Incentives, Contracts and Markets:

A General Equilibrium Theory of Firms*

William R. Zame

University of California, Los Angeles

zame@econ.ucla.edu

First version: July 13, 2005
This version: October 15, 2006

*I am deeply indebted to Bryan Ellickson for many stimulating conversations, to John Nachbar, Max Stinchcombe and Jeroen Swinkels for encouragement and pointed questions, to Jernej Copic, John Riley, Chris Shannon and Leo Simon for helpful conversations, to seminar audiences at Caltech, UC Berkeley, Universitat Pompeu Fabra, Barcelona JOCS, UC San Diego, the Banff International Research Station, Cornell University, NYU, Princeton University, Northwestern University, the University of Chicago, the Southwest Economic Theory Conference, the University of Arizona, USC, UC Davis, and three referees and the editor for comments. I am grateful to the Division of Humanities and Social Sciences at Caltech for sabbatical hospitality, and to the John Simon Guggenheim Foundation, the National Science Foundation (under grants SES-0317752 and SES-0518936), the Caltech Social and Information Sciences Laboratory, and the UCLA Academic Senate Committee on Research for financial support. Views expressed here are those of the author and do not necessarily reflect the views of any funding agency.

Incentives, Contracts and Markets: Toward A General Equilibrium Theory of Firms

Abstract

This paper takes steps toward integrating firm theory in the spirit of Alchian & Demsetz (1972) and Grossman & Hart (1986), contract theory in the spirit of Holmstrom (1979), and general equilibrium theory in the spirit of Arrow & Debreu (1954) and McKenzie (1959). In the model presented here, the set of firms that form and the contractual arrangements that appear, the assignments of agents to firms, the prices faced by firms for inputs and outputs, and the incentives to agents are all determined endogenously at equilibrium. Agents choose consumption — but they also choose which firms to join, which roles to occupy in those firms, and which actions to take in those roles. Agents interact anonymously with the (large) market, but strategically within the (small) firms they join. The model accommodates moral hazard, adverse selection, signaling and insurance. Equilibria may be Pareto ranked.

Keywords general equilibrium, incentives, contracts, firms, organizations, teams

JEL Classification Numbers D2, D5, D71, D8, L2

1 Introduction

Incentives and contracts are at the heart of much of the modern theory of the firm. Alchian & Demsetz (1972) focuses on the implications of monitoring and incentives for the structure of the firm. Holmstrom (1979) formulates the incentive problem more broadly as a principal-agent problem, and explores the extent to which contracting on observable/verifiable events can or cannot overcome incentive problems. Grossman & Hart (1986) focuses on the incompleteness of contracts and the consequent role of control rights. These papers, and the enormous literature they have inspired — much too vast to survey here — have provided a great deal of insight into the nature of the firm. However, this literature largely ignores the interaction between the firm and the market, taking as given the set of firms that form, the organizational structure of these firms, the assignments of agents to firms, the prices faced by firms for inputs and outputs, and the incentives for agents to take particular actions within a firm or even to participate in a firm at all.

The present paper begins an integration of firm theory in the spirit of Alchian & Demsetz (1972) and Grossman & Hart (1986), contract theory in the spirit of Holmstrom (1979), and general equilibrium theory in the spirit of Arrow & Debreu (1954) and McKenzie (1959). In the model presented here, the set of firms that form and the contractual arrangements that appear, the assignments of agents to firms, the prices faced by firms for inputs and outputs, and the incentives to agents are all determined endogenously at equilibrium. Agents choose consumption — but they also choose which firms to join, which roles to occupy in those firms, and which actions to take in those roles. Agents interact anonymously with the (large) market, but strategically within the (small) firms they join. Firms and the market are inextricably tied together: strategic choices within firms affect output, output affects market prices, market prices affect agents' budget sets and hence utilities, utilities affect incentives within firms, and incentives within firms affect strategic choices within firms. At equilibrium, competitive forces determine the prices of goods and which agents consume those goods, the firms that form and the wages and contractual arrangements that prevail, the assignment of agents to jobs and the efforts exerted by those agents — in short, every detail of the economy.

The central notions of the model are a *technology*, a *contract*, and a *firm type*. A technology consists of a set of roles (jobs), a space of actions for each role, a space of skills that may be possessed by individuals who fill each role, a set of verifiable consequences, and a stochastic production process. A contract is a system of payments to the members of the firm that sum to the net value of output. A firm type is a pair consisting of a technology and a contract. The set of possible firm types is given exogenously; one might view possible technologies as reflecting scientific possibilities and possible contracts as reflecting legal enforceability; alternatively, technologies and contracts might be the outcomes of a design process. Many firm types — involving many different technologies and/or many different contracts — might be potentially available — but many fewer firm types might actually be realized at equilibrium. Thus, the observed set of firm types — both technologies and contracts — is determined *endogenously* at equilibrium.[1,2]

Agents who choose to join a firm (or several firms) are exposed to risk: from the exogenously specified stochastic nature of the production process (or processes) and from the endogenously determined behavior of other agents with whom they are matched in the various firms. The model treats these risks as idiosyncratic, in the sense that the realizations of stochastic uncertainty across firms and the matching of agents within firms are assumed independent and random, and assumes that the Law of Large Numbers applies. Hence aggregate consumption and production are deterministic, rather than stochastic. Individuals cope with idiosyncratic risk by choosing consumption plans that are contingent on the realizations of the individual risks to

[1] This approach is parallel to that taken in Dubey & Geanakoplos (2002), Dubey, Geanakoplos & Shubik (2005) and Geanakoplos & Zame (2002), in which the set of possible financial instruments is viewed as given exogenously but the set of observed financial instruments is determined at equilibrium.

[2] Because the set of firm types might encompass a fine grid approximation to all possible firm types, the restriction to a finite set is not terribly restrictive. Allowing for an infinite set of firm types would lead to many technical and conceptual difficulties without adding much insight.

which they are exposed, but there are no spot markets for idiosyncratic risk. This does not mean that idiosyncratic risk is not insurable — only that such insurance, if offered, is offered within the structure of the various firm types.

A *common-beliefs equilibrium* of the model consists of commodity prices, wages (prices for roles/jobs), and, for each agent, consumption plans, firm-role-action choices, and beliefs (about the actions of others) such that: commodity markets clear, job markets clear, agents optimize in their budget sets given beliefs, and beliefs are correct and identical across agents. Beliefs about the actions of others enter into equilibrium because agents' wealth and utility depend on the actions of others. For active firm types (i.e., firm types that are formed), requiring that agents optimize given beliefs and that beliefs be correct simply amounts to requiring (as in Nash equilibrium) that agents' actions are optimal given the actions of others. However, if agents do not take the supply of jobs into account in determining their employment decisions — just as agents in the standard Arrow–Debreu–McKenzie consumption/exchange model do not take the supply of commodities into account in determining their own consumption decisions, and agents in the familiar security markets models do not take the supply of securities into account in determining their own portfolio choices — then beliefs about the actions of others in inactive firm types (i.e., firm types that are not formed) are an essential part of equilibrium, because they enter individual decisions *not* to choose roles in those firm types.

For active firm types, correctness of beliefs implies commonality. For inactive firm types, correctness of beliefs by itself has no bite. Requiring common beliefs imposes some discipline on individual beliefs but it does not rule out the possibility that agents beliefs are common but absurd — for instance, that others use dominated strategies. A refinement, *population perfect equilibrium* — an analog of trembling-hand perfect equilibrium — does rule out (some) absurd beliefs. Under natural assumptions, both common-beliefs equilibrium and population perfect equilibrium always exist.

Throughout, I use language (firms, jobs, workers, etc.) that is suggestive of the main interpretation, and most of the emphasis is on skills and actions that affect output, but the formal structure accommodates other interpretations as well. In particular, clubs, in the sense of Buchanan (1965) or Ellickson, Grodal, Scotchmer and Zame (1999; 2001), and groups in the sense of Ellickson, Grodal, Scotchmer & Zame (2006) are firms. The central characteristic of the firms treated here is that they are small (formally: they are finite sets in a continuum economy). Thus the treatment here excludes large publicly-owned firms (Microsoft) and political jurisdictions in the sense of Tiebout (1956).

A number of examples illustrate the main themes. (More realistic applications are intended for later papers.) Example 1 focuses on moral hazard in the spirit of Holmstrom (1979), and illustrates the interaction between market prices and incentives within firms. Example 2 focuses on adverse selection in the spirit of Akerlof (1970), and illustrates the interaction between market prices and the incentives to join firms. Example 3 illustrates how contracts screen agents. Example 4 illustrates the provision of insurance and market choice of contracts. Examples 5, 6 and illustrate the possibilities of robust mis-coordination and Pareto ranked equilibria; Theorem 2 shows that Pareto optimality prevails in a restricted — but important — context. Example 8, which contrasts with Bennardo & Chiappori (2003), illustrates the implications of population perfection and how the Harsanyi (1973) notion of perturbed payoffs explains randomization (in the population).

There is by now a substantial literature that seeks to incorporate asymmetric information (including moral hazard and adverse selection) in a general equilibrium context. The seminal work in this literature is Prescott & Townsend (1984a, 1984b); more recent contributions include Prescott & Townsend (2001) and Rustichini & Siconolfi (2002). Because the present paper follows a very different approach from this literature, it may be useful to note explicitly some of the differences from Prescott & Townsend:

- In Prescott & Townsend, production is either in large firms that make zero profit or by single households. In either case, there are no issues of matching or coordination and hence no possibility of mis-coordination. In this paper, production takes place in teams, in the sense of Marshack & Radner (1972). In these teams, each agent's wealth and utility may depend on his/her own actions and on the actions of others in the team(s) to which the agent belongs. Matching and coordination play an important role, and mis-coordination is possible.

- In Prescott & Townsend, agents choose lotteries over physical bundles, and prices are linear in probabilities, but need not be linear in physical consumption. In this paper, agents choose physical bundles and prices are linear in physical consumption.
- In Prescott & Townsend, agents choose lotteries over physical bundles but are not able to trade in commodity spot markets contingent on the realizations of these lotteries. In this paper, agents choose physical bundles contingent on the realization of uncertainty that they experience as a consequence of their memberships in firms, but cannot trade contracts contingent on the realization of that uncertainty.
- In Prescott & Townsend, contracts are exclusive (in particular, agents can belong to at most one firm) and specify (lotteries over) consumption. In this paper, contracts are non-exclusive (in particular, agents can belong to many firms) and specify income; consumption is chosen at market prices.

A different literature, including Ghosal & Polemarchakis (1997), Minelli & Polemarchakis (2000), Bisin, Geanakoplos, Gottardi, Minelli & Polemarchakis (forthcoming), Dubey & Geanakoplos (2002) and Dubey, Geanakoplos & Shubik (2005), also seeks to integrate trade in commodities and strategic interactions, and to accommodate moral hazard and adverse selection. Central to this literature is the assumption that deliveries are pooled and that the effects of individual actions are felt economy-wide. This seems a reasonable description of the market for mortgage pools (in which agents do not hold individual mortgages, but rather shares in large pools of mortgages) and perhaps of the interactions of agents in enormous public firms, but not of the interaction between a farmer and her hired hand (the farmer may employ a hired hand who is competent and works hard or one who is incompetent and shirks — but not an average of all hired hands) or more generally of the interactions of agents in small private firms, which are the thrust of the present paper.

Closer in spirit, although more limited in scope, is McAfee (1993), which treats an environment in which the competitive market determines the choice of selling mechanisms.

The framework of the present paper builds on Makowski (1976) and Ellickson, Grodal, Scotchmer & Zame (1999, 2001, 2006). As does the present paper, these papers construct frameworks in which agents interact with the market and with each other in small groups, determined endogenously in equilibrium.[3] However, these papers do not allow for hidden actions (moral hazard) or unobserved skills (adverse selection) or stochastic production. The focus in those papers, and in Makowski & Ostroy (2004), which uses a similar framework, is on competition and efficiency when agents consume and/or produce in small groups and actions are contractible. Rahman (2005) and Song (2006) expand this framework to accommodate moral hazard (but not adverse selection), but in a different direction, and with a different purpose, than here. Those papers focus on the possibility of decentralizing solutions to the planner's problem (incentive efficient allocations); they show that such decentralization requires coordinating agents (entrepreneurs), coordinated lotteries, exclusive contracts, and personalized (Lindahl) prices.

Following this Introduction, Section 2 uses Example 1 to motivate the model and illustrate some of the ideas and implications. Section 3 presents the bones of the formal model. Section 4 defines common beliefs equilibrium and asserts existence (Theorem 1). Section 5 illustrates how the model accommodates adverse selection (Example 2), screening (Example 3) and endogenous contracting and the provision of insurance (Example 4). Section 6 addresses welfare: mis-coordination can be a robust phenomenon and can lead to Pareto ranked equilibrium (Examples 5, 6, 7), but equilibria are Pareto optimal under some conditions (Theorem 2). Section 7 defines population perfect equilibrium, asserts existence (Theorem 3) and illustrates (Example 8) its implications. Section 8 concludes with a discussion of some of the issues involved in adapting the static model presented here to a dynamic environment. Proofs are collected in the Appendix.

[3] Cole & Prescott (1997) takes a rather different approach to the same problem, more in the spirit of Prescott & Townsend, and is considerably further from the present work.

2 A Motivating Example

To introduce and motivate the model, this Section presents an informal description and analysis of a simple moral hazard example in the spirit of Holmstrom (1979) to illustrate how trade with the market affects incentives within a firm.

Example 1 (Moral Hazard) Consider a world with two goods and a single kind of productive activity. Production requires the participation of two agents. The input to the production process is 2 units of the first good; the outcome of the production process depends on the effort exerted by the members of the firm, as shown in Figure 1 below.

	H	L
H	450	200
L	200	50

Figure 1

All agents are identical. Endowments are $e = (2, 0)$. Utility for consumption is Cobb-Douglas $u(c_1, c_2) = (c_1 c_2)^{1/2}$. Exerting High effort is unpleasant: an agent who exerts High effort experiences disutility of $d = 6$. Finally, output is observable and contractible but effort is not, so each member of the firm provides half the input and receives half the output.

As a benchmark, consider first the setting in which there are only two agents. Combining the assumption that ROW and COL each provide half the input and receive half the output with the assumptions about endowments and utility functions leads to Figure 2, which expresses the utility of ROW as a function of the choices of ROW and COL.

	H	L
H	9	4
L	10	5

Figure 2

(Utility for COL is obtained by interchanging roles.) Because $10 > 9$ and $5 > 4$, exerting Low effort is a dominant strategy. Because the situation for COL is symmetrical, we conclude that both agents will choose to exert Low effort, and obtain utility $\overline{u} = 5$. (Note that this exceeds the utility of consuming the endowment.)

Now suppose there are *many* agents, each of whom is free to enter or not enter a firm, to choose an effort level and *also* to trade with other agents in an anonymous competitive market. I claim that it *cannot* be an equilibrium for every agent to enter a firm and exert Low effort. To see this, notice that if that were an equilibrium then each agent's consumption would be $(1, 25)$, whence market prices would be $(1, 1/25)$. However, at those prices, an agent who deviated to High effort and obtained output of 100 units of good 2 could *trade* with the market at these prices. The optimal choice for such an agent would be $(5/2, 125/2)$; taking the utility penalty $d = 6$ into account, such an agent would obtain net utility $(25/2) - 6 = 13/2$, which is greater than 5. Thus, deviation from Low effort would be profitable, so it cannot be an equilibrium for all agents to enter a firm and exert Low effort, as claimed.

Assuming that agents are free to enter or not enter a firm, we can compute the market equilibrium. We have already seen that it cannot be an equilibrium for all agents to enter a firm and exert High effort and that it cannot be an equilibrium for all agents to enter a firm and exert Low effort. Moreover, it cannot be an equilibrium for no agents to work (enter a firm).[4] Assuming that jobs and actions in the firm are indivisible, it follows that equilibrium must be mixed in the population; that is, equilibrium must have the property that *ex ante* identical agents make different choices.

