

Michael Mandler

Policy discrimination with and without interpersonal comparisons of utility

Received: 20 April 2004 / Accepted: 17 April 2006
© Springer-Verlag 2006

Abstract Can the Pareto criterion guide policymakers who do not know the true model of the economy? If policymakers specify *ex ante* preferences for agents, then Pareto improvements from a distorted status quo are usually possible, and with more commodities than states, one can implement almost every Pareto optimum. Unlike the standard second welfare theorem, planners cannot dictate allocations: agents must trade. Unfortunately *ex ante* preferences impose interpersonal comparisons. If policymakers merely aim to maximize some social welfare function then optimal policies form an open set; hence small changes in the environment do not necessitate any policy response. Planners with symmetric information about agents can sometimes intervene without making interpersonal comparisons.

Keywords Welfare theorems · Pareto optimality · Optimal taxation · Policy paralysis · Ignorance priors

JEL Classification Numbers D50 · D63 · D81

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

This paper incorporates some much revised portions of Harvard Institute of Economic Research Discussion Paper 1729. I have discussed this paper with many people; Oliver Hart and John Geanakoplos in particular helped sharpen my thinking, and an anonymous referee suggested several major improvements.

M. Mandler
Department of Economics, Royal Holloway College, University of London,
Egham, Surrey TW20 0EX, UK
E-mail: m.mandler@rhul.ac.uk

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 implement them. So in what sense does the second welfare theorem recommend markets as an allocation mechanism? This puzzle bolsters the widespread doubts that the Pareto criterion can deliver practical advice. To address these questions, we make explicit policymakers' lack of information about primitives and ask when policymakers can recommend policies that correct a preexisting distortion, which for concreteness we assume to be taxes on net trades. As we will see, if a policymaker can posit a hypothetical *ex ante* stage at which each agent shares the policymaker's uncertainty and can make interpersonal comparisons between the potential preferences any given agent 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. 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 *ex ante* Pareto improvement, and again markets are indispensable. So there *is* a framework that makes rigorous the second welfare theorem's endorsement markets.

We use a general equilibrium model that is standard except that net commodity purchases are taxed. This distortion ensures that the status quo appears to call for policy intervention – externalities could serve just as well. 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.

When a policymaker can posit *ex ante* preferences for agents in the presence of policymaking uncertainty, call a policy x an *ex ante* improvement over y if x Pareto dominates y in terms of these *ex ante* preferences. If there are at least as many goods as states, then the second welfare theorem generically holds: 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. If the number of states is larger than the number of goods, then generically there is at least some policy that in the face of the preexisting distortion achieves an *ex ante* Pareto improvement (Theorem 2). Thus, despite doubts about the Pareto criterion's practicality, one can both acknowledge a policymaker's uncertainty and decree active policy advice.

But difficulties with the Pareto criterion remain. In the absence of 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 welfare-maximizing planner, no matter how the planner weights individual utility functions, would weakly recommend x over y . With policymaking uncertainty, this equivalence breaks down. The *ex ante* approach,

which we also call *agent-based*, extends method (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 method (1) to policymaking uncertainty no longer shares the advantage of method (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 raise the danger that very large numbers of policies will be declared optimal, in which case *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 related results lead us to identify agnosticism about interpersonal comparisons as the source of the impracticality of the Pareto criterion.

A countervailing consideration sometimes permits policy discrimination even when interpersonal comparisons are prohibited. We will require social welfare to be a function only of agents' ordinal and cardinal characteristics and not their indices. 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. But 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, when information is symmetric the concavity of utility implies that a utility-independent ranking will recommend that each agent consume the same bundle. Lerner (1944) famously used a similar logic to conclude that equalizing the distribution of endowments in a one-commodity model necessarily increases 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 when a planner foregoes problematic interpersonal comparisons. Lerner's work 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, suppose that one agent i 's consumption of good x imposes an externality on some other agent j . If a planner were maximizing some fixed social welfare function, a small upward shift in i 's externality effect on j would normally call for a small policy response, usually an increase in the tax on x . But if a policymaker instead wants simply to maximize *some* social welfare function (because to stay neutral, all welfare functions must be viewed as equally legitimate), the increase in the externality will not lead to any policy response. With the larger externality and the same tax rates the policymaker would still be maximizing a welfare function, perhaps one that assigns a lesser weight to j 's utility function. That social-welfare maximization can require nonzero taxes 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.

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.

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 closely related. Our model confronts each agent with the same choice set of net trades, similar to 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. Moreover, any implementation game must be anonymous if the policymaker begins with symmetric information about agents' characteristics. Finally if agents anticipate that, following the play of the revelation game, they can 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 model becomes similar to an implementation problem. Hammond used dominant strategies in his paper, but see Guesnerie (1995) and the references cited there for similar Bayesian arguments.

The contrast between the present paper and the implementation approach is misleading in a second and 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 there are no traditional assets, but this wrinkle does not interfere with the analytical machinery of the incomplete markets literature (see, e.g., Geanakoplos 1990; Magill and Quinzii 1996). 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 policy-maker. 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 $e \equiv (e_1, \dots, e_J)$ 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.¹ An *economy* is a $(e_j, \bar{u}_j)_{j=1}^J$ and an *allocation* is a $x \equiv (x_1, \dots, x_J) \in R_+^{LJ}$ such that $\sum_{j=1}^J (x_j - e_j) = 0$.

The economy begins with arbitrary ad valorem taxes $\tau = (\tau_1, \dots, \tau_L) \geq 0$ 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:

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

Definition 1 *An equilibrium with taxes τ is a (p, x) such that (1) x is an allocation, (2) $t = \sum_{j=1}^J \sum_{i=1}^L \tau_i p_i \max[0, x_{ij} - e_{ij}]$, (3) for each agent j , $x_j \in B_j(p, \tau, e_j, t)$, and (4) $x'_j \in B_j(p, \tau, e_j, t) \Rightarrow \bar{u}_j(x_j) \geq \bar{u}_j(x'_j)$.*

¹ We use the notation: $x \geq y \Leftrightarrow x_i \geq y_i$, all i ; $x > y \Leftrightarrow x \geq y$ and $x \neq y$; 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 $x_j \gg 0$, $\{z \in R_+^L : u_j(z) = u_j(x_j)\} \cap (R_+^L \setminus R_{++}^L) = \emptyset$.

Under our assumptions, an equilibrium for the model exists for any τ .² 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 $\Delta e \equiv (\Delta e_1, \dots, \Delta e_j) \in R^{LJ}$ such that $\sum_{j=1}^J \Delta e_j = 0$. We require that Δe be chosen so that an equilibrium exists, e.g., by supposing that $e + \Delta e \gg 0$. 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 allocation obtains with $(\tau, \Delta e)$. Let $f \equiv (f_1, \dots, f_j)$ indicate an equilibrium allocation that can occur with $(\tau, \Delta e)$, and 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 modifier "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 view a method for making interpersonal comparisons of utility as a set of 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 assign any importance to or even mention 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, but 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 keep the restriction here to ease comparison with the rest of the paper. For now, one may justify both the restriction and our use of welfare functions that sum 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.

