventually (based on numerical explorations) agent two will consume all of the aggregate endowment. But, if the initial weight on the more risk-sensitive agent is 0.5 then eventually (again based on numerical explorations) agent two will consume nothing.

Figure 4 graphs the evolution of Pareto weights for two examples. In the right graph we verify the result in Theorem 13 that for positive powers interior steady states are stable. In the left graph we investigate what happens for a negative power. We see that agent two eventually receives a Pareto weight of one.

When agents have identical power rewards, we do not have analytic results when there are more than two (types of) agents or when the state is not i.i.d. Such results should be feasible but the proofs are very involved. Neither assumption is essential for the results. From numerical explorations we find that the same results hold if the aggregate endowment follows an AR(1) process and is positively serially correlated. Numerically, we find when there are more than two agents, if all agents have the same positive power then interior steady states will be stable. Numerically, if all agents have the same negative power then, provided no two agents are initially in a private steady state, a *single* agent will eventually have a Pareto weight of one and all other agents will have a weight of zero. If some agents are in a private steady state then it is possible that multiple agents will have positive weights in the long-run.

Some of the results in this section have the same flavor as results in Dumas (1989) but there are important differences. Dumas (1989) considers a continuous time heterogeneous agent economy with

two agents who have heterogeneous power reward functions. The economy is more general than the economies considered in this paper in that it includes production. It is less general in that its analysis is restricted to time-additive utility. With time-additive utility the wealth distribution only varies over time in interesting ways if there is growth. In Dumas (1989) there can be (endogenous) growth. The interesting aspect of the results in this paper is that the long run results in risk-sensitive economies happen when there is no growth. If there was growth then there would be an additional mechanism that could cause consumption distributions to vary over time, independent of changes in Pareto weights.

Like some of the results of earlier sections, the results of this section are related to results in Sandroni (2000). Sandroni (2000) shows in the context of heterogeneous belief economies that agents who have the most accurate beliefs (as measured by entropy) survive and all other agents vanish.[15] As we know from Section 7.1, a social planning problem with risk-sensitive agents is equivalent to a social planning problem problem with agents who have heterogeneous beliefs. In this equivalent economy when all agents have positive powers, all agents eventually have the same beliefs and thus survive. In this equivalent economy when all agents have negative powers, agents do not eventually have the same beliefs and some agents do not survive.

The mechanism driving the results in the equivalent distorted beliefs economy is as follows. For positive powers, the beliefs of richer agents are farther from the true probabilities than the beliefs of poorer agents. Richer agents pay short and long run penalties for this: their Pareto weights tend to decrease over time. Over time the gap between the rich and poor decreases so that eventually all agents have the same beliefs and we are at an interior steady state. For negative powers, the beliefs of richer agents are closer to the true probabilities than the beliefs of poorer agents. Now poorer agents agents pay short and long run penalties for this: their Pareto weights tend to decrease over time. Hence the gap between rich and poor agents widens over time until eventually some agents consume nothing.

9.2 Heterogeneous power rewards

In this subsection we let agents have heterogeneous power reward functions of the form $u_i(c) = \frac{c^{1-\gamma_i}}{1-\gamma_i}$ with $\gamma_i > 0$ and $\gamma_i \neq 1$ and let all agents be strictly risk-sensitive with $\sigma_i < 0$. Assume for some pair of agents i and j it is the case that $\gamma_i \neq \gamma_j$. The power reward function satisfies the properties listed in Assumption four. We will assume throughout that Assumptions one thru three and Assumptions five thru six hold.

---

[15]Anderson (1998) also explored the use of entropy in determining the long run behavior of allocations when agents have recursive utility and/or heterogeneous beliefs. Also see Blume and Easley (1992) for a discussion of entropy.
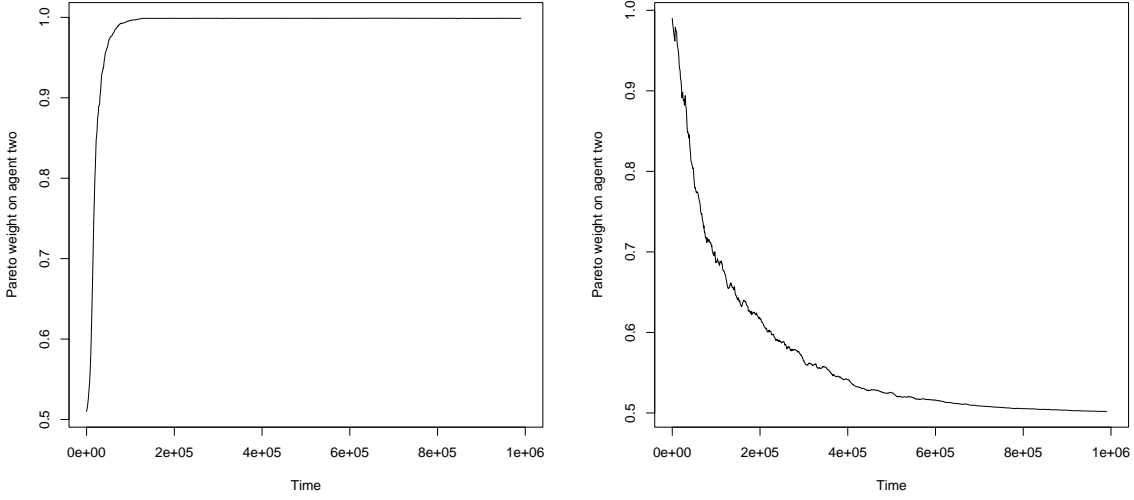
Figure 4: Evolution with identical power preferences

Evolution of the Pareto weight on agent of two when $\beta = 0.95$, $\sigma = -1$ and agents have the same γ. In the left graph γ is 2.0 for both agents and in the right graph γ is 0.5 for both agents. The probabilities are given by specification B. The initial Pareto weight on agent two in the left graph is 0.51 and in the right graph 0.99. The Pareto weight is graphed for 1×10^6 time periods. (On the graph one point is plotted for every 1000 periods so that 1000 points are plotted.)

We show there does not exist any interior regular steady states. If θ is an interior regular steady state then we must have[16]

$$\left(\frac{\theta_i}{\theta_j}\right)^{\frac{1-\gamma_i}{\gamma_i}} \frac{\sigma_i \left[c_j\left(e,\theta\right)\right]^{\frac{\gamma_j}{\gamma_i}\left(1-\gamma_i\right)}}{1-\gamma_i} - \frac{\sigma_j c_j\left(e,\theta\right)^{1-\gamma_j}}{1-\gamma_j} = d \tag{66}$$

for some constant d and all possible e. Differentiating this equation with respect to e yields

$$\left[\frac{\gamma_j}{\gamma_i}\left(\frac{\theta_i}{\theta_j}\right)^{\frac{1-\gamma_i}{\gamma_i}} \sigma_i \left[c_j\left(e,\theta\right)\right]^{\frac{\gamma_j}{\gamma_i}\left(1-\gamma_i\right)-1} - \sigma_j c_j\left(e,\theta\right)^{-\gamma_j}\right] \frac{\partial c_j\left(e,\theta\right)}{\partial e} = 0. \tag{67}$$

Dividing by $c_j^{-\gamma_j} \partial c_j/\partial e$ and rearranging terms yields the requirement that

$$\frac{\gamma_j}{\gamma_i}\left(\frac{\theta_i}{\theta_j}\right)^{\frac{1-\gamma_i}{\gamma_i}} \sigma_i \left[c_j\left(e,\theta\right)\right]^{\frac{\gamma_j}{\gamma_i}-1} = \sigma_j. \tag{68}$$

---

[16]This follows because the optimal one-period consumption allocation rules satisfy the first-order condition

$$\theta_i \left[c_i\left(e,\theta\right)\right]^{-\gamma_i} = \theta_j \left[c_j\left(e,\theta\right)\right]^{-\gamma_j}. \tag{65}$$

Solving this equation for $c_i\left(e,\theta\right)$ and placing the result in $\sigma_i u_i \left[c_i\left(e,\theta\right)\right] - \sigma_j u_j \left[c_j\left(e,\theta\right)\right]$ yields the left hand side of equation 66.