To solve for equilibrium it is convenient to analyze each of the possible configurations separately. To begin, look for an equilibrium in which some agents enter a firm and exert High effort and the remaining agents do not work. Write α for the fraction of agents who enter a firm and exert High effort, and normalize so that commodity prices are $(1, q)$. To determine α and q, equate demand and supply for good 1 and then equate utility for

[4] If no agents enter a firm then every agent would want a job in a firm, so the demand for jobs would exceed the supply of jobs.

agents who enter a firm and agents who do not:

$$\begin{aligned}
\alpha\left(\frac{1+225q}{2}\right)+(1-\alpha)(1) &= (\alpha)(1)+(1-\alpha)(2) \\
\left[\left(\frac{1+225q}{2}\right)\left(\frac{1+225q}{2q}\right)\right]^{1/2}-6 &= \left[\frac{1}{q}\right]^{1/2}
\end{aligned}$$

Solving these equations yields the equilibrium price of good 2 and the fraction of agents who enter the firm; approximately we have

$$q \approx .01\ ,\ \ \alpha \approx .60$$

To see that this is an equilibrium, it remains to show that agents who enter the firm and exert High effort would not prefer to deviate and exert Low effort; this will be true exactly when

$$\left[\left(\frac{1+225q}{2}\right)\left(\frac{1+225q}{2}\right)\right]^{1/2}-6 \ \geq\ \left[\left(\frac{1+100q}{2}\right)\left(\frac{1+100q}{2}\right)\right]^{1/2}$$

which is easily verified. Hence this is an equilibrium. Note that all agents obtain the same utility, but that agents who enter the firm and exert High effort consume more of each commodity than agents who do not work, as compensation for bearing the disutility of effort.

I leave it to the reader to show that this is the unique equilibrium. $\diamondsuit$

3 The Formal Model

Example 1 embodies many of the features of the general model formalized below, but several important features are missing. The general model allows for the possibility that production depends stochastically on both the actions of agents and their skills (characteristics), and allows for *many* potential types of firms, embodying many different technologies and contracts. These features, together with the features already present in Example 1, allow the model to capture moral hazard, adverse selection, screening, and endogenous contracting.

The data of the model consists of a finite set of perfectly divisible commodities (private goods), a finite set of firm types, and a space of agents.

3.1 Commodities

There are $L \geq 1$ perfectly divisible commodities (private goods), traded on competitive markets, so the commodity space is $\mathbb{R}^L_+$. Commodity prices are required to be strictly positive so the price space is $\mathbb{R}^L_{++}$.

3.2 Firm Types

A firm type is a pair consisting of a technology and a contract. In turn, a technology consists of a finite set of roles (jobs), compact metric spaces of skills and actions for each role, a finite set of observable/contractible (firm-specific) consequences, a specification of inputs and outputs in each consequence, a consequence-dependent distribution of income, and a stochastic mapping from actions and skills to consequences; a contract is a specification of the consequence-dependnent distribution of income.

Formally, a *technology* is a tuple $T = (R, \{S_r\}, \{A_r\}, \Omega, y, \pi)$ where :

- R is a finite set of *roles* (or *jobs*)

- for each role $r \in R$, S_r is a compact metric space of *skills* and A_r is a compact metric space of *actions*
- Ω is a finite set of *consequences*
- $y : \Omega \to \mathbb{R}^L$ is the *input/output mapping*
- π is a continuous family of *conditional probabilities*

$$\pi : (S_1 \times A_1) \times \ldots \times (S_R \times A_R) \to \mathbf{P}(\Omega)$$

  (where $\mathbf{P}(\Omega)$ is the space of probability measures on Ω)

The interpretation intended is that the consequence ω summarizes everything that is observable or verifiable to an outside agency — hence what is contractible — and that $\pi(\omega|\mathbf{s}, \mathbf{a})$ is the conditional probability that ω will occur if the various roles are filled by individuals whose skill profile is $\mathbf{s} = (s_1, \ldots, s_R)$ and who choose action profile $\mathbf{a} = (a_1, \ldots, a_R)$.

Given a particular technology, a *contract* is a bounded, continuous map

$$C : R \times \Omega \times \mathbb{R}^L_{++} \to \mathbb{R}$$

that is homogeneous of degree 1 in prices and has the property that

$$\sum_r C(r, \omega, p) = p \cdot y(\omega)$$

for each $\omega \in \Omega$, $p \in \mathbb{R}^L_{++}$.[5] $C(r, \omega, p)$ is the payment to the agent filling role r if the realized consequence is ω and market prices for commodities are p; note that payments $C(r, \omega, p)$ may be positive, negative, or zero. The requirement that payments sum to net income of the firm (i.e., that the budget balances) is just an accounting identity.

A *firm type* consists of a technology and a contract. A firm type delineates an array of possibilities for a set of agents; a *firm* is particular instance of a firm type, arising from the choices of individual agents.

[5] Homogeneity of degree 1 guarantees that nothing depends in an essential way on the price normalization.

There are $J \geq 0$ exogenously specified firm types; I use superscripts for the parameters in firm type j. It is frequently convenient to abuse notation and view J as either the number of firms or as the set of firms, to view as R^j as the either the number of roles in firm type j or the set of roles in firm type j, and so forth.[6]

3.3 Agents

An agent is described as a tuple consisting of a choice set, an endowment, a utility function, and an array of skills.

3.3.1 Choices

Each element of the choice set specifies a consequence-contingent consumption plan, the firm types to which the agent chooses to belong, the role (job) chosen in each firm, and the action chosen in each role.

Although agents choose consumption contingent on the realization of the uncertainty they face, it is convenient to adopt notation in which consumption is contingent on the realization of all uncertainty and later constrain the choice to be independent of the realization of uncertainty that is not faced by the agent. Hence, write

$$\Omega = \Omega^1 \times \ldots \times \Omega^J$$

and define a *consumption plan* to be a random variable

$$\tilde{x} : \Omega \to \mathbb{R}^L_+$$

[6]By definition, each role in a firm is filled by a single agent and each individual agent may choose to belong to at most one firm of each type — but perhaps to many different firms of different types. These modeling choices involve no loss of generality because otherwise identical roles or identical firm types could always be distinguished by giving them different names. In addition to being a bit simpler than the alternatives, these modeling choices have the advantage that they allow for the possibility that names can serve as coordinating devices; see Example 3 in Section 5, for instance.

(I follow the familiar convention of using tildes to denote random choices and Roman letters to denote non-random choices.) Write $\mathbb{R}_+^{L\Omega}$ for the space of all consumption plans.

To formalize firm-role-action choices, it is convenient to define a dummy role-action choice in each firm, so write:

$$\begin{aligned}
\mathbb{F}_0^j &= \{(r,a) : 1 \le r \le R^j, a \in A_r^j\} \\
\mathbb{F}^j &= \mathbb{F}_0^j \cup \{0\} \\
\mathbb{F}_0 &= \mathbb{F}_0^1 \times \ldots \times \mathbb{F}_0^J \\
\mathbb{F} &= \mathbb{F}^1 \times \ldots \times \mathbb{F}^J
\end{aligned}$$

An element $\phi = (\phi^1, \ldots, \phi^J) \in \mathbb{F}$ is a tuple of role-action choices in each firm type; by convention, $\phi^j = 0$ represents the choice *not to belong* to firm type j. For each j, define

$$\begin{aligned}
\rho^j &: \mathbb{F}^j \to R^1 \cup \ldots \cup R^J \cup \{0\} \\
\alpha^j &: \mathbb{F}^j \to A_1^j \cup \ldots \cup A_{R^j}^j \cup \{0\}
\end{aligned}$$

by $\phi^j = (\rho^j(\phi), \alpha^j(\phi))$. (I abuse notation in the obvious way so that $\rho^j(\phi) = \alpha^j(\phi) = 0$ if $\phi^j = 0$.)

For $\Phi \subset \mathbb{F}$ a non-empty set of firm-role-action choices, define the corresponding *choice set* to be:

$$X(\Phi) = \{(\tilde{x}, \phi) \in \mathbb{R}_+^{L\Omega} \times \Phi : \phi^j = 0 \Rightarrow \tilde{x} \text{ is independent of } \omega^j\}$$

Defining consumption sets in this way carries out the earlier promise that individual consumption choices depend formally on the realization of all uncertainty, but are independent of the realization of uncertainty not faced. Note that an agent may choose *many* jobs in *many* firms.

3.3.2 Endowments

An *endowment* $e \in \mathbb{R}_+^L$ specifies an initial claim to consumption. I identify $\mathbb{R}_+^L \subset \mathbb{R}_+^{L\Omega}$ as the constant functions $\Omega \to \mathbb{R}_+^L$, so endowments are *not ran-*

dom. Agents are not endowed with firm-role choices.[7] Whenever convenient, we abuse notation and view the endowment as the pair $(e, 0) \in X(\Phi)$ (so endowments are in consumption sets).

3.3.3 Utilities

Agents care about consumption, about their own choices, about the actions and skills of others in the firm to which they belong, and about the realized outcome. For simplicity, agents are assumed to be expected utility maximizers, so it suffices to specificy utility for *non-random* consumption. As above, it is convenient to view utility as defined over all tuples but require that it be independent of irrelevant components.

To formalize this, fix a consumption set $X(\Phi)$. Write

$$\begin{aligned} \mathbf{S}^j &= S_1^j \times \ldots \times S_{R^j}^j \\ \mathbf{S} &= \mathbf{S}^1 \times \ldots \times \mathbf{S}^J \\ \mathbf{A}^j &= A_1^j \times \ldots \times A_{R^j}^j \\ \mathbf{A} &= \mathbf{A}^1 \times \ldots \times \mathbf{A}^J \end{aligned}$$

Utility is a mapping

$$u : \mathbb{R}_+^L \times \mathbb{F} \times \mathbf{S} \times \mathbf{A} \times \Omega \to \mathbb{R}$$

so that $u(x, \phi, \mathbf{s}, \mathbf{a}, \omega)$ is the utility obtained if the agent consumes x, chooses the firm-role-action profile ϕ, faces firm members with skill-action profile $\mathbf{s}, \mathbf{a}$, and the profile of consequences ω occurs. I require that

$$\begin{aligned} \phi^j = 0 \quad &\Rightarrow \quad u \text{ is independent of } \omega^j \\ \phi^j = (r, a) \quad &\Rightarrow \quad u \text{ is independent of } \mathbf{s}_r^j, \mathbf{a}_r^j \end{aligned}$$

That is, utility is independent of skills and actions in firms to which the agent does not belong, and utility is independent of skills and actions of

[7]There would be no difficulty in allowing for the possibility that agents *are* endowed with firm-role choices; see Zame (2005).

others in roles which the given agent fills. I assume utility is continuous in its arguments and strictly increasing in each commodity. In addition, I require the following three additional assumptions.[8]

A1 $u(0, \phi, \mathbf{s}, \mathbf{a}, \omega) < u(e, \epsilon, \mathbf{s}, \mathbf{a}, \omega)$ for all $\phi, \mathbf{s}, \mathbf{a}, \omega$

A2 $\lim_{|x| \to \infty} u(x, \epsilon, \mathbf{s}, \mathbf{a}, \omega) = \infty$ for every $\mathbf{s}, \mathbf{a}, \omega$

A3 there is a $k > 0$ such that $u(x, \epsilon, \mathbf{s}, \mathbf{a}, \omega) \leq k(1 + |x|)$ for every $\mathbf{s}, \mathbf{a}, \omega$

3.3.4 Skills

Each agent is endowed with an array of *skills*: $s \in \mathbf{S}$.[9]

3.3.5 The Space of Agent Characteristics

Write Θ for the space of agent characteristics (Φ, u, e, s). If $\theta \in \Theta$ is a typical characteristic, I sometimes abuse notation to write $\theta = (X_\theta, u_\theta, e_\theta, s_\theta)$. Let $\mathcal{V}$ be the space of continuous functions $v : \mathbb{R}^L_+ \times \mathbb{F} \times \mathbf{S} \times \mathbf{A} \times \Omega \to \mathbb{R}$, equipped with the topology of uniform convergence on compact sets. In the topology Θ inherits as a closed subspace of $2^{\mathbb{F}} \times \mathbb{R}^L_+ \times \mathbb{F} \times \mathcal{V} \times \mathbf{S}$, the space of agent characteristics is a complete metric space.

3.4 The Economy

An *economy* consists of a finite set of commodities, a finite set of firms, and a Borel probability measure λ on Θ.

[8]Because the choice to belong to a firm is indivisible, some assumption such as **A1** is necessary to guarantee that optimal choices are upper-hemi-continuous in prices, and some assumptions such as **A2**, **A3** are necessary to guarantee that expected demand blows up as commodity prices approach the boundary of the price simplex.

[9]It a richer dynamic model, it would be natural to model skills as acquired; see Ellickson, Grodal, Scotchmer & Zame (2006) for example.

4 Common Beliefs Equilibrium

An equilibrium of the model will be defined as a set of market prices and individual choices with the property that markets clear and individuals optimize. Formalizing this requires careful specification of the objects to be priced, the meaning of market clearing, and the nature of individual optimization. Because these issues are central, and their formalization is complicated, a preliminary discussion may be useful.

In the model, agents choose (consequence-dependent) consumption of private goods, and also firms, roles and actions. The point of view taken here is that consumption of private goods and choices of firm and role are observable, verifiable and contractible, and hence are priced, but that action choices are not observable — or at least not verifiable – and hence are not contractible and are not priced.[10] Similarly, skills are not contractible and are not priced. Of course, the non-contractibility of actions seems the essence of moral hazard and the non-contractibility of skills seems the essence of adverse selection. Prices paid for roles in a given firm, which it is natural to term *wages*, represent transfers within the firm, and so sum to 0. Requiring wages to be specified as part of equilibrium seems very natural and probably needs little justification, but it is also necessary to guarantee that equilibrium exists.[11]

In the model, there are markets for jobs and markets for private goods; at equilibrium, both markets must clear. Market clearing for jobs means

[10]Observable skills and actions can be accommodated as in Ellickson, Grodal, Scotchmer & Zame (2006), by appropriately coding skills into consumption sets and actions into the description of firms. For instance, consider Example 1. To model ROW's actions as observable, posit two firm types, rather than one: in the first, ROW will be constrained to exert High effort; in the second, ROW will be constrained to exert Low effort. To model that output depends on the skill of COL, and that skill is observable, posit as many firm types as there are possible skills of COL, and restrict consumption sets so that only agents who have a particular skill can choose the role of COL in the corresponding firm type.

[11]In Example 1, equilibrium wages are 0, but if ROW were entitled to the entire output then equilibrium wages could not be 0, else all agents would prefer the role of ROW to the role of COL.

simply that, for each firm an equal mass of agents (perhaps 0) must choose each role. However because individual consumption and production may be random, market clearing for commodities requires some explanation. Note that individuals face two sources of uncertainty: i) the stochastic nature of the production process, and ii) the skills and behavior of other agents with whom they might be matched in a particular firm. (Recall that a firm is a particular instance of a firm type, formed by matching individual agents.) I assume here that matching of agents is random and uniform, that the realizations of uncertainty within firms (i.e., the realized consequences) are drawn independently from the induced distribution — in the language of Prescott & Townsend (1984a), *shocks are private* — and that the Law of Large Numbers applies.[12] The implication of these assumptions is that, although each individual agent faces risk (and accommodates these risks by choosing consequence-contingent consumption plans), these risks are purely idiosyncratic and wash out in the aggregate. Thus, *there is no aggregate risk*: aggregate input, output and consumption are deterministic (given the choices of agents).[13] I require prices to be deterministic as well.

To understand the nature of individual optimization, it is useful to think first about an Arrow–Debreu–McKenzie production economy in which no agent is endowed with good 1. By definition, an equilibrium for such an economy is a collection of prices and choices for which agents and firms optimize and markets clear. Note that agents and firms do *not* optimize subject to availability of good 1 (or any other good) — they optimize subject only to prices, even if it happens that, at equilibrium, good 1 is not produced

[12]The applicability of the Law of Large Numbers can be justified as in Duffie & Sun (2004a, 2004b).

[13]For instance, consider Example 1 of Section 2. Suppose 1/3 of the agents enter a firm in the role of ROW and exert High effort, 2/9 of the agents enter a firm in the role of COL and exert High effort, 1/9 of the agents enter a firm in the role of COL and exert Low effort, and the remaining 1/3 of the agents do not enter the firm. From the point of view of an agent in the role of ROW, this means that the probability of being matched with an agent (in the role of COL) who exerts High effort (so that output will be 450 units) is 2/3 and the probability of being matched with an agent who exerts Low effort (so that output will be 200 units) is 1/3. Similarly, aggregate production of good 2 will be $(1/3)(2/9)(450) + (1/3)(1/9)(200) = 1100/27$ units.

and hence is entirely unavailable. To paraphrase Debreu (1959): agents do not take aggregate supply into account when they make their optimization decisions. Similarly, in the standard security market model — in which securities are in zero net supply — investors may choose to buy or not to buy a security at a given price, but their portfolio choice does not depend on whether other agents choose to sell at that price (which is the only way in which a security in zero supply would be available). In either case, at an equilibrium at which some good is not produced or some security is not sold, prices must be such that that no one would demand the good or security *even if it were available.*

In the present model, it will be possible for an agent to obtain a job in a particular firm type only if other agents choose the other jobs in that particular firm type. In this respect, jobs are analogous to produced goods (in the Arrow–Debreu–McKenzie production model) or to securities in zero net supply (in the standard security market model). I therefore insist that agents optimize without taking the supply of jobs (i.e., the decisions of others to enter or not enter firms) into account. If consequences were independent of actions and skills, there would be little more to say, and we would be in the setting of Ellickson, Grodal, Scotchmer & Zame (1999,2001,2006). In this present case, however, consequences are dependent on actions and skills of others, so agents must take into account the actions and skills of others. The formal requirement is that agents form beliefs about the skills and action choices of others in each firm type, and optimize given those beliefs. But what discipline should be imposed on those beliefs? For firm types that are formed in equilibrium, the answer seems obvious: beliefs should be correct. For firm types that are *not* formed in equilibrium, imposing no discipline would admit equilibria which are viable only because different agents hold contradictory beliefs about behavior in those firms. I therefore require that all agents hold the same beliefs. This requirement is precisely parallel to the requirement made in Dubey & Geanakoplos (2002) (that agents make a common forecast about deliveries on pools, even those which are empty) and in Dubey, Geanakoplos & Shubik (2005) (that agents make a common forecast about the default rate on securities, even those for which there are no transactions). If some firm type is not formed at equilibrium, then prices, wages and beliefs must be such that that no one would demand a job in that firm type *even if it were available*.[14]

Thus, a *common-beliefs equilibrium* consists of commodity prices p, wages w, beliefs β and a probability distribution on $\Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F}$ such that: the equilibrium distribution is consistent with the population distribution, commodity markets clear, job markets clear, agents optimize in their budget sets given beliefs, and beliefs are correct. Formal definitions are given below.