Second, in any welfare function agents with identical sets of cardinal utility functions must be assigned the same utility function.

² 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.

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 $u = (u_1, \dots, u_J)$ 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 blocks any dependence of the social welfare function on agent indices, and therefore fits with the utility-independent rather than the agent-based rationale for the Pareto criterion. Anonymity plays 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. The requirement also 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. Even a planner agnostic about how to make interpersonal comparisons would not normally hold that one agent is more capable of experiencing satisfaction than another apparently identical agent based solely on the agents having different indices. 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 distraction that we can 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 are no dominating allocations. In addition, we say an allocation x is *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 $\bar{\tau}$ and $t > 0$ there must be a good i and agent j with $\bar{x}_{ij} - e_{ij} > 0$ and $\tau_i > 0$ and some good $k \neq i$ and agent $h \neq j$ with $\bar{x}_{ih} - e_{ih} < 0$ and $\bar{x}_{kh} - e_{kh} > 0$. Hence h 's marginal rate of substitution between i and k must equal $\bar{p}_i / (1 + \bar{\tau}_k) \bar{p}_k$ while j 's marginal rate of substitution between i and k must lie between $(1 + \bar{\tau}_i) \bar{p}_i / (1 + \bar{\tau}_k) \bar{p}_k$ and $(1 + \bar{\tau}_i) \bar{p}_i / \bar{p}_k$. 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 economy's characteristics will face a finite state space $\Omega = \{\omega_1, \dots, \omega_S\}$, $S \geq 2$, with associated probabilities $\pi = (\pi_1, \dots, \pi_S) \in \Delta_{++}^{S-1}$. 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) . In section 4, we could let each agent j have a distinct subjective probability $\pi(j)$. In section 5, the restriction to a single probability distribution strengthens our policy paralysis 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 s th row is $p(\omega_s)$. We also set the following notation for the remainder of the paper: $u_j = (u_j(\cdot, \omega_1), \dots, u_j(\cdot, \omega_S))$, $x_j = (x_j(\omega_1), \dots, x_j(\omega_S))$, $e_j = (e_j(\omega_1), \dots, e_j(\omega_S))$, $x(\omega_s) = (x_1(\omega_s), \dots, x_J(\omega_s))$, $e(\omega_s) = (e_1(\omega_s), \dots, e_J(\omega_s))$, $x = (x(\omega_1), \dots, x(\omega_S))$.

There may but does not have to be further uncertainty (above and beyond the policymaking uncertainty) at any ω_s . 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 an expected 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 $(f_j(\omega_1), \dots, f_j(\omega_S))$. Given an allocation x and taxes τ , the tax revenue at ω_s is given by $t(\omega_s) = \sum_{j=1}^J \sum_{i=1}^L \tau_i p_i(\omega_s) \max[0, x_{ij}(\omega_s) - e_{ij}(\omega_s)]$.

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 $p(\cdot)$ is invertible, the state could be inferred from the equilibrium price vector. But such inferences do

not affect market behavior; agents 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 define a parameter space of economies Q by letting $e \in R_{++}^{SLJ}$ be parameters, and by assuming for any agent j , any state ω_s , and any linear $h : R_+^L \rightarrow R$ that if $\bar{u}_j(\cdot, \omega_s) + h$ satisfies our assumptions on utility functions on $\{x_j(\omega_s) \in R_+^L : x_j(\omega_s) \leq \sum_{j=1}^J e_j(\omega_s)\}$ then $\bar{u}_j(\cdot, \omega_s) + h$ is a possible utility function for j at ω_s . If some goods at ω_s are contingent due to additional uncertainty at ω_s , leading $\bar{u}_j(\cdot, \omega_s)$ 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 state-contingent bundles over which j would choose if he or she faced the planner's state space. The simplest way to proceed is to endow each agent j with vNM preferences \succsim_j over lotteries (probability measures) on the prize space $R_+^L \times \Omega$ that has the consumption vector $x_j(\omega_s)$ as a typical element. For each j , \succsim_j determines a preference relation over the subset of lotteries where the probability of $x_j(\omega_s)$ is π_s (the probability held by the planner), and this subrelation will be representable by a function $Eu_j : R_+^{SL} \rightarrow R$ that is separable across states. That is, there exist functions $\{u_j(\cdot, \omega_s)\}_{s=1}^S$ such that $Eu_j(x_j) = \sum_{s=1}^S \pi_s u_j(x_j(\omega_s), \omega_s)$ for all $x_j \in R_+^{SL}$. Letting $U_j(\omega_s)$ denote the set of all increasing affine transformations of $\bar{u}_j(\cdot, \omega_s)$, we assume that $u_j(\cdot, \omega_s) \in U_j(\omega_s)$ for all ω_s . In words, j 's ex ante preferences over goods at any ω_s coincide with the preferences that j will in fact have at ω_s . Given that the ex ante stage is hypothetical, there are no further a priori restrictions to place on the \succsim_j . The selection of $u_j(\cdot, \omega_s)$ from $U_j(\omega_s)$ is the mathematical step at which the planner imposes interpersonal comparisons of welfare. We withdraw this step in the next section.

The vNM hypothesis (or a subjective expected utility variant) leads j 's ex ante utility to have the expected utility form. But in this section, all that matters is that each agent index j is assigned a single preference ordering, not its separability across states.³

We now label an allocation x to be *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 when $S = 1$. Allocation x is *strictly* ex ante superior to x' if strict inequalities hold for all j . Policies $(\tau, \Delta e, f)$

³ The present model allows agents to have diverse subjective probabilities, say in the sense of the Anscombe and Aumann (1963) theory. If j 's preferences are given by a $u_j(\cdot, \omega_s)$ at each ω_s and by a measure $\pi(j) \gg 0$, then by selecting the function $(\pi_s/\pi_s(j))u_j(\cdot, \omega_s) \in U_j(\omega_s)$ rather than $u_j(\cdot, \omega_s)$ and using the probability π , we arrive at a Eu_j of the form assumed above. This extension would not be applicable to section 5, where the Lerner example hinges on a specific and common specification of π .

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).

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. A result that there is scope for policy adjustment can therefore at best hold for a generic set of economies.

The ex ante suboptimality of an economy beginning at a status quo equilibrium (\bar{P}, \bar{x}) with taxes $\bar{\tau}$ can be attributed to two factors. First, if $\bar{\tau}$ is nonzero, $\bar{x}(\omega_s)$ 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 $\bar{u}_j(\cdot, \omega_s)$. 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 in 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 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 $(p(\omega_1), \dots, p(\omega_S))$ that support the allocation. And typically, if $L \geq S$, the price vectors $p(\omega_1), \dots, p(\omega_S)$ that rule at the S states will be linearly independent. Thus, for each agent j , the equations

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

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 vary 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, 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: unlike the policymaker, agents collectively know which state obtains and trading allows the economy to utilize this information. Moreover, since agents are

trading at each state, 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, most arbitrary status quo policies 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.