This equation can only be satisfied if $\gamma_i = \gamma_j$ or if $c_j(e, \theta)$ does not vary with e. We have assumed $\gamma_i \neq \gamma_j$ and $c_j(e, \theta)$ varies with the aggregate endowment because of the reasons given in footnote 12. By Theorem two there are no interior regular steady states. If Assumptions eight and nine hold then Corollary one guarantees that there also are no interior steady states.

What do the dynamics of allocations look like in this case? The no-crossing property does not apply so there are no interior bounds on the evolution of Pareto weights. One can numerically show that if γ for all agents is between zero and one then there exists a non-degenerate invariant distribution among agents. If γ for all agents is greater than one then with positive probability, one agent will have a Pareto weight of one. Which agent, if any, receives a weight of one in the limit depends upon the realizations of the endowment.

10 Conclusions

Although the allocation rules discussed in this paper were derived in an endowment economy they apply more generally when the optimal law of motion for a state vector follows a Markov process. Given the optimal law of motion and the optimal choice of aggregate consumption, the results of this paper can be used to show how consumption will be shared if optimal aggregate consumption satisfies the assumptions placed on the exogenous aggregate endowment in this paper.[17]

---

[17]A number of papers that were written at about the same time as this paper discuss related issues. Liu (2002) discusses heterogeneous agent economies with Knightian uncertainty. Hansen and Sargent (2003) discuss an environment with multiple agents who have a preference for robustness. Epstein (2001) and Epstein and Miao (2003) discuss sharing ambiguity.

Appendix

Value function iteration

This part of the appendix follows up on an observation made in Section 3.3. Following Lucas and Stokey (1984) and Kan (1995) we describe how one can iterate on an initial guess of the social planner's value function Q to compute the the social planner's value function. Given a guess at the social planner's value function Q^k, consider the following problem:

$$Q^{k+1}(x,\theta) = \max_{\{c_i \geq 0, q_i \geq \underline{q}_i\}_{i=1}^n} \sum_{i=1}^n \theta_i W_i(x, c_i, q_i) \tag{69a}$$

where

$$\sum_{i=1}^n c_i = c(x) \tag{69b}$$

and for each possible value, y, of next period's state

$$0 \leq \min_{\phi_y \in \Gamma^n} \left[Q^k(y, \phi_y) - \sum_{i=1}^n \phi_{iy} q_{iy} \right]. \tag{69c}$$

To compute the fixed point one can choose an arbitrary function $Q^0(x,\theta)$ which is continuous; increasing and concave in x; convex in θ; and finite. Then the social planner's value function can be computed by iteration. It will be the case that

$$Q(x,\theta) = \lim_{k \to \infty} Q^k(x,\theta) \tag{70}$$

where Q is the social planner's value function defined in equation 22 and Q^k is the kth iteration on the dynamic program.

Proof of Theorem one

I will first prove part one and then prove part two. Throughout we *only* discuss agents who have strictly positive Pareto weight. If $\bar{\theta}_i = 0$ for some agent i then his weight next period will always be zero. Let

$$q_{iy} = Q_i(y, \bar{\theta}) \tag{71}$$

be the utility remaining for agent i when the current period state is y and the current period Pareto weight vector is $\bar{\theta}$. Throughout y is an arbitrary value of the state and $\bar{\theta}$ is a fixed particular value for the Pareto weight vector.

Let the Pareto weight vector be $\bar{\theta}$ this period. If $\Omega_{ij}(x, \bar{\theta})$ is a singleton for agents i and j then there exists a real number g_{ij} such that

$$\sigma_i q_{iy} = g_{ij} + \sigma_j q_{jy}$$

for all possible values of the exogenous state y that can occur next period. g_{ij} can vary with i and j but cannot vary with y. If this condition is met for agents i and j then agents i and j will have the same marginal value of utility next period. More formally agent i's marginal value of utility next period when next period's state is y, is equal to agent j's:

$$
\begin{aligned}
\mathcal{M}_i(x, y, q_i) &= \frac{\beta \pi_i(x,y) \exp(\sigma_i q_{iy})}{\sum_{z \in \mathcal{X}} \pi(x,z) \exp(\sigma_i q_{iz})} \\
&= \frac{\beta \pi_j(x,y) \exp(\sigma_j q_{jy} + g_{ij})}{\sum_{z \in \mathcal{X}} \pi(x,z) \exp(\sigma_j q_{jz} + g_{ij})} \\
&= \mathcal{M}_j(x, y, q_j).
\end{aligned}
$$

If q_{iy} is given by equation 71 and

$$\frac{\phi_{iy}}{\phi_{jy}} = \frac{\bar{\theta}_i}{\bar{\theta}_j}$$

for all i and j the first-order conditions 29b and 29c are satisfied. Since there is a unique solution to the first-order conditions, next period's Pareto weight vector is $\bar{\theta}$. Since $\Omega_{ij}(x, \bar{\theta})$ is a singleton for all x this argument applies at every date in the future (with possibly a different value of g_{ij}). Hence the Pareto weight vector remains at $\bar{\theta}$ every date and state in the future.

Now for part two. If the Pareto weight vector equals $\bar{\theta}$ this period and in every state which can occur next period it must be the case, by first-order condition 29b, that $\mathcal{M}_i\left(x, y, q_i\right) = \mathcal{M}_j\left(x, y, q_j\right)$. Now,

$$\mathcal{M}_i\left(x, y, q_i\right) = \mathcal{M}_j\left(x, y, q_j\right) \tag{72a}$$

$$\frac{\beta \pi_i\left(x, y\right) \exp\left(\sigma_i q_{iy}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right)} = \frac{\beta \pi_j\left(x, y\right) \exp\left(\sigma_j q_{jy}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right)} \tag{72b}$$

$$\exp\left(\sigma_i q_{iy} - \sigma_j q_{jy}\right) = \frac{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right)} \tag{72c}$$

$$\sigma_i q_{iy} - \sigma_j q_{jy} = \log\left[\frac{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right)}\right]. \tag{72d}$$

Equation 72b follows from the definition of $\mathcal{M}$. Equation 72c follows by rearranging terms using the assumption that agents have the same beliefs. Equation 72d follows by taking logarithms.

We have that equation 72d must hold for all states y that can occur next period with strictly positive probability. However the right hand side cannot vary with y, so the left hand side must not vary with y. It follows that $\Omega_{ij}(x, \bar{\theta})$ is a singleton at all states that can occur next period. Since $\mathcal{X}$ is an ergodic set, all states, x, will be the current state at some future date with strictly positive probability. So $\Omega_{ij}(x, \bar{\theta})$ is a singleton for all states x.

Proof of corollary one

I will first prove part two and then prove part one. Throughout we *only* discuss agents who have strictly positive Pareto weight. If $\bar{\theta}_i = 0$ for some agent i then his weight next period will always be zero. Let

$$q_{ix} = Q_i\left(x, \bar{\theta}\right) \tag{73}$$

$$r_{ix} = u_i\left(c_i\left[c(x), \bar{\theta}\right]\right) \tag{74}$$

be shorthand for the lifetime and one-period utility remaining respectively for agent i when the current state is x and the current period Pareto weight vector is $\bar{\theta}$. Throughout x is an arbitrarily value of the state and $\bar{\theta}$ is a fixed particular value for the Pareto weight vector.