4.1 Wages

In addition to private goods, roles in firms are priced. It is natural to interpret a role in a firm as a job and its price as a *wage*, and to adopt the convention that wages are paid *to* the members of the firm.[15] To formalize this, write

$$\mathbb{M} = \{(j, r) : j \in J, r \in R^j\}$$

for the set of *memberships* or *firm-role pairs*. A *wage structure*, or *wage* for short, is a function $w : \mathbb{M} \to \mathbb{R}$ for which

$$\sum_{r \in R^j} w(j, r) = 0 \text{ for each } j$$

(The requirement that wages in a firm sum to 0 is an accounting identity, reflecting the definition of wages as lump sum transfers.) Write $\mathbb{W} \subset \mathbb{R}^{\mathbb{M}}$ for the space of wage structures. I emphasize that, in keeping with the interpretation that skills and actions are not contractible, skills and actions are *not* priced.

[14]In Example 1, for instance, this rules out the trivial equilibrium in which no agent demands a job because each assumes there will be no supply of partners.

[15]This is the opposite sign convention from that adopted in Ellickson, Grodal, Scotchmer and Zame (1999, 2001, 2006).

4.2 Budget Sets

If w is a wage structure and $\phi \in \mathbb{F}$ is a profile of role-action choices, it is convenient to abuse notation and write

$$w \cdot \phi = \sum_{\phi^j \neq 0} w(j, \rho^j(\phi))$$

for total wage income, given the wage structure w and the choices ϕ. (Of course agents do not receive wages in firms to which they do not belong.)

Given private good prices p and wages w, the choice $(\tilde{x}, \phi)$ is *budget feasible* for an agent with endowment e if consumption choices are budget feasible in each profile of consequences $\omega = (\omega^1, \ldots, \omega^J) \in \Omega$:

$$p \cdot \tilde{x}(\omega) \leq p \cdot e + w \cdot \phi + \sum_{\phi^j \neq 0} C^j(\rho^j(\phi), \omega^j, p)$$

The left side is expenditure on private goods; the right side is the value of the endowment of private goods plus wages for each job chosen plus contractual payments for each of these jobs. Write $B(e, p, w)$ for the set of budget feasible choices. (Keep in mind that income and consumption choices are random but that endowments, prices and wages are not random.)

4.3 Beliefs and Expected Utility

A system of *beliefs* is a probability measure β on $\mathbf{S} \times \mathbf{A}$, the space of all skill-action profiles for all firm types. Given beliefs β, write β^j for the marginal of β on $\mathbf{S}^j \times \mathbf{A}^j$ and β^j_{-r} for the marginal of β on $\mathbf{S}^j_{-r} \times \mathbf{A}^j_{-r}$.[16]

Consider an agent with characteristics $\theta = (\Phi, u, e, s)$, who holds beliefs β and chooses $(\tilde{x}, \phi) \in X(\Phi)$. Fix $\omega = (\omega^1, \ldots, \omega^J) \in \Omega$ and a skill-action profile $(\mathbf{s}, \mathbf{a}) \in \mathbf{S} \times \mathbf{A}$. If $\phi^j \neq 0$, the probability that this agent assigns to observing consequence ω^j in firm type j when complementary agents have skill-action profile $(\mathbf{s}, \mathbf{a})$ is

$$\pi^j(\omega^j | \phi, \mathbf{s}, \mathbf{a}) = \pi(\omega^j | s^j, \alpha^j(\phi), \mathbf{s}^j_{-r}, \mathbf{a}^j_{-r})$$

[16] As usual, I write $\mathbf{S}^j_{-r}$ for the profiles of skills in all roles *except* role r, etc.

If $\phi^j = 0$, define

$$\pi^j(\omega^j|\phi, \mathbf{s}, \mathbf{a}) = \pi^j(\omega^j|\mathbf{s}^j, \mathbf{a}^j)$$

Hence the probability that this agent assigns to observing the profile ω when complementary agents have skill-action profile $(\mathbf{s}, \mathbf{a})$ is

$$\pi(\omega|\phi, \mathbf{s}, \mathbf{a}) = \prod_{j \in J} \pi^j(\omega^j|\phi, \mathbf{s}, \mathbf{a})$$

and this agent's expected utility is:

$$Eu(\tilde{x}, \phi|\beta) = \int_{\mathbf{S}\times\mathbf{A}} u(\tilde{x}(\omega), \phi, \mathbf{s}, \mathbf{a}, \omega)\pi(\omega|\phi, \mathbf{s}, \mathbf{a})\, d\beta(\mathbf{s}, \mathbf{a})$$

4.4 Job Market Clearing

For each j, r, let

$$T^j_r = \{(\theta, \tilde{x}, \varphi) \in \Theta \times \mathbb{R}^{L\Omega}_+ \times \mathbb{F} : \rho^j(\phi) = r\}$$

This is the set of characteristics of agents who choose role r in firm type j. Say that the probability measure μ on $\Theta \times \mathbb{R}^{L\Omega}_+ \times \mathbb{F}$ is *consistent*, or that *the job market clears* if for each j and each $r, r' \in R^j$

$$\mu(T^j_r) = \mu(T^j_{r'})$$

That is: for each firm, the same number of agents choose each of the roles.

4.5 Aggregate Output

A choice of firms, roles and actions for each agent in the economy induces a distribution of skills and behaviors within each firm. From this distribution, the assumption of random matching and Law of Large Numbers determines aggregate output. Formally, let μ be a consistent probability measure on $\Theta \times \mathbb{R}^{L\Omega}_+ \times \mathbb{F}$. For each j, r for which $\mu(T^j_r) \neq 0$, define $h : T^j_r \to S^j_r \times A^j_r$ by

$$h(\theta, \tilde{x}, \phi) = (s^j, \alpha^j(\phi))$$

Let $\zeta_r^j = h_*\mu$ be the direct image measure and set $\gamma_r^j = \zeta_r^j/\|\zeta_r^j\|$; note that γ_r^j is the distribution of skills and actions in role r in firm type j, given the choices μ. Because matching is random, the distribution of skills and actions over all roles in firm type j is $\gamma^j(\cdot|\mu) = \gamma_1^j \times \ldots \times \gamma_{R^j}^j$ and the distribution of outcomes across firms of type j is given by

$$\Gamma^j(\omega^j|\mu) = \int_{\mathbf{S}^j \times \mathbf{A}^j} \pi^j(\omega^j|\mathbf{s}, \mathbf{a})\, d\gamma^j(\cdot|\mu)$$

Consistency requires that $\mu(T_r^j) = \mu(T_{r'}^j)$ for all $r, r' \in R^j$, so each of these — and their average — represents the number of firms that form. The Law of Large Numbers implies that the total output of all firms of type j is

$$Y^j(\mu) = \left(\frac{1}{R^j} \sum_{r \in R^j} \mu(T_r^j) \right) \sum_{\omega^j \in \Omega^j} y^j(\omega^j) \Gamma^j(\omega^j)$$

Hence aggregate output is

$$Y(\mu) = \sum_{j \in J} Y^j(\mu)$$

4.6 Aggregate Consumption

Because individual consumption is random (state-dependent), aggregate consumption cannot be defined simply as the integral of individual consumption with respect to the population measure. However, it follows from the Law of Large Numbers that aggregate consumption can be defined as the integral of individual *expected* consumption with respect to the population measure. Because the risks individuals face are correlated with their consumption decisions, some care must be taken in formalizing this definition.

Fix a consistent probability measure μ on $\Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F}$. For $\phi \in \mathbb{F}$, write

$$\begin{aligned} J(\phi) &= \{j \in J : \phi^j \neq 0\} \\ T(\phi) &= \{(\theta, \tilde{x}, \psi) \in \Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F} : \psi = \phi\} \end{aligned}$$

For $j \in J(\phi)$, define

$$\begin{aligned}
S^j(\phi) &= \begin{cases} S^j & \text{if} \quad j \notin J(\phi) \\ S^j_{-\rho^j(\phi)} & \text{if} \quad j \in J(\phi) \end{cases} \\
A^j(\phi) &= \begin{cases} A^j & \text{if} \quad j \notin J(\phi) \\ A^j_{-\rho^j(\phi)} & \text{if} \quad j \in J(\phi) \end{cases} \\
\mathbf{S}(\phi) &= S^1(\phi) \times \ldots \times S^J(\phi) \\
\mathbf{A}(\phi) &= A^1(\phi) \times \ldots \times A^J(\phi)
\end{aligned}$$

For $\mathbf{s} \in \mathbf{S}(\phi), \mathbf{a} \in \mathbf{A}(\phi)$, $j \in J$ and $\omega = (\omega^1, \ldots, \omega^J) \in \Omega$, define $\pi^j(\omega^j|\phi, \mathbf{s}, \mathbf{a})$ and $\pi(\omega|\phi, \mathbf{s}, \mathbf{a})$ as in Subsection 4.3 and $\gamma^j(\cdot|\mu)$ as in Subsection 4.5. Set $\gamma(\cdot|\mu) = \gamma^1(\cdot|\mu) \times \ldots \times \gamma^J(\cdot|\mu)$ and let $\gamma_\phi(\cdot|\mu)$ be the marginal of $\gamma(\cdot|\mu)$ on $\mathbf{S}(\phi) \times \mathbf{A}(\phi)$; $\gamma_\phi(\cdot|\mu)$ is the distribution of skills and actions in all roles in all firms except those in ϕ. Expected consumption of $(\Phi, u, e, s, \tilde{x}, \phi)$ is

$$E(\tilde{x}|\mu) = \sum_{\omega \in \Omega} \int_{\mathbf{S}(\phi) \times \mathbf{A}(\phi)} \tilde{x}(\omega) \pi(\omega|\phi, \mathbf{s}, \mathbf{a}) \, d\gamma_\phi(\mathbf{s}, \mathbf{a}|\mu)$$

The total expected consumption of agents in $T(\phi)$ is $\int_{T(\phi)} E(\tilde{x}|\mu) \, d\mu$, so the total expected consumption of all agents, which (using the Law of Large Numbers) I identify with aggregate consumption, is

$$\begin{aligned}
X(\mu) &= \sum_{\phi \in \mathbb{F}} \int_{T(\phi)} E(\tilde{x}) \, d\mu \\
&= \sum_{\phi \in \mathbb{F}} \int_{T(\phi)} \sum_{\omega \in \Omega} \int_{\mathbf{S}(\phi) \times \mathbf{A}(\phi)} \tilde{x}(\omega) \pi(\omega|\phi, \mathbf{s}, \mathbf{a}) \, d\gamma_\phi(\mathbf{s}, \mathbf{a}|\mu) \, d\mu
\end{aligned}$$

(Of course, some of the sets $T(\phi)$ may have measure 0, in which case the corresponding contributions to aggregate consumption will be 0 as well.)

4.7 Equilibrium

Given an economy λ, a *feasible configuration* for λ is a probability measure μ on $\Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F}$ for which

- the marginal of μ on Θ is λ
- μ is consistent (the job market clears)
- almost all individual choices are physically feasible

$$\mu\{(\Phi, u, e, s, \tilde{x}, \phi) : (\tilde{x}, \phi) \notin X(\Phi)\} = 0$$

- commodity markets clear: $X(\mu) = Y(\mu) + \int e d\lambda$

A *common beliefs equilibrium* for λ consists of prices $p \in \Delta$, wages $w \in \mathbb{R}^{\mathbb{M}}$, beliefs β, and a feasible configuration μ on $\Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F}$ such that

- almost all choices are budget feasible:

$$\mu\{(\Phi, u, e, s, \tilde{x}, \phi) : (\tilde{x}, \phi) \notin B(e, p, w)\} = 0$$

- almost all choices are optimal given prices and beliefs:

$$\mu\{(\Phi, u, e, s, \tilde{x}, \phi) : \\ \exists(\tilde{x}', \phi') \in X(\Phi) \cap B(e, p, w), Eu(\tilde{x}', \phi'|\beta) > Eu(\tilde{x}, \phi|\beta)\} = 0$$

- beliefs are correct for firms that form: $\mu(T_r^j) \neq 0 \Rightarrow \beta^j = \gamma^j(\cdot|\mu)$

Subject to boundedness of individual endowments (a condition shown necessary in Ellickson, Grodal, Scotchmer & Zame (1999)) and to the usual caveat that all goods are available in the aggregate, common beliefs equilibrium exists; the proof is deferred to the Appendix.

Theorem 1 *If individual endowments are uniformly bounded and all goods are available in the aggregate (i.e., $\int e\, d\lambda \gg 0$),then a common beliefs equilibrium exists.*

5 Examples

The examples below, together with Example 1, illustrate the working of the model. Example 1 illustrated how moral hazard interacts with the market; Example 2 illustrates how adverse selection interacts with the market. Example 3 illustrates market screening and Example 4 illustrates insurance and endogenous contracting. All of these examples have been selected for ease of analysis rather than realism.

Example 2 (Adverse Selection) In this example, output depends on Worker skill, which is hidden, giving rise to adverse selection.

Formally: There are two goods. In the single firm type, there are two roles, Investor and Worker (so $R = \{I, W\}$), each with a single action. The set of Investor skills is a singleton, but the set of Worker skills is an interval $S_W = [\varepsilon, 1]$. (I choose $\varepsilon > 0$ below.) There are two possible consequences $\Omega = \{G, B\}$; consequence-dependent output is $y(G) = (0, 1)$, $y(B) = (0, 0)$. Conditional probabilities depend only on the skill of the Worker: $\pi(G|s) = s$. The Investor owns all the output of the firm: $C(I, \omega, p) = p \cdot y(\omega)$. (Here and elsewhere, I suppress degenerate/irrelevant variables.)

There are two types of agents: Type I agents can choose to become an Investor or not to enter a firm, and have endowment $(1, A)$; $4/5 < A < 1$. Utility of Type I agents depends only on consumption: $u_I(c_1, c_2) = (c_1 c_2)^{1/2}$. Type II agents can choose to be a Worker or not to enter a firm, and have endowment $(1, 0)$. Utility of Type II agents depends on consumption, on firm-role-action choice, and on own skill s:

$$\begin{aligned} u_s(c_1, c_2; 0) &= c_1 \\ u_s(c_1, c_2; W) &= c_1 - s \end{aligned}$$

Note that more skilled agents have higher disutility for work.[17]

[17]Nothing would change if Type II agents' marginal utility for the second good were strictly positive, but small.

To choose ε, define $h : [0,1] \times [0,1] \to [0,1]$ by

$$h(\varepsilon, w) = \left[\frac{\varepsilon + w}{2}\right] [(1-w)(A+1)]^{1/2} + \left[1 - \frac{\varepsilon + w}{2}\right] [(1-w)A]^{1/2}$$

Note that

$$h(0,0) = A^{1/2} \ ; \ \frac{\partial h}{\partial \varepsilon}(0,0) < 0 \ ; \ \frac{\partial h}{\partial w}(\varepsilon, w) < 0 \ \text{ if } \ w < 1$$

Hence choosing choose $\varepsilon > 0$ sufficiently small guarantees that $h(\varepsilon, w) < A^{1/2}$ for $\varepsilon \leq w \leq 1$.