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. Let the entire model be (Ω, π) with the initial certainty model's economy assigned to ω_1 . If we are given ex ante utilities Eu_1, \dots, Eu_j for (Ω, π) , then 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 strictly ex ante superior to any status quo policy $(\bar{\tau}, \Delta e = 0, f)$ with $f(\omega_1) = \bar{x}$ if π_1 is bigger than some threshold $\hat{\pi}_1 < 1$. 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. Observe that the probability threshold $\hat{\pi}_1$ is 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 $f(\omega_1) = \bar{x}$. 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, $\bar{u}_j(\cdot, \omega_1), \dots, \bar{u}_j(\cdot, \omega_S)$, 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. Consider x_j and x'_j such that for some pair of states ω_k and ω_s , $\bar{u}_j(x_j(\omega_k), \omega_k) > \bar{u}_j(x'_j(\omega_k), \omega_k)$ and $\bar{u}_j(x_j(\omega_s), \omega_s) < \bar{u}_j(x'_j(\omega_s), \omega_s)$. If $U_j(\omega_k) \neq U_j(\omega_s)$, there is no neutral way to decide if j 's welfare is greater with

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

x_j than with x'_j . Yet the policymaker must specify a preference for j between x_j and x'_j . If say $Eu_j(x_j) > Eu_j(x'_j)$, the policymaker must hold that the gains of those ex post preferences of j 's that rank x_j ahead of x'_j outweigh the losses of the ex post preferences of j 's that rank x'_j ahead of x_j . Given that the actual agent j never faced this uncertainty, this claim amounts to an interpersonal comparison of welfare.

To see the similarity between these comparisons for a single j and traditional social welfare comparisons, consider the Harsanyi (1953) theory of social welfare, in which an agent ranks social choices while pretending there is an equal chance that he or she will be any of the agents in society. Whatever the merits of Harsanyi's proposal, it requires comparison of different types of satisfaction and will therefore reproduce the disputes that vex social decision-making: when parties disagree on who derives the more intense satisfaction from the allocation of some good, they will also disagree on how to rank social choices in the Harsanyi's set-up. The construction of ex ante preferences for an agent j is strikingly similar to the Harsanyi problem: the policymaker must rank allocations for j in ignorance of which preferences j will have; the only distinction of the ex ante problem for j is that rather than being an equal probability of being each of society's agents the probability that j has utility $\bar{u}_j(\cdot, \omega_s)$ is π_s . Indeed the ex ante problem for j and the Harsanyi problem are so similar that a policymaker pursuing the ex ante Pareto program would be well-advised to use the Harsanyi thought experiment to choose weights on the $\bar{u}_j(\cdot, \omega_s)$. Whether by the Harsanyi device or some other route, a policymaker may be able to construct ex ante preferences for j but those preferences will not be controversy-free; any dispute about interpersonal comparisons will reappear in a disagreement about which ex ante ordering should be employed.

To avoid interpersonal comparisons of utility, we consider alternatives to the ex ante ordering. The utility-independent ordering does not deem one set of weights on ex post utilities to be a better way to assemble a social welfare function than any other set of weights. On the other hand, the ordering ignores the index attached to each utility, which potentially makes the ordering more discriminating. 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 $u = (u_1, \dots, u_J)$ 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 the allocations they induce. 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. The rationales given earlier for these restrictions still apply except now uncertainty is intrinsic to the policymaker's problem and hence to apply the Harsanyi (1955) theorem on additive social welfare functions we no longer have to resort to the contrivance that further uncertainty obtains at some state. The restrictions in any event make our policy paralysis conclusions stronger: any policy that is utility-independent or maximization optimal with the restrictions remains optimal without the restrictions, but the reverse implication does not hold.

Some other social-welfare criteria are similar 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 is 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 $\pi(j)$ is used to calculate j 's expected utility, which again ensures 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 restriction: a change in allocations can harm some potential agent and still be a utility-independent improvement if some other potential agent with the same set of cardinal utility functions enjoys sufficient utility gains. As mentioned in section 2, the social welfare literature has made a compelling case for anonymity. In addition, we wish to identify the policy consequences of agnosticism about interpersonal welfare comparisons alone, without requiring that every potential agent must be left unharmed. Such a requirement would apply the agent-based approach to the entire set of potential agents, and is not entailed by a prohibition of interpersonal comparisons per se.

Finally, anonymity makes possible important exceptions that relieve the bleak landscape of policy paralysis. A policymaker's ignorance of the primitives of an economy can in some highly symmetric situations make it easier to discriminate among policies by a utility-independent ranking. If we were to hamper utility-independence by requiring that policy changes leave every potential agent unharmed, these interesting cases would be excluded.

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 Theorem 3 does not restrict the probabilities of the states in $\Omega \cup \Omega'$, the added states in Ω' can have arbitrarily small probability. Theorem 3 treats policy changes such that $\Delta e \neq 0$, which arise, for instance, when compensation for a change in τ is attempted. It is not difficult, using 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 utility functions for some agents at the added and base states were to 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 evaluate policy changes looking only at the welfare of individuals changes at a subset of states: the entire state space matters. Indeed, given any set of additional states and any 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.⁵

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 change, but only when two or more agents have identical sets of cardinal utility functions. Although in the certainty case it is plausible to dismiss as a fluke any such violation of the diversity condition, 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 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.⁶

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 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 1 We suppose that the policymaker believes that all agents are equally likely to have any given set of cardinal utilities. To make this precise, define $\pi_{U_j} = \sum_{\omega_s \in \Omega: U_j(\omega_s) = U} \pi_s$ for any set of utilities U and agent j . Our assumption is then

⁵ 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.

⁶ 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.

that the policymaker satisfies the *ignorance priors* (or symmetry) condition that for any pair of agents i and j and set of utilities U , $\pi_{U_i} = \pi_{U_j}$. Let π_U denote this common probability.

Suppose $\sum_{j=1}^J e_j(\omega_s)$ does not vary by state and consider any allocation x that is also state-invariant: for each j , there is a \bar{x}_j with $x_j(\omega_s) = \bar{x}_j$ for all ω_s . Then, for any U and $u \in U$,

$$\begin{aligned} \sum_{(\omega_s, j): U_j(\omega_s)=U} \pi_s u(x_j(\omega_s)) &= \sum_{j=1}^J \sum_{\omega_s: U_j(\omega_s)=U} \pi_s u(\bar{x}_j) = \pi_s u(\bar{x}_j) \\ &= \sum_{j=1}^J \pi_{U_j} u(\bar{x}_j) = \pi_U \sum_{j=1}^J u(\bar{x}_j). \end{aligned}$$

Now let $\psi = \frac{1}{J} \sum_{j=1}^J e_j(\omega_s)$, let Ψ be the state-invariant allocation where each j at each ω_s consumes ψ , and suppose y is another state-invariant allocation where $\bar{y}_j = y_j(\omega_s) \neq \psi$ for at least one j . So

$$\begin{aligned} \sum_{(\omega_s, j): U_j(\omega_s)=U} \pi_s u(y_j(\omega_s)) &= \pi_U \sum_{j=1}^J u(\bar{y}_j) \text{ and} \\ \sum_{(\omega_s, j): U_j(\omega_s)=U} \pi_s u(\Psi(\omega_s)) &= \pi_U \sum_{j=1}^J u(\psi). \end{aligned}$$

For U with $\pi_U > 0$ and $u \in U$, the strict concavity of u and $\frac{1}{J} \sum_{j=1}^J \bar{y}_j = \psi$ therefore imply

$$\sum_{(\omega_s, j): U_j(\omega_s)=U} \pi_s u(\psi) > \sum_{(\omega_s, j): U_j(\omega_s)=U} \pi_s u(y_j(\omega_s)).$$

If we sum this inequality over the U with $\pi_{U_j} > 0$, we conclude that for any assignment u ,

$$\sum_{j=1}^J E u_j(\psi) > \sum_{j=1}^J E u_j(y_j). \quad (5.1)$$

Hence the allocation Ψ is utility-independent superior to any state-invariant $y \neq \Psi$.