If $\bar{\theta}$ is a steady state then by Theorem one and Assumption eight there exists a real number g_{ij} such that

$$\sigma_i q_{iy} = g_{ij} + \sigma_j q_{jy} \tag{75}$$

for all states y. Since q_i satisfies the recursion

$$q_{ix} = r_{ix} + \frac{\beta}{\sigma_i} \log \sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right) \tag{76}$$

it follows that (after multiplying both sides by σ_i)

$$\sigma_i q_{ix} = \sigma_i r_{ix} + \beta \log \sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right) \tag{77a}$$

$$= \sigma_i r_{ix} + \beta \log \sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz} + g_{ij}\right) \tag{77b}$$

$$= \sigma_i r_{ix} + \beta\left[g_{ij} + \log \sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right)\right] \tag{77c}$$

where equation 77b uses equation 75, with $y = z$, to write q_{iz} in terms of q_{jz}. Multiplying the recursion 76 for agent j by σ_j yields

$$\sigma_j q_{jx} = \sigma_j r_{jx} + \beta \log \sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right). \tag{78}$$

Subtracting 78 from 77c yields

$$\sigma_i q_{ix} - \sigma_j q_{jx} = \sigma_i r_{ix} - \sigma_j r_{jx} + \beta g_{ij} \tag{79}$$

$$g_{ij} = \sigma_i r_{ix} - \sigma_j r_{jx} + \beta g_{ij} \tag{80}$$

$$(1 - \beta) g_{ij} = \sigma_i r_{ix} - \sigma_j r_{jx}. \tag{81}$$

Equation 80 uses equation 75, with $y = x$. The right hand side of equation 81 can not vary with x since the left hand side does not depend upon x.

Now for part one. If $\omega_{ij}(x, \bar{\theta})$ is a singleton for all agents i and j and for all states x then there exists a real number d_{ij} such that

$$\sigma_i r_{ix} = \sigma_j r_{jx} + d_{ij} \tag{82}$$

50

for all possible values of the exogenous state x that can occur next period. The real numbers d_{ij} can vary with i and j but can not vary with x. Again, we assume all agents have strictly positive weights. If any agent has a weight of zero, they can be ignored since they will consume nothing and contribute nothing to the social planner's value function.

Consider the following guess at the social planner's value function

$$Q^0(x, \theta) = \sum_{i=1}^{n} \theta_i u_i (c_i(x, \theta)). \tag{83}$$

We will iterate, starting from Q^0, to obtain a sequence of value functions which converge to Q, the solution of the dynamic programming problem described in Section 3.3. On iteration k, let Q^k be the social planner's value function and let Q_i^k be the value function for agent i. We have (using arguments similar to the ones which justified equation 27 in the text)

$$Q_i^k(x, \theta) = \begin{cases} Q^k(x, \theta) & \theta_i = 1 \\ \frac{\partial Q^k(x, \theta)}{\partial \theta_i} & 1 > \theta_i > 0 \\ \underline{q}_i & \theta_i = 0. \end{cases} \tag{84}$$

Define

$$q_{ix}^k = Q_i^k(x, \bar{\theta}). \tag{85}$$

On each iteration we will show that

$$\sigma_i q_{ix}^k - \sigma_j q_{jx}^k = g_{ij}^k \tag{86}$$

where g_{ij}^k is a constant which does not vary with x but may vary with i, j and k. Equation 86 holds when $k = 0$ since

$$q_{ix}^0 = \frac{\partial Q^0(x, \theta)}{\partial \theta_i}|_{\theta=\bar{\theta}} = u_i (c_i(x, \bar{\theta})). \tag{87}$$

and

$$\sigma_i q_{ix}^0 - \sigma_j q_{jx}^0 = \sigma_i r_{ix} - \sigma_j r_{jx} = d_{ij}. \tag{88}$$

The last equation follows from equation 82. Set $g_{ij}^0 = d_{ij}$.

On iteration $k + 1$, if equation 86 holds on iteration k then the optimal choice of next period's Pareto weights are $\phi_y = \bar{\theta}$ for all y when this period's Pareto weights are $\bar{\theta}$. To see this note that the first-order conditions on iteration $k + 1$ include

$$\left(\frac{\phi_{iy}}{\phi_{jy}} \right) \mathcal{M}_j \left(x, y, q_j^k \right) = \left(\frac{\theta_i}{\theta_j} \right) \mathcal{M}_i \left(x, y, q_i^k \right)$$

for all y where q_i^k is given by equation 85. These are the analogs to the first-order conditions 29b and 29c when the value function is Q^k from next period on. If next period's Pareto weights are $\bar{\theta}$ then equation 86 guarantees that the first-order condition is satisfied with $\mathcal{M}_i = \mathcal{M}_j$.

Now, we show that if equation 86 holds on iteration k then it must hold on iteration $k + 1$. We have at $\bar{\theta}$

$$Q^{k+1}(x, \bar{\theta}) = \sum_{i=1}^{n} \bar{\theta}_i \left(r_{ix} + \frac{\beta}{\sigma_i} \log \sum_{z \in \mathcal{X}} \pi(x, z) \exp \left(\sigma_i q_{iz}^k \right) \right)$$

$$= \sum_{i=1}^{n} \bar{\theta}_i \left(r_{ix} + \frac{\beta}{\sigma_i} \log \sum_{z \in \mathcal{X}} \pi(x, z) \exp \left(\sigma_j q_{jz}^k + g_{ij}^k \right) \right)$$

$$= \sum_{i=1}^{n} \bar{\theta}_i \left(r_{ix} + \frac{\beta}{\sigma_i} \left[g_{ij}^k + \log \sum_{z \in \mathcal{X}} \pi(x, z) \exp \left(\sigma_j q_{jz}^k \right) \right] \right).$$

It follows that

$$q_{ix}^{k+1} = r_{ix} + \frac{\beta}{\sigma_i} \left[g_{ij}^k + \log \sum_{z \in \mathcal{X}} \pi(x, z) \exp \left(\sigma_j q_{jz}^k \right) \right]$$

and

$$\sigma_i q_{ix}^{k+1} - \sigma_j q_{jx}^{k+1} = \sigma_i r_{ix} - \sigma_j r_{jx} + \beta g_{ij}^k \tag{89}$$

$$= d_{ij} + \beta g_{ij}^k. \tag{90}$$

The right hand side of equation 90 does not vary with x so neither does the left hand side. We can set

$$g_{ij}^{k+1} = d_{ij} + \beta g_{ij}^k. \tag{91}$$

In the limit

$$\sigma_i Q_i \left(x, \bar{\theta}\right) - \sigma_j Q_j \left(x, \bar{\theta}\right) = \lim_{k \to \infty} \left[\sigma_i q_{ix}^k - \sigma_j q_{jx}^k\right] \tag{92}$$

$$= \lim_{k \to \infty} g_{ij}^k \tag{93}$$

$$= \frac{d_{ij}}{1 - \beta}. \tag{94}$$

Since the right hand side of equation 94 does not vary x it must be that

$$\sigma_i Q_i \left(x, \bar{\theta}\right) - \sigma_j Q_j \left(x, \bar{\theta}\right) \tag{95}$$

does not vary with x. Hence by Theorem one, since $\Omega_{ij}(x, \bar{\theta})$ is a singleton for all states x, $\bar{\theta}$ is a steady state.

Proof of Theorem two

First we show that if $\bar{\theta}$ is a regular steady state for all agents then for each pair of agents i and j who have strictly positive weights

$$\sigma_i u_i \left(c_i \left(e, \bar{\theta}\right)\right) - \sigma_j u_j \left(c_j \left(e, \bar{\theta}\right)\right) \tag{96}$$

does not vary with e for all positive real numbers e, though it may vary with i and j. Then we show the converse.

