Policy discrimination with and without interpersonal comparisons of utility*

Michael Mandler
Department of Economics
Royal Holloway, University of London
Egham, Surrey TW20 0EX
United Kingdom

m.mandler@rhul.ac.uk

August 2004

Abstract:

Does the Pareto criterion discriminate among policy choices when the policymaker does not know the correct model of the economy? If the policymaker specifies ex ante preferences for each agent, there will typically be some policy change that improves the welfare of each agent relative to a status quo that suffers from a preexisting distortion. And if there are at least as many commodities as states, the second welfare theorem generically obtains: for almost every Pareto optimum, there is a policy that attains this allocation. Moreover, agents must trade under these policies; optimal allocations cannot be instituted by government fiat as they could be in the standard formulation of the second welfare theorem. The drawback is that ex ante preferences impose interpersonal welfare comparisons. If we instead require that policy changes increase all possible social welfare functions, and we are allowed to perturb a base model with additional states, then all policies including the distorted status quo are optimal. Furthermore, the set of policies that maximize some welfare function is open; consequently, small changes in the environment usually do not call for any policy response. Certain cases where a policymaker has symmetrical information about agents (following Lerner's classical treatment) do permit policy intervention.

Keywords: welfare theorems, Pareto optimality, taxes, policy paralysis

* This paper incorporates some (much revised) portions of Harvard Institute of Economic Research Discussion Paper 1729. I have discussed the ideas underlying this paper with many people over many years; let me thank Oliver Hart and John Geanakoplos in particular. An anonymous referee suggested several major improvements.

1. Introduction

There is a well-known puzzle about the second welfare theorem: if a policymaker knows the preferences and endowments of all agents, then it might as well act like a central planner and just assign agents the Pareto optimal allocation that it wants them to consume. If on the other hand the policymaker is uncertain about the economy's primitives it will be unable even to identify Pareto optima, let alone design transfers that achieve them. So in what sense does the second welfare theorem recommend markets as an allocation mechanism? This puzzle bolsters a common suspicion that the Pareto criterion is impractical for real-world policymaking. To address both the puzzle and the suspicion, we make explicit policymakers' lack of information about primitives and ask when policymakers can recommend policies that correct a preexisting distortion, namely taxes on net trades. As we will see, if a policymaker can posit a hypothetical ex ante stage at which agents share the policymaker's uncertainty and can make interpersonal comparisons between the potential preferences agents might have, then in some cases almost any first-best ex ante Pareto optimum can be achieved, and with policies that are just as sweeping as second welfare theorem policies: all tax distortions should be removed. Furthermore and in contrast to the puzzle, a policymaker cannot obtain these optima by dictating allocations; markets have to be used. In the remaining cases where the first best cannot be reached, policymakers can generically achieve at least some an ex ante Pareto improvement, and again markets are indispensable. So there *is* a framework that makes rigorous the second welfare theorem's endorsement of markets.

To examine the scope for policy adjustment, we distort a standard general equilibrium model by taxing net commodity purchases. The distortion ensures that the status quo appears to call for policy intervention; externalities could serve just as well, but taxes are tractable and have a long theoretical history. When policymakers know the primitives of the model, the welfare theorems imply that any policy (a tax rate for each good and a redistribution of endowments) that collects positive tax revenue is Pareto dominated by some zero-tax policy. We suppose instead that although each agent knows his or her own characteristics, the policymaker has only a probability distribution over the primitives of the economy, and say that *policymaking uncertainty* then obtains.

If a policymaker can posit ex ante preferences for agents in the presence of policymaking uncertainty, then call a policy *x* an *ex ante* improvement over *y* if *x* Pareto dominates *y* in terms of these ex ante preferences. The policies recommended by such a rule are similar to second welfare theorem recommendations if the number of states is no larger than the number of goods: almost every first-best allocation can then be reached by some policy (Theorem 1). But in contrast to the standard presentation of the second welfare theorem, in which the government knows the model and could therefore institute optima by fiat instead, under policymaking uncertainty individuals and markets have an indispensable role to play. Agents collectively know which state has occurred, and markets harness that information. When the number of states is larger than the number of goods, then generically there is at least some policy response to the preexisting distortion that achieves an ex ante Pareto improvement (Theorem 2). Thus, despite suspicions that the Pareto criterion cannot discriminate effectively among policies, there is a framework that recognizes a policymaker's uncertainty and decrees active policy intervention.

But valid criticisms of the Pareto criterion remain. When there is no policymaking uncertainty, the Pareto criterion may be viewed in two essentially equivalent ways: policy *x* weakly Pareto dominates *y* if (1) no agent is made worse off at *x* compared to *y*, or (2) any social-welfare-maximizing planner, no matter how the planner's welfare function weights individual utility functions, would recommend *x* over *y*. Under policymaking uncertainty, however, this equivalence breaks down. The ex ante approach, which we also call *agent-based*, follows (1) by identifying each agent with an ex ante preference relation. But these ex ante preferences must weight the potential utility functions that an agent might have. Since the agents themselves never face any uncertainty about what their preferences are, the ex ante preferences must rely on the policymaker's judgments about how to interpersonally compare welfare. Thus, our extension of (1) to policymaking uncertainty no longer shares the advantage of (2) of not committing to a particular weighting of individual utility functions.

To stay free from interpersonal comparisons, we define a second ordering that follows (2) in remaining neutral regarding weights on utility functions. We say policy *x* is *utility-independent* superior to *y* if, for all sum-of-expected-utilities social welfare functions, *x* is recommended over *y*. We also label a policy *x* to be *maximization-optimal* if there are utility functions for the potential agents such that *x* maximizes the resulting sum-of-expected-utilities welfare function. This leeway to choose utility representations means that utility-independence and maximization optimality are agnostic about how to compare the welfare of different preferences. Utility-independence or maximization optimality can lead to very large numbers of policies to be declared optimal, in which case we say that *policy paralysis* occurs. Our first policy paralysis result states that if a sufficient number of states (which can have arbitrarily small probability) are added to a base model, then any policy is utility-independent optimal (Theorem 3). This and associated results lead us to identify agnosticism about interpersonal comparisons as the source of the impracticality of the Pareto criterion.

A countervailing consideration however permits at least some policy discrimination in the face of policymaking uncertainty. Throughout the paper, we require that the social welfare attributed to an agent depends only on ordinal and cardinal information about the agent and not on the agent's index. This restriction complies with a long tradition of anonymity postulates in social choice theory, and since it limits the set of admissible welfare functions it strengthens our policy paralysis results. The restriction also implies that when a policymaker has symmetric information about two agents the policymaker should assign the same expected welfare to either agent receiving any given consumption bundle – no matter how the policymaker believes interpersonal welfare comparisons should be made. As we will see in the *Example* of section 5, the concavity of utility then implies that a utility-independent ranking will recommend that each agent should consume the same bundle. Lerner (1944) famously used a similar ignorance-based argument to conclude that equalizing the distribution of endowments in a one-commodity model will necessarily increase expected social welfare. According to Lerner, even policymakers who disagree about which types of agents more efficiently translate utility into social welfare can agree on an equal income distribution if they have symmetric information about which agent is of which type. Lerner's reasoning implies that agnosticism about interpersonal welfare comparisons can sometimes facilitate policy discrimination and that utility-independence does not always lead to complete paralysis. Our *Example* shows that Lerner's argument extends to multicommodity settings and can allow policy intervention even a planner foregoes problematic interpersonal comparisons. Lerner-style exceptions to policy paralysis arise only in classical economic environments (private consumption must be present), not in pure social choice settings (Mandler (1999)). Lerner's work nevertheless remains strangely ignored in general equilibrium theory, a lapse we hope to remedy.

Our second policy paralysis result, while not contradicting Lerner's argument, shows that plausible conditions will rule out the symmetry it presupposes. We show that under these conditions the maximization-optimal policies form an open set (Theorem 4). Consequently, if some tax and transfer policy is maximization optimal and the parameters of the model change slightly, that policy will remain optimal: local policy paralysis obtains. For example, the parameters might specify that one agent's consumption of good *x* imposes a consumption externality on some other agent. If we were maximizing some fixed social welfare function, a small increase in the size of the externality would normally call for a small policy response, perhaps an increase in the tax on *x*. But if a policymaker instead wants simply to maximize *some* social welfare function (in order to stay neutral regarding welfare comparisons), the increase in the externality will not lead to any policy response. With the larger externality and the same tax rates the policymaker would perhaps be maximizing a welfare function that assigns a lesser weight to the agent who is the victim of the externality. The fact that nonzero tax vectors can be optimal is hardly news (see, e.g., Mirrlees (1986)). Our point is that a rule that says policymakers should maximize some welfare function leads to a very large number of tax vectors being optimal. It can be that locally *every* tax and transfer policy is optimal; each policy serves as an efficient way to serve some classical social welfare goal.[1]

The Pareto criterion can be justified by different rationales. The ex ante ordering appeals to the principle that ex ante no individual should be made worse off by a policy change. The utility-independence ordering (or maximization optimality) argues that no particular way of making interpersonal welfare comparisons should be granted privileged status. Our purpose is

[1] The paralysis conclusion is subject to the caveat that, as our model stands, a policymaker can achieve ex post optimality simply by setting taxes equal to zero. But as we point out in the concluding discussion, this is an artifact of using taxes as a distortion; if some externality were present, for example, even ex post optimality would be unattainable.

not to judge which rationale is the right one: they are geared to different purposes. Rather our aim is to identify when each vantage point leads to effective discrimination among policy choices. To summarize, the Pareto criterion is workable if policymakers posit an ex ante stage at which agents experience the policymaker's lack information; without ex ante preferences, policy adjustment is problematic. And since this ex ante stage is hypothetical, the preferences that hold at this stage impose interpersonal comparisons of welfare. The role played by a hypothetical ex ante stage recalls the literature on Bayesian games (cf. Aumann (1998) and Gul (1998)), in which agents play correlated equilibrium actions only if there is a hypothetical ex ante stage at which agents have symmetric information and common priors. Here, it is the usefulness of the Pareto criterion that depends on an ex ante stage though not on common priors.

We take the policymaker's information to be fixed in this paper; the implementation and mechanism design literatures in contrast consider policies that induce agents to reveal their private information. Our modeling strategy is partly guided by our focus on the traditional policy tools of competitive markets. But the two approaches are in any event closely related. Our model confronts each agent with the same choice set of net trades; outside of some details that stem from the absence of production, these are the net trades that arise with Diamond-Mirrlees taxes. As Hammond (1979) points out, if a large number of agents play an anonymous revelation game in which agents announce their characteristics, each agent could equivalently be confronted with a common choice set of net trades: each agent who announces his characteristics in a revelation game will be assigned some net trade vector and so one could instead let the agent choose from the set of all net trade vectors selected by some agent in the distribution of possible characteristics. Anonymity, moreover, will be a necessary feature of any implementation scheme when the policymaker's information about agents' characteristics is symmetric across agents. Finally if agents anticipate that, following the play of the revelation game, they will have the opportunity to trade further on competitive markets, the only final allocations that can occur in equilibrium are those that could arise if agents chose from Diamond-Mirrlees choice sets of net trades. Thus, with a large number of agents, our setting is similar to an implementation setting. Hammond used dominant strategies in his paper, but see

Guesnerie (1995) and the references cited there for similar Bayesian arguments. Notice that our policy paralysis conclusions complement the result in the "limits to redistribution" literature that any dominant-strategy mechanism that implements efficient allocations in a large economy must lead to an undistorted Walrasian equilibrium with no transfers (see Champsaur and Laroque (1981), and earlier Varian (1976), Hammond (1979)). If only Walrasian outcomes are possible, we cannot require additionally that outcomes are Pareto-improving relative to an arbitrary status quo: achieving both first-best efficiency and a Pareto improvement will usually be impossible. Our paralysis results are in the same spirit: a preexisting distortion that blocks first-best efficiency cannot be removed if policy changes must be utility-independent improving. Given the Hammond (1979) connection between implementation and tax equilibria, this parallelism makes sense.

The contrast between the present paper and the implementation approach is misleading in a second and perhaps more important respect. We reach policy paralysis conclusions even when a policymaker is virtually certain about agent characteristics. Hence these results apply to any mechanism that does not reveal agent characteristics with complete certainty. When choosing economy-wide policy instruments, such as tax rates, governments inevitably have to come to policy decisions in the presence of at least some residual uncertainty about agents' characteristics, and our results hold in that setting.

The ex ante or agent-based approach specifies ex ante preferences for agents and is therefore formally a model of incomplete markets. It so happens in the present setting that the information asymmetry that agents know their own preferences but not others' preferences blocks the existence of markets for assets with state-dependent payoffs. But this is a minor point that does not interfere with the analytical machinery of the incomplete markets literature (see Geanakoplos (1990) and Magill and Quinzii (1996) for overviews). Indeed, it is a pleasant surprise that the techniques of incomplete markets are so well-suited to explaining seemingly distant social choice issues. Conversely, we argue in the conclusion that our results shed light on the dilemmas of policy design that have appeared in the incomplete markets literature, and on the theory of the second best as well.