As a benchmark, consider first the setting in which there is a single agent of each type, but the skill of the Type II agent is drawn from the uniform distribution on $[0,1]$ (the realization known to the Type II agent but not to the Type I agent). I claim that autarky is the only equilibrium. To see this, note that if a wage $w > \varepsilon$ is offered to Workers, then a Type II agent having skill s will be willing to be a Worker if and only if $w \geq s$. Keeping in mind that output depends on skill, it follows that a Type I agent who pays the wage w and becomes an Investor will enjoy expected utility:

$$\overline{u} \ = \ \left[\frac{\varepsilon + w}{2}\right] [(1-w)(A+1)]^{1/2} + \left[1 - \frac{\varepsilon + w}{2}\right] [(1-w)A]^{1/2}$$

By construction, $\overline{u} < A^{1/2}$, which is the utility a Type I agent would obtain from consuming her endowment $(1, A)$. Hence no Type I agent would ever wish to become an Investor. It follows that $w \leq \varepsilon$, so no Type II agents become Workers, and the equilibrium is autarkic. (This deliberately parallels Akerlof (1970).)

Now consider a setting with many agents, with more agents of Type I than of Type II, and with skills of Type II agents uniformly distributed on $[\varepsilon, 1]$. Note first that autarky cannot be an equilibrium. For if it were, then no Type II agent would demand to be a Worker, whence $w \leq \varepsilon$, and commodity prices would be $(1, 1/A)$. An Investor who pays the wage w will obtain $(1-\varepsilon, A+1)$ when she experiences the G realization and $(1-\varepsilon, A)$ when she experiences the B realization. Keeping in mind that the probability of a G realization is at least ε and that the Investor will be able to trade at the market prices of $(1, 1/A)$, a simple computation shows that expected utility for an Investor will be $\widehat{u} = A^{1/2}(2 - \varepsilon + \varepsilon A^{-1})/2$. By assumption, $A < 1$, so $\widehat{u} > A^{1/2}$. Hence Type I agents would demand to be Investors and the market for jobs would not clear. This is a contradiction, so we conclude that autarky cannot be an equilibrium.

To solve for equilibrium, take good 1 as numeraire, so commodity prices are $(1, q)$, and posit a wage $w > \varepsilon$. Assign the Type II agents who have skill $s \leq w$ to be Workers and an equal number of Type I agents to be Investors. Given stochastic production, calculate the price q for good 2 at which supply equals demand. Keeping in mind the fraction of Investors who obtain output and the fraction who do not, as well as the fraction of Type II agents who do not enter a firm, it follows that q is the solution to

$$\begin{aligned} \frac{w^2}{2} + A &= [1 - w]\left[\frac{1+q}{2q}\right] + w\left[\frac{\varepsilon + w}{2}\right]\left[\frac{1 - w + (A+1)q}{2q}\right] \\ &\quad + w\left[1 - \frac{\varepsilon + w}{2}\right]\left[\frac{1 - w + Aq}{2q}\right] \end{aligned}$$

This equation is linear in q; solving expresses q as a quadratic function of w. Now compute the imputed (indirect) expected utility $V_I(w)$ for Investors and $V_0(w)$ for Type I agents who do not enter a firm

$$\begin{aligned} V_I(w) &= \left[\frac{\varepsilon + w}{2}\right]\left[\frac{1 - w + (A+1)q}{2}\right] q^{-1/2} \\ &\quad + \left[1 - \frac{\varepsilon + w}{2}\right]\left[\frac{1 - w + qA}{2}\right] q^{-1/2} \\ V_0(w) &= \left[\frac{1+q}{2}\right] q^{-1/2} \end{aligned}$$

Note that $V_I(\varepsilon) > V_0(\varepsilon)$. A little calculus shows that there is a *unique* w^*, $\varepsilon < w^* < 1$, with $V_S(w^*) = V_0(w^*)$. For this wage w^* and the corresponding commodity prices $(1, q^*)$, all agents optimize and markets clear, so this defines the unique common beliefs equilibrium — and at this equilibrium there is trade. Because they are *ex ante* identical, all Type I agents obtain the same *ex ante* expected utility. However, some Type I agents become Investors and are lucky (obtain a unit of good 1), some Type I agents become Investors and are unlucky, and the remaining Type I agents do not enter a firm; these various Type I agents obtain different realized consumption and different *ex post* realized utility.

Note that if ε were 0 then there would be a second, autarkic equilibrium, supported by the wage $w = 0$ and by the (correct) belief that all Type II agents who choose to become Workers have skill 0. Evidently, this autarkic equilibrium is unstable. $\diamondsuit$

Example 3 (Screening) In this example, there is a single technology but two contracts; at equilibrium, contracts screen agents into their more efficient roles.

Formally, the single technology T has two roles ($R = \{E, W\}$), each with a single action; there are two consequences ($\Omega = \{G, B\}$). Output is $y(G) = 1$, $y(B) = 0$, and probabilities depend on skills:

$$\pi^{\mathcal{P}}(G|s_E, s_W) = \min\{1, s_E + s_E\}$$

There are two contracts; because the sum of contractual payments is the value of output, it is enough to specify payments to Entrepreneurs:

$$C^1(E, \omega) = \frac{1}{2}y(\omega)\ ,\ C^2(E, \omega) = y(\omega)$$

(In the first contract, participants share equally; in the second contract Entrepreneurs own the entire output.) All agents have endowment $e = 1$, care only about consumption, and are risk neutral: $u(x) = x$. Agent skills are uniformly distributed on $[0, 1]$. Finally, there are two types of firm: $F^1 = (T, C^1)$, $F^2 = (T, C^2)$.

It is easily checked that there are only two (kinds of) equilibria:

> **Non-screening equilibrium** Only firm type/contract 2 is active. Agents in $[0, 1/2]$ choose to become Workers in firms of type 2, agents in $(1/2, 1]$ choose to become Entrepreneurs in firms of type 2; wages are $w^2(W) = 3/8$. Firm type 1 is inactive; wages are $w^1 = 0$; all agents believe that firms of type 1 are populated entirely by agents of skill $s = 0$. (Shadow wages and beliefs for firm type 1 are indeterminate.)

Screening equilibrium Both firm types/contracts are active. Agents in $[0, 1/4]$ choose to become Workers in firms of type 2, agents in $[3/4, 1]$ choose to become Entrepreneurs in firms of type 2; $1/2$ the agents in $(1/4, 3/4)$ choose to become Workers in firms of type 1 and $1/2$ the agents in $(1/4, 3/4)$ choose to become Entrepreneurs. Wages are $w^1 = 0$, $w^2(W) = 3/8$.

It is enlightening to compare the social gains for the equilibrium in which only the equal-participation contract $C^1$is available, for the equilibria above, and for the full-information outcome in which skills are public information (so agents of skill s are matched with agents of skill $1 - s$) :

- **Only C^1 available** social gain = 80/192
- **Non-screening equilibrium** social gain = 88/192
- **Screening equilibrium** social gain = 91/192
- **Full-information equilibrium** social gain = 96/192

(Of course there is nothing special about *two* contracts; I leave it to the reader to see that as the number of contracts tends to infinity, the screening equilibrium converges to full efficiency.)

The non-screening equilibrium is supported by implausible beliefs. Such beliefs might be ruled out by an appropriate notion of stability, but I have no plausible notion to offer. $\diamondsuit$

The approach taken in this paper views technologies and contracts as part of the description of the economy and wages (lump sum transfers) as determined at equilibrium. An alternative approach would have been to view technologies as part of the description of the economy and both wages and contracts as determined at equilibrium. This approach would certainly lead to a consistent model, but it would be seem less satisfactory than it might appear, because equilibrium would usually be highly indeterminate: Theorem

1 guarantees that *any* contractual arrangement can be supported at equilibrium by some lump-sum wage, and there would seem to be no obvious way of choosing among equilibrium contracts. By contrast, the approach taken here makes it possible to endogenize contracts because equilibrium selects a contract or set of contracts among many possible contracts — even among all possible contracts. The example below, set in the context of the provision of insurance, illustrates this point.

Example 4 (Insurance and Endogenous Contracting) In this example, Workers can use either of two production processes, one safe and one risky. Risky production can be insured by forming a partnership with an Insurer; the market selects the optimal insurance contract.

Formally: there is a single good. There are two kinds of agents, Workers and Insurers, and three technologies. The first technology T_1 involves a single Worker, a single action, a single consequence; output is η (a parameter), $\sqrt{3} < \eta < 2$. The second technology T_2 involves a single Worker, a single action, two equally likely consequences G, B; output is $y(G) = 3$, $y(B) = -1$. The third technology T_3 involves a single Worker and a single Insurer, each with a single action, two equally likely consequences G, B; output is $y(G) = 3$, $y(B) = -1$. Note that T_2 and T_3 are identical except that T_3 involves an Insurer; the interpretation intended is that the Worker does the production and the Insurer provides insurance against the bad outcome. Workers are endowed with 1 unit of the good, Insurers are endowed with $K > 0$ units of the good (K is a parameter of the model). Both Workers and Insurers derive utility solely from consumption, $u(c) = \log c$, and can enter at most a single firm. Finally, there are more Insurers than Workers. (This is analogous to requiring free entry of Insurers.)

We want to analyze the situation in which many contracts are available. For T_1 and T_2, contracts are degenerate, hence suppressed; the issue is the nature of contracts for T_3. Note first that if C, C' are contracts which trade at wages w, w' (respectively), then C, w and C', w' are equivalent — in the sense of yielding the same total payments — exactly when $C(W,G) - w(I) = C'(W,G) - w'(I)$ and $C(W,B) - w(I) = C'(W,B) - w'(I)$. Hence there is no loss in considering only contracts C_b having the form $C_b(W,G) = -b$, $C_b(W,B) = b$. Write $q(b) = w(I)$ for the wage paid to the Insurer, which we view as the price of C_b. Suppose that all possible contracts C_b are available; we want to know the contracts traded at equilibrium, and their prices.[18]

Some preliminary calculations will be useful. Define the *Insurer neutral price* $q(b)$ for the contract C_b to be the price at which an Insurer would be indifferent between participating or consuming his endowment; $q(b)$ is determined by the equality:

$$\frac{1}{2}\log(K + b + q(b)) + \frac{1}{2}\log(K - b + q(b)) = \log(K)$$

A worker who accepts the contract C_b and pays $q(b)$ will enjoy expected utility:

$$U(b) = \frac{1}{2}\log(3 - b + q(b)) + \frac{1}{2}\log(-1 + b - q(b))$$

If all contracts C_b are offered at the Insurer neutral price $q(b)$, then workers choose a contract C_b to maximize expected utility. Straightforward calculations, left to the reader, establish the following facts:

- If $K \leq 4$ then there is *no* contract which, at the price $q(b)$, yields the Worker expected utility as great as $\sqrt{3}$.
- If $K > 4$ then among all contracts C_b offered at prices $q(b)$, the unique contract which yields the Worker the highest utility $U(K)$ has

$$\begin{aligned} b(K) &= 2K(K^2 - 4K)^{-1/2} \\ q(b(K)) &= K(K-2)(K^2 - 4K)^{-1/2} - K \end{aligned}$$

$U((b(K))$ is a strictly increasing function of $K > 4$, and

$$\lim_{K\to\infty} b(K) = 2 \ , \ \lim_{K\to\infty} q(K) = 0 \ , \ \lim_{K\to\infty} U(K) = \log 2$$

Hence there is a unique $K(\eta) > 4$ such that $U(K(\eta)) = \log(\eta)$.

[18] Formally: for each $b \in \mathbb{R}$ there is a firm (T_3, C_b). To be rigorous, the set of firms should be finite, but this abuse should cause no confusion, and at most two contracts will be active at equilibrium.

We can now solve for equilibrium as a function of insurer wealth K. Note first that there are more Insurers than Workers, so in equilibrium some Insurers consume their endowment; hence in equilibrium, all Insurers obtain utility $\log K$. There are three cases:

- $K < K(\eta)$: At equilibrium, all Insurers obtain utility $\log K$, so no Insurer would accept any contract at a price less than the Insurer neutral price. However, at the Insurer neutral price, Workers would obtain utility less than they could obtain in the safe technology T_1. Hence no contract C_b is active. (A little more formally: no firm type (T_3, C_b) is active.) Shadow prices for contracts C_b are $q(b)$. Because $\eta > 1$, Workers prefer working in the safe technology T_1 to consuming their endowments; hence all Workers work in the safe technology. (No insurance is provided.) Note that shadow prices for some contracts C_b are indeterminate: for some contracts C_b, $U(b) < \eta$, so a shadow price slightly below $q(b)$ will not be enough to induce any Worker to enter.

- $K > K(\eta)$: By construction, it is an equilibrium for all contracts C_b to be offered at the Insurer neutral price $q(b)$, for all Workers and an equal mass of Insurers to choose the contract $C_{b(K)}$ at the price $q(b(K))$, and for all other Insurers to be inactive. To see that there are no other equilibria (modulo the same indeterminacy as above), note that the price for any active contract must be the Insurer neutral price (else Insurers who enter the contract would obtain utility different from Insurers who do not enter any contract), and that the price for any any inactive contract must be at least as high as the Insurer neutral price (else Insurers could enter that contract and obtain utility greater than Insurers who do not enter any contract). This implies that $C_{b(K)}$ is the unique active contract.

- $K = K(\eta)$: Now, equilibrium is indeterminate: some Workers choose the safe technology T_1; some Workers and an equal mass of Insurers choose the contract $C_{b(K(\eta))}$; the price for C_b is $q(b)$.

In the limit as insurers become infinitely wealthy — hence risk neutral — workers obtain full insurance at the actuarially fair price. $\diamondsuit$

6 Welfare

Because team production requires coordination, mis-coordination is possible. The following examples illustrates that there are many possibilities for mis-coordination and that mis-coordination can lead to Pareto inferior — even Pareto ranked — equilibria. Example 5 observes that any normal form game can be imbedded in an economy, so that all the welfare ranking possible for Nash equilibria are possible for economies as well. Example 6 uses the same idea to show how agents can mis-coordinate on the "wrong" production plan. For these examples, the lower-ranked Pareto equilibrium is not immune to deviations by the members of a single firm, and prices are irrelevant. Example 7 illustrates the possibility of Pareto ranked equilibria that *are* immune to deviations by the members of a single firm and for which prices are relevant.

Example 5 (Nash Miscoordination) Let $\Gamma = (N, (A_n), (v_n))$ be a game in normal form (with player set N, action sets A_n, utility functions v_n). Imbed this game in an economy $\mathcal{E}_\Gamma$ as follows. There is a single good. There is a single firm type, with N roles; role n has action set A_n. There is a single consequence and no production, and the contract is degenerate $C \equiv 0$. There are N types of agents in the economy, with an equal mass of agents of each type. Agents of type n are endowed with 1 unit of the consumption good and can choose not to enter a firm or to enter in role n. Utility of agents of type n depends on consumption, on job choice, and on the profile of actions in the firm:

$$u_n(x; 0) = x \ , \ u_n(x; n, \mathbf{a}) = x + v_n(\mathbf{a})$$

If $\sigma = (\sigma_1, \ldots, \sigma_N)$ is a Nash equilibrium of Γ with the property that each agent enjoys strictly positive expected utility, then we can construct a common beliefs equilibrium for $\mathcal{E}_\Gamma$ by setting the commodity price equal to 1, wages equal to 0, requiring all agents to choose 1 unit of consumption and requiring the distribution of action choices for agents of type n to coincide with σ_n.

Hence if Γ admits Nash equilibria in which each agent enjoys strictly positive expected utility that are Pareto ranked, then $\mathcal{E}_\Gamma$ enjoys common beliefs equilibria that are Pareto ranked. ◇

Example 6 (Output Miscoordination) There is a single good. There is a single firm, with two roles ROW, COL. ROW has two actions H, L; COL has two actions H, L. There are two consequences: G, B; output is $y(G) = 4$, $y(B) = 0$. Consequences depend stochastically on actions; $\pi(G|\cdot, \cdot)$ is given by the matrix in Figure 2:

	H	L
H	1	0
L	0	$\frac{1}{2}$

Figure 2

Firm participants share equally: $C(\cdot, \omega) = \frac{1}{2}y(\omega)$. All agents are identical; they are endowed with one unit of the consumption good, and may choose not to enter the firm, or to enter the firm in either role and choose either action. Utility depends only on consumption: $u(x) = x$.

There are two pure equilibria and a "mixed" equilibrium:

E1 1/2 of all agents choose ROW, H; 1/2 of all agents choose COL, H; wages are $w \equiv 0$

E2 1/2 of all agents choose ROW, L; 1/2 of all agents choose COL, L; wages are $w \equiv 0$

E3 1/2 of all agents choose ROW; of these, 1/3 choose H, 2/3 choose L; 1/2 of all agents choose COL; of these, 1/3 choose H, 2/3 choose L; wages are $w \equiv 0$

In **E1** the common utility of all agents is 2, in **E2** it is 1, in **E3** it is 2/3, so **E1** Pareto dominates **E2** which in turn Pareto dominates **E3**.