We show the first part by contradiction. Assume the Pareto weight vector $\bar{\theta}$ assigns positive weights to agents i and j. Assume there exist hypothetical non-negative values of the aggregate endowment e_a and e_b such that

$$\sigma_i u_i \left(c_i \left(e_a, \bar{\theta}\right)\right) - \sigma_j u_j \left(c_j \left(e_a, \bar{\theta}\right)\right) \neq \sigma_i u_i \left(c_i \left(e_b, \bar{\theta}\right)\right) - \sigma_j u_j \left(c_j \left(e_b, \bar{\theta}\right)\right). \tag{97}$$

Construct an economy in which the state space is $\mathcal{X} = \{e_a, e_b\}$ where e_a and e_b are the possible values for the state vector. Specify aggregate consumption to be the identity function of the current state so that $c(e_a) = e_a$ and $c(e_b) = e_b$. Let

$$\pi(e_a, e_a) = \pi(e_a, e_b) = \pi(e_b, e_a) = \pi(e_b, e_b) = 0.5. \tag{98}$$

It follows by Corollary one that $\bar{\theta}$ is not a steady state in this economy. Hence $\bar{\theta}$ cannot be a regular state because there exists a specification of the aggregate state space, aggregate consumption and transition probabilities in which $\bar{\theta}$ is not a steady state.

Now for the converse. If equation 96 does not vary with e then it follows that, regardless of the specification of the state space,

$$\sigma_i u_i \left(c_i \left[c\left(x\right), \bar{\theta}\right]\right) - \sigma_j u_j \left(c_j \left[c\left(x\right), \bar{\theta}\right]\right)$$

does not vary with the state x, since $c(x)$ is a positive real number by Assumption Assumption two. Hence $\omega_{ij}(x, \bar{\theta})$ is a singleton for all states x and Corollary one guarantees that $\bar{\theta}$ is a steady state, no matter what the state space $\mathcal{X}$, transition probabilities, and consumption are.

Proof of corollary two

Let $\bar{\theta}$ be a steady state Pareto weight vector and assume the one-period consumption allocation rules are linear in the aggregate endowment. Write the one-period consumption allocation rules as

$$c_i(e, \theta) = s_i(\theta)e \tag{99}$$

where s_i is a function which can depend upon the Pareto weights, but not the aggregate endowment.

By Corollary one, given the current value of the state x, for all values of the state y that can occur next period with strictly positive probability

$$\sigma_i u_i \left(c_i \left(e, \bar{\theta}\right)\right) - \sigma_j u_j \left(c_j \left(e, \bar{\theta}\right)\right) \tag{100}$$

does not vary with $e = c(y)$ for any pair of agents i and j who have strictly positive weight. The derivative of line 100 with respect to e is

$$\sigma_i u_i' \left[s_i \left(\bar{\theta}\right) e\right] s_i(\bar{\theta}) - \sigma_j u_j' \left[s_j \left(\bar{\theta}\right) e\right] s_j(\bar{\theta}). \tag{101}$$

Now the first-order condition for current consumption allocations

$$\bar{\theta}_i u_i' \left[s_i \left(\bar{\theta}\right) e\right] = \bar{\theta}_j u_j' \left[s_j \left(\bar{\theta}\right) e\right] \tag{102}$$

implies that line 101 can be written as

$$\bar{\theta}_j \alpha(\bar{\theta}) u_j' \left[s_j \left(\bar{\theta}\right) e\right] \tag{103}$$

where we define

$$\alpha(\bar{\theta}) = \frac{\sigma_i s_i(\bar{\theta})}{\bar{\theta}_i} - \frac{\sigma_j s_j(\bar{\theta})}{\bar{\theta}_j}. \tag{104}$$

Since $\bar{\theta}_j \alpha$ is a constant for a fixed $\bar{\theta}$ and u'_j is strictly greater than zero then either the expression in line 103 (and hence also line 101) is always exactly zero, always strictly negative, or always strictly positive. It can not be always strictly negative or always strictly positive since if it was equation 100 would have to vary with y since $e = c(y)$ because Assumption nine guarantees that there are at least two distinct values of $e(y)$ occur with strictly positive probability. So α must always be exactly zero and equation 100 must be equal to a constant for all values of the aggregate endowment e such that $e > 0$. Hence $\bar{\theta}$ is a regular steady state.

Proof of Theorem three

We will prove the result by showing that if $\bar{\lambda}_{ij}$ is a steady state Pareto ratio for agents i and j then for all Pareto weight vectors θ, such that

$$\frac{\theta_i}{\theta_j} = \bar{\lambda}_{ij} \tag{105}$$

it is the case that for all $x \in \mathcal{X}$ and all $y \in \mathcal{X}$:

$$\sigma_i u_i \left(c_i \left[c(x), \theta \right] \right) - \sigma_j u_j \left(c_j \left[c(x), \theta \right] \right) = \sigma_i u_i \left(c_i \left[c(y), \theta \right] \right) - \sigma_j u_j \left(c_j \left[c(y), \theta \right] \right). \tag{106}$$

Let $\bar{\lambda}_{ij}$ be a steady state Pareto ratio for agents i and j. Define

$$\xi_{ij}(x, \theta) = c_i \left[c(x), \theta \right] + c_j \left[c(x), \theta \right] \tag{107}$$

to be the sum of the optimal consumption of agents i and j. Consider an auxiliary economy in which only agents i and j participate. Let the exogenous state in this economy comprise of the vectors x and θ. The law of motion for x is the same as in the original economy. The law of motion for θ is given by the optimal law of motion for the Pareto weights in the original economy. In the auxiliary economy the evolution of θ is treated as exogenous. Denote the state in the auxiliary economy as

$$\hat{x} = \begin{bmatrix} x \\ \theta \end{bmatrix}. \tag{108}$$

The auxiliary economy is itself a risk-sensitive heterogeneous agent economy. In the auxiliary economy let the two dimensional Pareto weight vector be $\hat{\theta}$ and set aggregate consumption, $\hat{c}$, so that

$$\hat{c}(\hat{x}) = \xi_{ij}(x, \theta). \tag{109}$$

In the auxiliary economy agents i and j will receive the same utility and the same consumption as in the original economy. The Pareto weights will be constant and such that

$$\frac{\hat{\theta}_i}{\hat{\theta}_j} = \bar{\lambda}_{ij} \tag{110}$$

$$\hat{\theta}_i + \hat{\theta}_j = 1. \tag{111}$$

We will be in a steady state for all agents in the auxiliary economy. Since the one-period consumption allocation rules are linear, by Corollary two this steady state is regular. It follows that the weighted difference of one-period utility allocations is constant for all values of the state that can occur with positive probability in the future. This guarantees that for all $x \in \mathcal{X}$ and $y \in \mathcal{X}$ equation 106 holds.

Proof of Theorem four

Let $\bar{\theta}$ be a regular steady state. For any i and j let θ be a Pareto weight vector such that

$$\frac{\theta_i}{\theta_j} = \frac{\bar{\theta}_i}{\bar{\theta}_j}. \tag{112}$$

The solution to the social planner's problem gives us formulas for the optimal choice of consumption for agents i and j which we denote as $c_i \left[c(x), \theta \right]$ and $c_j \left[c(x), \theta \right]$. Assume $\bar{\theta}_i > 0$. If it is not then the proof is obvious.