2. Welfare criteria with policymaking certainty

We first lay out a benchmark model that we suppose is known to the policymaker. There are L commodities and J agents. Each agent j has an endowment $e_j \in R^L_{++}$ and a utility function $\bar{u}_j$ defined on consumption bundles $x_j \in R^L_+$. Let $e \equiv (e_1, ..., e_J)$ and let x_{ij} and e_{ij} refer, respectively, to agent j's consumption and endowment of good i. We assume that each $\bar{u}_j$ is twice continuously differentiable, differentiably strictly concave, and differentiably strictly increasing, and that the indifference curves of $\bar{u}_j$ that intersect R^L_{++} do not also intersect the coordinate axes.[2] An *economy* is a $(e_j, \bar{u}_j)_{j=1}^J$ and an *allocation* is a $x \equiv (x_1, ..., x_J) \in R^{LJ}_+$ such that $\Sigma_{j=1}^J (x_j - e_j) = 0$.

The economy begins with arbitrary ad valorem taxes $\tau = (\tau_1, ..., \tau_L) \geq 0$ that (to ensure that the taxes are in fact distorting) are imposed only on the value of net purchases. The revenue that results is for simplicity distributed in equal parts to the J individuals. Letting $p \in R^L_+ \setminus \{0\}$ indicate the before-tax price vector and $t \geq 0$ the government's tax revenue, the budget set facing agent j is:

$$B_j(p, \tau, e_j, t) = \{x_j \colon \Sigma_{i=1}^L ((1 + \tau_i)\, p_i \max [0, x_{ij} - e_{ij}] + p_i \min [0, x_{ij} - e_{ij}]) \leq (1/J)\, t\}.$$

Definition 1. *An equilibrium with taxes τ is a* (p, x) *such that* (1) *x is an allocation*, (2) *for each agent j, $x_j \in B_j(p, \tau, e_j, t)$, where* $t = \Sigma_{j=1}^J \Sigma_{i=1}^L \tau_i\, p_i \max [0, x_{ij} - e_{ij}]$, *and* (3) $x_j' \in B_j(p, \tau, e_j, t) \Rightarrow \bar{u}_j(x_j) \geq \bar{u}_j(x_j')$.

Under our assumptions, an equilibrium for the model exists for any τ.[3] Observe that if τ is sufficiently high in all coordinates, agents do not trade, they consume their endowment.

In addition to setting τ, the government can also transfer endowments by choosing a $\Delta e \equiv (\Delta e_1, ..., \Delta e_J) \in R^{LJ}$ such that $\Sigma_{j=1}^J \Delta e_j = 0$. We require that Δe be chosen so that an

---

[2] We use the notation: $x \geq y \Leftrightarrow x_i \geq y_i$, all i; $x > y \Leftrightarrow x \geq y, x \neq y$; and $x \gg y \Leftrightarrow x_i > y_i$, all i. Formally, $\bar{u}_j$ being differentiably strictly concave and differentiably strictly increasing means that, for all x_j, $D^2\bar{u}_j(x_j)$ is negative definite and $D\bar{u}_j(x_j) \gg 0$. The indifference curve condition is that, for all $x_j \gg 0$, $\{z \in R^L_+ : u_j(z) = u_j(x_j)\} \cap (R^L_+ \setminus R^L_{++}) = \varnothing$.

[3] See Shafer and Sonnenschein (1976), particularly note 4.1, and observe that $e_j \gg 0$ is always an element of B_j. Consequently, B_j, seen as a correspondence of x (via the effect of x on t) and p, is, in addition to being convex-valued, also continuous and nonempty-valued.

equilibrium still exists, e.g., by supposing $e + \Delta e \gg 0$. Multiple equilibria may arise for a given $(\tau, \Delta e)$, but since we want to give the policymaker as much latitude as possible we assume that the policymaker can choose which equilibrium price vector and allocation obtains with $(\tau, \Delta e)$. Letting $f \equiv (f_1, ..., f_J)$ indicate an equilibrium allocation that can occur with $(\tau, \Delta e)$, call $(\tau, \Delta e, f)$ a *policy*. Also, $(\tau, \Delta e)$ are *policy instruments*, and we say that a policy $(\tau, \Delta e, f)$ *reaches* the allocation f. Beginning at a status quo equilibrium $(\bar{p}, \bar{x})$ with taxes $\bar{\tau}$, the policy of maintaining the status quo is simply $(\bar{\tau}, \Delta e = 0, f = \bar{x})$.

The Pareto ordering may be characterized in two different ways under policymaking certainty. One way is to define an allocation x to be *ex ante* or *agent-based superior* to x' if for all agents j, $\bar{u}_j(x_j) \geq \bar{u}_j(x_j')$, and for some j, $\bar{u}_j(x_j) > \bar{u}_j(x_j')$. The agent-based criterion assigns a welfare significance to each agent and thus to each agent's index. The term "ex ante" will be self-explanatory once we introduce policymaking uncertainty.

An alternative *utility-independent* characterization of the Pareto ordering declares one allocation superior to another if it is recommended by all possible methods of making interpersonal comparisons of utility. We equate a method for making interpersonal comparisons of utility with the weights given to utilities when they are summed to a social welfare function. In contrast to the agent-based approach, the principle that a policymaker should remain agnostic about how to make interpersonal comparisons does not mention or assign any importance to agent indices.

We impose two restrictions on which utility functions can be admitted into social welfare sums. First, for each j, any two admissible utility representations for j can differ only by an increasing affine transformation. We will use this restriction repeatedly. In this section, it is dispensable since the essential equivalence of our two characterizations of the Pareto ordering would continue to hold if we admitted all increasing transformations; we maintain the restriction to ease comparison with the rest of the paper. For now, one may justify both this restriction and our use of welfare functions that are sums of individual utilities by supposing that the goods in the model are contingent commodities (say because there is objective, non-policymaking uncertainty in the background) and that each agent j's preferences have a von Neumann-Morgenstern (vNM) utility representation given by $\bar{u}_j$; Harsanyi (1955) then implies that every vNM social welfare function that obeys the Pareto principle can be represented as a sum of increasing affine transformations of the $\bar{u}_j$. In subsequent sections, uncertainty will be intrinsic to the environment and we will not need to implant it separately into the background.

Second, in any welfare function agents with identical sets of cardinal utility functions must be assigned the same utility function. The effect of this *anonymity requirement* is to make the indices assigned to agents irrelevant for social welfare functions; it is thus aligns with the utility-independent rather than the agent-based rationale for the Pareto criterion. Anonymity will play two important roles: it makes our policy paralysis results in section 5 stronger, and it makes possible important exceptions to policy paralysis when the policymaker has symmetric information about some agents.

Before considering the anonymity requirement further, we state formally our restrictions on admissible utility functions and the associated definition of utility independence.

Definition 2. *For each j, let* U_j *denote the set of all increasing affine transformations of* $\bar{u}_j$. *A utility assignment is a* $u = (u_1, ..., u_J)$ *such that for all j,* $u_j \in U_j$, *and the following anonymity requirement obtains: for any pair of agents* (j, h), *if* $U_j = U_h$ *then* $u_j = u_h$. *The allocation x is utility-independent superior to* x' *if, for all assignments u,* $\sum_{j=1}^{J} u_j(x_j) > \sum_{j=1}^{J} u_j(x_j')$.

The anonymity requirement follows a long social choice tradition (dating to May (1952), Suppes (1966), and Sen (1970)) that maintains that an agent's index is irrelevant for social decision-making; only information about preference or cardinal strength of preference should matter. In addition, one of our main purposes is to compare the standard agent-based Pareto criterion with the ranking that would be made by a social welfare-maximizing planner who is agnostic about how to make interpersonal comparisons of utility. But even a planner agnostic about how to make interpersonal comparisons would not postulate that one agent is more capable of experiencing satisfaction than another apparently identical agent based solely on the agents having different indices. Moreover, if we were to let utility weights in social welfare functions depend on agent indices, we would import an agent-based emphasis on indices into the utility-independent criterion, thus preventing us from seeing when the two approaches give genuinely different results. The anonymity requirement is nevertheless weak.

It has bite only when a policymaker encounters a set of agent who are both ordinally and cardinally identical; ; when there is no policymaking uncertainty, these cases are dismissible.

Here and subsequently, we define a policy $(\tau, \Delta e, f)$ to be superior to $(\tau, \Delta e, f)'$ in either an ex ante/agent-based or utility-independent sense if f is superior to f' by the corresponding ordering of allocations. But the distinction between policies and allocations is irrelevant in the certainty model: any allocation x can be reached by a policy that sets $\Delta e = x - e$ and sets τ high enough to induce agents not to trade.

The ex ante/agent-based and utility-independent orderings usually coincide under policymaking certainty, but our anonymity requirement permits an exception. If x is agent-based superior to x' then x is also superior to x' by the utility-independent definition, but the reverse implication need not hold. For instance, if $J = 2$, $U_1 = U_2$, and U_1 contains only strictly concave functions, then an allocation x such that $u_1(x_1) > u_1(x_2)$ is utility-independent inferior to a x' with $x_1' = x_2' = (1/2)x_1 + (1/2)x_2$. Yet clearly x' is not superior to x by the agent-based ordering. In the absence of policymaking uncertainty, such cases are a minor wrinkle that we may exclude with a *diversity condition* stating that no pair of agents has the same set of cardinal utilities; the agent-based and utility-independent orderings then will rank allocations in the same way. We will see that a comparable diversity condition is inappropriate under policymaking uncertainty.

The agent-based and utility-independent orderings automatically generate definitions of optimality by the requirement that there is no dominating allocations. In addition, we define an allocation x to be *maximization optimal* if there is an assignment u such that $\sum_{j=1}^{J} u_j(x_j) \geq \sum_{j=1}^{J} u_j(x_j')$ for all other allocations x'. A maximization-optimal allocation must also be utility-independent and agent-based optimal, but the reverse implications need not hold. Thus, as well as being more important in the welfare economics literature, maximization optimality is in principle more restrictive. But given our convexity assumptions the three definitions of optimality do coincide at interior optima if the diversity condition holds.

These orderings and optimality concepts give familiar and decisive advice. If the economy begins at a status quo equilibrium $(\bar{p}, \bar{x})$ with tax vector $\bar{\tau}$ such that for some good i and agent j, $\bar{x}_{ij} - \bar{e}_{ij} > 0$ and $\bar{\tau}_i > 0$, there must then be some other agent $h \neq j$ with $\bar{x}_{ih} - \bar{e}_{ih} < 0$

and h must be a net purchaser of some good, say k. Hence h's marginal rate of substitution between i and k must equal $\frac{\bar{p}_i}{\bar{p}_k + \bar{\tau}_k}$ while j's marginal rate of substitution between i and k must be greater than or equal to $\frac{\bar{p}_i + \bar{\tau}_i}{\bar{p}_k + \bar{\tau}_k}$. The marginal rates of substitution of the two agents therefore differ and the equilibrium allocation will be neither agent-based or utility-independent optimal. Under either ordering, there exist allocations x^* that are both optimal and superior to $\bar{x}$ and there are policies $(\tau, \Delta e, f)$ such that $x^* = f$, e.g., set $\Delta e = x^* - e$ and let τ be arbitrary. The welfare theorems thus give strong advice when the policymaker knows the model of the economy.

3. Policymaking uncertainty

A policymaker who is uncertain about the model faces a state space $\Omega = \{\omega_1, ..., \omega_S\}$, $S \geq 2$, with associated probability $\pi = (\pi_1, ..., \pi_S) \in \Delta_{++}^{S-1}$. Each state ω_s specifies for each agent j an ex post utility and an endowment, denoted $\bar{u}_j(\cdot, \omega_s)$ and $e_j(\omega_s)$ respectively, that satisfy the assumptions of the certainty model of section 2. A *model* is a pair (Ω, π). Notice that we specify a single π that applies to any expected utility calculation. When in the next section we consider ex ante orderings, this amounts to a notational convenience; we could let each agent j have a distinct subjective probability $\pi(j)$ and the same theorems would hold. When we consider policy paralysis in section 5, however, the restriction to a single probability strengthens our conclusions; a diversity of individual probabilities would only expand the set of tests that a utility-independent improvement would have to satisfy and hence make policy paralysis more likely.

Consumption by agent j at ω_s is denoted $x_j(\omega_s)$. Let $p(\omega_s)$ denote an equilibrium price vector at state ω_s, and P the $S \times L$ matrix whose sth row is $p(\omega_s)$. We also set the following notation for the remainder of the paper:

$$u_j = (u_j(\cdot, \omega_1), ..., u_j(\cdot, \omega_S)),$$

$$x_j = (x_j(\omega_1), ..., x_j(\omega_S)),$$

$$e_j = (e_j(\omega_1), ..., e_j(\omega_S)),$$

$$x(\omega_s) = (x_1(\omega_s), ..., x_J(\omega_s)),$$

$$e(\omega_s) = (e_1(\omega_s), ..., e_J(\omega_s)),$$

$$x = (x(\omega_1), ..., (x(\omega_S)).$$

There may but does not have to be further uncertainty (above and beyond the policymaking uncertainty) at any ωs. Consumption for j then consists of commodities contingent on the resolution of the additional uncertainty and each $\bar{u}_j(\cdot, \omega_s)$ is assumed to be a vNM utility representation of j's preferences at ω_s.