If $L = 1$, there must be a j such that $y_j(\omega_s) > \psi$ for all ω_s and hence j must be worse off with ψ than with y at every state. The allocation giving each agent ψ therefore cannot be superior to y according to any of the possible ex ante orderings. So 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 that, for each j , $e_j(\omega_s)$ also does not vary by state. 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 $(\bar{\tau}, 0, \bar{f})$ if $\bar{\tau}$ is high enough to prevent trade from occurring at all ω_s .

Since 5.1 is a strict 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 $\bar{\tau}$ in the previous paragraph does not have to be set high enough 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 our ignorance priors assumption applies to disparate preferences over many commodities.⁷

Our requirements that cardinally and ordinally identical potential agents are represented by the same utility function and that each $U_j(\omega_s)$ contains only increasing affine (rather than all monotonic) transformations of some strictly concave utility are both crucial for the non-paralysis conclusion.

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. Simpler examples would suffice to show this (e.g., let 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, nontrivial policy advice is possible in some plausible cases.⁸

The Lerner construction notwithstanding, there is a limit to what symmetry can accomplish. One might speculate that a planner with sufficiently symmetric information about how different goods enter into agent utility functions would want to lower any one of the τ_i , but this is not correct. For example, a low τ_1 and a high τ_2 can serve as an effective tool to transfer wealth to agents who consume good 1 intensively. If agents' indices convey no information that would allow lump-sum transfers to the intensive good 1 consumers, then taxes can serve as the optimal way to redistribute income. Policy paralysis can then set in: any vector of taxes can maximize social welfare for some choice of welfare weights.

In the local policy paralysis result below, we do not assume as in Theorem 3 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 for the maximization definition are stronger.

To avoid the redundancy in our definition of an endowment transfer, we now let Δe denote $(\Delta e_2, \dots, \Delta e_J)$ and, for any policy $(\tau, \Delta e, f)$, set $e_1(\omega_s) = -\sum_{j=2}^J \Delta e_j(\omega_s)$.

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 $g : \Pi \rightarrow R_+^{SLJ}$, where $\Pi \subset R_+^L \times R^{L(J-1)}$ is open and contains $(\tau, \Delta e)$, such that $g(\tau', \Delta e')$ is an equilibrium allocation for any $(\tau', \Delta e') \in \Pi$ and $g(\tau, \Delta e) = f$. Although policies are generically differentiable (see the proof of Theorem 2),

⁷ 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.

⁸ Other, more trifling policy recommendations can also be made. For example, if $\bar{\tau}$ is high enough to prevent trade at all ω_s , then any $(\tau, \Delta e = 0)$ that allows some trade at some ω_s is a utility-independent improvement.

welfare maximization need not always select one of these generic policies; a non-generic, nondifferentiable policy can be dictated. We thus incur a small loss in generality in considering only differentiable policies.

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_{j_s} where any pair of identical ex post utilities is multiplied by the same weight.⁹ Let B denote this set of weights, $\{b \in R_{++}^{SJ} : U_j(\omega_s) = U_{j'}(\omega_{s'}) \Rightarrow b_{j_s} = b_{j'_s}\}$, which has dimension equal to the number of distinct utilities in Ω . Given a differentiable policy $(\tau, \Delta e, f)$ and an assignment u , and letting g be the function specified above, we parameterize welfare functions by defining $w_u : B \times \Pi \rightarrow R$ by $w_u(b, (\tau, \Delta e)) = \sum_{j=1}^J E \hat{u}_j(g_j(\tau, \Delta e))$, 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 $\Delta E = \{\Delta e : e_j(\omega_s) + \Delta e_j(\omega_s) \geq 0 \text{ for all } j \text{ and } \omega_s\}$ and assuming for all $(\Delta e \in \Delta E, \tau)$ that an equilibrium exists at each ω_s .

Definition 6 A differentiable policy $(\tau, \Delta e \in \Delta E, f)$ 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) $D_{\tau, \Delta e}^2 w_u(1^{JS}, (\tau, \Delta e))$ is negative definite.¹⁰

Definition 7 A differentiable policy $(\tau, \Delta e, f)$ satisfies the rank condition for the assignment u if $D_{(\tau, \Delta e), b}^2 w_u(1^{JS}, (\tau, \Delta e))$ 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 $D_{(\tau, \Delta e), b}^2 w_u(1^{JS}, (\tau, \Delta e))$ when there is a sufficient diversity of utilities. But the rank condition is substantive and its meaning is important. It says that there are enough utility functions in the model that for each policy instrument there is an independent linear combination of changes in welfare weights that will alter the marginal social welfare of that instrument. Thus each policy instrument has a *distinctive* effect on social welfare; it affects the welfare of a different combination of ex post utilities. This requirement is not onerous: one would expect, for example, that a change in some τ_i will have a different impact on intensive buyers and sellers of good i compared to its impact on other potential agents. Following the policy paralysis theorem, we show that we can perturb a model in such a way that 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.

⁹ 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.

¹⁰ For any positive integer m , 1^m denotes the vector of m 1's.

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. So Theorem 4 indicates that when a policymaker aims to maximize some welfare function, 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; 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 all policies are utility-independent optimal, a paralysis result using the more stringent maximization definition of optimality faces the hurdle that the set of agent utilities reachable through some policy is not convex, in which case utility-independent policies need not be maximization optimal.

We now show that the rank condition is weak in that the inclusion of small-probability states can ensure the condition is satisfied. Note that as we add more ex post agents to the model, more columns are added to $D_{(\tau, \Delta e), b}^2 w_u(1^{JS}, (\tau, \Delta e))$ 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 u^* such that w_{u^*} satisfies the rank condition.*

6 Discussion

Our results are both positive and negative. The ex ante Pareto ordering will recommend a move from most status quo policies, but the ordering incorporates a system of interpersonal welfare comparisons. On the other hand, a thorough avoidance of interpersonal comparisons can lead a vast number policies to be optimal. Our purpose is not to judge if one approach is better than the other; they are geared to different purposes. The ex ante ordering appeals to the principle that ex ante no individual should be made worse off by a policy change. The utility-independent ordering (or maximization optimality) assumes that no system for making interpersonal welfare comparisons should be granted privileged status.

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 dictating transfers of 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

some distortions are uncorrectable, optimal policies can be counter-intuitive and depend on unobtainable information.