Now construct a value for the aggregate endowment e, such that if the Pareto weights were $\bar{\theta}$, not θ, agents i and j would still consume $c_i \left[c(x), \theta \right]$ and $c_j \left[c(x), \theta \right]$. For agent k, let the function $\hat{a}_k(x, \theta)$ solve the equation

$$\bar{\theta}_k u'_k \left[\hat{a}_k(x, \theta) \right] = \bar{\theta}_i u'_i \left(c_i \left[c(x), \theta \right] \right) \tag{113}$$

if $\bar{\theta}_k > 0$. If $\bar{\theta}_k = 0$ then set $\hat{a}_k(x, \theta) = 0$. $\hat{a}_k$ is the amount that agent k would consume at the steady state $\bar{\theta}$, if agent i consumes $c_i \left[c(x), \theta \right]$. Let

$$e(x, \theta) = \sum_{k=1}^{n} \hat{a}_k(x, \theta). \tag{114}$$

Now it follows that

$$c_i\left[e(x,\theta),\bar{\theta}\right] = c_i\left[c(x),\theta\right] \tag{115}$$

$$c_j\left[e(x,\theta),\bar{\theta}\right] = c_j\left[c(x),\theta\right]. \tag{116}$$

Since $\bar{\theta}$ is a regular steady state

$$\sigma_i u_i\left(c_i\left[e(x,\theta),\bar{\theta}\right]\right) - \sigma_j u_j\left(c_j\left[e(x,\theta),\bar{\theta}\right]\right) \tag{117}$$

does not vary with x and θ. It follows that

$$\sigma_i u_i\left(c_i\left[c(x),\theta\right]\right) - \sigma_j u_j\left(c_j\left[c(x),\theta\right]\right) \tag{118}$$

does not vary with x and θ.

Now the proof of Corollary one can be imitated to show that if equation 118 holds for all x and θ then the Pareto ratios for agents i and j remain constant over time. The key step is to show that if we start from the guess in equation 83 then on iteration k

$$\sigma_i Q_i^k\left(x,\theta\right) - \sigma_j Q_j^k\left(x,\theta\right) = g_{ij}^k \tag{119}$$

is satisfied whenever θ satisfies equation 112. The constant g_{ij}^k does not vary with x and θ but may vary with i and j, The previous paragraph shows that this holds when $k = 0$. It is easy to show that it holds for all k.

Proof of Theorem five

Let **E** be values of the aggregate endowment at all dates and states. Let $F_i\left(\mathbf{E}\right)$ be the optimal consumption for agent i at all dates and states if the Pareto weight vector remains at $\bar{\theta}$ in all dates and states. At the steady state $\bar{\theta}$ the social planner's preferences over the aggregate endowment can be written as

$$\sum_{i=1}^n \bar{\theta}_i U_i\left[x_0, F_i\left(\mathbf{E}\right)\right] \tag{120}$$

where U_i is defined in equation eight. Since the steady state is regular for any agents i and j who have positive weight

$$U_i\left[x_0, F_i\left(\mathbf{E}\right)\right] = \left(-\frac{\sigma_j}{\sigma_i}\right) U_j\left[x_0, F_j\left(\mathbf{E}\right)\right] + \frac{g_{ij}}{\sigma_i} \tag{121}$$

where g_{ij} is a constant (see the proof of Theorem one for more discussion) and the social planner's preferences can be written as

$$\Lambda_j U_j\left[x_0, F_j\left(\mathbf{E}\right)\right] + \Phi_j \tag{122}$$

where

$$\Lambda_j = \left[\sum_{i=1}^n \bar{\theta}_i\left(-\frac{\sigma_j}{\sigma_i}\right)\right] \tag{123}$$

$$\Phi_j = \left[\sum_{i=1}^n \bar{\theta}_i \frac{g_{ij}}{\sigma_i}\right]. \tag{124}$$

The constant Φ_j does not affect the social planner's preferences and can be disregarded.

Consider the case where all agents have the same risk-sensitivity factors. In this case $\Lambda_j = 1$ and the social planner's preferences at the steady state can be captured by

$$U_j\left[x_0, F_j\left(\mathbf{E}\right)\right]. \tag{125}$$

Since j is arbitrary this argument works for any j. When agents have different risk-sensitivity factors it is always possible to normalize the risk-sensitivity factors, to say σ_j for some particular j, by scaling the agent's reward functions (see the discussion in Section 2.3). To maintain the same steady state Pareto optimum the Pareto weights will have to be adjusted also. After they are adjusted the argument at the beginning of this paragraph can be invoked to show that the social planner's preferences at the steady can be captured by equation 125 with j arbitrary.

Proof of Theorem six

Let θ be a Pareto weight vector and $\bar{\theta}$ a regular steady state. We will show if

$$\frac{\theta_i}{\theta_j} > \frac{\bar{\theta}_i}{\bar{\theta}_j} \tag{126}$$

then the optimal choice of next period's Pareto weight vector, ϕ_y, satisfies

$$\frac{\phi_{iy}}{\phi_{jy}} > \frac{\bar{\theta}_i}{\bar{\theta}_j} \tag{127}$$

in all states y.

Define $\hat{\theta}$ to be a modified version of θ which assigns the same weight to all agents as θ except that the weights assigned to agents i and j are set to the steady state Pareto ratio for agents i and j:

$$\frac{\hat{\theta}_i}{\hat{\theta}_j} = \frac{\bar{\theta}_i}{\bar{\theta}_j} \tag{128}$$

$$\hat{\theta}_k = \theta_k \qquad k \neq i, j. \tag{129}$$

When next period's state is y, let ϕ_y and $\hat{\phi}_y$ be the optimal choice of next period's Pareto weights when this period's Pareto weights are θ and $\hat{\theta}$ respectively. Also define

$$q_{iy} = Q_i\left(y, \phi_y\right) \tag{130}$$

$$\hat{q}_{iy} = Q_i\left(y, \hat{\phi}_y\right) \tag{131}$$

for all states y and all agents i. Let this period's state be x.

Since $\bar{\theta}$ is a regular steady state and $\bar{\theta}_i/\bar{\theta}_j = \hat{\theta}_i/\hat{\theta}_j$ is a steady state for agents i and j it must be that

$$\exp(\sigma_i \hat{q}_{iy} - \sigma_j \hat{q}_{jy}) = \exp(g_{ij}) \tag{132}$$

$$\frac{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j \hat{q}_{jz}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i \hat{q}_{iz}\right)} = \exp(-g_{ij}) \tag{133}$$

for some constant g_{ij}. (This follows from Theorem two and arguments similar to arguments in the proof of Corollary one.)

When the current Pareto weights are θ the first-order conditions for the social planner's problems are

$$\left(\frac{\phi_{iy}}{\phi_{jy}}\right) = \left(\frac{\theta_i}{\theta_j}\right) \exp(\sigma_i q_{iy} - \sigma_j q_{jy}) \left(\frac{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right)}\right) \tag{134}$$

for all y. By optimality in the social planning problem it must be that

$$\frac{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_j q_{jz}\right)}{\sum_{z \in \mathcal{X}} \pi\left(x, z\right) \exp\left(\sigma_i q_{iz}\right)} \geq \exp(-g_{ij}). \tag{135}$$

Now assume for some state y

$$\frac{\phi_{iy}}{\phi_{jy}} < \frac{\bar{\theta}_i}{\bar{\theta}_j}. \tag{136}$$

It follows that

$$\exp(\sigma_i q_{iy} - \sigma_j q_{jy}) \geq \exp(g_{ij}) \tag{137}$$

and hence by equation 134

$$\frac{\phi_{iy}}{\phi_{jy}} \geq \frac{\theta_i}{\theta_j} > \frac{\bar{\theta}_i}{\bar{\theta}_j}. \tag{138}$$

This contradicts the assumption in equation 136. Conclude that equation 127 holds for all states y.