An allocation under policymaking uncertainty is a x such that each $x(\omega_s)$ is an allocation at ω_s. An equilibrium with taxes $\tau \geq 0$ is now a (P, x) such that, for each ω_s, $(p(\omega_s), x(\omega_s))$ is an equilibrium for the economy that occurs at ω_s when taxes are τ. A policy is a $(\tau, \Delta e, f) \in R^L_+ \times R^{LJ} \times R^{SLJ}_+$ such that each $f(\omega_s)$ is an equilibrium allocation at ω_s when endowments equal $e(\omega_s) + \Delta e$ and taxes are τ. Since the policymaker chooses tax rates and redistributions before agents interact on the market, τ and Δe are not state-contingent and therefore retain their previous dimensionality but f now specifies consumption at each ω_s. Let f_j now denote $(f_j(\omega_1), ..., f_j(\omega_S))$. Given an allocation x and taxes τ, the tax revenue at ω_s, $\sum_{j=1}^{J} \sum_{i=1}^{L} \tau_i\, p_i(\omega_s) \max\,[0, x_{ij}(\omega_s) - e_{ij}(\omega_s)]$, is $t(\omega_s)$.

After the policymaker selects $(\tau, \Delta e, f)$, markets equilibrate and $p(\omega_s)$, $x(\omega_s)$, and $t(\omega_s)$ are simultaneously determined. If the function p is invertible, the state could be inferred from the equilibrium price vector. But since agents already know their own preferences, this information has no value to agents; they simply choose utility-maximizing trades given the observed price vector. The policymaker does care what the true state is, but $(\tau, \Delta e, f)$ is set before $p(\omega_s)$ is observed. We suppose implicitly that each agent knows only his or her own preferences. Information is therefore asymmetric, preventing trade in assets with state-dependent payoffs.

We define a parameter space of economies Q by letting $e \in R^{SLJ}_{++}$ be parameters, and by assuming for any agent j that if h is any small linear function on R^{SL}_+, then $\bar{u}_j + h$ is a possible ex ante utility function for j. More precisely, h must have the form $\sum_{s=1}^{S} \sum_{i=1}^{L} a_i(\omega_s)\, x_{ij}(\omega_s)$, where each $a_i(\omega_s) \in R$, and where, for some $\varepsilon > 0$, $|h(x_j)| < \varepsilon$ for all $x_j \in R^{SL}$ with $\|x_j\| \leq 1$. We choose ε to be small enough that our assumptions on utilities continue to hold on a rectangle in R^{LS} that contains 0 and $\sum_{j=1}^{J} e_j$. If goods at some ω_s are contingent because of additional uncertainty at ω_s, which would lead $\bar{u}_j(\cdot, \omega_s)$ to be separable across those goods, the linearity of h ensures that this separability is retained.

The set Q has a finite number of dimensions and we denote a typical element of Q as (e, h). For any set A in a finite-dimensional Euclidean space, a *generic subset* refers to an open subset of A whose complement has Lebesgue measure 0.

4. Effective policy discrimination with the ex ante ordering

In the presence of policymaking uncertainty, the ex ante/agent-based approach begins with an ex ante preference ordering for each agent j over the hypothetical choices j would make if he or she faced the policymaker's state space. In principle, we should posit for each agent j vNM preferences $\succeq_j$ defined on lotteries where the typical prize is a consumption vector $x_j(\omega_s)$. But since we will need to consider only lotteries in which the probability of $x_j(\omega_s)$ is π_s, we instead just directly suppose that $\succeq_j$ induces preferences over state-contingent commodity bundles x_j that have a vNM utility representation using the probability π. We assume that j's vNM utility at ω_s is an affine transformation of $\bar{u}_j(\cdot, \omega_s)$. Then, letting $U_j(\omega_s)$ denote the set of all increasing affine transformations of $\bar{u}_j(\cdot, \omega_s)$, j's ex ante utility over the lotteries with probability π will be a function $Eu_j\colon R_+^{SL} \to R$ given by $Eu_j(x_j) \equiv \sum_{s=1}^{S} \pi_s u_j(x_j(\omega_s), \omega_s)$, where $u_j(\cdot, \omega_s) \in U_j(\omega_s)$ for all ω_s.

The assumption of a π common to all agents is in the present section a matter of notation. If we instead were to characterize an agent j not just with functions $u_j(\cdot, \omega_s)$ representing j's vNM utility for consumption vectors but with a separate $\pi(j)$ as well, we could then select the function $\dfrac{\pi_s}{\pi_s(j)} u_j(\cdot, \omega_s) \in U_j(\omega_s)$ and probability π and thus arrive at the same expected utility function for j given above.

The vNM hypothesis applied to the policymaking uncertainty alone, not any additional nonpolicymaking uncertainty at some ω_s, leads j's ex ante utility to be an expectation of functions that in each state differ only by an affine transformation. But in this section, the vNM hypothesis on preferences is made for concreteness. What matters is that each agent index j is assigned a single preference ordering, not additive separability across states.

We now say that an allocation x is *ex ante* (or *agent-based*) *superior to* x' if, for all j, $Eu_j(x_j) \geq Eu_j(x_j')$, and, for some j, $Eu_j(x_j) > Eu_j(x_j')$. This ordering coincides with the ex ante/agent-based of section 2 arises when $S = 1$. Allocation x is *strictly* ex ante superior to x' if strict inequalities hold for all j. Policies (τ, Δe, f) are ex ante ranked according to the ex ante ordering of their allocations f. In contrast to the certainty model, there can now be ex ante optimal allocations that cannot be reached by any policy (since Δe must be constant across states).

Our conclusions in this section will hold only for typical configurations of the primitives of the model. By fluke it might happen that the status quo τ and $\Delta e = 0$ lead to an ex ante optimal allocation, in which case no policy adjustment would be called for. Results on the scope for policy adjustment can therefore at best hold only for a generic set of models or economies.

The ex ante suboptimality of an economy beginning at a status quo equilibrium ($\overline{P}$, $\overline{x}$) with taxes $\overline{\tau}$ can be attributed to two factors. First, if $\overline{\tau}$ is nonzero, $\overline{x}(\omega_s)$ will normally be suboptimal for the economy at ω_s. Second, no agent who actually trades possesses the ex ante utility Eu_j; the trading agents have the ex post utilities $\overline{u}_j(\cdot, \omega_s)$. Consequently, relative to the hypothetical agents with the ex ante utilities, markets are incomplete and agents cannot insure themselves against the uncertainty in Ω. Allocations will therefore normally be ex ante suboptimal even when $\tau = 0$. As we will now see, the policy instruments τ and Δe will typically allow the policymaker to engineer an ex ante improvement as a response to this suboptimality: the status quo will typically be ex ante suboptimal relative to what can be reached by some policy. Most dramatically, if there are at least as many goods as states, the ex ante/agent-based approach usually recommends policy changes just as sweeping as the second welfare theorem: virtually any first best allocation (including ex ante improvements on the status quo) can be reached and with taxes set to 0.

Theorem 1. If $L \geq S$, there is a generic subset of economies G such that for any economy in G there is a generic subset of ex ante optimal allocations each of which can be reached by some policy with $\tau = 0$.

The logic underlying the proof of Theorem 1 (in the appendix along with all other proofs) is simple. Since each agent shares the same marginal rate of substitution at an ex ante optimal allocation x, there are prices ($p(\omega_1)$, ..., $p(\omega_S)$) that support the allocation. And typically, if $L \geq S$, the price vectors $p(\omega_1)$, ..., $p(\omega_S)$ that rule at the S states will be linearly independent. Thus, for each agent j, the equations

$$p(\omega_s) \cdot \Delta e_j = p(\omega_s) \cdot (x_j(\omega_s) - e_j(\omega_s)),\ s = 1, ..., S,$$

have a solution Δe_j, and so if the policymaker sets $\tau = 0$ and each j's transfer equal to this Δe_j then j can exactly afford the bundle $x_j(\omega_s)$ at ω_s when prices equal $p(\omega_s)$.

The optimal allocations identified by Theorem 1 cannot be achieved by direct command decision; the policymaker does not know which ω_s obtains, and usually the target allocation $x(\omega_s)$ will differ by state. Thus, although Theorem 1 is akin to the second welfare theorem, it assigns markets a more fundamental role. In the standard presentation of the second welfare theorem, there is no policymaking uncertainty ($S = 1$), and so optimality could always be achieved instead with taxes left at the status quo levels: let Δe move agents' endowments directly to an optimal allocation, thus making trade unnecessary. But when $S \geq 2$ agents must typically trade at all states since the post-transfer endowments $e_j(\omega_s) + \Delta e_j$ typically will not equal the target $x_j(\omega_s)$ at any ω_s. Markets and trade therefore have an indispensable function in the presence of policymaking uncertainty: unlike the policymaker, agents collectively know which state obtains and trading allows the economy to utilize this information. Moreover, since agents are trading, reaching a first best allocation requires that tax rates be set to zero.

What can be said when the number of states is greater than the number of goods, $S > L$? Generically at least some policy adjustment relative to an arbitrary status quo is possible:

Theorem 2. If $S \geq 2$, then for any τ there is a generic subset of economies G such that for each equilibrium allocation x with taxes τ of each economy in G there is a policy that reaches an allocation that is a strict ex ante improvement over x.[4]

Thus, typically an arbitrary status quo policy will not be ex ante optimal. And although there may not be a policy with $\tau = 0$ that is ex ante superior to the status quo, it follows from the proof of Theorem 2 there will be at least some policy in which τ differs from the status quo τ that is ex ante superior to the status quo: policymakers can adjust arbitrarily given tax rates.

[4] If $S = 1$ and $\tau > 0$, the conclusion of the theorem continues to hold.

Policies that achieve strict ex ante improvements are also robust to the addition of a small amount of uncertainty. Suppose, in the $S = 1$ certainty model, that we begin with a status quo equilibrium $(\bar{p}, \bar{x})$ with taxes $\bar{\tau}$ and find a $(\tau', \Delta e', f')$ that leads to a strict ex ante Pareto improvement. We can add a small amount of uncertainty by introducing an arbitrary set of new states and assigning the new states small probability. The entire model is then (Ω, π), where we assign the initial certainty model's economy to ω_1. If we are given ex ante utilities Eu_1, ..., Eu_J for (Ω, π), then, for π_1 sufficiently near 1, a policy $(\tau', \Delta e', f'')$ such that $f''(\omega_1) = f'$ and where the $f''(\omega_s)$, $s \neq 1$, are set arbitrarily will be is strictly ex ante superior to any status quo policy $(\bar{\tau}, \Delta e = 0, f)$ with $f(\omega_1) = \bar{x}$. So if a policymaker has access to ex ante utilities, then the addition of a sufficiently small amount of uncertainty will not lead to the reversal of a proposed policy change. But observe that the probabilities for the uncertainty perturbation that will preserve policy recommendations are a function of the ex ante utilities. For a given (Ω, π) – even if π_1 is near 1 – there may well be ex ante utilities such that $(\tau', \Delta e', f'')$ does not lead to an ex ante improvement over a status quo policy $(\bar{\tau}, 0, f)$ at which $f(\omega_1) = \bar{x}$ and where, say, $f(\omega_s) = f''(\omega_s)$ for $s \neq 1$. All that is necessary is that at some ω_s some j is worse off with $(\tau', \Delta e', f'')$ than with $(\bar{\tau}, 0, f)$ and that $u_j(\cdot, \omega_s)$ is a sufficiently large multiple of $\bar{u}_j(\cdot, \omega_s)$.

5. Policy recommendations without interpersonal comparisons of utility

The ex ante or agent-based approach to social decision-making prescribes for each agent j an ex ante utility Eu_j. Each Eu_j imposes a weighting of ex post utilities: given the base set of utilities, $\bar{u}_j(\cdot, \omega_1)$, ..., $\bar{u}_j(\cdot, \omega_S)$, each $u_j(\cdot, \omega_s)$ in Eu_j is an affine transformation of $\bar{u}_j(\cdot, \omega_s)$. Since the policymaker's uncertainty about agents' potential preferences does not correspond to any uncertainty experienced by the agents themselves, the weights on the $\bar{u}_j$ must reflect the policymaker's judgments about which potential preferences experience the greater satisfaction and thus deserve greater priority. As long as j's ex post preferences differ, there will be x_j and x_j' such that x_j is preferred to x_j' by one of j's ex post preference relations but where the reverse judgment is held by another of j's ex post preference relations. But the policymaker must specify a preference for j between x and x'. If, say, $Eu_j(x_j) > Eu_j(x_j')$, the policymaker in effect maintains that those of j's ex post preferences that rank x ahead of x' gain more satisfaction than what is lost by those of j's ex post preferences that hold the reverse judgment. Given that the actual agent j never faced this uncertainty – the uncertainty is entirely the policymaker's – this claim amounts to an interpersonal comparison of welfare.

These interpersonal comparisons are embodied in the utility weights used to assemble Eu_j. To see this, consider the Harsanyi (1953) model of social welfare, in which agents rank social choices pretending that they are ignorant of which person in society they will be. Harsanyi's decision problem requires a chooser to gauge the intensity of satisfaction that different agents experience, and thus involves the same type of judgments needed to specify a social welfare ordering. If, for instance, two parties disagree about what ordering should guide social decision-making due to an underlying dispute about who would experience the deepest happiness from the allocation of some good, then they will also disagree on how to rank social choices in the Harsanyi problem, and vice versa. The construction of ex ante preferences for any agent j in turn involves the same type of interpersonal comparisons of satisfaction as Harsanyi's problem. Indeed, the only difference between j's ex ante problem and the Harsanyi problem is that the latter assigns the same probability to assuming the identity of each of society's agents whereas the former uses the probability π and some identities may not be possible for any given j. Agent j's ex ante problem and the Harsanyi problem are so similar that one could advise a policymaker pursuing the ex ante Pareto program to use the Harsanyi thought experiment as a way to choose weights on the $\overline{u}_j(\cdot, \omega_s)$.