Moreover, the equilibria above are robust to the introduction of other contracts. To see this, let $\mathcal{C}$ be any set of contracts containing the contract C.

	H	L
H	$(\alpha^j, 4-\alpha^j)$	$(\beta^j, -\beta^j)$
L	$(\beta^j, -\beta^j)$	$(\frac{1}{2}(\alpha^j+\beta^j), 2-\frac{1}{2}(\alpha^j+\beta^j))$

Figure 3

Fix a contract $C^j \in \mathcal{C}$, $C^j \neq C$, and write $\alpha^j = C^j(\text{ROW}, G)$, $\beta^j = C^j(\text{ROW}, B)$. The matrix of expected contractual deliveries (to ROW) is shown in Figure 3. This matrix defines a 2×2 game. A simple calculation shows that for each α, β this game has a totally mixed equilibrium $(\frac{1}{3}H + \frac{2}{3}L, \frac{1}{3}H + \frac{2}{3}L)$, yielding expected payments to ROW and COL of $(\alpha^j + 2\beta^j)/3$ and $(4-(\alpha^j+2\beta^j))/3$ respectively. Define

$$w^j(\text{ROW}) = \frac{2-(\alpha^j+2\beta^j)}{3}$$

Now we can define equilibria for the economy as follows:

E1′ *for the contract* C:
1/2 of all agents choose ROW, H; 1/2 of all agents choose COL, H; $w \equiv 0$

for all other contracts C^j:
no agents choose C^j; $w^j(\text{ROW}) = \frac{2}{3} - \frac{\alpha^j+2\beta^j}{3}$; all agents believe that the distribution of action choices will be $(\frac{1}{3}H + \frac{2}{3}L, \frac{1}{3}H + \frac{2}{3}L)$

E2′ *for the contract* C:
1/2 of all agents choose ROW, L 1/2 of all agents choose COL, L; wages are $w \equiv 0$

for all other contracts C^j:
no agents choose C^j; $w^j(\text{ROW}) = \frac{2}{3} - \frac{\alpha^j+2\beta^j}{3}$; all agents believe that the distribution of action choices will be $(\frac{1}{3}H + \frac{2}{3}L, \frac{1}{3}H + \frac{2}{3}L)$

E3′ *for the contract* C:
1/2 of all agents choose ROW; of these, 1/3 choose H, 2/3 choose L; 1/2 of all agents choose COL; of these, 1/3 choose H, 2/3 choose L; wages are $w \equiv 0$

for all other contracts C^j:
no agents choose C^j; $w^j(\text{ROW}) = \frac{2}{3} - \frac{\alpha^j + 2\beta^j}{3}$; all agents believe that the distribution of action choices will be $(\frac{1}{3}H + \frac{2}{3}L, \frac{1}{3}H + \frac{2}{3}L)$

Hence Pareto ranking of equilibrium is robust to the exogenous specification of possible contractual arrangements. $\diamondsuit$

Example 7 (Price Miscoordination) There are two goods. There is a single firm, with two roles ROW, COL. Each of ROW, COL can choose to exert High or Low effort. There are three consequences: G, M, B; production is $y(G) = (-2, +4)$, $y(M) = (-2, +8/3)$, $y(B) = (-2, +2)$. Consequences depend deterministically on actions as given by the matrix in Figure 4:

	H	L
H	G	M
L	M	B

Figure 4

The contract is equal sharing: $C(\cdot, \omega) = p \cdot y(\omega)$.

All agents are identical, and can choose either role and either action in a single firm, or not enter a firm. Endowments are $e = (1, 0)$; utility depends on consumption and individual effort level:

$$
\begin{aligned}
u(x_1, x_2; 0) &= (x_1 x_2)^{1/2} \\
u(x_1, x_2; L) &= (x_1 x_2)^{1/2} \\
u(x_1, x_2; H) &= H[(x_1 x_2)^{1/2}]
\end{aligned}
$$

where

$$
H(t) = \begin{cases} \frac{2}{3}t & \text{if} \quad 0 \le t \le \frac{\sqrt{3}}{2} \\ \frac{2}{3}\left(\frac{\sqrt{3}}{2}\right) + \frac{1}{15}\left(t - \frac{\sqrt{3}}{2}\right) & \text{if} \quad \frac{\sqrt{3}}{2} \le t < \infty \end{cases}
$$

Straightforward calculations (as in Example 1) show that there are two equilibria:

E1 1/5 of all agents choose ROW, High effort; 1/5 of all agents choose COL, High effort; 3/5 of all agents do not enter a firm; commodity prices are $(1, 3/4)$; wages are $w \equiv 0$

E2 1/4 of all agents choose ROW, Low effort; 1/4 of all agents choose COL, Low effort; 1/2 of all agents do not enter a firm; commodity prices are $(1, 1)$; wages are $w \equiv 0$

The common utility of all agents in **E1** is $\sqrt{3}/3$, while the common utility of all agents in **E2** is $1/2$, so **E1** Pareto dominates **E2**.

I claim that these equilibria are immune to deviations by members of a single firm. To see this, consider first equilibrium **E1**. If the members of a single firm jointly deviate to Low effort, they do not affect prices, which remain $(1, 3/4)$. Hence their income after production is $3/4$, whence optimal consumption is $(3/8, 1/2)$ and utility is $\sqrt{3}/4 < \sqrt{3}/3$. Hence the members of the firm do not wish to deviate. Now consider equilibrium **E2**. If the members of a firm jointly deviate to High effort, they do not affect prices, which remain $(1, 1)$. Hence their income after production is 2, whence optimal consumption is $(1, 1)$ and utility is $H(1) = (9\sqrt{3} + 2)/30 > \sqrt{3}/3$. However, this joint deviation is itself subject to individual deviations: either agent could deviate back to L, obtaining income after production of $4/3$, leading to optimal consumption of $(2/3, 8/9)$ and utility of $4\sqrt{3}/9 > (9\sqrt{3} + 2)/30$. Hence deviation is not a self-enforcing agreement. $\diamondsuit$

6.1 Pareto Optimality

As the examples above show, mis-coordination and consequent Pareto ranking of equilibria can be robust phenomena. However, full Pareto optimality does obtain when moral hazard, adverse selection and idiosyncratic risk are all absent. To formalize this, let μ be an economy, and let μ^1, μ^2 be feasible configurations for the economy μ. Say μ^1 *weakly Pareto dominates* μ^2 if for every $(\Phi, u, (e, \epsilon), s) \in \mathcal{C}$ and every $(\phi^i, x^i) \in \mathbb{F} \times \mathbb{R}_+^{L\Omega}$ for which $(\Phi, u, (e, \epsilon), s, \phi^i, x^i) \in \operatorname{supp} \mu^i$ (each i), we have

$$E^1 u(x^1, \phi^1, \cdot, \cdot) \geq E^2 u(x^2, \phi^2, \cdot, \cdot) \tag{1}$$

(where E^i denotes expectations according to μ^i), with strict inequality for at least one pair $(\phi^1, x^1), (\phi^2, x^2)$. Say μ^1 *strictly Pareto dominates* μ^2 if (1) holds with strict inequality for every pair $(\phi^1, x^1), (\phi^2, x^2)$. Say that μ is *fully Pareto optimal* if it is not weakly Pareto dominated by any feasible configuration.

THEOREM 3 *If there is no moral hazard (i.e., each A^j_r is a singleton), no adverse selection (i.e.,* $\mathbf{S}$ *is a singleton), and no idiosyncratic risk (i.e., each Ω^j is a singleton) then equilibrium states are fully Pareto optimal.*

The proof (deferred to the Appendix) is simple, but does not parallel the familiar proof for the Arrow–Debreu–McKenzie model because firms do not maximize profit.

7 An Equilibrium Refinement

The definition of common beliefs equilibrium rules out contradictory beliefs, but not absurd beliefs — for example that others use dominated strategies. Here I offer a refinement that does rule out (at least some) absurd beliefs.

The refinement is suggested by extensive-form trembling hand perfection. In an extensive form game, trembling hand perfection is a particular way of imposing discipline on beliefs about behavior (equivalently, plans) at nodes that are not reached; it does so by allowing only beliefs that are limits of optimal behavior in slight perturbations for which those nodes *are* reached. Similarly, the refinement offered here imposes discipline on beliefs about behavior in firm types that are not formed; it does so by allowing only beliefs that are limits of optimal behavior in slight perturbations for which those firm types *are* formed. (Because formation of firms is limited by budget constraints as well as by individual choice, this notion may be applicable to beliefs only in some of the firms that do not form.)

Dubey & Geanakoplos (2002) and Dubey, Geanakoplos & Shubik (2005) offer refinements that also use perturbations to rule out equilibria supported by "unreasonable" beliefs. However, their perturbations require a small set of agents (a government, perhaps) to behave in a particular, suboptimal, way: in Dubey & Geanakoplos (2002) the government enters every pool and makes full deliveries; in Dubey, Geanakoplos & Shubik (2005) the government borrows in every market and never defaults. By contrast, the refinement offered here requires entry decisions that may be sub-optimal, but behavior within firms that is fully optimal.

7.1 Population Perfect Equilibrium

Trembling hand perfection is characterized by a number of equivalent conditions, including the following: A behavior strategy profile σ for an extensive form game is trembling hand perfect if there is a sequence $\{\tau_n\}$ of strictly positive real numbers and a sequence $\{\sigma_n\}$ of totally mixed behavior strategy profiles such that $\tau_n \to 0$, $\sigma_n \to \sigma$ and in each profile σ_n, each agent puts weight at most τ_n on actions that are not part of a best response. I replace the requirement that all agents play totally mixed strategies, which guarantees that all nodes are reached with positive probability, by the requirement that all feasible firms form; I replace the requirement that each agent puts small weight on actions that are not part of a best response by the requirement that the set of agents who do not optimize is small. In addition I require that all agents optimize conditional on their firm-role choices.

To formalize this, fix commodities, firm types and a population distribution λ. Let $B(p, w, e)$ be the budget set, given prices p, wages w, and endowment e. For each firm $j \in J$, role $r \in R^j$ and firm-action profile $\psi \in \mathbb{F}$, set

$$\begin{aligned} B_{jr}(p, w, e) &= \{(\tilde{x}, \phi) \in B(p, w, e) : \rho^j(\phi) = r\} \\ B_{\psi}(p, w, e) &= \{(\tilde{x}, \phi) \in B(p, w, e) : \rho^j(\phi) = \rho^j(\psi), \text{ all } j\} \end{aligned}$$

$B_{jr}(p, w, e)$ is the set of budget-feasible choices that include choosing role r in firm j; $B_{\psi}(p, w, e)$ is the set of budget-feasible choices for which firm-role choices agree with those in ψ. Given prices p^* and wages w^*, say that firm j is *budget feasible in the large* at p^*, w^* if

$$\lambda\Big(\{\theta \in \Theta : B_{jr}(p, w, e_\theta) \neq \emptyset\}\Big) > 0$$

for each $r \in R^j$. (That is, formation of the firm j is consistent with budget constraints for some non-null set of agents.) Say j is *locally feasible* at p^*, w^* if it is budget feasible in the large for all p, w in some neighborhood of p^*, w^*.

The equilibrium (p^*, w^*, μ, β) for λ is *population perfect* if for each $\tau > 0$ there is a tuple $(p_\tau, w_\tau, \mu_\tau, \beta_\tau)$ — a *τ-tremble of* (p^*, w^*, μ, β) — of prices, wages, configurations and beliefs such that

- μ_τ is a feasible configuration
- if firm j is locally feasible at p^*, w^* then firm j is active in μ_τ
- $\mu_\tau\Big(\{(\theta, \tilde{x}, \phi) : (\theta, \tilde{x}, \phi) \notin B(p_\tau, w_\tau, e_\theta)\}\Big) = 0$

- $\mu_\tau\Big(\{(\theta,\tilde{x},\phi) : (\theta,\tilde{x},\phi) \text{ is not optimal in } B(p_\tau,w_\tau,e_\theta)\}\Big) < \tau$
- $\mu_\tau\Big(\{(\theta,\tilde{x},\phi) : (\theta,\tilde{x},\phi) \text{ is not optimal in } B_\phi(p_\tau,w_\tau,e_\theta)\}\Big) = 0$
- $|p_\tau - p| < \tau$, $|w_\tau - w| < \tau$, $\|\mu_\tau - \mu\| < \tau$, $\|\beta_\tau - \beta\| < \tau$

Roughly: (p^*, w^*, μ, β) is population perfect if it is the limit of approximate equilibria in which all firms that are locally feasible at p^*, w^* are active. The notion of approximate equilibrium is most agents optimize in their budget sets and all agents optimize subject to their firm-role choices. Note that any equilibrium in which all firms are active is population perfect.

Theorem 3 *If endowments are bounded and all goods are represented in the aggregate, then a population perfect equilibrium exists.*

7.2 Bertrand Competition and Zero Profits

A simple example will illustrate the impact of population perfection. The example is suggested by Bennardo & Chiappori (2003), which examines the effects of Bertrand competition in a world of moral hazard, when utility is not separable in consumption and effort. Their conclusion is that it is possible to have a competitive equilibrium with free entry at which firms make positive profits; the model presented here suggests a different conclusion.

Example 6 There is a single good. There are two types of agents, Workers and Investors. There is a single technology T, which requires the participation of a single Worker and a single Investor; the Worker can choose to exert Low effort or High effort; the Investor has a single action. Consequences $\Omega = \{G, B\}$ depend stochastically on effort; output depends on consequences:

$$\begin{array}{rclcrcl} \pi(G|H) &=& \frac{3}{5} &,& \pi(G|L) &=& \frac{2}{5} \\ y(G) &=& 20 &,& y(B) &=& -20 \end{array}$$

Workers care about consumption and effort; Investors care only about consumption:

$$\begin{aligned} u_W(x;0) &= x \\ u_W(x;L) &= x \\ u_W(x;H) &= \frac{5}{2}\sqrt{x} \\ u_I(x) &= x \end{aligned}$$

Worker endowments are $e_W = 4$; Investor endowments are $e_I = 30$. Workers and Investors can belong to at most one firm.

As in Example 4, we may "normalize" contracts so that $C(W,B) \equiv 0$. For $a \in \mathbb{R}$, define the contract C^a by $C^a(W,H) = a$. I claim that if $\mathcal{C}$ is *any* set of contracts containing C^5 (which, as we shall see below, is, among the contracts consistent with High effort, the one most preferred by Workers) then there is a unique (up to determinacy of shadow prices for inactive contracts) population perfect equilibrium; in this equilibrium, only the contract C^5 is active, $w^5(W) = w^5(I) = 0$, all Workers and an equal mass of Investors belong to the firm (T, C^5), 5/7 of Workers exert High effort and 2/7 of Workers exert Low effort, Worker utility is 13/2 (above their autarkic utility of 4) and Investor utility is 30 (equal to their autarkic utility).

To see this, fix a (finite) set of contracts $\mathcal{C}$ containing C^5. Let $\mathcal{H}$ be the contingent consumptions (x_G, x_B) for which it is weakly optimal for the Worker to exert High effort. This is the set of solutions to the inequality

$$\frac{3}{5}\left(\frac{5}{2}\sqrt{x_G}\right) + \frac{2}{5}\left(\frac{5}{2}\sqrt{x_B}\right) \geq \frac{1}{2}x_G + \frac{1}{2}x_B \tag{2}$$

A straightforward argument shows that $\mathcal{H}$ is compact. On the boundary Γ of $\mathcal{H}$, the Worker is indifferent between High and Low effort, and expected utility along Γ is $Eu_W(x_G, x_B) = \frac{1}{2}x_G + \frac{1}{2}x_B$. Calculation shows that expected utility is maximized at $x_G = 9, y_G = 4$, and $Eu_W(9,4) = \frac{13}{2}$; see Figure 5.