Proof of Theorem seven

I will prove part (2). The proof of part (1) is analogous. Note that

$$\Delta_{21}\left(y, \phi\right) = \delta_{21}\left(y, \phi\right) + F\left(\phi\right) \tag{139}$$

where we have defined

$$F\left(\phi\right) = \beta \log \sum_{z \in \mathcal{X}} \pi(\cdot, z) \exp\left[\sigma_2 Q_2\left[z, \phi_z(\cdot, \phi)\right]\right] - \beta \log \sum_{z \in \mathcal{X}} \pi(\cdot, z) \exp\left[\sigma_1 Q_1\left[z, \phi_z(\cdot, \phi)\right]\right]. \tag{140}$$

Because the state is i.i.d., F is a function of only ϕ and not the exogenous state. Also the functions $\pi(\cdot, z)$ and $\phi_z(\cdot, \phi)$ do not depend on their first argument. Note well that the state is a scalar in this proof and that aggregate consumption is equal to the state.

Note the following two properties of Δ_{21} :

(i) For any fixed $\phi \in \Phi$, $\Delta_{21}(y, \phi)$ is strictly decreasing in y. This follows since $\delta_{21}(y, \phi)$ is strictly decreasing in y and F does not depend upon y.

(ii) For a fixed y, $\Delta_{21}(y, \phi)$ is strictly decreasing in ϕ_2 (where $\phi = \begin{bmatrix} 1 - \phi_2, & \phi_2 \end{bmatrix}'$.) This follows since, by Assumption four and optimality in the social planning problem, $Q_2(y, \phi)$ is strictly increasing in ϕ_2 and $Q_1(y, \phi)$ is strictly decreasing in ϕ_2.

Consider any two possible realizations of the state, a and b such that $b > a$. Let ϕ_a and ϕ_b be the optimal choices of the next period's Pareto weights when next period's state are a and b respectively. Define

$$\lambda_a = \frac{\phi_{2a}}{\phi_{1a}} \qquad \lambda_b = \frac{\phi_{2b}}{\phi_{1b}}$$

to be the ratio of the optimal choices of Pareto weights. The first-order conditions include equations of the form

$$\lambda_c = h \exp[\Delta_{21}(c, \phi_c)]$$

for $c = a$ and for $c = b$. h is a constant which is the same if $c = a$ or $c = b$.

Dividing the left and right sides of the first-order condition for b by the left and right sides of the first-order condition for a yields the necessary condition:

$$\frac{\lambda_b}{\lambda_a} = \exp[\Delta_{21}(b, \phi_b) - \Delta_{21}(a, \phi_a)]. \tag{141}$$

By property (i) above and the fact $b > a$,

$$1 > \exp[\Delta_{21}(b, \phi_a) - \Delta_{21}(a, \phi_a)]. \tag{142}$$

Hence it could not be optimal to set $\lambda_b = \lambda_a$. (If it were optimal to set $\lambda_b = \lambda_a$ then it follows from equation (141) that equation (142) would have to be satisfied with equality.) If $\lambda_b > \lambda_a$ then equation (141) could not be satisfied since the left side is greater than one and the right side is less than one since by property (ii) above

$$1 > \exp[\Delta_{21}(b, \phi_a) - \Delta_{21}(a, \phi_a)] > \exp[\Delta_{21}(b, \phi_b) - \Delta_{21}(a, \phi_a)].$$

[This follows since $\lambda_b > \lambda_a$ implies $\phi_{2b} > \phi_{2a}$ which implies $\Delta_{21}(b, \phi_b) < \Delta_{21}(b, \phi_a)$.] Contradiction. Conclude $\lambda_b < \lambda_a$ and hence $\phi_{2b} < \phi_{2a}$.

Proof of Theorem nine

(In this proof we recycle notation. In particular $\omega, \Omega, \delta, F, \Delta_t$, and p_t are used differently than in other sections.)

Since $E_t \lambda_{t+1} \leq \lambda_t$ and $\lambda_t \geq 0$, the martingale convergence theorem guarantees that $\{\lambda_t\}_{t=0}^{\infty}$ converges with probability one to a finite random variable, λ^*.

Let $\lambda_t(\omega)$ denote the sample path for the Pareto ratios for a given realization of the exogenous state for all period's $t = 0, 1, \ldots$ (here ω indexes realizations of values for the exogenous state at all dates). Let Ω be the set of all sample paths. By Egoroff's theorem, see Davidson (1994) page 283, for every $\delta > 0$ there exists a set $F(\delta) \subseteq \Omega$ with prob$[F(\delta)] \geq 1 - \delta$ such that for $\omega \in F(\delta)$

$$\lim_{t \to \infty} \lambda_t(\omega) = \lambda^*(\omega)$$

uniformly on $F(\delta)$. For every $\delta > 0$ and every $\epsilon > 0$ there exists a $T(\delta, \epsilon)$ such that for $t > T(\delta, \epsilon)$

$$|\lambda_t(\omega) - \lambda^*(\omega)| < \epsilon \tag{143}$$

for all $\omega \in F(\delta)$. For given δ and ϵ it follows that

$$|\lambda_{t+1}(\omega) - \lambda_t(\omega)| < 2\epsilon \tag{144}$$

for $t > T(\delta, \epsilon)$, for all $\omega \in F(\delta)$.

Choose some $z > 0$ and some $\delta > 0$. Let d be the number defined in hypothesis one of this theorem corresponding to value of $\underline{\lambda}$ given by z and the value of $\overline{\lambda}$ given by some finite upper bound for λ^*. Choose $\epsilon > 0$ such that $\epsilon < d/2$ and $\epsilon < z$. Define

$$S_t = \left\{ \omega \in F(\delta) : \overline{\lambda} > \lambda_t(\omega) > z \right\}, \tag{145a}$$

$$N_t = \left\{ \omega \in \Omega : \overline{\lambda} > \lambda_t(\omega) > z \right\}, \tag{145b}$$

$$\Delta_t = \left\{ \omega \in \Omega : |\lambda_{t+1} - \lambda_t| < 2\epsilon \right\}. \tag{145c}$$

Let prob$[S_t] = p_t$. By equation (144), we know that for $t > T(\delta, \epsilon)$

$$\text{prob}[N_t \text{ and } \Delta_t] \geq p_t. \tag{146}$$

56

But for $t > T(\delta, \epsilon)$ we also know

$$\text{prob}\left[N_t \text{ and } \Delta_t\right] \leq (p_t + \delta)(1 - \underline{\pi}). \tag{147}$$

This follows since $\text{prob}\left[N_t\right] \leq p_t + \delta$ and conditioned on $\lambda_t > z$ the probability that $|\lambda_{t+1} - \lambda_t| < 2\epsilon$ is less than or equal to $1 - \underline{\pi}$ (recall that $\underline{\pi}$ was defined in the statement of this theorem). Equations (146) and (147) imply that

$$p_t \leq \delta \left(\frac{1 - \underline{\pi}}{\underline{\pi}}\right). \tag{148}$$

The arguments of this paragraph can be repeated for any $\delta > 0$. Hence equation (148) holds for all $\delta > 0$. Egoroff's Theorem can again be invoked to show that $\{\lambda_t\}_{t=0}^{\infty}$ converges with probability one to a random variable which is less than or equal to z. Since the arguments of this paragraph can be repeated for all $z > 0$ it must be that $\{\lambda_t\}_{t=0}^{\infty}$ converges with probability one to zero.

Proof of lemma one

The first-order conditions for the social planner's problem when agents are risk-sensitive with the same (correct) beliefs can be written as

$$\theta_{it+1}\theta_{jt}\mathcal{D}_j(x_t, \theta_t, x_{t+1}) = \theta_{it}\theta_{jt+1}\mathcal{D}_i(x_t, \theta_t, x_{t+1}) \tag{149}$$

where

$$\mathcal{D}_i(x, \theta, y) = \frac{\exp\left[\sigma_i Q_i\left[y, \phi_y(x, \theta)\right]\right]}{\sum_{z \in \mathcal{X}} \pi(x, z)\exp\left[\sigma_i Q_i\left(z, \phi_z(x, \theta)\right)\right]}$$

for all i. Equation (48) can be obtained by dividing both sides of equation (149) by $\theta_{jt}\theta_{jt+1}$ and taking expected values conditioned on t of both sides of the resulting equation. Now use the identity that for any random variables y and z it is the case that $E(yz) = \text{cov}(y, z) + E(y)E(z)$ and rearrange terms.