Since a main purpose of Paretian welfare economics is to avoid interpersonal comparisons of utility, we now turn to utility-independence as a way to avoid all such comparisons. Utility-independence asserts that no set of weights on ex post utilities in social welfare functions is more legitimate than another. On the other hand, utility-independence relinquishes any use of the fact that an ex post utility is attached to some agent's index. We begin by specifying the utilities that can be admitted into social welfare functions in the presence of policymaking uncertainty.

Definition 5. *A utility assignment under policymaking uncertainty is a* $u = (u_1, ..., u_J)$ *such that for all agents j and h and all states* ω_s *and* ω_l, $u_j(\cdot, \omega_s) \in U_j(\omega_s)$ *and (anonymity)* $U_j(\omega_s) = U_h(\omega_l)$ *implies* $u_j(\cdot, \omega_s) = u_h(\cdot, \omega_l)$.

A welfare function now assigns a welfare number in R to each allocation in R_+^{SLJ}. Given a utility assignment u, $\sum_{j=1}^J Eu_j$ is the welfare function that assigns welfare level $\sum_{j=1}^J Eu_j(x_j)$ to allocation x. The definition of utility-independence remains as in section 2: allocation x is utility-independent superior to x' if, for all assignments u, $\sum_{j=1}^J Eu_j(x_j) > \sum_{j=1}^J Eu_j(x_j')$. An allocation x is utility-independent optimal if there is no utility-independent superior allocation, and is *maximization optimal* if there is an assignment u such that, for all allocations x', $\sum_{j=1}^J Eu_j(x_j) \geq \sum_{j=1}^J Eu_j(x_j')$. Policies are again ranked or are optimal based on how the allocations they induce are ranked. When $S = 1$, these definitions coincide with those given in section 2.

As in section 2, welfare functions use the same utility function to represent all potential agents with the same set of cardinal utilities and are additively separable in agents' ex post utilities. These restrictions on welfare functions are justified, respectively, by the arguments given earlier in favor of anonymity and by the Harsanyi (1955) theorem on additive social welfare functions. Observe that unlike section 2 uncertainty is now intrinsic to the policymaker's problem and hence we no longer have to resort to the contrivance that further uncertainty obtains at some state. We emphasize that these restrictions make policy paralysis conclusions stronger: by limiting the number of welfare functions, it becomes *easier* to conclude that one policy is utility-independent superior to another and hence *harder* to conclude that a policy is optimal (undominated).

Some other social-welfare criteria are similar to but not identical to utility-independence. For instance an interim or ex post definition of the Pareto ordering (Holmström and Myerson (1983)) would label one allocation superior to another if each agent in the distribution of possible agents is at least weakly better off (and one potential agent strictly better off). Equivalently, we could retain the ex ante ordering but require additionally that, for allocations to be ranked, every agent j is better off no matter what subjective probability $\pi(j)$ we used to calculate j's expected utility. This would again ensure that when one allocation is ranked superior to another, no potential agent is worse off. These criteria differ from utility-independence because of the latter's anonymity requirement: a change in allocations can harm some potential agent $j(\omega_s)$ and still be a utility-independent improvement if some other potential agent with the same set of cardinal utility functions as $j(\omega_s)$ enjoys sufficient utility gains. As mentioned in section 2, the rationales for anonymity in the social welfare literature argue in favor of our restrictions on utility assignments. Moreover, we wish to identify the policy consequences of agnosticism about interpersonal welfare comparisons alone, without additionally dictating that every potential agent must be left unharmed. Such an additional requirement would in effect apply the agent-based approach to the entire set of potential agents; it does not follow from the utility-independent principle that policymakers should remain agnostic about how to make interpersonal comparisons.

Finally, anonymity opens the door to important exceptions that relieve a bleak conclusion that policy paralysis always obtains. A policymaker's ignorance of the primitives of an economy can in certain highly symmetric situations make it *easier* to discriminate among policies by a utility-independent ranking. We will see that these symmetric cases become plausible under policymaking uncertainty whereas earlier they were sensibly ruled out by the diversity condition. If we were to hamper utility-independence further by requiring that policy changes leave every potential agent unharmed, these interesting cases would be excluded.[5]

Before turning to the exceptions, we record that policy paralysis obtains when any base model is perturbed through the addition of further states. Specifically, no policy is utility-independent superior to an arbitrary status quo policy if L states can be added to the base model, thus contrasting sharply with the scope for policy change allowed by the ex ante Pareto criterion.

Theorem 3. For each base set of states Ω, there is a set of L states Ω' such that in any model with state space $\Omega \cup \Omega'$, no policy $(\tau, \Delta e \neq 0, f)$ is utility-independent superior to any status quo policy $(\bar{\tau}, 0, \bar{f})$.

Since the probabilities of the states in $\Omega \cup \Omega'$ can be set arbitrarily, the added states in Ω' can have arbitrarily small probability. Theorem 3 treats policy changes such that $\Delta e \neq 0$, which

---

[5] As we will see, the cases that permit policy discrimination under a utility-independent ordering hinge on the use of equal probabilities to model a policymaker's ignorance. This connection between ignorance and equal probabilities (the principle of insufficient reason) argues against the use of arbitrary subjective probabilities to calculate an agent's ex ante utility.

arise, for instance, when compensation for a change in τ is attempted. It is not difficult, by adding more additional states, to cover policy changes that involve only a change in τ.

Theorem 3 suffers from the drawback that the added states can vary as a function of the base model Ω and can therefore omit agents with the same utilities as agents in the base model. Consequently, to prove Theorem 3, it is sufficient to show that some agent at some added state is harmed by any proposed policy change. If some utility functions for agents at the added and base states coincide, then even if some j at some additional state $\hat{\omega}$ were harmed by a change from $(\tau, \Delta e, f)$ to $(\tau, \Delta e, f)'$, $(\tau, \Delta e, f)'$ could still be ranked utility-independent superior: other potential agents with identical utility representations might collectively gain more utility in expectation from the policy change than j's expected loss at $\hat{\omega}$. It is therefore impossible to infer the overall consequences of policy changes from how the welfare of individuals changes at a subset of states: it is the overall distribution of characteristics that matters. It should be clear, moreover, that for any set of additional states and any given policy change, there exists an accompanying base model such that the policy change is a utility-independent improvement for the model combining the base and additional states.[6]

In the certainty ($S = 1$) model as well, some agent can be made worse off by a policy change even though the utility-independent ordering recommends the policy change. But while in the certainty case it is plausible to dismiss as irrelevant any example that does not obey the diversity condition (i.e., an example where different agents have identical sets of cardinal utility functions), it is the norm for the same potential utility functions to arise at multiple states and for multiple agents. If, for example, a base model specifies that agent j either has the ex post utility u_j or u_j', it is reasonable to allow j to have each of these utilities with non-negligible probability at some of the additional states (e.g., when the probability of j having any given utility is independent of what preferences the other agents have). Similarly, if the policymaker has identical information about a pair of agents, then the support of the distribution of those

[6] In models of social choice, policy paralysis requires only that preference relations in certain open sets are elements of the state space, regardless of the preferences that appear at other states (see Mandler (1999), Theorem 4). Since agents with identical preferences have the same preferences over policies in pure social choice settings, a policy that harms one potential agent harms all potential agents with the same utility function. For the same reason, the Lerner exception to policy paralysis that we now consider cannot arise in pure social choice settings.

agents' utility functions should be the same. Thus, the methodology permitted by Theorem 3 of adding idiosyncratic states to a fixed base model can sometimes be suspect.[7]

Indeed, the following example shows that a highly symmetric model can allow some allocations and policies to be ranked by the utility-independent ordering. The example is inspired by the Lerner's (1944) argument that equalizing the distribution of income will increase social welfare, even when individuals derive utility from income at different rates, so long as the policymaker is ignorant about which agents are the more efficient producers of utility. Both for Lerner and in the example below, ignorance can make a policy criterion *more* discriminating. The example again illustrates that the utility-independent ordering can recommend policy changes that are rejected by any ex ante ordering and therefore that the utility-independent ordering is neither weaker nor stronger than any given ex ante ordering.

Example. Suppose that $\Sigma_{j=1}^{J} e_j(\omega_s)$ does not vary as a function of the state ω_s and that the policymaker has "ignorance" priors over the agents' utilities. That is, for each pair of agents *i* and *j* and each state ω_s, assume that the following *symmetry condition* holds:

(5.1) $$\Sigma_{\omega_l \in \Omega : U_i(\omega_l) = U_j(\omega_s)} \pi_l = \Sigma_{\omega_l \in \Omega : U_j(\omega_l) = U_j(\omega_s)} \pi_l .$$

That is, the likelihood that agent *i* has a set of cardinal utility functions *U* is equal to the likelihood that any other *j* has the same *U*. Let $\psi = \frac{1}{J} \Sigma_{j=1}^{J} e_j(\omega_s)$ and *x* be an allocation such that $x(\omega_s)$ does not vary as a function of ω_s and that $x_j(\omega_s) \neq \psi$ for at least one *j*. The symmetry condition implies that any distinct utility *u* that appears in some $U_k(\omega_s)$ consumes $x_j(\omega_s)$, $j = 1, ..., J$, each with probability $\frac{1}{J}$. Since $\Sigma_{j=1}^{J} \frac{1}{J} x_j(\omega_s) = \psi$, the strict concavity of *u* implies $u(\psi) > \Sigma_{(\omega_l, j) : U_j(\omega_l) = U_k(\omega_s)} \pi_l \, u(x_j(\omega_l))$. So, letting ψ also denote the allocation where every agent at every state consumes ψ, it follows that for any assignment *u*,

(5.2) $$\Sigma_{j=1}^{J} E u_j(\psi) > \Sigma_{j=1}^{J} E u_j(x_j).$$

---

[7] It is worth noting, however, that a proof for Theorem 3 need not use additional states with utility representations that do not occur at $\omega_s \in \Omega$. What is necessary is that the probability of any $\omega_s \in \Omega$ that has one or more agents with utilities that appear in an additional state $\hat{\omega}$ is sufficiently small.

Hence the allocation giving each agent ψ is utility-independent superior to any x that is constant across ω_s.

If $L = 1$, there must be a j such that $x_j(\omega_s) > \psi$ for all ω_s. Such agents are worse off with ψ at every state. The allocation giving each agent ψ therefore cannot be superior to x according to any of the possible ex ante orderings. Once again we see that the utility-independent ordering can endorse a change in allocations rejected by any ex ante ordering.

Some policies can be ranked as well. Assume now in addition for each j that $e_j(\omega_s)$ also does not vary across states. If, for some j, $e_j(\omega_s) \neq \psi$, then any $(\tau, (\Delta e_j = \psi - e_j(\omega_s))_{j=1}^{J}, f)$ is utility-independent superior to any status quo policy $(\overline{\tau}, 0, \overline{f})$ if τ and $\overline{\tau}$ are both high enough to prevent trade from occurring at all ω_s.

Since (5.2) is an inequality, the example is robust in the sense that small changes in the primitives of the model – in $U_j(\omega_s)$, the $e_j(\omega_s)$, and π – will still allow some allocations and policies to be ranked. For the same reason, the τ and $\overline{\tau}$ in the policies of the previous paragraph do not have to be set so high as to prevent all trade, just high enough that only a small amount of trade occurs.

Although our conclusions are similar in spirit to Lerner (1944), Lerner's agents consume just one good – income – and thus all share the same ordinal preferences (if not the same cardinal utility) whereas (5.1) applies to disparate preferences over many commodities.[8]

Our anonymity requirement that cardinally and ordinally identical potential agents are represented by the same utility function and that only increasing affine (rather than all monotonic) transformations of some strictly concave utility function are elements of $U_j(\omega_s)$ are both crucial for the non-paralysis conclusion. We could in principle allow further information about agents beyond their preferences and their cardinal utilities to affect the weighting of utilities in social welfare functions. But then we would have to consider cases where a policymaker's information is symmetric with regard to this additional information as well, and once again utility-independence in conjunction with anonymity would lead to nontrivial policy

[8] For other formalizations of Lerner's argument, see McManus et al. (1972), McCain (1972), and Sen (1973), which all suppose that each agent's utility is a function of one good.

recommendations.

The significance of the example is not that there can be allocations and policies that are suboptimal according to the utility-independent or maximization definitions. Even simpler examples would suffice to show this (e.g., suppose that all agents in all states have the same cardinal utility function). What the example underscores is that even with no restriction on the number and diversity of preference orderings, some nontrivial policy advice is possible in some plausible cases.[9] ■

In the local policy paralysis result below, we do not assume that certain utilities appear with non-negligible probability only at certain carefully constructed states. We cast the result in terms of the historically more important maximization definition of optimality. Since maximization-optimal policies are also utility-independent optimal but the converse need not hold, results that apply to the maximization definition are stronger.

We will say that a policy $(\tau, \Delta e, f)$ is *differentiable* if the allocation induced by the policy is locally a continuously differentiable function of the policy instruments $(\tau, \Delta e)$. That is, there must be a continuously differentiable function g from an open $\Pi \subset R_+^L \times R^{L(J-1)}$ to allocations R_+^{SLJ} such that $g(\tau', \Delta e')$ is an equilibrium allocation for any $(\tau', \Delta e') \in \Pi$, $(\tau, \Delta e) \in \Pi$, and $g(\tau, \Delta e) = f$. Most policies are differentiable; the lemma in the proof of Theorem 2 in fact shows that policies are generically differentiable. But welfare maximization need not always select one of these generic policies; sometimes a nongeneric policy at which the equilibrium allocation is not a differentiable function of $(\tau, \Delta e)$ may be dictated. We nevertheless restrict ourselves to maximization-optimal policies that are differentiable, thus incurring a small loss of generality.