Now consider a population perfect equilibrium (guaranteed to exist by Theorem 3); let $w^5 = w^5(W)$ be the equilibrium wage paid to the Worker in the firm (T, C^5). The argument is in several steps:

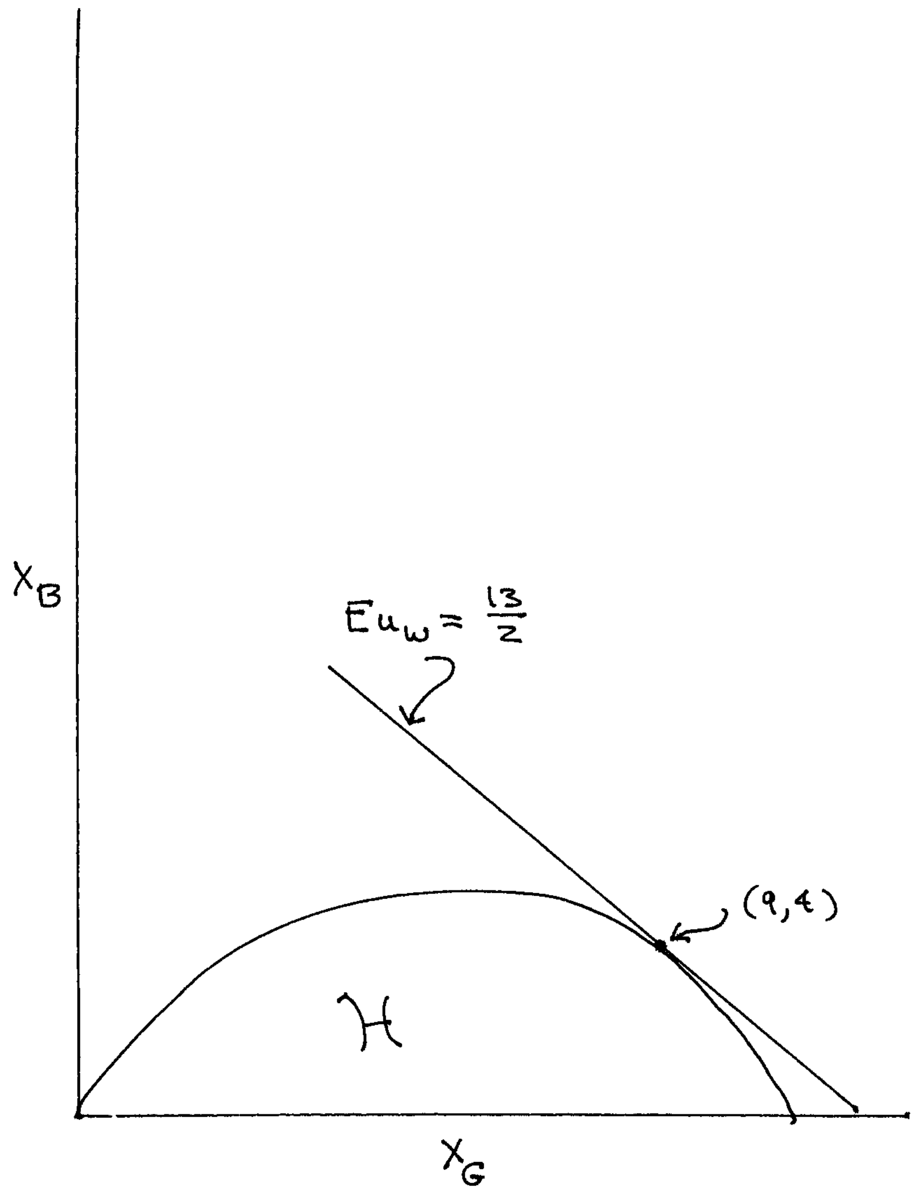

Figure 5: The High-effort Region

Step 1 Because there are more Investors than Workers, some Investors do not enter a firm and consume their endowment. Hence, all Investors must obtain the utility of consuming their own endowment. Hence equilibrium wages must have the property that all contracts yield Investors expected returns at most 0, and all active contracts yield Investors expected returns exactly 0.

Step 2 The wage $w^5 = 0$. To see this, consider the other possibilities:

- $w^5 > 0$: A Worker who accepted the contract C^5 would enjoy contingent consumption that is the sum of endowment, wage, and the contractual transfer:

$$(4,4) + (w^5, w^5) + (5,0) = (9 + w^5, 4 + w^5)$$

  Because $w^5 > 0$, this contingent consumption would yield expected utility greater than 13/2. This need not be an optimal choice for a Worker, but any optimal choice must yield expected utility at least this large, and hence must yield contingent consumption strictly above the indifference curve $Eu_W = 13/2$, and hence outside the region $\mathcal{H}$. This means that the optimal choice of Workers would be to accept *some* contract and exert Low effort. If supply and demand for this contract (or contracts) are equal, then Investors who accept this contract obtain strictly negative returns, which is worse than they could do by not accepting any contract. This is a contradiction.

- $-4 < w^5 < 0$: If this firm were active, Workers would exert High effort; because Investor beliefs are correct, Investors would necessarily believe Workers would exert High effort. However, in any sufficiently small tremble, this firm *would* be active and the wage would be in the open interval $(-4, 0)$, so that Investors *would* necessarily believe that Workers would exert High effort. Population perfection thus guarantees that in the *given* equilibrium, Investors believe that Workers would exert High effort in the firm (T, C^5). (If we did not impose population perfection, there would be equilibria in which (T, C^5) is inactive because Investors believe that Workers would exert Low effort, even though that would be a dominated strategy.) But then Investors would expect positive returns from entering the firm (T, C^5), and all Investors would demand *some* firm in which they expect positive returns. Because there are more Investors than Workers, this is a contradiction.

- $w^5 \leq -4$: Whatever their beliefs, Investors' expected returns from this contract would be at least $\frac{1}{2}(15)+\frac{1}{2}(-20)-w^5 > 0$. As before, this would imply that all Investors demand some firm in which they expect positive returns, which is again a contradiction.

Hence $w^5 = 0$, as claimed.

Step 3 Workers who accept C^5 at the wage $w^5 = 0$ can obtain utility $\frac{13}{2}$ at the contingent consumption $(9, 4)$. In any contract other than C^5, workers can only obtain utility this large if their contingent consumption is above the indifference curve $Eu_W = \frac{13}{2}$, which would induce Low effort, so that Investor expected returns would be negative. Hence no contract other than C^5 can be active. Since $\frac{13}{2}$ is above their autarkic utility, all Workers accept the contract C^5. In order that Investor expected returns be 0, it is necessary that $\frac{5}{7}$ of Workers exert High effort and $\frac{2}{7}$ of Workers exert Low effort.

Bennardo & Chiappori (2003) model such environments as two-stage games: in the first stage, Investors agents offer contracts and wages; in the second stage, Workers choose contracts and effort. Bennardo & Chiappori focus on the (subgame perfect) equilibrium in which all Investors offer the contract C^5 and the wage $w^5 = 0$, and all Workers accept and exert High effort. In this equilibrium, those Investors who attract a Worker make positive profit, and those Investors who do not attract a Worker make zero profit. The point of Bennardo & Chiappori (2003) is that the familiar Bertrand undercutting story does not work: Investors who do not attract a Worker could only attract one by offering a better contract/wage, but a better contract/wage would induce Low effort and thus leads to a loss, so such an offer would not be an improving deviation.

To understand why population perfect equilibrium is different from the Bennardo-Chiappori equilibrium, two points should be made. The first is that Bennardo & Chiappori (2003) follow a common convention in *assuming* that Workers who are indifferent choose to exert High effort. However, there are many other subgame perfect equilibria as well, including one in which 5/7 of Workers exert High effort and 2/7 of Workers exert Low effort; at *this* equilibrium — which corresponds to the population perfect equilibrium — all Investors make 0 expected profit.

The second point is that the Bennardo-Chiappori conclusion of positive profits at equilibrium is an artifact of the assumption that all Workers are exactly identical. To make the point most simply, suppose that only the contract C^5 can be offered and, following Harsanyi (1973), consider a perturbation in which the Workers are not precisely identical, but are distinguished by endowments that are uniformly distributed on some small interval $[a, b]$, where $a < 4 < b$ and $|b - a|$ is small.

First look for a population perfect equilibrium. If the wage w is offered, Workers with endowment e will accept and exert High effort if $w + e < 4$ and accept and exert Low effort if $w + e > 4$. Because there are more Investors than Workers, expected returns to Investors must be 0, so the equilibrium wage w^* solves the equation:

$$0 = \frac{4 - a - w^*}{b - a}(1 - w^*) + \frac{b - 4 + w^*}{b - a}(-5 - w^*) = 0 \qquad (3)$$

($1 - w^*$ is the expected net return to the Investor, conditional on obtaining a Worker who exerts High effort; $(4 - a - w^*)/(b - a)$ is the fraction of Workers who exert High effort, etc.) Solving, we find that

$$w^* = \frac{24 - a - 5b}{7 + a - b}$$

(At equilibrium, only Workers who have endowment exactly $e^* = 4 - w^*$ are indifferent between exerting High effort and Low effort; of course this is a set of measure 0.) The fraction of Workers who exert High effort is $(4 - a - w^*)/(b - a)$ and the fraction who exert Low effort is $(b - 4 + w^*)/(b - a)$. A little algebra shows that as $a, b \to 4$, this equilibrium converges to that identified above: $w^* = 0$, 5/7 of the Workers exert High effort and 2/7 of the Workers exert Low effort.

Now look for a subgame perfect equilibrium for the two-stage game. Say the equilibrium wage proposal is $\overline{w}$. If $\overline{w} > w^*$ then Investors make losses, which cannot occur at equilibrium. On the other hand, if $\overline{w} < w^*$ so Investors make profits, then the usual Bertrand undercutting argument *does* apply: Investors who have not attracted a Worker could offer a wage w between $\overline{w}$ and w^*; such a wage would attract a Worker who would exert High effort, and hence would make positive profit for the Investor. It follows that the equilibrium wage must be $\overline{w} = w^*$. As noted above, at that wage, Workers who have endowment $e < 4 - w^*$ exert High effort and Workers who have endowment $e > 4 - w^*$ exert Low effort; only workers who have endowment $e = 4 - w^*$ are indifferent between exerting High and Low effort, and this is a set of measure 0. Hence the equilibrium of the two-stage game coincides with the population perfect equilibrium — and as $a, b \to 4$, this equilibrium converges to the zero-profit equilibrium, which is thus the unique subgame perfect equilibrium that is "stable" with respect to small perturbations in preferences. $\diamond$

8 Conclusion

This paper offers a model in which agents trade anonymously with the market but interact strategically in small productive groups (firms). Contracts are determined endogenously at equilibrium. The model accommodates moral hazard, adverse selection, and screening. Mis-coordination and Pareto ranked equilibria are robust phenomena.

Several extensions of the basic model seem natural and worthwhile. The most obvious would be to relax the finiteness requirements to allow allow for a continuum of types of firm, for a continuum of possible production states in each firm, and for a continuum of actions in each role. Such extensions seems conceptually straightforward, although the modeling and technical details involved (especially wages in a continuum of firms) seem daunting.

A more important extension involves time. The model constructed in this paper is static/atemporal; it seems important to understand a properly dynamic version of the model. Among the important issues in such a dynamic model are the acquisition of skills, the possibility that choices are conditioned on previous actions and on the realization of previous uncertainty, the possibility of both short-term and long-term relationships, and the availability of intertemporal financial markets.

Appendix: Proofs

Before beginning the proofs of the Theorems, I collect some technical material. The proofs of the following two lemmas are straightforward and omitted.

LEMMA 1 *Θ is a Borel subset of $2^{\mathbb{F}} \times \mathbb{R}^L_+ \times \mathbb{F} \times \mathcal{V} \times \mathbf{S}^{\mathbb{F}}$, which is a complete separable metric space.*

LEMMA 2 *Let $K \subset \Theta$ be a compact set of agent characteristics.*

(i) For every consumption level $c^* \in \mathbb{R}_+$ there is a utility level $u^* < \infty$ such that if $\|\tilde{x}\| = \sup_\omega |\tilde{x}(\omega)| \leq c^*$ then

$$u_\theta(\tilde{x}, \phi, \mathbf{s}, \mathbf{a}, \omega) \leq u^*$$

for every $\theta \in K$, $\phi \in \Phi_\theta$, $\mathbf{s} \in \mathbf{S}$, $\mathbf{a} \in \mathbf{A}$, $\omega \in \Omega$

(ii) For every utility level $u^* < \infty$ there is a consumption level $c^* \in \mathbb{R}_+$ such that if $|x| \geq c^*$ then:

$$u_\theta(x, 0, \mathbf{s}, \mathbf{a}, \omega) \geq u^*$$

for every $\theta \in K$, $\mathbf{s} \in \mathbf{S}$, $\mathbf{a} \in \mathbf{A}$, $\omega \in \Omega$.

Now fix a set of firms. Define $\rho^* : \mathbb{F} \to \mathbb{R}^{\mathbb{M}}$ by

$$\rho^*(\phi)(j, r) = \begin{cases} 1 \text{ if } r = \rho^j(\phi) \\ 0 \text{ if } r \neq \rho^j(\phi) \end{cases}$$

Set

$$\nabla = \{\eta \in \mathbb{R}^{\mathbb{M}} : \eta^j_r = \eta^j_{r'} \text{ for all } j \in J, r, r' \in R^j\}$$

The following characterization of consistent distributions is straightforward and left to the reader.

LEMMA 3 *The distribution μ on $\Theta \times \mathbb{R}^{L\Omega}_+ \times \mathbb{F}$ is consistent if and only if $\int \rho^*(\phi)\, d\mu \in \nabla$.*

The next two lemmas are derived from Ellickson, Grodal, Scotchmer & Zame (1999): the first translates Lemma 7.2 to the present language (so I omit the proof); the second improves on Step 7 in the proof of Theorem 6.1.

LEMMA 4 *There is a constant* $b > 0$ *such that if* $w \in W$, $\phi_1, \ldots, \phi_I \in \mathbb{F}$ *and there are* $\alpha_1, \ldots, \alpha_I > 0$ *for which* $\sum_i \alpha_i \rho^*(\phi_i) \in \nabla$ *then*

$$\min_i w \cdot \phi_i \leq -b \max_i w \cdot \phi_i$$

LEMMA 5 *For every* $E > 0$ *there is a constant* b_E *such that if* ν *is an economy in which individual endowments are bounded by* E *and* (p, w, β, μ) *is an equilibrium for* ν *then there is a wage* $w^* \in \mathbb{W}$ *such that* (p, w^*, β, μ) *is also an equilibrium for* ν *and* $\|w^*\| = \sup_{j,r} |w^*(j, r)| \leq b_E$.

Proof The proof is by contradiction. Suppose the assertion is not true; then for every n there is an economy λ_n in which individual endowments are bounded by E and an equilibrium $(p_n, w_n, \beta_n, \mu_n)$ for λ_n with the property that there does not exist a wage w_n^* such that $(p_n, w_n^*, \beta_n, \mu_n)$ is an equilibrium for λ_n and $\|w_n^*\| \leq n$. Passing to a subsequence if necessary, assume that for each $\phi \in \mathbb{F}$ the sequence $\{w_n \cdot \phi\}$ has a limit — perhaps infinite. Set

$$\begin{aligned} \mathbb{F}^0 &= \{\phi \in \mathbb{F} : w_n \cdot \phi \to v(\phi) \text{ for some } v(\phi) \in \mathbb{R}\} \\ \mathbb{F}^+ &= \{\phi \in \mathbb{F} : w_n \cdot \phi \to +\infty\} \\ \mathbb{F}^- &= \{\phi \in \mathbb{F} : w_n \cdot \phi \to -\infty\} \end{aligned}$$

Define a linear map $F : \mathbb{W} \to \mathbb{R}^{\mathbb{F}^0}$ by $F(w)(\phi) = w \cdot \phi$; let $\operatorname{ran} F$ be the range of F. Use the fundamental theorem of linear algebra to find a linear map $G : \operatorname{ran} F \to \mathbb{W}$ such that the composition FG is the identity on $\operatorname{ran} F$. Because G is continuous, there is a constant $\|G\|$ so that $|G(f)| \leq \|G\| \, |f|$ for every $f \in \mathbb{R}^{\mathbb{F}^0}$.

Let b be the constant given by Lemma 4. Set

$$\overline{D} = \sup\{C^j(r, \omega^j, p) : j \in J, r \in R^j, \omega^j \in \Omega^j, p \in \Delta\}$$

Choose n_0 sufficiently large so that if $n \geq n_0$ then

$$\begin{array}{rcll} |w_n \cdot \phi - v(\phi)| & \leq & +1 & \text{for every} \quad \phi \in \mathbb{F}^0 \\ w_n \cdot \phi & \geq & +(1+b^{-1})(1+\|G\|)(E+J\overline{D}) & \text{for every} \quad \phi \in \mathbb{F}^+ \\ w_n \cdot \phi & \leq & -(1+b^{-1})(1+\|G\|)(E+J\overline{D}) & \text{for every} \quad \phi \in \mathbb{F}^- \end{array}$$

For each $n \geq n_0$ set

$$w_n^* = GF(w_n) - GF(w_{n_0}) + w_{n_0}$$

I claim that $(p_n, w_n^*, \beta_n, \mu_n)$ is an equilibrium for λ_n if $n \geq n_0$. To see this, it is only necessary to show that optimal choices given p_n, w_n, β_n are optimal choices when wages are w_n^*.