Proof of Theorem 10

Lets begin with part (1a) of the theorem. Let this period's Pareto weight vector θ be such that $1 > \theta_2 > \bar{\theta}_2$. Since $\bar{\theta}$ is a regular steady state next period's Pareto weight vector ϕ will be such that $\phi_{2y} > \bar{\theta}_2$ for any value of next period's state, y. The proof will rely on the following three facts:

(i) The optimal choice of next period's Pareto ratio ϕ_{2y}/ϕ_{1y} is a strictly decreasing function of y. (This follows from Theorem seven and the assumption that $\delta_{21}(y, \phi)$ is a strictly decreasing function of y.)

(ii) For a fixed ϕ, $Q_1(y, \phi)$ is a strictly increasing function of y. Since $\sigma_1 \leq 0$ it follows that either $\exp\left[\sigma_1 Q_1(y, \phi)\right]$ is a strictly decreasing function of y (when $\sigma_1 < 0$) or it is always equal to one (when $\sigma_1 = 0$).

(iii) For a fixed y, $Q_1(y, \phi)$ is a strictly decreasing function of ϕ_2. Since $\sigma_1 \leq 0$ it follows that either $\exp\left[\sigma_1 Q_1(y, \phi)\right]$ is a strictly increasing function of ϕ_2 (when $\sigma_1 < 0$) or it is always equal to one (when $\sigma_1 = 0$).

It follows from (i) thru (iii) that either $\exp\left[Q_1(y, \phi_y)\right]$ is a strictly decreasing of y (when $\sigma_1 < 0$) or it is always equal to one (when $\sigma_1 = 0$). (Here $Q_1(y, \phi_y)$ is evaluated at the optimal choice of Pareto weights ϕ_y when next period's state is y.) Since part (i) showed that ϕ_{2y}/ϕ_{1y} is a strictly decreasing function of y it must be the case that

$$\frac{\text{cov}_t\left(e^{\sigma_1 Q_1\left(x_{t+1}, \theta_{t+1}\right)}, \frac{\theta_{2t+1}}{\theta_{1t+1}}\right)}{E_t\left[e^{\sigma_1 Q_1\left(x_{t+1}, \theta_{t+1}\right)}\right]} \geq 0 \tag{150}$$

for all x_t and all finite $\frac{\theta_{2t}}{\theta_{1t}} > \bar{\theta}_2/\bar{\theta}_1$.

Now we have from Lemma one that the Pareto ratios are a non-negative martingale. Paralleling the proof of Theorem 11 we can define suitable d and $\underline{\pi}$ so that Theorem nine can be invoked to show that the ratio of agent two's Pareto weight to agent one's Pareto weight converges to $\bar{\theta}_2/\bar{\theta}_1$. (The theorem can be used to show that $\theta_{2t}/\theta_{1t} - \bar{\theta}_2/\bar{\theta}_1$ converges to zero.) It is possible to find a d and $\underline{\pi}$ since there are no other interior regular steady states.

The proof of part (1b) when $0 < \theta_2 < \bar{\theta}_2$ is analogous and can be formulated by switching the roles of the agents in the proof of (1a) above.

I now sketch the proof of part (2b) when $0 < \theta_2 < \bar{\theta}_2$. In this case using arguments similar to the proof above for case (1a) it can be shown that

$$\frac{\text{cov}_t\left(e^{\sigma_1 Q_1\left(x_{t+1}, \theta_{t+1}\right)}, \frac{\theta_{2t+1}}{\theta_{1t+1}}\right)}{E_t\left[e^{\sigma_1 Q_1\left(x_{t+1}, \theta_{t+1}\right)}\right]} \geq 0 \tag{151}$$

and

$$E_t\left[\frac{\theta_{2t+1}}{\theta_{1t+1}}\right] \leq \frac{\theta_{2t}}{\theta_{1t}}. \tag{152}$$

This suggests that the Pareto ratio (of agent two's Pareto weight to agent one's Pareto weight) converges to zero. However we do not have enough to guarantee that the Pareto ratio will converge to zero with probability one. There is another

steady state, $\bar{\theta}$, which is interior. That prevents us from finding a d and $\underline{\pi}$ which satisfy the conditions of Theorem nine, because in any interval which includes the interior steady state d will be zero.

However, $\{\theta_{2t}/\theta_{1t}\}_{t=0}^{\infty}$ is bounded below by zero and bounded above by $\bar{\theta}_2/\bar{\theta}_1$. Equation 151 and the martingale convergence theorem guarantee that the Pareto ratios converge to a finite random variable λ^* with probability one. It is straightforward (the formal argument is similar to the argument in Theorem nine) to show that λ^*, with probability one, can take on at most two values, zero and $\bar{\theta}_2/\bar{\theta}_1$. It must be that

$$\text{prob}_t \left[\lambda^* = 0\right] \geq 1 - \frac{\theta_{2t}/\theta_{1t}}{\bar{\theta}_2/\bar{\theta}_1} > 0.$$

(This follows from equation (152).) This provides a strictly positive lower bound for the probability that the Pareto ratios converge to zero. Since the Pareto ratios could converge to $\bar{\theta}_2/\bar{\theta}_1$ with probability of at most $\frac{\theta_{2t}/\theta_{1t}}{\bar{\theta}_2/\bar{\theta}_1}$ (which is strictly less than one unless this period's Pareto ratio is $\bar{\theta}_2/\bar{\theta}_1$), the steady state Pareto weight vector $\bar{\theta}$ is not stable.

The proof of part (1a) when $1 > \theta_2 > \bar{\theta}_2$ is analogous and can be formulated by switching the roles of the agents in the proof of (2b) above.

Proof of Theorem 11

We already know that any two time-additive agents are always in a private steady state in which their Pareto ratios are constant over time. Let agent j have time-additive preferences with $\sigma_j = 0$ and agent i have risk-sensitive preferences with $\sigma_i < 0$. Recall the first-order condition from Section 3.4 when this period's exogenous state is x and this period's Pareto weight vector is θ :

$$\frac{\phi_{iy}(x,\theta)}{\phi_{jy}(x,\theta)} = \left(\frac{\theta_i}{\theta_j}\right) \frac{\exp(\sigma_i q_{iy})}{\sum_{z \in \mathcal{X}} \pi(x,z) \exp(\sigma_i q_{iz})} \tag{153}$$

where $q_{iy} = Q_i\left[y, \phi_y(x,\theta)\right]$. For $\underline{\lambda}$ and $\overline{\lambda}$ such that $+\infty > \overline{\lambda} > \underline{\lambda} > 0$ set

$$d = \min_{\underline{\lambda} \leq \frac{\theta_i}{\theta_j} \leq \overline{\lambda}} \ \min_{x \in \mathcal{X}} \ \max_{y \in \mathcal{X}} \ \left| \frac{\theta_i}{\theta_j} - \left(\frac{\theta_i}{\theta_j}\right) \frac{\exp(\sigma_i q_{iy})}{\sum_{z \in \mathcal{X}} \pi(x,z) \exp(\sigma_i q_{iz})} \right|. \tag{154}$$

Because of Assumption nine and the fact that there are no interior steady state Pareto ratios for agents i and j it must be that $d > 0$. The minimum is obtained since the function minimized is continuous (in θ_i/θ_j) and the constraint set is compact.