In the incomplete-markets model, a planner with policymaking uncertainty faces two sources of uncertainty: the uncertainty facing agents and an additional uncertainty about the parameters of the model. If the planner can make interpersonal comparisons of welfare and 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-improvements are present 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 an ex ante Pareto improvement is possible. 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 markets are complete. 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. Finally, with incomplete markets and no policymaking uncertainty, we return to the puzzle that opened this paper: a planner who knows the model could simply mandate a *first*-best optimum of its choosing. It is therefore difficult to name a scenario where market incompleteness by itself calls for a second-best Pareto improvement reached by a forced transfer of initial-period assets.

As for the theory of the second best, consider an economy with some uncorrectable distortions but where some instruments remain in the policymaker's toolkit. 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 some policy change 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 alone 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 again 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. An allocation x ex post dominates x' if, for all j and ω_s , $\bar{u}_j(x_j, \omega_s) \geq \bar{u}_j(x'_j, \omega_s)$, with strict inequality for some j and ω_s . So an allocation x is ex post optimal if there is no feasible x' that ex post dominates x , and a policy $(\tau, \Delta e, f)$ is ex post optimal if there is no $(\tau', \Delta e', f')$ such that f' ex post dominates f . Since ex post optimality does not weigh the gains of one potential preference relation against the losses of another, it does not make interpersonal comparisons of utility. As our framework now stands, a policymaker can reach an ex post optimal allocation by setting $\tau = 0$. Since any status quo policy with $\tau > 0$ will usually not reach an ex post optimal allocation, policy paralysis would seem to disappear. But this reasoning is unpersuasive. First, although a policy that sets $\tau = 0$ is indeed unusual

in reaching *allocations* that are ex post optimal, that does not mean ex post optimal *policies* are rare. Indeed, under the assumptions of Theorem 3, every policy change from the status quo harms some potential agent and hence arbitrary status quo policies are ex post optimal as well. Bear in mind that for an allocation x to be ex post optimal, it must be that, for every ω_s , $x(\omega_s)$ is Pareto undominated by another feasible allocation at ω_s , and this can be hard to achieve [these x form a $S(J - 1)$ -dimensional subset of the $SL(J - 1)$ -dimensional set of allocations]. 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 an *allocation* x achieved by some policy is of dubious relevance; the allocations that x ex post dominates will typically not be reached by any 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 solve the policy paralysis problem.

Appendix

Proof of Theorem 1 The set of ex ante optimal allocations is a manifold of dimension $J - 1$ (see, e.g., Mas-Colell 1985, Proposition 4.6.1), 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 $P(x)$. We normalize $p(x)$ and hence $P(x)$ by requiring $p(x)$ to lie in the SL dimensional unit simplex.

Since $L \geq S$, we can define for each $(x, (e, h)) \in Y \times Q$ the square matrices P_s , $s = 1, \dots, S$, by setting, for $k \leq s$, the k th 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 for arbitrary $s < S$ that if there is a generic subset $G_s \subset Y \times Q$ at which P_s with rank s , then there is a generic subset $G_{s+1} \subset Y \times Q$ at which P_{s+1} with rank $s + 1$. We define the function $g_{s+1} : G_s \rightarrow R$ 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, 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 then 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 , $b_j = (p_{\omega_1}(x) \cdot x_j(\omega_1), \dots, p_{\omega_S}(x) \cdot x_j(\omega_S))$. Since $P(x)$ has rank S , there is for any j a solution Δe_j to $(p_{\omega_1}(x) \cdot (\Delta e_j + e_j(\omega_1)) + e_j(\omega_1)), \dots, p_{\omega_S}(x) \cdot (\Delta e_j + e_j(\omega_S))) = b_j$, that is, a Δe_j such that

$$P(x)\Delta e_j = b_j - (p_{\omega_1}(x) \cdot e_j(\omega_1), \dots, p_{\omega_S}(x) \cdot e_j(\omega_S)). \quad (\text{A.1})$$

For $j = 2, \dots, J$, set Δe_j as a solution to A.1, and set $\Delta e_1 = -\sum_{j=2}^J \Delta e_j$; it is readily confirmed that Δe_1 also solves A.1 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 . \square

Proof of Theorem 2 We first briefly sketch a proof that, for any $\bar{\tau}$, 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 χ_{ω_s} from an open set $\Pi_O \subset R_+^L \times R^{L(J-1)}$ of policy instruments that contains $(\bar{\tau}, \Delta e = 0)$ to allocations such that (1) $\chi_{\omega_s}(\bar{\tau}, \Delta e = 0) = x(\omega_s)$ and (2) for any $(\tau, \Delta e) \in \Pi_O$, $\chi_{\omega_s}(\tau, \Delta e)$ is a locally unique equilibrium allocation of (e, h) in state s when the policy instruments are $(\tau, \Delta e)$.

Due to the kinks in the budget set as agents switch from buying to not buying or from selling to not selling a good, an equilibrium is locally characterized by an array of market-clearing and first-order conditions that depends on the set of goods each agent buys, sells, or neither buys nor sells. Specifically a first order condition (FOC) on some j 's marginal utility for k must hold with equality only if j buys or sells k , not when j consumes his endowment. Putting aside the unique no-trade allocation, we can restrict attention to arrays of equilibrium conditions where each good that is bought by some agent is sold by some agent.

Fixing a state and array, let \hat{p} denote the prices of the goods that are traded with one price set equal to 1, let \hat{x} and \hat{e} respectively denote profiles of consumption and endowments on the coordinate subspace in R^{LJ} where each coordinate is a good i traded by some agent j . Let F denote the C^1 function whose domain has typical element $((e, h), t, \hat{p}, \hat{x}, \lambda)$, where λ is the profile of the agents' Lagrange multipliers, that is defined by the agents' FOCs for all goods the agents buy or sell, the agents' budget constraints, the definition of t , and the market-clearing conditions of the traded goods (agents consume their endowment of untraded goods). If (p, x) is an equilibrium for the economy (e, h) with taxes $\bar{\tau}$, there is an array of equilibrium conditions such that, for the resulting F and λ , $F((e, h), t, \hat{p}, \hat{x}, \lambda) = 0$. We omit the largely routine calculation that $DF((e, h), t, \hat{p}, \hat{x}, \lambda)$ has full row rank whenever $F((e, h), t, \hat{p}, \hat{x}, \lambda) = 0$. Iterating this argument over all arrays of equilibrium conditions, the transversality theorem implies that generically equilibrium allocations are locally isolated. To exclude the nondifferentiability that occurs when simultaneously a j satisfies the FOC for buying or selling i and $x_{ij} - e_{ij} = 0$, we add to the range of each F the additional term $x_{ij} - e_{ij}$, thus defining a function F^* . We again omit the calculation showing that DF^* has full row rank, which shows that generically if $F^*(t, \hat{p}, \hat{x}, \lambda) = 0$ then $D_{t, \hat{p}, \hat{x}, \lambda} F^*(t, \hat{p}, \hat{x}, \lambda)$ has full row rank, and hence since $D_{t, \hat{p}, \hat{x}, \lambda} F^*(t, \hat{p}, \hat{x}, \lambda)$ has fewer columns than rows that there is no $(t, \hat{p}, \hat{x}, \lambda)$ such that $F^*(t, \hat{p}, \hat{x}, \lambda) = 0$. Thus at some generic set G equilibrium allocations are locally unique and the implicit function theorem gives us the function $\chi = (\dots, \chi_{\omega_s}, \dots)$.