One way to generate the welfare functions that can arise with (Ω, π) is to pick an arbitrary assignment u and then multiply each ex post utility function $u_j(\cdot, \omega_s)$ by some positive weight b_{js} where any pair of identical ex post utilities is multiplied by the same

[9] Other, more trifling policy recommendations can also be made. For example, if $\overline{\tau}$ is high enough to prevent trade at all ω_s, then any $(\tau, \Delta e = 0)$ such that τ allows some trade at some ω_s is a utility-independent improvement.

weight.[10] Let B denote this set of weights, $\{b \in R^{SJ}_{++}: U_j(\omega_s) = U_{j'}(\omega_{s'}) \Rightarrow b_{js} = b_{j's'}\}$, which has dimension equal to the number of distinct utilities in Ω. Given a differentiable policy (τ, Δe, f) and an assignment u, and letting g be the function specified above, we define the welfare functions parameterized by B, $w_u: B \times \Pi \to R$ by setting $w_u(b, (\tau, \Delta e)) = \sum_{j=1}^{J} E\,\hat{u}_j(g_j(\tau, \Delta e))$, where $\hat{u}$ is the assignment $\hat{u} = b \cdot u$.

We put aside the question of whether equilibria exist at boundary policies by now requiring that endowment redistributions are in the set $\Delta E = \{\Delta e: e_j(\omega_s) + \Delta e_j \geq 0$ for all j and $\omega_s\}$ and assuming for all ($\Delta e \in \Delta E$, τ) that an equilibrium exists at each ω_s.

Definition 6. *A differentiable policy* (τ, $\Delta e \in \Delta E$, f) *is* a regular maximum for the assignment u *if* (1) *whenever* $(\tau, \Delta e, f)'$ *has* $f' \neq f$ *and* $\Delta e' \in \Delta E$, $\sum_{j=1}^{J} E u_j(f_j) > \sum_{j=1}^{J} E u_j(f_j')$, *and* (2) $D^2_{\tau,\Delta e} w_u(1^{JS}, (\tau, \Delta e))$ is negative definite.[11]

Definition 7. *A differentiable policy* (τ, Δe, f) *satisfies the rank condition for the assignment u if* $D^2_{(\tau,\Delta e), b} w_u(1^{JS}, (\tau, \Delta e))$ *has rank* LJ.

Differentiability and regularity of a policy are the traditional conditions that guarantee a maximum is well-behaved; they ensure that calculus can be applied, that a strict second order condition obtains, and that two or more policies do not simultaneously maximize the same welfare function. The assumptions are also "open" properties that continue to hold if the model is smoothly perturbed. The rank condition is an open property as well since LJ is the maximal rank of $D^2_{(\tau,\Delta e), b} w_u(1^{JS}, (\tau, \Delta e))$. But the rank condition is also substantive and its meaning is important. It says that there are enough utility functions in the model so that for every policy instrument we can find an independent combination of changes in welfare weights that will alter the marginal social welfare of that instrument. This means that each policy instrument has a *distinctive* effect on social welfare in that it affects the welfare of a different combination of ex post utilities. For example, a change in some τ_i will have a different impact on intensive

[10] We do not by this method generate all possible welfare functions since we are restricted to linear transformations of the $u_j(\cdot, \omega_s)$. But the excluded constant terms permitted by affine transformations never change any ranking of policies determined by a sum of utilities.

[11] For any positive integer m, 1^m denotes the vector of m 1's.

buyers and sellers of good *i* compared to its impact on other potential agents. (Thus, the rank condition is not satisfied in the Example above.) Following the policy paralysis theorem, we show that we can add additional states to a model to ensure that the rank condition is satisfied; these states guarantee that the model is sufficiently rich in agents so that precisely these distinctive effects of different policy instruments are present.

Theorem 4. *The policy instruments* $(\tau, \Delta e)$ *such that some differentiable* $(\tau, \Delta e, f)$ *is a regular maximum for some u and where* w_u *satisfies the rank condition form an open set*.

Suppose that the entire uncertainty model is perturbed slightly – say by the addition of a small consumption externality – in such a way that the primitives of the model change smoothly as a function of the perturbation. If the status quo policy is differentiable and a regular maximum and the corresponding rank condition is satisfied, it will remain so after a small enough perturbation. Theorem 4 then indicates that if policymakers aim to maximize some welfare function, then a small externality will induce no policy response.

The proof of Theorem 4, in the appendix, reverses the standard implicit function procedure of solving the first order conditions of a welfare maximization problem for an optimal allocation as the parameters of the problem change slightly; instead we solve the first order conditions for welfare weights as the allocation changes. A global version of Theorem 4, for either the utility-independent or maximization definitions of optimality, faces difficulties. The Example above is a sign that there is no general condition that rules out models in which many policies are suboptimal in the utility-independent sense. And even when many or all policies are utility-independent optimal, a paralysis result using the more stringent maximization definition of optimality faces the added hurdle that the set of agent utilities reachable through some policy is not convex, which implies that utility-independent policies need not be maximization optimal.

The rank condition is weak in that if we add states with diverse ex post agents to a model, the condition is necessarily satisfied; the added states moreover can have arbitrarily small relative probability. Note that as we add more ex post agents to the model, more columns are added to $D^2_{(\tau,\Delta e),\,b} w_u(1^{JS}, (\tau, \Delta e))$ but not more rows; the number of rows always equals the number of policy instruments LJ.

Theorem 5. For any differentiable policy $(\tau, \Delta e, f)$ *for* (Ω, π) *that is a regular maximum for some u, there exists a* $(\hat{\Omega}, \hat{\pi})$ *such that, for every* $\lambda \in [0, 1]$, *the model* $(\Omega \cup \hat{\Omega}, (\lambda\pi, (1-\lambda)\hat{\pi}))$ *has a differentiable policy that is a regular maximum for some* $\hat{u}$ *such that* $w_{\hat{u}}$ *satisfies the rank condition.*

6. Discussion

Our results are both positive and negative. The ex ante Pareto criterion will recommend a move from most status quo policies, but this criterion incorporates a system for making interpersonal welfare comparisons. On the other hand, a thorough-going avoidance of interpersonal comparisons can lead a large number policies to be optimal.

Our results illuminate some common complaints about the usefulness of the Pareto criterion. When markets are incomplete, it is well-known that a policymaker can institute Pareto improvements by redistributing initial-period asset holdings. The necessary transfers require detailed information, however, and so it is tempting to conclude that such policy interventions are impractical (see, e.g., Geanakoplos and Polemarchakis (1990)). Similar observations were made in the wake of the theorem of the second best (Lipsey and Lancaster (1956)): when there are distortions that the policymaker cannot correct, optimal policies can be counter-intuitive and depend on unobtainable information about the parameters of the model.

Our explicit modeling of policymaking uncertainty allows us to evaluate this type of reasoning. Consider a model of incomplete markets in which the policymaker is uncertain about the attributes of the underlying economy. (If there were no policymaking uncertainty, policy analysis would proceed as it does in the incomplete markets literature.) The policymaker then faces two sources of uncertainty: the uncertainty that confronts the agents and an additional uncertainty about the parameters of the model. If the policymaker can make interpersonal comparisons of welfare and therefore construct an ex ante ordering, he or she could then devise policies that are improving relative to the status quo policy of letting agents choose their asset portfolios without government intervention. Pareto-improving reallocations of assets exist in the incomplete markets model in the absence of policymaking uncertainty; the inclusion of policymaking uncertainty simply adds new dimensions of market incompleteness for the hypothetical ex ante agents. But section 4 shows that policymaking uncertainty alone, even when the ex post agents face *no* market incompleteness, is typically enough to guarantee that some policy changes are called for. On the other hand, utility-independent or maximization welfare rules argue against any change of policies and this conclusion does not hinge on market incompleteness. The set of optimal policies, as section 5 shows, is sizable even when the (ex post) agents face complete markets. Thus, market incompleteness does not introduce any special or additional problem of policy paralysis: the difficulty lies in not knowing the model with certainty and simultaneously trying to avoid interpersonal comparisons of welfare.

Analogous observations apply to the difficulty, described by the theorem of the second best, of devising policy recommendations in the presence of multiple uncorrectable distortions in the economy. If a policymaker can formulate an explicit state space to describe his or her uncertainty and can furnish ex ante preferences, the ex ante ordering will typically recommend that policy be changed from an arbitrary status quo. In this paper, for example, one could suppose that some or all of the taxes on net trades are uncorrectable; the proof of Theorem 2 indicates that the endowment transfers can still engineer an ex ante improvement. On the other hand, if a utility-independent or maximization welfare rule is in effect, then policy paralysis will occur even when the policymaker has the freedom to set all tax rates equal to zero. It is the difficulty of specifying ex ante preferences that makes policy adjustment problematic, not the presence of uncorrectable distortions.

A final word is necessary on allocations and policies that are optimal ex post. Let us suppose that the diversity condition of section 2 is satisfied at each state; in the present setting, this allows to ignore the distinction between the agent-based and utility-independent orderings. Define an *allocation* x to be *ex post optimal* if there is no other (feasible) allocation x' such that $x'(\omega_s)$ Pareto dominates $x(\omega_s)$ at every ω_s. A *policy* $(\tau, \Delta e, f)$ is *ex post optimal* if there is no other policy $(\tau, \Delta e, f)'$ such that f' Pareto dominates f at every ω_s. Since ex post optimality does not make hypothetical comparisons between the different preferences an agent might have, it does not make interpersonal comparisons of utility. As our framework now stands, a policymaker can achieve an ex post optimal allocation by setting $\tau = 0$. Since any status quo policy with $\tau \gg 0$ will not achieve an ex post optimal allocation (assuming that in some state agents are endowed with a Pareto inefficient allocation), the dilemma of policy paralysis would seem to disappear. But this reasoning is unpersuasive. First, although a policy that sets $\tau = 0$ is indeed unusual in achieving an *allocation* that is optimal ex post, that does not mean ex post optimal *policies* are rare. Indeed, on the conditions of Theorem 3, since every policy change relative to the status quo harms some ex post agent, arbitrary status quo policies are ex post optimal as well. Bear in mind that for an allocation to be ex post optimal, it must be undominated at *every* state by any other feasible allocation, and this can be hard to achieve. But for a policy to be ex post optimal, it must merely be that every alternative policy harms some agent at some state. Since policymakers are in the business of selecting policies, the ex post optimality of the allocation achieved by a policy is of dubious relevance; it usually involves a comparison with allocations achieved by no policy. Second, and just as importantly, the ex post optimality of $\tau = 0$ is an artifact of the way we have modeled distortions. Had there been externalities, for instance, in addition to taxes, and if the policymaker were uncertain about the parameters of the externalities, there would usually be *no* tax policy that achieves an ex post optimal allocation. Ex post optimality therefore does not provide a general resolution of the policy paralysis problem.

Appendix

Proof of Theorem 1. The set of interior ex ante optimal allocations is a manifold of dimension $J - 1$ (see, e.g., Mas-Colell (1985, Proposition 4.6.9)), which we denote Y, and thus, generic subsets of Y are well-defined. For any ex ante optimal allocation $x \gg 0$ (we can ignore boundary optima as nongeneric), there is a supporting $p(x) \in R^{SL}_{++}$ such that each $DEu_j(x_j)$ is proportional to $p(x)$, and we arrange $p(x)$ as the $S \times L$ matrix $P(x)$ with the sth row given by the coordinates of $p(x)$ that are proportional to $D_{x_j(\omega_s)}Eu_j$. We normalize $p(x)$ and hence $P(x)$ by requiring $p(x)$ to lie in the LS dimensional unit simplex.

Since $L \geq S$, we can define the square matrices P_s, $s = 1, ..., S$, by setting, for $k \leq s$, the kth row of P_s equal to the first s coordinates of $p_{\omega_k}(x)$. We now show that there is a generic subset of $Y \times Q$ such that P_S has rank S. Since for any $(x, (e, h)) \in Y \times Q$, P_1 trivially has rank 1, it is sufficient to show that, for $s = 1, ..., S - 1$, if there is a generic subset $G_s \subset Y \times Q$ at which P_s with rank s, then there is another generic subset $G_{s+1} \subset Y \times Q$ at which P_{s+1} with rank $s+1$. Given the induction assumption, we define the function g_{s+1}: $G_s \to R$ by setting $g_{s+1}(x, (e, h))$ equal to the determinant of P_{s+1}. Calculating det P_{s+1} by cofactor expansion along row $s + 1$, the derivative of det P_{s+1} with respect to the $(s+1)$st entry of $p_{\omega_{s+1}}(x)$ must be nonzero given the induction assumption that P_s has rank s. Moreover, we can change this coordinate of $p(x)$ without changing any other coordinate by increasing $D_{x_{s+1}(\omega_{s+1})}Eu_j$ for all j. Thus $Dg_{s+1} \neq 0$, and so by the implicit function theorem (for manifolds, sometimes called the preimage theorem, see, e.g., Guillemin and Pollack (1974)), the subset of G_s such that det P_s+1 = 0, say Z_{s+1}, is a manifold of dimension equal to $\dim(Y \times Q) - 1$ and hence a closed and measure-0 subset of $Y \times Q$. We therefore set $G_{s+1} = G_s \setminus Z_{s+1}$. Hence on G_S, $P(x)$ has rank S. Moreover, by Fubini's theorem, there must be a generic subset $G \subset Q$ such that, for all $(e, h) \in G$, $P(x)$ has rank S for all x in a generic subset of the ex ante optimal allocations of (e, h).