Let

$$\underline{\pi} = \min_{x,y \in \mathcal{X}} \pi(x,y)$$

be the smallest transition probability between any two states. By Assumption eight this probability must be strictly greater than zero. Thus the probability that the difference on the right hand side of equation 154 is at least d is greater than or equal to $\underline{\pi}$.

Now since, as explained in the text, the ratio of agent i's Pareto weight to agent j's Pareto weight is a martingale which is bounded below by zero, Theorem nine can be invoked to show that the ratio converges to zero with probability one.

Proof of Theorem 12

First note that the one-period consumption allocation rules are

$$c_i(e,\theta) = s_i e$$

where

$$s_i = \frac{\theta_i^{\frac{1}{\gamma}}}{\sum_{j=1}^{n} \theta_j^{\frac{1}{\gamma}}}.$$

For

$$\sigma_i u_i\left[c_i(e,\theta)\right] - \sigma_j u_j\left[c_j(e,\theta)\right] = \left(\sigma_i s_i^{1-\gamma} - \sigma_j s_j^{1-\gamma}\right) \frac{e^{1-\gamma}}{1-\gamma}$$

to be constant for any e it is necessary that the term in parenthesis on the right hand side be zero. This will only happen if

$$\frac{s_i}{s_j} = \left(\frac{\sigma_j}{\sigma_i}\right)^{\frac{1}{1-\gamma}}$$

which is equivalent to requiring that $\theta = \bar{\theta}$. By Theorem two $\bar{\theta}$ is the unique interior regular steady state.

Proof of Theorem 13

For this case
$$\delta_{21}\left(y,\phi\right)=\frac{y^{1-\gamma}Z(\phi)}{1-\gamma}.$$
where
$$Z(\phi)=\left(\frac{1}{\sum_{j=1}^{2}\phi_{j}^{\frac{1}{\gamma}}}\right)^{1-\gamma}\left(\sigma_{2}\phi_{2}^{\frac{1-\gamma}{\gamma}}-\sigma_{1}\phi_{1}^{\frac{1-\gamma}{\gamma}}\right).$$

First lets consider the case in which $0<\gamma<1$ and let this period's Pareto weight vector, θ, be such that $\theta_2 > \bar{\theta}_2$. From the no-crossing property next period's Pareto weight vector ϕ must be such that $\phi_2 > \bar{\theta}_2$ which implies $Z(\phi) < 0$ and δ_{21} is strictly decreasing in y. By symmetry when $\theta_2 < \bar{\theta}_2$ it must be the case that $Z(\phi) > 0$ and δ_{21} is strictly increasing in y. Now Theorem seven and Theorem 10 can be applied to the show the results in the theorem for when $0 < \gamma < 1$.

Now lets consider the case in which $\gamma > 1$ and let this period's Pareto weight vector, θ, be such that $\theta_2 > \bar{\theta}_2$. From the no-crossing property next period's Pareto weight vector ϕ must be such that $\phi_2 > \bar{\theta}_2$ which implies $Z(\phi) > 0$ and δ_{21} is strictly increasing in y. By symmetry when $\theta_2 < \bar{\theta}_2$ it must be the case that $Z(\phi) < 0$ and δ_{21} is strictly decreasing in y. Now Theorem seven and Theorem 10 can be applied to the show the results in the theorem for when $\gamma > 1$.

References

Alvarez, F. and U. J. Jermann (2000). Efficiency, equilibrium, and asset pricing with risk of default. *Econometrica 68*, 775–797.

Anderson, E. W. (1998). *Uncertainty and the Dynamics of Pareto Optimal Allocations*. Ph. D. thesis, University of Chicago, Chicago, IL.

Atkeson, A. and R. E. Lucas, Jr. (1995). Efficiency and equality in a simple model of unemployment insurance. *Journal of Economic Theory 66*(1), 64–88.

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

Benhabib, J., S. Jafarey, and K. Nishimura (1988). The dynamics of efficient intertemporal allocations with many agents, recursive preferences, and production. *Journal of Economic Theory 44*, 301–320.

Blume, L. and D. Easley (1992). Evolution and market behavior. *Journal of Economic Theory 9*(40), 9–40.

Browning, M., L. P. Hansen, and J. J. Heckman (2000). Micro data and general equilibrium models. In J. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*. Amsterdam: North Holland.

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

Davidson, J. (1994). *Stochastic Limit Theory*. Oxford: Oxford.

Doob, J. (1953). *Stochastic Processes*. New York: Wiley.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica 60*(2), 353–394.

Duffie, D., P.-Y. Geoffard, and C. Skiadas (1994). Efficient and equilibrium allocations with stochastic differential utility. *Journal of Mathematical Economics 23*, 133–146.

Dumas, B. (1989). Two-person dynamic equilibrium in the capital market. *Review of Financial Studies 2*, 157–188.

Dumas, B., R. Uppal, and T. Wang (2000). Efficient intertemporal allocations with recursive utility. *Journal of Economic Theory 99*, 240–259.

Epstein, L. (2001). Sharing ambiguity. *American Economic Review 91*(2), 45–50.

Epstein, L. and J. Miao (2003). A two-person dynamic equilibrium under ambiguity. *Journal of Economic Dynamics and Control 27*, 1253–1288.

Epstein, L. G. (1987). The global stability of efficient intertemporal allocations. *Econometrica 55*, 329–355.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–969.

Hansen, L. P. (1987). Calculating asset prices in three exchange economies. In *Advances in Econometrics, Fifth World Congress*. Cambridge: Cambridge University Press.

Hansen, L. P. and T. J. Sargent (1995). Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic Control 40*(5), 968–971.

Hansen, L. P. and T. J. Sargent (1997). Recursive linear models of dynamic economies. Manuscript.

Hansen, L. P. and T. J. Sargent (2003). Robust control of forward-looking models. *Journal of Monetary Economics 50*, 581–604.

Hansen, L. P., T. J. Sargent, and T. D. Tallarini, Jr. (1999). Robust permanent income and pricing. *Review of Economic Studies 66*(4), 873–907.

Hayashi, F., J. Altonji, and L. Kotlikoff (1996). Risk-sharing between and within families. *Econometrica 64*, 261–294.

Heaton, J. and D. Lucas (1995). The importance of investor heterogeneity and financial market imperfections for the behavior of asset prices. *Carnegie-Rochester Series on Public Policy 42*, 1–32.

Jacobson, D. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions on Automatic Control AC-18*(2), 124–131.

Kan, R. (1995). Structure of Pareto optima when agents have stochastic recursive preferences. *Journal of Economic Theory 66*(2), 626–631.

Liu, W.-F. (2002). Heterogeneous agent economies with knightian uncertainty. Manuscript, University of Washington.

Lucas, Jr., R. E. and N. L. Stokey (1984). Optimal growth with many consumers. *Journal of Economic Theory 32*, 139–171.

Ma, C. (1993). Market equilibrium with heterogeneous recursive-utility-maximizing agents. *Economic Theory 3*(2), 243–266.

Ozaki, H. and P. A. Streufert (1996). Dynamic programming for non-additive stochastic objectives. *Journal of Mathematical Economics 25*, 391–442.

Phelan, C. and R. M. Townsend (1991). Computing multi-period, information-constrained optima. *Review of Economic Studies 58*(5), 853–881.

Sandroni, A. (2000). Do markets favor agents able to make accurate predictions? *Econometrica 68*, 1303–1341.

Schroder, M. and C. Skiadas (1999). Optimal consumption and portfolio selection with stochastic differential utility. *Journal of Economic Theory 68*, 68–126.

Tallarini, Jr., T. D. (2000). Risk-sensitive real business cycles. *Journal of Monetary Economics 45*(3), 507–532.