Now consider an arbitrary selection of S of the functions F defined above, $\mathbf{F} = (F_{\omega_1}, \dots, F_{\omega_S})$, one F_{ω_s} chosen from each state. We restrict ourselves to an open subset of the domain of \mathbf{F} , say Y with typical element y , that contains all 0's of \mathbf{F} and such that each $D_{t(\omega_s), \hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s)} F_{\omega_s}$ is nonsingular and hence that χ is C^1 on some open $T \subset R^L \times R^{L(J-1)}$. Let $\mu : T \times Y \rightarrow R^J$ defined by $\mu_j(\tau, \Delta e, y) = Eu_j(\chi_j(\tau, \Delta e)) = \sum_{s=1}^S \pi_s u_j(\chi_{j, \omega_s}(\tau, \Delta e, y))$ indicate the ex ante utility levels that occur with χ as a function of the policy instruments. The proof is complete if we can show for a generic subset of economies that, at any \bar{y} with $\mathbf{F}(\bar{y}) = 0$, $D_{\tau, \Delta e} \mu(\bar{\tau}, 0, \bar{y})$ has rank J : since then the linear map $D_{\tau, \Delta e} \mu(\bar{\tau}, 0, \bar{y})$ is onto, there is a $(\tau', \Delta e')$ such that $D_{\tau, \Delta e} \mu(\bar{\tau}, 0, \bar{y})(\tau', \Delta e') \gg 0$ and hence for some $\varepsilon > 0$ some allocation 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 , consider the derivatives of μ with respect to $\varepsilon_{22}, \varepsilon_{12}, \varepsilon_{13}, \dots, \varepsilon_{1J}$. We define the functions \mathbf{F}_i , $i = 1, \dots, J$, by appending to \mathbf{F} 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 j th row consists of the derivatives of μ_j with respect to the first i of the variables $\varepsilon_{22}, \varepsilon_{12}, \varepsilon_{13}, \dots, \varepsilon_{1J}$. Thus, each M_i , $i \geq 2$ is M_{i-1} with an additional row and column added. We show that there is a generic subset of economies for which each \mathbf{F}_i has no 0 and hence that $D_{\tau, \Delta e} \mu(\bar{\tau}, 0, \bar{y})$ has rank J when $\mathbf{F}(\bar{y}) = 0$.

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 (Geanakoplos and Polemarchakis (1986)). The matrix of the direct effects of $\varepsilon_{22}, \varepsilon_{13}, \dots, \varepsilon_{1J}$ on μ , which we call *DIR*,

$$\begin{pmatrix} \varepsilon_{22} & \varepsilon_{12} & \varepsilon_{13} & \dots & \varepsilon_{1J} \\ -\sum_{s=1}^S \pi_s D_{x_{21}(\omega_s)} u_1(x_1, \omega_s) & -\sum_{s=1}^S \pi_s D_{x_{11}(\omega_s)} u_1(x_1, \omega_s) & -\sum_{s=1}^S \pi_s D_{x_{11}(\omega_s)} u_1(x_1, \omega_s) & \dots & -\sum_{s=1}^S \pi_s D_{x_{11}(\omega_s)} u_1(x_1, \omega_s) \\ \sum_{s=1}^S \pi_s D_{x_{22}(\omega_s)} u_2(x_2, \omega_s) & \sum_{s=1}^S \pi_s D_{x_{12}(\omega_s)} u_2(x_2, \omega_s) & 0 & \dots & 0 \\ 0 & 0 & \sum_{s=1}^S \pi_s D_{x_{13}(\omega_s)} u_3(x_3, \omega_s) & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \sum_{s=1}^S \pi_s D_{x_{1J}(\omega_s)} u_J(x_J, \omega_s) \end{pmatrix}.$$

The function \mathbf{F}_1 is transverse to 0 (i.e., $D\mathbf{F}_1$ has full row rank whenever $\mathbf{F}_1 = 0$) 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 *DIR* while leaving the value of \mathbf{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 $t(\omega_s), \hat{p}(\omega_s), \hat{x}(\omega_s), \lambda(\omega_s)$, $s = 1, \dots, S$, would have full row rank, which is impossible since this derivative has more rows than columns. We henceforth remove the closed 0-measure set of parameters such that $\mathbf{F}_1 = 0$ from the domain of the \mathbf{F}_i . To show that \mathbf{F}_2 is transverse to 0 requires an initial argument that

$$\frac{D_{x_{11}(\omega_1)} u_1(x_1, \omega_1)}{D_{x_{21}(\omega_1)} u_1(x_1, \omega_1)} = \frac{D_{x_{11}(\omega_2)} u_1(x_1, \omega_2)}{D_{x_{21}(\omega_2)} u_1(x_1, \omega_2)}$$

is not satisfied at any 0 of \mathbf{F} for a generic subset of economies. This is readily established with a separate transversality argument that shows that if we 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 \mathbf{F} . Given that this equality is not satisfied, we may by independently rescaling $(Du_2(\cdot, \omega_1), \lambda_{2,\omega_1})$ and $(Du_2(\cdot, \omega_2), \lambda_{2,\omega_2})$ perturb the row 2-column 2 entry of the DIR , without changing the other entries of DIR or the value of \mathbf{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 see that \mathbf{F}_2 is transverse to 0. We then proceed by induction, restricting the domain of each $\mathbf{F}_i, i = 3, \dots, J$, to exclude the points at which $\mathbf{F}_{i-1} = 0$ has a solution: simply by rescaling $(Du_i(\cdot, \omega_i), \lambda_{i,\omega_i})$, each of the remaining \mathbf{F}_i is seen to be transverse to 0, using the cofactor expansion of $\det M_i$ along row i . Thus generically $D_{\tau, \Delta e} \mu(\bar{\tau}, 0, \bar{y})$ has rank J when $\mathbf{F}(\bar{y}) = 0$, as desired. \square

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 $h \neq j$ 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 $D_{x_j(\omega^l)}(e_j(\omega^l), \omega^l), l = 1, \dots, L$, 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 for, any status quo policy $(\bar{\tau}, 0, \bar{f})$, that $\bar{f}_j(\omega^l) = e_j(\omega^l)$ 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 $\sum_{j=1}^J \Delta e_j = 0$, $p(\omega^l) \cdot \Delta e_j = 0$. Since, for all j and ω^l , there is some $\lambda_j^l > 0$ such that $p_i(\omega^l) = \lambda_j^l D_{x_j(\omega^l)} u_j(e_j(\omega^l), \omega^l)$ for all goods i such that $\Delta e_{ij} \neq 0$, $\lambda_j^l D_{x_j(\omega^l)} u_j(e_j(\omega^l), \omega^l) \cdot \Delta e_j = 0$. Condition (2) then implies that $\Delta e_j = 0$ for all j , a contradiction. \square