For any such x, define for each j, $c_j = (p_{\omega_1}(x) \cdot x_j(\omega_1), ..., p_{\omega_S}(x) \cdot x_j(\omega_S))$. Since $P(x)$ has rank S, there is for any j a solution Δe_j to $(p_{\omega_1}(x) \cdot (\Delta e_j + e_j(\omega_1)), ..., p_{\omega_S}(x) \cdot (\Delta e_j + e_j(\omega_S))) = c_j$, that is, a Δe_j such that

(A1) $$P(x)\Delta e_j = c_j - (p_{\omega_1}(x) \cdot e_j(\omega_1), ..., p_{\omega_S}(x) \cdot e\ (\omega_S)).$$

For $j = 2, ..., J$, set Δe_j as a solution to A1, and set $\Delta e_1 = -\sum_{j=2}^{J} \Delta e_j$; it is readily confirmed that Δe_1 also solves (A1) for $j = 1$. Since, for each ω_s, $p_{\omega_s}(x)$ is an equilibrium price vector for the economy at ω_s when $\tau = 0$ and Δe is specified as above, setting $f = x$ reaches the ex ante optimal allocation x. ■

Proof of Theorem 2. Step I. We first show that, for any τ, there is a generic subset of economies $G \subset Q$ such that for any $(e, h) \in G$ and any equilibrium allocation $x(\omega_s)$ in state s of (e, h) there exists a C^1 function g_{ω_s} from an open set $\Pi_O \subset R_+^L \times R^{L(J-1)}$ of policy instruments that contains $(\tau, \Delta e = 0)$ to allocations such that (1) $g_{\omega_s}(\tau, \Delta e = 0) = x(\omega_s)$ and (2) for any $(\tau', \Delta e') \in \Pi_O$, $g_{\omega_s}(\tau', \Delta e'))$ is a locally unique equilibrium allocation of (e, h) in state s when the policy instruments are $(\tau', \Delta e')$.

To establish this preliminary result, we fix the state and omit any notation of it. Let a *labeling* be a pair of nonempty disjoint subsets of $\{1, ..., L\} \times \{1, ..., J\}$, denoted B and S, such that $(i, j) \in B$ if and only if there exists a $j' \neq j$ such that $(i, j') \in S$ and if $(i, j) \in S$ then there exists a $i' \neq i$ such that $(i', j) \in B$. That is, some agent j buys a good i if and only if some other agent j' sells i, and if some agent sells some good i then that agent buys some other good i'. Reset commodity indices so that the first $\iota = \#\{i: \exists\, j \text{ such that } (i, j) \in B \cup S\}$ goods are the goods in $B \cup S$. Let $\hat{p} = (1, p_2, ..., p_\iota)$, let $\hat{x}$ denote the projection of a consumption profile $x \in R^{LJ}$ onto the $\#(B \cup S)$-dimensional coordinate subspace of the consumption bundles listed in $B \cup S$, let $\hat{e}$ denote the projection of the endowment profile e onto the same subspace, and let $\kappa = \#\{j: \exists\, i \text{ such that } (i, j) \in B \cup S\}$.

For any of the finite number of labelings, let $F: Q \times R^L \times R^{\iota-1}_{++} \times R^{\#B\cup S}_{++} \times R^{\kappa}_{++} \times R_{+} \to R^{\#(B\cup S)+\kappa+\iota}$ denote the C^1 function given by $F((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = \langle \{ [D_{x_{ij}} u_j(x_j) - \lambda_j p_i]_{i:(i,j)\in S}, [D_{x_{ij}} u_j(x_j) - \lambda_j p_i(1+\tau_i)]_{i:(i,j)\in B}, \sum_{i:(i,j)\in B}(1+\tau_i)p_i(x_{ij} - e_{ij}) + \sum_{i:(i,j)\in S} p_i(x_{ij} - e_{ij}) - t/J \}_{j:(i,j)\in B\cup S}, t - \sum_{(i,j)\in B} p_i \tau_i (x_{ij} - e_{ij}), \sum_{j=1}^{J}(x_{2j} - e_{2j}), ..., \sum_{j=1}^{J}(x_{Lj} - e_{Lj}) \rangle$, where each $x_{ij} = \hat{x}_{ij}$ if $(i, j) \in B \cup S$, and $x_{ij} = e_{ij}$ otherwise. If (p, x) is an equilibrium for the economy (e, h) with taxes τ, there is a Lagrange multiplier λ_j for each agent j such that (x_j, λ_j) solves j's maximization problem and $F((e, h), \tau, p_2/p_1, ..., p_\iota/p_1, \hat{x}, \lambda, t) = 0$ for some labeling. Note that F sets any x_{ij} such that $(i, j) \notin B \cup S$ equal to e_{ij}. When a 0 of F describes an equilibrium, these are goods for which agents optimally consume exactly their endowment; there is no condition setting the marginal utility of these goods equal to $\lambda_j p_i$ or $\lambda_j(1 + \tau_i)p_i$ due to the kinks in agents' budget sets at endowment points. If $D_{\hat{p},\hat{x},\lambda,t} F((e,h), \tau, \hat{p}, \hat{x}, \lambda, t)$ is nonsingular whenever $F((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = 0$, the inverse function theorem implies that, for each (e, h), the $(\hat{p}, \hat{x}, \lambda, t)$ such that $F((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = 0$ are locally isolated. Hence the equilibrium allocations such that every agent consumes nonendowment bundles in the same coordinate subspace are also locally isolated. The transversality theorem (see, e.g., Guillemin and Pollack (1974)) therefore implies that local uniqueness in this sense obtains for a full measure set of economies if $DF((e, h), \tau, \hat{p}, \hat{x}, \lambda, t)$ has full row rank whenever $F((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = 0$. We omit the largely routine calculation that full row rank does in fact obtain.

We have considered only labelings that consist of nonempty sets of indices, thus excluding the no-trade equilibria where each agent consumes a nonendowment bundle in no coordinate subspace; but for any economy there is at most one no-trade allocation. Hence, using the intersection of the finite number of full measure sets defined above, one for each labeling, we conclude that there is a full measure set of economies such that each equilibrium at which every agent consumes nonendowment bundles in the same coordinate subspace is locally unique. We next show that generically each equilibrium allocation x is also locally isolated from equilibrium allocations in which nonendowment consumptions lie in a different coordinate subspace. If $x = (\hat{x}, x_{ij} = e_{ij}$ for $(i, j) \notin B \cup S)$ is an equilibrium allocation for some economy (e, h) and x fails to be locally isolated, there must be a labeling, a corresponding F, and a $(\tilde{i}, \tilde{j}) \in B \cup S$ such that $F((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = 0$ and $x_{\tilde{i}\tilde{j}} = e_{\tilde{i}\tilde{j}}$. To exclude this possibility, we add to the range of each F the additional term $x_{\tilde{i}\tilde{j}} - e_{\tilde{i}\tilde{j}}$, where $(\tilde{i}, \tilde{j}) \in B \cup S$, thus defining a function F^*. We again omit the calculation showing that DF^* has full row rank, i.e., rank $\#(B \cup S) + \kappa + \iota + 1$. Hence, for a full measure set of economies, if $F^*((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = 0$ then $D_{\hat{p}, \hat{x}, \lambda, t} F^*((e, h), \tau, \hat{p}, \hat{x}, \lambda, t)$ has rank $\#(B \cup S) + \kappa + \iota + 1$. But since $D_{\hat{p}, \hat{x}, \lambda, t} F^*((e, h), \tau, \hat{p}, \hat{x}, \lambda, t)$ has only $\#(B \cup S) + \kappa + \iota$ columns, it must be that at this set of economies there exists no $(\hat{p}, \hat{x}, \lambda, t)$ such that $F^*((e, h), \tau, \hat{p}, \hat{x}, \lambda, t) = 0$. Constructing such a full measure set for each $(\tilde{i}, \tilde{j}) \in B \cup S$, and taking the intersection of all of the full-measure sets of economies defined so far, we conclude that for a full-measure set G of economies, equilibrium allocations are generically locally unique. (Normalized equilibrium prices are not locally unique since $\hat{p}$ does not specify prices for goods that no one buys or sells.) That these equilibrium allocations are C^1 functions of (e, h) and τ then follows from the implicit function theorem, and the openness of the set of economies we have identified follows from the fact that for any F we may place the endogenous variables $(\hat{p}, \hat{x}, \lambda, t)$ in a compact set. When a 0 of F specifies equilibrium values for the variables, the specification of g is completed by setting $x_{ij} = e_{ij}$ when $(i, j) \notin B \cup S$.

Step II. Turning to the theorem itself, consider an arbitrary selection of S of the functions F defined above, $\mathbf{F} = (F_{\omega_1}, ..., F_{\omega_S})$, one F_{ω_s} chosen from each state. If we express G as the product $G_{\omega_1} \times ... \times G_{\omega_S}$, then $\mathbf{F}$ is defined on the S-fold product of the sets $G_{\omega_s} \times R^L \times R_{++}^{\iota-1} \times R_{++}^{\#B \cup S} \times R_{++}^{\kappa} \times R_{+}$. We restrict ourselves to an open subset of this domain, say Y, that contains all 0's of **F** and such that each $D_{\hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s), t(\omega_s)} F_{\omega_s}(y)$ is nonsingular, where $y = (y_{\omega_1}, ..., y_{\omega_S})$ denotes a typical element of Y. Also, $\hat{p}(y_{\omega_s})$, $\hat{x}(y_{\omega_s})$, etc., will denote the indicated coordinates of y_{ω_s}. We now extend the g given in step I by defining a C^1 function χ: $Z \to R^{JSL}$, where Z is an open subset of $Y \times R_+^L \times R^{L(J-1)}$ that contains $(y, \tau, \Delta e)$ whenever $y \in Y$, $\tau = \tau(y_{\omega_s})$, $\Delta e = \Delta e(y_{\omega_s})$, and $D_{\hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s), t(\omega_s)} F_{\omega_s}$ evaluated at $(h(y_{\omega_s}), e(y_{\omega_s}) + \Delta e, \tau, \hat{p}(y_{\omega_s}), \hat{x}(y_{\omega_s}), \lambda(y_{\omega_s}), t(y_{\omega_s}))$ is nonsingular for each ω_s. Given $(y, \tau', \Delta e) \in Z$, consider the implicit function theorem solution values of $(\hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s), t(\omega_s))$ for the equation

$$F_{\omega_s}(h(y_{\omega_s}), e(y_{\omega_s}) + \Delta e, \tau', \hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s), t(\omega_s)) = F_{\omega_s}(y_{\omega_s}),$$

where if $\Delta e = 0$, $\tau' = \tau(y_{\omega_s})$, and $F_{\omega_s}(y_{\omega_s}) = 0$, then we set the solution $(\hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s), t(\omega_s))$ to equal the corresponding coordinates of y_{ω_s}. To define χ, set the $\hat{x}(\omega_s)$ coordinates of $\chi_{\omega_s}(\tau', \Delta e, y)$ to equal the $\hat{x}(\omega_s)$ coordinates of this solution and the remaining coordinates, as before, to equal agents' endowments. Notice that when $F_{\omega_s}(y_{\omega_s}) = 0$ and equilibrium values of the variables are specified, $\chi_{\omega_s}(\tau', \Delta e, y)$ will coincide with the g_{ω_s} given in step I. If $\mathbf{F}(y) = 0$ and y specifies equilibrium values for $(\hat{p}, \hat{x}, \lambda, t)$, then χ gives, as a function of τ' near τ and Δe near 0, the economy's unique nearby equilibrium consumption profile. Let μ: $R^L \times R^{L(J-1)} \times Y \to R^J$ denote the ex ante utilities of these consumption profiles, i.e., $\mu_j(\tau', \Delta e, y) = Eu_j(\chi_j(\tau', \Delta e, y)) = \sum_{s=1}^{S} \pi_s\, u_j(\chi_{j,\omega_s}(\tau', \Delta e, y))$. The proof is complete if we can show for a generic subset of economies that, for any 0 of any such **F**, $D_{\tau,\Delta e}\,\mu(\tau, \Delta e = 0, y)$ has rank J: since then the linear map $D_{\tau,\Delta e}\,\mu(\tau, \Delta e, y)$ is onto, there is a $(\tau', \Delta e')$ such that $D_{\tau,\Delta e}\,\mu(\tau, \Delta e, q)\,(\tau', \Delta e') \gg 0$ and hence for $\varepsilon > 0$ sufficiently small, one of the allocations reached by $(\tau, \Delta e) + \varepsilon\,(\tau', \Delta e')$ increases each Eu_j.

Letting ε_{ij} denote a transfer from agent 1 to agent j of good i, we will need to consider only the derivatives of μ with respect to $\varepsilon_{22}, \varepsilon_{12}, \varepsilon_{13}, ..., \varepsilon_{1J}$. We define the functions $\mathbf{F}_i$, $i = 1, ..., J$, by supplementing **F** with an additional term equal to the determinant of a matrix M_i of derivatives of μ. For $\mathbf{F}_1$, M_1 is just the 1×1 matrix $D_{\varepsilon_{22}}\mu_1$. Each M_i, $i \geq 2$, is an $i \times i$ matrix whose columns consist of derivatives of coordinates of μ with respect to the first i of the variables $\varepsilon_{22}, \varepsilon_{12}, \varepsilon_{13}, ..., \varepsilon_{1J}$ and whose rows consists of the derivatives of the first i of the variables $\mu_1, ..., \mu_J$. Thus, each M_i, $i \geq 2$, is M_{i-1} with an additional row and column added. We now show that there is a generic subset of economies for which each $\mathbf{F}_i$ has no 0; since we can repeat this argument for any **F**, this shows that any $D_{\tau,\Delta e}\,\mu(\tau,\,\Delta e = 0, y)$ has rank J.