The first task is to show that (almost all) optimal choices given p_n, w_n, β_n are feasible when wages are w_n^*. To this end, for each index $n \geq n_0$ and each $\phi \in \mathbb{F}$, set

$$\begin{array}{rcl} T(\phi) & = & \{(\xi, \tilde{x}, \psi) \in \Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F} : \psi = \phi\} \\ \mathbb{F}(\mu_n) & = & \{\phi \in \mathbb{F} : \mu_n(T(\phi)) > 0\} \end{array}$$

I assert that $\mathbb{F}(\mu_n) \subset \mathbb{F}^0$. To see this, let $\phi \in \mathbb{F}(\mu_n)$. If $\phi \in \mathbb{F}^-$ then $w_n \cdot \phi < -E - J\overline{D}$. However, in each of the economies λ_n, individual endowments are bounded by E and profit shares in each firm are bounded by $\overline{D}$, so no agent can have income from endowment and profit distributions greater than $E + J\overline{D}$ (at any prices). Hence, no agent can afford to pay wages greater than $E + J\overline{D}$. Put differently: $w_n \cdot \phi \geq -E - J\overline{D}$, which means $\phi \notin \mathbb{F}^-$. On the other hand, if $\phi \in \mathbb{F}^+$ then Lemma 3 and consistency of μ_n imply

$$\sum_{\phi \in \mathbb{F}(\mu_n)} \mu_n(T(\phi, \mu_n)) = \sum_{\phi \in \mathbb{F}} \mu_n(T(\phi, \mu_n)) = \int \rho^*(\phi)\, d\mu_n \in \nabla$$

and Lemma 4 implies that there is some $\phi' \in \mathbb{F}(\mu_n)$ such that $w_n \cdot \phi' < -E - J\overline{D}$, which again is a contradiction. It follows that $\mathbb{F}(\mu_n) \subset \mathbb{F}^0$, as asserted. However, because FG is the identity on $\operatorname{ran} F$,

$$F(w_n^*) = FGF(w_n) - FGF(w_{n_0}) + F(w_{n_0}) = F(w_n)$$

Equivalently, $w_n^* \cdot \phi = w_n \cdot \phi$ for every $\phi \in \mathbb{F}^0$. Hence, choices that are optimal given p_n, w_n, β_n have the same cost when wages are w_n^*, and in particular are feasible when wages are w_n^*.

The second task is to show that optimal choices given p_n, w_n^*, β_n are feasible when wages are w_n. To see this, let $(\tilde{x}, \phi)$ be optimal (for some characteristics) given p_n, w_n^*, β_n. Observe that ϕ cannot belong to $\mathbb{F}^-$, because if it were then the construction would guarantee that $w_n^* \cdot \phi < E - J\overline{D}$, and again ϕ could not be part of a feasible choice at at prices p_n and wages w_n^*. If $\phi \in \mathbb{F}^0$ then $w_n \cdot \phi = w_n^* \cdot \phi$, so budget feasibility of $(\tilde{x}, \phi)$ at p_n, w_n^* implies budget feasibility at p_n, w_n. Finally, if $\phi \in \mathbb{F}^+$ then the construction guarantees that $w_n^* \cdot \phi \leq w_n \cdot \phi$, so if $(\tilde{x}, \phi)$ is budget feasible given p, w_n^* it is certainly budget feasible given p, w_n, as desired.

Thus, $(p_n, w_n^*, \beta_n, \mu_n)$ is an equilibrium for λ_n. On the other hand, the definition of $\mathbb{F}^0$ implies that $\|F(w_n)\| \leq \max_\phi v(\phi)$, so

$$\|w_n^*\| \leq \|G\| \max_\phi v(\phi) + \|w_{n_0}\| + \|G\| \, \|F(w_{n_0})\|$$

Because the right hand side is fixed, it is less than n for n sufficiently large, which contradicts the supposition that there does not exist a wage structure w_n^* such that $(p_n, w_n^*, \beta_n, \mu_n)$ is an equilibrium for λ_n and $\|w_n^*\| \leq n$. This contradiction completes the proof. ■

With these preliminaries in hand I turn to the proof of Theorem 1.

Proof of Theorem 1 The idea of the proof is to construct artificial economies for which prices can be bounded away from 0 and wages can be bounded above and below, construct common beliefs equilibria for these artificial economies, and then construct a common beliefs equilibrium for the given economy as a limit of common beliefs equilibria for the artificial economies. The proof proper is in a number of steps. For each $\varepsilon > 0$, Step 1 constructs an artificial economy; Step 2 constructs a compact convex space of prices, wages, beliefs and choice distributions; Step 3 constructs a correspondence from this space to itself; Step 4 constructs a fixed point for this correspondence, and shows that this fixed point is an equilibrium for the artificial economy. To take a limit, it is necessary that equilibrium wages for these artificial economies remain bounded as $\varepsilon \to 0$. The equilibrium wages constructed in Step 4 may not have this property, but Step 5 shows that it is possible to modify the equilibria so that modified wages do remain bounded as $\varepsilon \to 0$. Step 6 shows that equilibrium prices in the artificial economies stay away from the boundary of the price simplex, and Step 7 shows that the limit of a subsequence of the modified equilibria for the artificial economies is an equilibrium for the given economy.

Step 1 The artificial economies are constructed to contain a few agents whose demands are easy to estimate and whose commodity demands are unsatisfiable when any prices are sufficiently close to 0 or any wages are sufficiently high or low. To accomplish this, write $\overline{e} = \int e \, d\lambda$ for the aggregate endowment; by assumption, $\overline{e} \gg 0$. For each $j \in J$, $r \in R^j$, choose and fix an arbitrary skill $s_{jr} \in S^j_r$ and action $a_{jr} \in A^j_r$. Define $\delta_{jr} \in \mathbb{F}$ by

$$\delta^{j'}_{jr} = \begin{cases} (r, a_{jr}) & \text{if} \quad j' = j \\ 0 & \text{if} \quad j' \neq j \end{cases}$$

For each j, r, define a characteristic $\theta_{jr} = (\Phi, e, u, \sigma) \in \Theta$ by

$$\Phi = \{0, \delta_{jr}\} \; , \; e = \overline{e}$$
$$u(x, \phi, \mathbf{s}, \mathbf{a}, \omega) = |x|$$

Now fix $\varepsilon > 0$. Write $\overline{R} = R^1 + \ldots + R^J$; this is the total number of roles in all firms. Set

$$\lambda_\varepsilon = (1 - \varepsilon)\lambda + \varepsilon \left(\frac{1}{\overline{R}} \right) \sum_{j,r} \delta_{jr}$$

Note that λ_ε is a probability measure, so defines an economy.

Step 2 The spaces of prices and wages are constructed so that the commodity demands of the artificial agents are impossibly large when prices or wages are on the boundary. To accomplish this, recall that contract payments are bounded, and set

$$\overline{D} = \sup\{|C^j(r, \omega^j, p)| : j \in J, r \in R^j, \omega \in \Omega^j, p \in \Delta\}$$

Write $M = \overline{e} + L\overline{D}\,\overline{R} + J \max_{j,\omega^j} |y^j(\omega^j)|$ (this will serve as an upper bound for the norm of supply). By assumption, $\overline{e} \gg 0$ so $\overline{e} \geq e_0(1, \ldots, 1) = e_o\mathbf{1}$ for some $e_0 > 0$. Set

$$t = \frac{\varepsilon e_0}{2LM + \overline{R}(M + \overline{D})}$$

Define trimmed price and wage spaces by

$$\begin{aligned} \Delta_\varepsilon &= \{p \in \Delta : p_\ell \geq t \text{ for all } \ell\} \\ \mathbb{W}_\varepsilon &= \{w \in \mathbb{W} : |w(j,r)| \leq \frac{1}{t} \text{ for all } j, r\} \end{aligned}$$

Note that $\Delta_\varepsilon, \mathbb{W}_\varepsilon$ are compact convex sets.

The compact space of beliefs and choice distributions is constructed so that choice distributions are consistent with the population measure λ_ε and with the bounds on prices and wages. To this end, let E be an upper bound for individual endowments, set

$$\overline{C} = \frac{1}{t}(E + J\overline{D}) + \frac{J}{t^2}$$

and let $\mathbb{M}_\varepsilon$ be the space of probability measures on $\Theta \times \mathbb{R}_+^{L\Omega} \times \mathbb{F}$ that are supported on $\Theta \times [0, \overline{C}]_+^{L\Omega} \times \mathbb{F}$ and have marginal on Θ equal to λ_ε. Write B for the space of beliefs; that is, probability measures on $\mathbf{S} \times \mathbf{A}$. Say that beliefs $\beta \in B$ and a choice distribution $\mu \in \mathbb{M}_\varepsilon$ are *compatible* if $\beta_r^j = \gamma_r^j(\cdot|\mu)$ whenever $\mu(T_r^j) \neq 0$. (See Subsection 4.5 and 4.6.) Finally, write

$$H_\varepsilon = \{(\beta, \mu) \in B \times \mathbb{M}_\varepsilon : \beta \text{ is compatible with } \mu\}$$

By construction, λ_ε is a Borel measure on Θ, which (by Lemma 1) is a Borel set in $2^{\mathbb{F}} \times \mathbb{R}_+^L \times \mathbb{F} \times \mathcal{V} \times \mathbf{S}^{\mathbb{F}}$. Hence λ_ε may be regarded as a probability measure on $2^{\mathbb{F}} \times \mathbb{R}_+^L \times \mathbb{F} \times \mathcal{V} \times \mathbf{S}^{\mathbb{F}}$ for which $\lambda_\varepsilon(\Theta) = 1$. Because $2^{\mathbb{F}} \times \mathbb{R}_+^L \times \mathbb{F} \times \mathcal{V} \times \mathbf{S}^{\mathbb{F}}$ is a complete separable metric space, λ_ε is tight and regular; it follows that H_ε is tight and hence relatively compact in the topology of weak convergence of measures. (See Billingsley (1968).). It is evident that H_ε is closed, hence compact, and it is convex and non-empty.

Step 3 The correspondence

$$\Psi : \Delta_\varepsilon \times \mathbb{M}_\varepsilon \times H_\varepsilon \to \Delta_\varepsilon \times \mathbb{M}_\varepsilon \times H_\varepsilon$$

is constructed so that, given (p, w, β, μ), the image set $\Psi(p, w, \beta, \mu)$ consists of all (p', w', β', μ') for which μ' is an optimal choice distribution given p, w, β, β' is compatible with μ', and p', w' maximize the value of aggregate excess demand at β, μ.

The first part of the construction is easy. Given (p, w, β, μ), the construction in Steps 1 and 2 guarantees that all budget feasible commodity choices for every characteristic θ in the support of λ_ε belong to $[0, \overline{C}]^{L\Omega}$. Hence I may define $\Psi_1(p, w, \beta, \mu)$ as the set of pairs $(\beta', \mu') \in H_\varepsilon$ such that the marginal of μ' on Θ is λ_ε, μ' is supported on tuples $(\theta, \tilde{x}, \phi)$ for which $(\tilde{x}, \phi)$ is optimal in θ's budget set, given prices p, wages w, and beliefs β, and β' is compatible with μ'.

Aggregate demand is defined exactly as in Subsection 4.6:

$$X(\mu) \quad = \quad \sum_{\phi \in \mathbb{F}} \int_{T(\phi)} \sum_{\omega \in \Omega} \int_{\mathbf{S}(\phi) \times \mathbf{A}(\phi)} \tilde{x}(\omega) \pi(\omega | \phi, \mathbf{s}, \mathbf{a}) \, d\gamma_\phi(\mathbf{s}, \mathbf{a} | \mu) \, d\mu \tag{4}$$

Because μ may not be consistent, the definition of aggregrate supply is more roundabout. Write $\mathbf{1} = (1, \ldots, 1) \in \mathbb{R}^L_+ \subset \mathbb{R}^{L\Omega}_+$. For each firm j, role r, price p define

$$y^j(\omega, r, p) = \begin{cases} y^j(\omega) + [C^j(1, \omega, p) - p \cdot y^j(\omega)] \, \mathbf{1} & \text{if} \quad r = 1 \\ \\ C^j(r, \omega, p) \mathbf{1} & \text{if} \quad r \neq 1 \end{cases} \tag{5}$$

Notice that

$$\begin{aligned} p \cdot y^j(\omega, r, p) \quad &= \quad C^j(r, \omega, p) \quad \text{for each } r \\ \sum_r y^j(\omega, r, p) \quad &= \quad y^j(\omega) \end{aligned} \tag{6}$$

Define aggregate output by

$$Y(\mu, p) = \sum_{\phi \in \mathbb{F}} \int_{T(\phi)} \sum_{\omega \in \Omega} \int_{\mathbf{S}(\phi) \times \mathbf{A}(\phi)} y^j(\omega, r, p) \pi(\omega | \phi, \mathbf{s}, \mathbf{a}) \, d\gamma_\phi(\mathbf{s}, \mathbf{a} | \mu) \, d\mu \tag{7}$$

and aggregate supply as $\overline{e} + Y(\mu, p)$. (This definition of aggregate output agrees with the definition in Subsection 4.5 when μ is consistent.)

Define $z(p, w, \beta, \mu) \in \mathbb{R}^{L\Omega}$ and $\zeta(\mu) \in \mathbb{R}^{\mathbb{M}}$ by

$$\begin{aligned} z(p, w, \beta, \mu) &= X(\mu) - \bar{e} - Y(\mu, p) \\ \zeta^j_r(p, w, \beta, \mu) &= \mu(T^j_r) - \frac{1}{R^j} \sum \mu(T^j_r) \end{aligned} \tag{8}$$

The quantity z plays the role of excess commodity demand, and the quantity ζ plays the role of excess firm-role demand.

Define $\Psi_2(p, w, \beta, \mu)$ to be the set of price-wage pairs that maximize the value of excess demand:

$$\Psi_2(p, w, \beta, \mu) = \arg\max \{p' \cdot z(p, w, \beta, \mu) - w' \cdot \zeta(p, w, \beta, \mu) : p' \in \Delta_\varepsilon, w' \in \mathbb{W}_\varepsilon\}$$

(By construction, both wages and excess firm-role demand belong to $\mathbb{R}^{\mathbb{M}}$; $w' \cdot \zeta(p, \beta, \mu)$ is the ordinary inner product. The quantity $w' \cdot \zeta(p, \beta, \mu)$ is subtracted, in order to be consistent with the sign convention for wages.)

Finally, the correspondence Ψ is defined as the product of Ψ_2 with Ψ_1:

$$\Psi(p, w, \beta, \mu) = \Psi_2(p, w, \beta, \mu) \times \Psi_1(p, w, \beta, \mu)$$

Step 4 It is straightforward to check that Ψ is an upper-hemi-continuous correspondence, and has compact, convex, non-empty values. (Continuity of Ψ relies on upper-hemi-continuity of individual demand, which is a consequence of Assumption A1.) Kakutani's fixed point theorem implies that Ψ has a fixed point $(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$; I claim that $(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$ is an equilibrium for λ_ε.

To see this, write $z_\varepsilon = z(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$, $\zeta_\varepsilon = \zeta(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$. It is immediate from the definitions of budget sets, aggregate consumption, the quantities $y^j(\omega, r, p)$, aggregate output and excess demands (equations (4), (5), (7), (8)) that Walras's Law holds in the aggregate:

$$p_\varepsilon \cdot z_\varepsilon - w_\varepsilon \cdot \zeta_\varepsilon = 0 \tag{9}$$

Next I show that p_ε does not belong to the boundary of the price simplex Δ_ε. To see this, note first that because $(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$ is a fixed point, Walras's Law, equation (9), implies that the maximum value of excess demand is

0. Because we can always take $w' = 0$ it follows that the maximum value of commodity excess demand is at most 0. If p_ε is on the boundary of Δ_ε then some commodity ℓ has price $(p_\varepsilon)_\ell = t$. Because the artificial agents have endowment $\overline{e}$, they have wealth at least $p_\varepsilon \cdot \overline{e} \geq e_0 > 0$, and hence can purchase at least e_0/t units of good ℓ, which would yield utility at least e_0/t. In view of the nature of their utility functions, their commodity demand must be at least e_0/t in expectation. Hence the norm of total demand of the artificial agents is at least $\varepsilon e_0/t$. Because every agent's demand is non-negative and the norm of aggregate supply is bounded by M, it follows that excess demand for at least one good is at least $\varepsilon e_0/tL - M$ and that excess commodity demand for other goods is at least $-M$. Hence the value of excess demand at the price $(1/L, \ldots, 1/L)$ is at least

$$\frac{1}{L}\left(\frac{\varepsilon e_0}{tL} - LM\right) > 0$$

This contradicts the fact that p_ε maximizes the value of commodity excess demand, so we conclude that p_ε does not belong to the boundary of the price simplex Δ_ε, as desired.

Because the value $p \cdot z_\varepsilon$ of commodity excess demand is linear in prices, and the maximum is attained for a price in the interior of the price simplex Δ_ε, the value of commodity excess demand must be independent of prices. Hence commodity excess demand must be $-c\mathbf{1}$ for some $c \geq 0$.