Proof of Theorem 4 Letting $(\hat{\tau}, \Delta \hat{e}, \hat{f})$ be differentiable and a regular maximum for u and such that $D_{(\tau, \Delta e), b}^2 w_u(1^{JS}, (\hat{\tau}, \Delta \hat{e}))$ has rank LJ , there must be a LJ dimensional coordinate subspace of B , say B^* , such that $D_{(\tau, \Delta e), b^*}^2 w_u(1^{JS}, (\hat{\tau}, \Delta \hat{e}))$ 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 $(\tau, \Delta e) \in \Pi_O$, $D_{\tau, \Delta e}^2 w_u((b^*((\tau, \Delta e))), 1^{JS-LJ})$, $(\tau, \Delta e)$ 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 $(\tau, \Delta e)_t \rightarrow (\hat{\tau}, \Delta \hat{e})$ 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 $b = (b^*((\tau, \Delta e)_t), 1^{JS-LJ})$ 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_u and the fact that $b^*((\tau, \Delta e)_t) \rightarrow 1^{LJ}$, $(\bar{\tau}, \Delta \bar{e})$ 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 $D_{(\tau, \Delta e), b}^2 w_u(1^{JS}, (\tau, \Delta e))$, the policies that satisfy the rank condition are also open, which completes the proof. \square

Proof of Theorem 5 Let $(\bar{\tau}, \Delta \bar{e}, \bar{f})$ denote the given differentiable policy for (Ω, π) . We let $\hat{\Omega}$ consist of two types of states and require for all $\omega \in \hat{\Omega}$ that each utility repeats no utility from any other state in Ω and no utility for any other agent at ω .

The first type is $\{\tilde{\omega}_1, \dots, \tilde{\omega}_L\}$. Choose the $\bar{u}_j(\cdot, \tilde{\omega}_i)$ so that (i) $D_{x_j(\tilde{\omega}_i)} \bar{u}_j(e_j(\tilde{\omega}_i) + \Delta \bar{e}_j, \tilde{\omega}_i) = D_{x_{j'}(\tilde{\omega}_i)} \bar{u}_{j'}(e_{j'}(\tilde{\omega}_i) + \Delta \bar{e}_{j'}, \tilde{\omega}_i)$ for all pairs (j, j') , (ii) the equilibrium allocation at $\tilde{\omega}_i$ is a C^1 function χ of $(\tau, \Delta e)$, (iii) each $\bar{u}_j(\cdot, \tilde{\omega}_i) \circ \chi_j$ is differentiable strictly concave, and (iv) $D_{x_1(\tilde{\omega}_1)} \bar{u}_1(e_1(\tilde{\omega}_1) + \Delta \bar{e}_1, \tilde{\omega}_1), \dots, D_{x_1(\tilde{\omega}_L)} \bar{u}_1(e_1(\tilde{\omega}_L) + \Delta \bar{e}_1, \tilde{\omega}_L)$ are linearly independent.

To assemble the second type, we first define preliminary states ω^i , $i = 1, \dots, L$. For $i = 1, \dots, L - 1$, define ω^i by letting each j have a utility $\bar{u}_j(\cdot, \omega^i)$ such that, for any $(x_{ij}, x_{Lj}) \in R_+^2$, $\bar{u}_j(\cdot, (x_{ij}, x_{Lj}), \omega^i)$ is a constant function and, for any $(x_{-ij}, x_{-Lj}) \in R_+^{L-2}$, $\bar{u}_j(\cdot, (x_{-ij}, x_{-Lj}), \omega^i)$ is differentiable strictly concave and differentiable strictly increasing. 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 $(\tau, \Delta e)$ in a neighborhood of $(\bar{\tau}, \Delta \bar{e})$, ω^i has a locally unique equilibrium allocation given by a C^1 function χ of $(\tau, \Delta e)$, (2) letting μ_{j, ω^i} denote the composition $\bar{u}_j(\cdot, \omega^i) \circ \chi_j$, then $D_{\tau_i} \mu_{1, \omega^i}(\bar{\tau}, \Delta \bar{e}) > 0$, $D_{\tau_L} \mu_{1, \omega^i}(\bar{\tau}, \Delta \bar{e}) < 0$, $D_{\tau_i} \mu_{2, \omega^i}(\bar{\tau}, \Delta \bar{e}) < 0$, and $D_{\tau_L} \mu_{2, \omega^i}(\bar{\tau}, \Delta \bar{e}) > 0$, and (3) each μ_{j, ω^i} is differentiable strictly concave. Define ω^L by letting all agents derive utility only from goods L and $L - 1$, letting conditions (1) and (3) be satisfied, and by requiring (2'): $D_{\tau_L} \mu_{1, \omega^L}(\bar{\tau}, \Delta \bar{e}) > 0$, $D_{\tau_{L-1}} \mu_{1, \omega^L}(\bar{\tau}, \Delta \bar{e}) < 0$, $D_{\tau_L} \mu_{2, \omega^L}(\bar{\tau}, \Delta \bar{e}) < 0$,

and $D_{\tau_{L-1}} \mu_{2, \omega^L}(\bar{\tau}, \Delta \bar{e}) > 0$. We now use $\omega^1, \dots, \omega^L$ to specify the second type of 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 . Finally, we set $\hat{\Omega} = \{\tilde{\omega}_1, \dots, \tilde{\omega}_L\} \cup \Omega^1 \cup \dots \cup \Omega^L$. Let $\hat{S} = \#\hat{\Omega}$.

Let $v(b, (\tau, \Delta e))$ denote $\sum_{j=1}^J \sum_{\omega_s \in \hat{\Omega}} \hat{\pi}_s b_{js} u_j(\chi_{j, \omega_s}(\tau, \Delta e), \omega_s)$, where $\hat{\pi}$ satisfies $\hat{\pi}_s = \hat{\pi}_{s'}$ when ω_s and $\omega_{s'}$ are elements of the same Ω^i . We define r_{js} for $\omega_s \in \hat{\Omega}$ and $j = 1, \dots, J$, so that, for the assignment $u = (\dots, r_{js} \bar{u}_j, \dots)$ for the utilities that appear in $\hat{\Omega}$, $v(1^{J\hat{S}}, \cdot)$ is maximized at $(\bar{\tau}, \Delta \bar{e})$. To see that this can be done, notice that any change in τ_i will neither affect the $x(\tilde{\omega}_i)$ (since, given (i), $x(\tilde{\omega}_i)$ is Pareto optimal at $\tilde{\omega}_i$) nor affect $x(\omega_s)$ for any ω_s derived from ω^k , $k \neq i$, L since agents at these ω_s neither buy nor sell i . Conditions (2) and (2') then allow r_{js} (for ω_s derived from ω^i and ω^L) to be set so that $D_{\tau_i} v(1^{J\hat{S}}, (\bar{\tau}, \Delta \bar{e})) = 0$. As for the Δe_j argument of v , our inclusion of all permutations of the agent indices and our restriction on $\hat{\pi}$ imply that for any utility function \hat{u} that appears at some ω_s in some Ω^i , $\sum_{\omega_s \in \Omega^i} \hat{\pi}_s D_{\Delta e_j} \mu_{\hat{u}(\omega_s)}(\bar{\tau}, \Delta \bar{e}) = 0$, where $\mu_{\hat{u}(\omega_s)}(\tau, \Delta e)$ gives the utility level of the agent that has \hat{u} at ω_s at policy $(\tau, \Delta e)$. Given condition (i), if we set each r_{js} for $\omega_s \in \{\tilde{\omega}_1, \dots, \tilde{\omega}_L\}$ equal to 1, then $(\bar{\tau}, \Delta \bar{e})$ must maximize $v(1^{J\hat{S}}, \cdot)$.