We can decompose the effects of changes in the ε_{ij} on μ into a sum of the direct utility effects of the transfers, which depend on Du_j, and the indirect effects via changes in the $p(\omega_s)$ and t, which depend on $D^2 u_j$ but not on Du_j (see Geanakoplos and Polemarchakis (1986) for more on this point). The matrix of the direct effects of the pertinent ε_{ij} on μ is given by DE =

$$\begin{array}{ccccc} \varepsilon_{22} & \varepsilon_{12} & \varepsilon_{13} & \cdots & \varepsilon_{1J} \end{array}$$

$$\begin{pmatrix} -\Sigma_{s=1}^{S}\pi_s D_{x_2}u_1(\chi_{1,\omega_s},\omega_s) & -\Sigma_{s=1}^{S}\pi_s D_{x_1}u_1(\chi_{1,\omega_s},\omega_s) & -\Sigma_{s=1}^{S}\pi_s D_{x_1}u_1(\chi_{1,\omega_s},\omega_s) & \cdots & -\Sigma_{s=1}^{S}\pi_s D_{x_1}u_1(\chi_{1,\omega_s},\omega_s) \\ \Sigma_{s=1}^{S}\pi_s D_{x_2}u_2(\chi_{2,\omega_s},\omega_s) & \Sigma_{s=1}^{S}\pi_s D_{x_1}u_2(\chi_{2,\omega_s},\omega_s) & 0 & \cdots & 0 \\ 0 & 0 & \Sigma_{s=1}^{S}\pi_s D_{x_1}u_3(\chi_{3,\omega_s},\omega_s) & & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & & \Sigma_{s=1}^{S}\pi_s D_{x_1}u_J(\chi_{J,\omega_s},\omega_s) \end{pmatrix}.$$

The function $\mathbf{F}_1$ is transverse to 0 (i.e., $D\mathbf{F}_i$ has full row rank whenever $\mathbf{F}_i = 0$) since we may simultaneously multiply $(Du_j(\cdot,\omega_s),\ \lambda_{j,\omega_s})$ for each j and some ω_s by the same constant, thus perturbing the upper left term of DE while leaving the value of **F** unchanged. It follows that for a generic subset of economies $\mathbf{F}_1 = 0$ has no solution: if it did then the matrix of derivatives of $\mathbf{F}_1$ with respect to the endogenous variables $\hat{p}(\omega_s)$, $\hat{x}(\omega_s)$, $\lambda(\omega_s)$, $t(\omega_s)$, $s = 1, ..., S$, would have full row rank at any solution, which is impossible since this matrix has more rows than columns. We henceforth remove the closed 0-measure set of parameters such that $\mathbf{F}_1 = 0$ from the range of the remaining $\mathbf{F}_i$. To show that $\mathbf{F}_2$ is transverse to 0 requires an initial argument that

$$\frac{D_{x_1}u_1(x_1,\omega_1)}{D_{x_2}u_1(x_1,\omega_1)} = \frac{D_{x_1}u_1(x_1,\omega_2)}{D_{x_2}u_1(x_1,\omega_2)}$$

is not satisfied at any 0 of **F** for a generic subset of economies. This is readily established with an separate transversality argument that shows that we can add this equation to an arbitrary pair of F's for the economies at states 1 and 2, then the resulting function is transverse to 0 (perturb, at one of the states, every j's marginal utility for one of the goods and that good's price) and hence this equation is generically not satisfied at a 0 of any **F**. Given that this equality is not satisfied, we may by independently rescaling $(Du_j(\cdot,\omega_1),\ \lambda_{j,\omega_1})$ and $(Du_j(\cdot,\omega_2),\ \lambda_{j,\omega_2})$ perturb the row 2-column 2 entry of the DE, without changing the other entries of DE or the value of **F**. If we calculate det M_2 by expansion of cofactors in the second row, and given our earlier restriction to parameters such that $\mathbf{F}_1 \neq 0$ and hence det $M_1 \neq 0$, we set that $\mathbf{F}_2$ is transverse to 0. We then proceed by induction, restricting the domain of each $\mathbf{F}_i$, $i = 3, ..., J$, to exclude the points at which $\mathbf{F}_{i-1} = 0$ has a solution: simply by rescaling $(Du_j(\cdot, \omega_1), \lambda_{j,\omega_1})$ for all j, each of the remaining $\mathbf{F}_i$ is seen to be transverse to 0, using the cofactor expansion of det $\mathbf{F}_i$ along row i. Thus generically $D_{\tau,\Delta e}\mu(\tau, \Delta e = 0, y)$ has rank J at any 0 of $\mathbf{F}$, as desired. ■

Proof of Theorem 3. Choose Ω' so that, for all j and ω^l, (1) $U_j(\omega^l) \neq U_h(\omega^l)$ for any agent $h \neq j$ and $U_j(\omega^l) \neq U_h(\hat{\omega})$ for any h and $\hat{\omega} \in \Omega \cup \Omega' \setminus \omega^l$, (2) the vectors $Du_j(e_j(\omega^l), \omega^l)$, $l = 1, ..., L$, are linearly independent, and (3) $e(\omega^l) \gg 0$ is a Pareto optimal allocation for the economy $(u_j(\cdot, \omega^l), e_j(\omega^l))_{j=1}^J$.

The strict concavity of the $u_j(\cdot, \omega^l)$ and (3) imply that for any status quo policy $(\bar{\tau}, 0, \bar{f})$, $\bar{f}_j(\omega^l) = e_j(\omega^l)$ for all j and ω^l. Given (1), it is sufficient to show that at any $(\tau, \Delta e \neq 0, f)$ there exists a ω^l and j such that $u_j(e_j(\omega^l), \omega^l) > u_j(f_j(\omega^l), \omega^l)$. Suppose, to the contrary, that $u_j(f_j(\omega^l), \omega^l) \geq u_j(e_j(\omega^l), \omega^l)$ for all ω^l and j, and hence (given strict concavity) that $f_j(\omega^l) = e_j(\omega^l)$ for all ω^l and j. Given the arguments in section 2 on the suboptimality of equilibria where traded goods have nonzero taxes, it follows that if $\Delta e_{ij} \neq 0$ for any agent j and any good i and $J \geq 2$ and $L \geq 2$, then $\tau_i = 0$. Therefore $t(\omega^l) = 0$ which also holds when $J = 1$ or $L = 1$ since then there is no trade. From the definition of the budget constraint, $p(\omega^l) \cdot (e_j(\omega^l) - (e_j(\omega^l) + \Delta e_j)) \leq 0$, where $p(\omega^l)$ is an equilibrium price vector corresponding to f at ω^l. Hence $p(\omega^l) \cdot \Delta e_j \leq 0$ and, since $\Sigma_{j=1}^J \Delta e_j = 0$, $p(\omega^l) \cdot \Delta e_j = 0$. Since, for all j and ω^l, there is some $\lambda_j^l > 0$ such that $p_i(\omega^l) = \lambda_j^l D_{x_{ij}} u_j(e_j(\omega^l), \omega^l)$ for all goods i such that $\Delta e_{ij} \neq 0$, $\lambda_j^l D_{x_j} u_j(e_j(\omega^l), \omega^l) \cdot \Delta e_j = 0$. Condition (2) then implies that $\Delta e_j = 0$ for all j, a contradiction. ■

Proof of Theorem 4. Letting $(\hat{\tau}, \Delta\hat{e}, \hat{f})$ be differentiable and a regular maximum for u and such that $D^2_{(\tau,\Delta e), b} w_u(1^{JS}, (\hat{\tau}, \Delta\hat{e}))$ has rank LJ, there must be a LJ dimensional coordinate subspace of B, say B^*, such that $D^2_{(\tau,\Delta e), b^*} w_u(1^{JS}, (\hat{\tau}, \Delta\hat{e}))$ is nonsingular, where b^* denotes a typical element of B^*. Label coordinates so that B^* is spanned by the first LJ coordinates of R^{JS}. By the implicit function theorem, there is a C^1 function, say b^*, from some open subset $\Pi' \subset R^{LJ}$ containing $(\hat{\tau}, \Delta\hat{e})$ to B^* such that $b^*(\hat{\tau}, \Delta\hat{e}) = 1^{LJ}$ and

$$D_{(\tau,\Delta e)}w_u((b^*((\tau,\Delta e)),1^{JS-LJ}),(\tau,\Delta e))=0$$

for all $(\tau, \Delta e) \subset \Pi'$. The fact that $(\hat{\tau}, \Delta\hat{e}, \hat{f})$ is a regular maximum implies that there are open sets $B_O \subset B^*$ and $\Pi_O \subset R^{LJ}$ containing 1^{LJ} and $(\hat{\tau}, \Delta\hat{e})$, respectively, such that for $b^* \in B_O$ and $(\tau, \Delta e) \in \Pi_O$, $D^2_{\tau,\Delta e}w_u((b^*((\tau, \Delta e)), 1^{JS-LJ}), (\tau, \Delta e))$ is negative definite.

The above establishes that all $(\tau, \Delta e) \in \Pi_O$ are strict maxima of w_u for some $b \in B_O$ if we constrain $(\tau, \Delta e)$ to be an element of Π_O. We now show that there is an open $\Pi^* \subset \Pi_O$ containing $(\hat{\tau}, \Delta\hat{e})$ such that for some b all $(\tau, \Delta e) \in \Pi^*$ are unconstrained strict maxima of w_u. Suppose, to the contrary, that there is a sequence $\{(\tau, \Delta e)_t\}$, where $(\tau, \Delta e)_t \neq (\hat{\tau}, \Delta\hat{e})$ for all t, such that $(\tau, \Delta e)_t \to (\hat{\tau}, \Delta\hat{e})$ and such that each $(\tau, \Delta e)_t$ is not a strict maximum of w_u. Let $\{(\tilde{\tau}, \Delta\tilde{e})_t\}$ be a sequence such that, for all t, $(\tilde{\tau}, \Delta\tilde{e})_t$ is a (possibly nonstrict) maximum of w_u when $b = (b^*((\tau, \Delta e)_t), 1^{JS-LJ})$ and where $\Delta\tilde{e}_t \in \Delta E$. Since each $(\tau, \Delta e)_t$ is not a strict maximum, we may choose $\{(\tilde{\tau}, \Delta\tilde{e})_t\}$ so that $\{(\tilde{\tau}, \Delta\tilde{e})_t\} \neq (\tau, \Delta e)_t$ for all t. Since each $(\tau, \Delta e)_t$ is a strict maximum of w_u when $b = (b^*((\tau, \Delta e)_t)$ and $(\tau, \Delta e)$ is restricted to Π_O, $(\tilde{\tau}, \Delta\tilde{e})_t \notin \Pi_O$ for all t. We have already restricted $\Delta\tilde{e}_t$ to be an element of the compact set ΔE; we may also assume that $\tilde{\tau}_t$ lies in a compact subset of R^L_+ since if τ is sufficiently large, no trade and hence the same f occurs. Since therefore we can restrict ourselves to a compact set of policy instruments, say $\bar{\Pi}$, and Π_O is open, there is a subsequence of $\{(\tilde{\tau}, \Delta\tilde{e})_t\}$ converging to a $(\bar{\tau}, \Delta\bar{e}) \in \bar{\Pi} \setminus \Pi_O$. Given the continuity of w and the fact that $b^*((\tau, \Delta e)_t) \to 1^{LJ}$, $(\bar{\tau}, \Delta\bar{e})$ is an unconstrained maximum of w_u when $b = 1$, contradicting $(\hat{\tau}, \Delta\hat{e})$ being a strict maximum.

The openness of the policies that satisfy Definition 6 (2) is self-evident. In addition, since LJ is the maximal rank of $D^2_{(\tau,\Delta e),\, b}w_u(1^{JS}, (\tau, \Delta e))$, the policies that satisfy the rank condition are also open, which completes the proof. ■

Proof of Theorem 5. Include in $\hat{\Omega}$ a set of L states $\{\tilde{\omega}_1, ..., \tilde{\omega}_L\}$ at which each $e(\tilde{\omega}_i)$ is Pareto optimal for some set of J utility functions utilities $\bar{u}_j(\cdot, \tilde{\omega}_i)$, $j = 1, ..., J$, and the total resources $\sum_{j=1}^{J} e_j(\tilde{\omega}_i)$. Choose the $\bar{u}_j(\cdot, \tilde{\omega}_i)$ so that (i) $\bar{u}_j(\cdot, \tilde{\omega}_i) \neq \bar{u}_{j'}(\cdot, \omega)$ but $D_{x_j(\tilde{\omega}_i)}\bar{u}_j(e_j(\tilde{\omega}_i), \tilde{\omega}_i) = D_{x_{j'}(\tilde{\omega}_i)}\bar{u}_{j'}(e_{j'}(\tilde{\omega}_i), \tilde{\omega}_i)$ for all pairs (j, j') and all $\omega \in \Omega \cup \{\tilde{\omega}_1, ..., \tilde{\omega}_L\}$, (ii) the equilibrium allocation at $\tilde{\omega}_i$ is a C^1 function g of $(\tau, \Delta e)$, (iii) each $\bar{u}_j(\cdot, \tilde{\omega}_i) \circ g_j$ is differentiably strictly concave, and (iv) the vectors $D_{x_1(\tilde{\omega}_1)}\bar{u}_1(e_1(\tilde{\omega}_1), \tilde{\omega}_1), ..., D_{x_1(\tilde{\omega}_L)}\bar{u}_1(e_1(\tilde{\omega}_L), \tilde{\omega}_L)$ are linearly independent.