I now show that w_ε does not belong to the boundary of the wage space $\mathbb{W}_\varepsilon$. If w_ε is on the boundary of $\mathbb{W}_\varepsilon$, then some wage has absolute value $1/t$; because wages in each firm type sum to 0, it follows that there is some j, r so that $w(j,r) \geq 1/t\overline{R}$. Each artificial agent with characteristics θ_{jr} could choose role r in firm j, obtain income at least $1/tL - \overline{D}$, and spend this income on the cheapest private good, obtaining utility at least $1/tL - \overline{D}$. The expected total commodity demand of each such artificial agent must be at least $1/tL - \overline{D}$. Because the total mass of such agents is $\varepsilon/\overline{R}$, total demand of these artificial agents is at least

$$\frac{\varepsilon}{\overline{\overline{R}}}\left(\frac{1}{tL} - \overline{D}\right)$$

Because the total commodity demand of all other agents is no less than $-M$, it follows that excess demand for at least one commodity is strictly positive, which contradicts the fact that commodity excess demand is $-c\mathbf{1}$. It follows that w_ε does not belong to the boundary of the price simplex $\mathbb{W}_\varepsilon$.

Write

$$\nabla = \{\eta \in \mathbb{R}^{\mathbb{M}} : \eta^j_r = \eta^j_{r'} \text{ for all } j \in J, r, r' \in R^j\}$$

Because the value $-w' \cdot \zeta_\varepsilon$ of firm-role excess demand is linear in w', and the maximum is attained for a wage in the interior of W_ε, the value of firm-role excess demand must be independent of wages. Because W_ε is closed under multiplication by -1, firm-role excess demand must actually be 0. Because wages are are arbitrary, subject to the bound $1/t$ and the constraint that wages in each firm type sums to 0, firm-role excess demand ζ_ε must belong to ∇. Because λ and hence λ_ε are consistent, $f \in \nabla$, so $g(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon) \in \nabla$ as well. This implies that μ_ε is consistent.

Because firm-role excess demand ζ_ε belongs to ∇, the value of firm-role excess demand is 0 at every wage $w \in \mathbb{W}$. Because the maximum value of excess demand is 0, the maximum value of commodity excess demand z_ε must also be 0. Because commodity excess demand is $-c\mathbf{1}$, this implies that commodity excess demand $z_\varepsilon = 0$.

Straightforward algebraic manipulations, using the consistency of μ_ε together with the definitions and properties of $y^j(\omega, r, p)$ (equations (5) and (6)) show that $Y(\mu_\varepsilon, p) = Y(\mu_\varepsilon)$, and hence that $(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$ is an equilibrium for λ_ε.

Step 5 Individual endowments for the economy λ_ε are bounded by E. Apply Lemma 5 to find wages w^*_ε such that $(p_\varepsilon, w^*_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$ is an equilibrium for λ_ε and $\|w^*_\varepsilon\| \leq b_E$.

Step 6 The next step is to show that prices $\{p_\varepsilon\}$ are bounded away from the boundary of the price simplex Δ. Suppose this were not so; then there is a subsequence $\{p_{\varepsilon_n}\}$ converging to some price $\overline{p} \in \operatorname{bdy} \Delta$. I find a set of agents whose total demand at p_{ε_n} blows up as $n \to \infty$; this will provide a contradiction.

Because $\overline{p} \neq 0$, there is some index m with $\overline{p}_m > 0$. Write $\Theta^m = \{\theta : (e_\theta)_m > 0\}$ and note that $\lambda(\Theta^m) > 0$. For each index k, let Θ^m_k be the set of characteristics $\theta \in \Theta^m$ such that for every $x, \phi, \mathbf{s}, \mathbf{a}, \omega, p$ we have:

$$u(x, \phi, \mathbf{s}, \mathbf{a}, \omega) \quad \leq \quad k(1 + |x|)$$

Because $\Theta^m = \bigcup \Theta^m_k$, there is some index k' for which $\lambda(\Theta^m_{k'}) > 0$. Use regularity and tightness of λ to find a compact subset $K \subset \Theta^m_{k'}$ such that $\lambda(K) > 0$. For each $p \in \Delta$, the budget set of an agent with characteristics $\theta \in K$ includes $((p_m(e_\theta)_m / \overline{p}_\ell)\delta_\ell; 0)$. Lemma 2 implies that as $n \to \infty$, and uniformly on K, agents with characteristics in K can afford arbitrarily large consumption of good ℓ, hence obtain arbitrarily high utility, hence demand arbitrarily high expected consumption. Because $\lambda(K) > 0$ this implies that the total demand of agents with characteristics in K tends to ∞ (as $n \to \infty$). Because individual endowments and aggregate supply are uniformly bounded, independently of n, this is impossible. This proves that the family $\{p_\varepsilon\}$ is bounded away from the boundary of Δ.

Step 7 Because prices $\{p_\varepsilon\}$ are bounded away from the boundary of the simplex Δ, the family $\{\mu_\varepsilon\}$ of choice distributions is tight. Because prices $\{p_\varepsilon\}$ and wages w^*_ε lie in bounded sets and beliefs β_ε lie in the compact set B, the family $\{(p_\varepsilon, w^*_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)\}$ lies in a compact set, so some subsequence converges: say $(p_{\varepsilon_n}, w^*_{\varepsilon_n}, \beta_{\varepsilon_n}, \mu_{\varepsilon_n}) \to (p, w, \beta, \mu)$, with $p \in \Delta$. It is easily checked that (p, w, β, μ) is an equilibrium for the economy λ, so the proof is complete. ■

Proof of Theorem 2 Because there is no moral hazard and no adverse selection, consequences beliefs are degenerate, so will be suppressed. Let (p, w, μ) be an equilibrium for the economy λ. If some feasible configuration weakly dominates μ then (because preferences are strictly monotone in private goods and consuming no private goods is worse than consuming) it is always possible to re-distribute consumption to obtain a feasible configuration that strongly dominates, μ, so for the purpose of obtaining a contradiction, suppose σ is any such feasible configuration.

For each $\phi \in \mathbb{F}$, write

$$T_\phi = \{(\theta, x, \psi) : \psi = \phi\}$$

For notational convenience, write production y^j in firm j as the sum of equal production for each individual:

$$y^j = \sum_{r \in R^j} \frac{y^j}{R^j}$$

Feasibility means that consumption is no greater than the sum of endowment and net production:

$$\sum_{\phi \in \mathbb{F}} \int_{T_\phi} x \, d\sigma \leq \sum_{\phi \in \mathbb{F}} \int_{T_\phi} e \, d\sigma + \sum_{\phi \in \mathbb{F}} \int_{T_\phi} \sum_{\phi^j \neq 0} \frac{y^j}{R^j} \, d\sigma$$

Hence the value at price p of consumption is no greater than the sum of the value of endowment and net production:

$$\sum_{\phi \in \mathbb{F}} \int_{T_\phi} p \cdot x \, d\sigma \leq \sum_{\phi \in \mathbb{F}} \int_{T_\phi} p \cdot e \, d\sigma + \sum_{\phi \in \mathbb{F}} \int_{T_\phi} \sum_{\phi^j \neq 0} p \cdot \frac{y^j}{R^j} \, d\sigma$$

Because σ is feasible, it is consistent; keeping in mind that wages in each firm sum to 0 and that contractual payments in each firm sum to the value of net output, it follows that:

$$\sum_{\phi \in \mathbb{F}} \int_{T_\phi} p \cdot x \, d\sigma \leq \sum_{\phi \in \mathbb{F}} \int_{T_\phi} \left[p \cdot e + \sum_{\phi^j \neq 0} w^j + \sum_{\phi^j \neq 0} \sum_{r \in R^j} C^j(r, p) \right] \, d\sigma$$

Hence there is a $\phi \in \mathbb{F}$ and a compact set $K \subset T_\phi$ such that $\sigma(K) > 0$ and if $(\theta, x, \phi) \in K$ then

$$p \cdot x \leq p \cdot e + \sum_{\phi^j \neq 0} w^j + \sum_{\phi^j \neq 0} \sum_{r \in R^j} C^j(r, p)$$

Hence (x, ϕ) is budget feasible at prices p and wages w. Let $\widehat{K}$ be the projection of K into Θ, and note that $\lambda(\widehat{K}) \geq \sigma(K) > 0$. In view of the definition of strong Pareto dominance, we see that if $\theta \in \widehat{K}$ then there is some (x, ϕ) which is budget feasible at prices p and wages w and which strictly preferred to the actual choice in μ. Because this violates the equilibrium condition for (p, w, μ), we have arrived at a contradiction. We conclude that μ is Pareto optimal, as asserted. ■

Proof of Theorem 3 The argument is a modification of the proof of Theorem 1. Fix $\tau > 0$. To define a τ-tremble, fix $\varepsilon > 0$ and construct an artificial economy as in the proof of Theorem 1. Given prices and wages p, w, define

$$M_0(p, w) = \{(\theta, \tilde{x}, \phi) : (\tilde{x}, \phi) \text{ is optimal in } B(p, w, e_\theta)\}$$

For each firm $j \in J$ and role $r \in R^j$, define

$$M_{jr}(p, w) = \{(\theta, \tilde{x}, \phi) : (\tilde{x}, \phi) \text{ is optimal in } B_{jr}(p, w, e_\theta)\}$$

(Note that $M_{jr}(p, w)$ might be empty.) Finally, define

$$\begin{aligned}\alpha_j(p, w) = \sup\{\delta \in [0, 1) : &\text{ firm } j \text{ is budget feasible at every} \\ &p', w' \text{ with } |p - p'| < \delta, |w - w'| < \delta\}\end{aligned}$$

It is easy to see that α_j is continuous in p, w. Now define a modified demand correspondence $\widehat{\Psi}_1$ by letting $\widehat{\Psi}_1(p, w, \beta, \mu)$ be the set of pairs (μ', β') for which

- the marginal of μ' on Θ is λ
- $\mu' = \mu_0 + \sum_{j \in J} \sum_{r \in R^j} \mu_{jr}$, where μ_0 is supported on $M_0(p, w)$, μ_{jr} is supported on $M_{jr}(p, w)$, and

$$\|\mu_{jr}\| = \frac{1}{JR^j} \alpha_j(p, w) \tau$$

- β' is compatible with μ'

(Note that if $\alpha_j(p, w) = 0$ for all j then $\widehat{\Psi}_1$ coincides with the demand correspondence Ψ_1 defined in the proof of Theorem 1.) Define the correspondence Ψ_2 exactly as in the proof of Theorem 1, and define $\widehat{\Psi}$ to be the product

$$\widehat{\Psi}(p, w, \beta, \mu) = \Psi_2(p, w, \beta, \mu) \times \widehat{\Psi}_1(p, w, \beta, \mu)$$

Proceeding exactly as in the proof of Theorem 1, it can be seen that $\widehat{\Psi}$ has a fixed point $(p_\varepsilon, w_\varepsilon, \beta_\varepsilon, \mu_\varepsilon)$. Taking the limit of a subsequence as $\varepsilon \to 0$ yields a tuple $(p^\tau, w^\tau, \beta^\tau, \mu^\tau)$. We can find a sequence $\{\tau_n\} \to 0$ for which the corresponding tuples converge; call the limit $(p^*, w^*, \beta^*, \mu^*)$. As in the proof of Theorem 1, it can be seen that $(p^*, w^*, \beta^*, \mu^*)$ is a common beliefs equilibrium for λ and each $(p^{\tau_n}, w^{\tau_n}, \beta^{\tau_n}, \mu^{\tau_n})$ is a τ_n-tremble, so $(p^*, w^*, \beta^*, \mu^*)$ is a population perfect equilibrium. ■

References

George Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* 84 (1970), 488-500.

Armen Alchian and Harold Demsetz, "Production, Information Costs, and Economic Organization," *American Economic Review* 62 (1972), 777-795.

Kenneth Arrow and Gerard Debreu, "Existence of an Equilibrium for a Competitive Economy," *Econometrica* 22 (1954), 265-290.

Alberto Bennardo and Pierre-Andre Chiappori, "Bertrand and Walras Equilibria under Moral Hazard," *Journal of Political Economy* 111 (2003), 785-817.

Patrick Billingsley, *Convergence of Probability Measures*, New York: John Wiley and Sons (1968).

Alberto Bisin, John Geanakoplos, Piero Gottardi, Enrico Minelli and Herakles Polemarchakis, "Markets for Contracts," *Journal of Mathematical Economics* (forthcoming).

James Buchanan, "An Economic Theory of Clubs," *Economica* 33 (1965), 1-14.

Harold Cole and Edward C. Prescott, "Valuation Equilibrium with Clubs," *Journal of Economic Theory* 74 (1997), 19-39.

Gerard Debreu, *Theory of Value*, New Haven: Yale University Press (1959).

Pradeep Dubey and John Geanakoplos, "Competitive Pooling: Rothschild-Stiglitz Reconsidered," *Quarterly Journal of Economics* 117 (2002), 1529-1570.

PRADEEP DUBEY, JOHN GEANAKOPLOS AND MARTIN SHUBIK, "Default and Punishment in General Equilibrium," *Econometrica* 73 (2005), 1-37.

DARRELL DUFFIE AND YENENG SUN, "The Existence of Independent Random Matching," Stanford GSB Working Paper (2004).

DARRELL DUFFIE AND YENENG SUN, "The Exact Law of Large Numbers for Independent Random Matching," Stanford GSB Working Paper (2004).

BRYAN ELLICKSON, BIRGIT GRODAL, SUZANNE SCOTCHMER AND WILLIAM R. ZAME, "Clubs and the Market," *Econometrica* 67 (1999), 1185-1217.

BRYAN ELLICKSON, BIRGIT GRODAL, SUZANNE SCOTCHMER AND WILLIAM R. ZAME, "Clubs and the Market: Large Finite Economies," *Journal of Economic Theory* 101 (2001), 40-77.

BRYAN ELLICKSON, BIRGIT GRODAL, SUZANNE SCOTCHMER AND WILLIAM R. ZAME, "The Organization of Production, Consumption and Learning," in K. Vind, ed., *The Birgit Grodal Symposium*. Berlin: Springer-Verlag (2006). .

JOHN GEANAKOPLOS AND WILLIAM R. ZAME, "Collateral and the Enforcement of Intertemporal Contracts," UCLA Working Paper (2002).

SAYANTAN GHOSAL AND HERAKLES POLEMARCHAKIS, "Nash-Walras Equilibria," *Ricerche Economiche* 51 (1997), 31-40.

SANFORD J. GROSSMAN AND OLIVER HART, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy* 94 (1986), 691-719.

JOHN HARSANYI, "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2 (1973), 1-23.

S. Hart, W. Hildenbrand and E. Kohlberg, "Equilibrium Allocations as Distributions on the Commodity Space," *Journal of Mathematical Economics* 1 (1974), 159-166.

Bengt Holmstrom, "Moral Hazard and Observability," *Bell Journal of Economics* 10 (1979), 74-91.

Louis Makowski, "A General Equilibrium Theory of Organizations," Mimeograph (1976).

Louis Makowski and Joseph Ostroy, "Competitive Contractual Pricing with Transparent Teams," UCLA Working Paper (2004).

Jacob Marshack and Roy Radner, *Economic Theory of Teams*, New Haven: Yale University Press (1972).

Andreu Mas-Colell, "On a Theorem of Schmeidler," *Journal of Mathematical Economics* 13 (1984), 201-206.

R. Preston McAfee, "Mechanism Design by Competing Sellers," *Econometrica* 61 (1963), 1281-1312.

Lionel McKenzie, "On the Existence of General Equilibrium for a Competitive Market," *Econometrica* 27 (1959), 54-71.

Enrico Minelli and Herakles Polemarchakis, "Nash-Walras Equilibria of a Large Economy," *Proceedings of the National Academy of Sciences (USA)* 97 (2000).

Edward C. Prescott and Robert M. Townsend, "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica* 52 (1984a), 21-45.

Edward C. Prescott and Robert M. Townsend, "General Competitive Analysis in an Economy with Private Information," *International Economic Review* 25 (1984b), 1-20.

Edward S. Prescott and Robert M. Townsend, "Clubs as Firms in Walrasian Economies with Private Information," University of Chicago Working Paper (2001).

David Rahman, "Contractual Pricing with Incentive Constraints," UCSD Working Paper (2005).

Aldo Rustichini and Paolo Siconolfi, "General Equilibrium in Economies with Adverse Selection," University of Minnesota Working Paper (2002).

Joon Song, "Contractual Matching: Limits of Decentralization," UCLA Dissertation (2006).

C. M. Tiebout, "A Pure Theory of Local Public Goods," *Journal of Political Economy* 64 (1956), 416-424.

William R. Zame, "Incentives, Contracts and Markets: A General Equilibrium Theory of Firms," UCLA Working Paper (2005).