Let $\hat{f}(\omega_s)$ for $\omega_s \in \hat{\Omega}$ be given by the χ functions we have defined, evaluated at $(\bar{\tau}, \Delta \bar{e})$. Our differentiability assumptions imply that $(\bar{\tau}, \Delta \bar{e}, \hat{f})$ is a differentiable policy for $(\hat{\Omega}, \hat{\pi})$, while the concavity assumptions on the $\bar{u}_j(\cdot, \tilde{\omega}_i) \circ \chi_j$ and the μ_{j, ω^i} and our choices for the r_{js} imply that $(\bar{\tau}, \Delta \bar{e}, \hat{f})$ is a regular maximum for $(\dots, r_{js} \bar{u}_j, \dots)$ in the model $(\hat{\Omega}, \hat{\pi})$. It follows that $(\bar{\tau}, \Delta \bar{e}, f^*)$, where f^* is \bar{f} for $\omega_s \in \Omega$ and \hat{f} for $\omega_s \in \hat{\Omega}$, is a differentiable and a regular maximum for the assignment u^* , consisting of the u given in the Theorem for the utilities in Ω and $(\dots, r_{js} \bar{u}_j, \dots)$ for the utilities in $\hat{\Omega}$, in the model $(\Omega \cup \hat{\Omega}, (\lambda \pi, (1 - \lambda) \hat{\pi}))$ generated by any $\lambda \in [0, 1]$.

It remains to show that $w_{\hat{u}}$ satisfies the rank condition. Consider the columns of the matrix $D_{(\tau, \Delta e), b}^2 v(1^{J\hat{S}}, (\tau, \Delta e))$ that correspond, respectively, to the b 's assigned to agents 2 through J at $\tilde{\omega}_1, \dots, \tilde{\omega}_L$ and the b 's assigned to $\bar{u}_1(\cdot, \omega^i)$, $i = 1, \dots, L$. 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

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

where \tilde{P}_i is the square matrix with columns $D_{x_j(\tilde{\omega}_i)} \bar{u}_j(e_j(\tilde{\omega}_i) + \Delta \bar{e}_j, \tilde{\omega}_i)$, $j = 2, \dots, J$, and $+$'s and $-$'s indicate the signs of entries. Since the linear indepen-

dence assumption (iv) implies that each \tilde{P}_i is nonsingular, the above matrix of columns has rank LJ . Since the submatrix of $D_{(\tau, \Delta e), b}^2 w_{u^*}(1^{JS}, (\tau, \Delta e))$ 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 $(1-\lambda)$), $D_{(\tau, \Delta e), b}^2 w_{u^*}(1^{JS}, (\tau, \Delta e))$ also has rank LJ .

To ensure that the utilities in the Ω^i meet the maintained assumptions of the model, perturb the utilities given above by adding a small multiple of a differentially strictly concave and differentially strictly increasing function of the remaining $L-2$ goods. Since $D_{(\tau, \Delta e), b}^2 v(1^{J\hat{S}}, (\tau, \Delta e))$ having rank LJ is a full rank condition, its rank will persist for a small perturbation. And given that $D_{(\tau, \Delta e), b}^2 v(1^{J\hat{S}}, (\tau, \Delta e))$ 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 $D_{\Delta e_j} v(1^{J\hat{S}}, (\tau, \Delta e)) = 0$. \square

References

1. Anscombe, F., Aumann, R.: A definition of subjective probability. *Ann. Math. Stat.* **34**, 199–205 (1963)
2. Geanakoplos, J.: An introduction to general equilibrium with incomplete markets. *J. Math. Econ.* **19**, 1–38 (1990)
3. Geanakoplos, J., Polemarchakis, H.: Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete. In: Heller, W., Starr, R., Starrett, D. (eds.) *Essays in Honor of Kenneth Arrow*, vol 3. Cambridge University Press, Cambridge (1986)
4. Geanakoplos, J., Polemarchakis, H.: Observability and optimality. *J. Math. Econ.* **19**, 153–165 (1990)
5. Guesnerie, R.: *A contribution to the pure theory of taxation*. Cambridge University Press, Cambridge (1995)
6. Hammond, P.: Straightforward incentive compatibility in large economies. *Rev. Econ. Stud.* **46**, 263–282 (1979)
7. Harsanyi, J.: Cardinal utility in welfare economics and in the theory of risk-taking. *J. Polit. Econ.* **61**, 434–435 (1953)
8. Harsanyi, J.: Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *J. Polit. Econ.* **63**, 309–321 (1955)
9. Holmström, B., Myerson, R.: Efficient and durable decision rules with incomplete information. *Econometrica* **51**, 1799–1819 (1983)
10. Lerner, A.: *The economics of control*. Macmillan, London (1944)
11. Lipsey, R., Lancaster, R.: The general theory of second best. *Rev. Econ. Stud.* **24**, 11–32 (1956)
12. Magill, M., Quinzii, M.: *Theory of incomplete markets*, vol. 1. MIT Press, Cambridge (1996)
13. Mandler, M.: Interpersonal comparisons of utility and the policy paralysis problem. *Soc. Choice Welf.* **17**, 95–115 (1999)
14. Mas-Colell, A.: *The theory of general economic equilibrium: a differentiable approach*. Cambridge University Press, Cambridge (1985)
15. May, K.: A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica* **20**, 680–684 (1952)
16. McCain, R.: Distributional equality and aggregate utility: further comment. *Am. Econ. Rev.* **62**, 497–500 (1972)
17. McManus, M., Walton, G., Coffman, R.: Distributional equality and aggregate utility: further comment. *Am. Econ. Rev.* **62**, 489–496 (1972)
18. Mirrlees, J.: The theory of optimal taxation. In: Arrow, K., Intriligator, M. (eds.) *Handbook of mathematical economics*, volume 3. North Holland, Amsterdam (1986)
19. Sen, A.: *Collective choice and social welfare*. Holden Day, San Francisco (1970)
20. Sen, A.: On ignorance and equal distribution. *Am. Econ. Rev.* **63**, 1022–1024 (1973)

21. Shafer, W., Sonnenschein, H.: Equilibrium with externalities, commodity taxation, and lump sum transfers. *Int. Econ. Rev.* **17**, 601–611 (1976)
22. Suppes, P.: Some formal models of grading principles. *Synthese* **6**, 284–306 (1966)