Next, for each of the L goods, construct a further set of states in $\hat{\Omega}$ as follows. As a preliminary, we first specify the states ω^i, $i = 1, ..., L$. For $i = 1, ..., L - 1$, define ω^i by letting each j have a utility $\bar{u}_j(\cdot, \omega)$ that is the sum of a C^2 differentiably strictly concave and differentiably strictly increasing function of goods i and L. Set $e(\omega^i)$ so that $e_j(\omega^i)$ is a constant function of j. Choose the J utility functions on goods i and L so that (1) for each distinct pair of agents j and j', $\bar{u}_j(\cdot, \omega^i) \neq \bar{u}_{j'}(\cdot, \omega^i)$, (2) for distinct pair of states ω^i and $\omega^{i'}$ and any pair of agents j and j', $\bar{u}_j(\cdot, \omega^i) \neq \bar{u}_{j'}(\cdot, \omega^{i'})$, (3) ω^i has a unique equilibrium allocation $f(\omega^i)$ for $(\tau, \Delta e) \in K$ given by a C^1 function g of $(\tau, \Delta e)$, and (4) letting μ_{j,ω^i} denote the composition $\bar{u}_j(\cdot, \omega^i) \circ g_j$, then, for every $(\tau, \Delta e) \in K$, $D_{\tau_i}\mu_{1,\omega^i}(\tau, \Delta e) > 0$, $D_{\tau_L}\mu_{1,\omega^i}(\tau, \Delta e) < 0$, $D_{\tau_i}\mu_{2,\omega^i}(\tau, \Delta e) < 0$, and $D_{\tau_L}\mu_{2,\omega^i}(\tau, \Delta e) > 0$, and (5) each μ_{j,ω^i} is differentiably strictly concave. Define ω^L by letting all agents derive utility only from goods L and $L - 1$, letting conditions (1) through (3) and (5) be satisfied, and by requiring $D_{\tau_L}\mu_{1,\omega^L}(\tau, \Delta e) > 0$, $D_{\tau_{L-1}}\mu_{1,\omega^L}(\tau, \Delta e) < 0$, $D_{\tau_L}\mu_{2,\omega^L}(\tau, \Delta e) < 0$, and $D_{\tau_{L-1}}\mu_{2,\omega^L}(\tau, \Delta e) > 0$. We now use $\omega^1, ..., \omega^L$ to specify the states in $\hat{\Omega}$: for each ω^i, let Ω^i denote the $J!$ states constructed by taking all possible permutations of the agent indices of the utilities in ω^i and set $\hat{\Omega} = \{\tilde{\omega}_1, ..., \tilde{\omega}_L\} \cup \Omega^1 \cup ... \cup \Omega^L$. Let $\hat{S} = \# \hat{\Omega}$.

Let $v(b, (\tau, \Delta e))$ denote $\sum_{j=1}^{J} \sum_{\omega_s \in \hat{\Omega}} \hat{\pi}_s b_{js} u_j(g_{j,\omega_s}(\tau, \Delta e), \omega_s)$, where $g(\tau, \Delta e)$ gives the unique equilibrium allocation for $(\tau, \Delta e)$ and the $\hat{\pi}_s$ can take any value such that $\hat{\pi}_s = \hat{\pi}_{s'}$ if ω_s and $\omega_{s'}$ are both elements of the same Ω^i. Set r_{js} for $\omega_s \in \hat{\Omega}$ and $j = 1, ..., J$, so that, for the assignment $u = (..., r_{js}\bar{u}_j, ...)$ for the utilities that appear in $\hat{\Omega}$, $D_{\tau_i}v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$ for each τ_i and $D_{\Delta e_{ij}}v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$ for each Δe_{ij} when evaluated at the $(\tau, \Delta e)$ that are the policy instruments of the regular maximum given by the assumptions of the Theorem. The r_{js} may be set so that $D_{\tau_i}v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$ since for each τ_i, $i = 1, ..., L$, the μ_{1,ω^i} and μ_{2,ω^i} defined in (4) above are respectively increasing and decreasing in τ_i; hence for any values of the r_{js} assigned to the utility functions not owned by agents 1 and 2 at ω^i, we can adjust either the r_{js} assigned to $\bar{u}_1(\cdot, \omega^i)$ or the r_{js} assigned to $\bar{u}_2(\cdot, \omega^i)$ to ensure that $D_{\tau_i}v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$. Notice that any change in τ will not alter $x(\tilde{\omega})$ since $x(\tilde{\omega})$ is Pareto optimal and will not change $x(\omega_s)$ for any ω_s derived from ω^k for $k \neq i$ since agents at these ω_s neither buy nor sell i. To ensure that $D_{\Delta e_j}v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$, observe that our inclusion of all permutations of the agent indices and our restriction on $\hat{\pi}_s$ imply that $\sum_{\omega_s \in \Omega^i} \hat{\pi}_s D_{\Delta e_j} \mu_{j'(\omega_s)}(\tau, \Delta e) = 0$, where μ_j is the composition $u_j(\cdot, \omega_s) \circ g_{j,\omega_s}(\tau, \Delta e)$ and where, for any fixed utility $\hat{u}$ that appears at some ω_s in some Ω^i, $j'(\omega_s)$ denotes the agent that has $\hat{u}$. So, if we set each r_{js} for $\omega_s \in \{\tilde{\omega}_1, ..., \tilde{\omega}_L\}$ equal to 1, the $(\tau, \Delta e)$ that are the policy instruments of the given regular maximum must maximize $v(1^{J\hat{S}}, (\tau, \Delta e))$. Hence $D_{\Delta e_j} v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$.

Our differentiability assumptions implies that the policy instruments of the given regular maximum $(\tau, \Delta e, f)$ are differentiable policy instruments for $(\hat{\Omega}, \hat{\pi})$, while the concavity assumptions on the μ_{j,ω^i} and our choices for the r_{js} imply that $(\tau, \Delta e)$, joined with the allocation that $(\tau, \Delta e)$ induces, is a regular maximum for $(..., r_{js} \bar{u}_j , ...)$ in the model $(\hat{\Omega}, \hat{\pi})$. It follows that $(\tau, \Delta e, \hat{f})$, where $\hat{f}$ is f joined with the allocation that $(\tau, \Delta e)$ induces at $\hat{\Omega}$, is differentiable and a regular maximum for the assignment $\hat{u}$, where $\hat{u}$ is the u given in the Theorem joined with $(..., r_{js} \bar{u}_j , ...)$ for the utilities in $\hat{\Omega}$, in the model $(\Omega \cup \hat{\Omega}, (\lambda\pi, (1-\lambda)\hat{\pi}))$ generated by any λ.

It remains to show that $w_{\hat{u}}$ satisfies the rank condition. Consider the columns of the matrix $D^2_{(\tau,\Delta e), b} v(1^{J\hat{S}}, (\tau, \Delta e))$ that correspond, respectively, to the b's assigned to agents 2 through J at $\tilde{\omega}_1, ..., \tilde{\omega}_L$ and the b's assigned to $\bar{u}_1(\cdot, \omega^i)$, $i = 1, ..., L$. Given our assumptions on the μ_{1,ω^i} and the Pareto optimality of the allocations at the $\tilde{\omega}$ states, these columns have the form

$$\begin{array}{c} \Delta e_2 \\ \vdots \\ \Delta e_J \\ \tau_1 \\ \vdots \\ \tau_{L-1} \\ \tau_L \end{array} \begin{bmatrix} \tilde{P}_2 & & 0 & 0 & \cdots & 0 & 0 \\ & \ddots & & \vdots & & \vdots & \vdots \\ 0 & & \tilde{P}_J & 0 & \cdots & 0 & 0 \\ 0 & & 0 & + & & 0 & 0 \\ \vdots & & \vdots & & \ddots & & \\ 0 & & 0 & 0 & & + & - \\ 0 & & 0 & - & \cdots & - & + \end{bmatrix}$$

where $\tilde{P}_j$ is the nonsingular square matrix whose ith column is $D_{x_j(\tilde{\omega}_i)} \bar{u}_j(e_j(\tilde{\omega}_i), \tilde{\omega}_i)$ and +'s and −'s indicate the signs of entries. Since the linear independence assumption (iv) implies that each $\tilde{P}_j$ is nonsingular, this matrix of columns has rank LJ. Since the submatrix of $D^2_{(\tau,\Delta e), b} w_{\hat{u}}(1^{JS}, (\tau, \Delta e))$ that consists of the columns that correspond to the same variables has the same rank as the above matrix (each column merely being rescaled by $(1-\lambda)$),

$D^2_{(\tau,\Delta e),\, b}\, w_{\hat{u}}(1^{JS}, (\tau, \Delta e))$ also has rank LJ.

The utility functions given in the definitions of the ω^i are increasing only in goods i and L. To ensure that the utilities in $\hat{\Omega}$ meet the maintained assumptions of the model, perturb the utilities given above by adding a small multiple of a C^2 differentiably strictly concave and differentiably strictly increasing function of the remaining $L-2$ goods. Since $D^2_{(\tau,\Delta e),\, b}\, v(1^{J\hat{S}}, (\tau, \Delta e))$ having rank LJ is a full rank condition, its rank will persist for a small perturbation. And given that $D^2_{(\tau,\Delta e),\, b}\, v(1^{J\hat{S}}, (\tau, \Delta e))$ has rank LJ, the implicit function theorem implies that we may adjust the b's so as to maintain the equalities $D_{\tau_i} v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$ and $D_{\Delta e_j} v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$. ■

References

Aumann, R., "Common priors: a reply to Gul," *Econometrica* 66: 929-938.

Champsaur, P. and G. Laroque, 1981, "Fair allocations in large economies" *Journal of Economic Theory* 25: 269–282.

Geanakoplos, J., 1990, "An introduction to general equilibrium with incomplete markets," *Journal of Mathematical Economics* 19: 1-38.

Geanakoplos, J. and H. Polemarchakis, 1986, "Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete," in *Essays in Honor of Kenneth Arrow*, Vol. 3, W. Heller, R. Starr, and D. Starrett, eds., Cambridge: Cambridge University Press.

Geanakoplos, J. and H. Polemarchakis, 1990, "Observability and optimality," *Journal of Mathematical Economics* 19: 153-65.

Guesnerie, R., 1995, *A Contribution to the Pure Theory of Taxation*, Cambridge: Cambridge University Press.

Guillemin, V. and A. Pollack, 1974, *Differential Topology*, Englewood Cliffs, NJ: Prentice-Hall.

Gul, F., 1998, "A comment on Aumann's Bayesian view," *Econometrica* 66: 923-927.

Hammond, P., 1979, "Straightforward incentive compatibility in large economies," *Review of Economic Studies* 46: 263-82.

Harsanyi, J., 1953, "Cardinal utility in welfare economics and in the theory of risk-taking," *Journal of Political Economy* 61: 434-5.

Harsanyi, J., 1955, "Cardinal welfare, individualistic ethics and interpersonal comparisons of utility," *Journal of Political Economy* 63: 309-21.

Holmström, B. and R. Myerson, 1983, "Efficient and durable decision rules with incomplete information," *Econometrica* 51: 1799-1819.

Lerner, A., 1944, *The Economics of Control*, London: Macmillan.

Lipsey, R. and R. Lancaster, 1956, "The general theory of second best," *Review of Economic Studies* 24: 11-32.

Mandler, M., 1999, "Interpersonal comparisons of utility and the policy paralysis problem," *Social Choice and Welfare* 17: 95-115.

Magill, M. and M. Quinzii, 1996, *Theory of Incomplete Markets*, Volume 1, Cambridge: MIT Press.

Mas-Colell, A., 1985, *The Theory of General Economic Equilibrium: A Differentiable Approach*, Cambridge: Cambridge University Press.

May, K., 1952, "A set of independent necessary and sufficient conditions for simple majority decision," *Econometrica* 20: 680-4.

McCain, R., 1972, "Distributional equality and aggregate utility: further comment," *American Economic Review*: 62: 497-500.

McManus, M., G. Walton, and R. Coffman, 1972, "Distributional equality and aggregate utility: further comment," *American Economic Review* 62: 489-496.

Mirrlees, J., 1986, "The theory of optimal taxation," in *Handbook of Mathematical Economics*, Volume 3, K. Arrow and M. Intriligator, eds., Amsterdam: North Holland.

Sen, A., 1970, *Collective Choice and Social Welfare*, San Francisco: Holden Day.

Sen, A., 1973, "On ignorance and equal distribution," *American Economic Review* 63: 1022-1024.

Shafer, W. and H. Sonnenschein, 1976, "Equilibrium with externalities, commodity taxation, and lump sum transfers," *International Economic Review* 17: 601-11.

Suppes, P., 1966, "Some formal models of grading principles," *Synthese* 6: 284-306.

Varian, H., 1976, "Two problems in the theory of fairness," *Journal of Public Economics* 5: 249-60.