ARTICLES

INEQUALITY AND UNCERTAINTY:
THEORY AND LEGAL APPLICATIONS

MATTHEW D. ADLER[†] & CHRIS WILLIAM SANCHIRICO[††]

"Welfarism" is the principle that social policy should be based solely on individual well-being, with no reference to "fairness" or "rights." The propriety of this approach has recently been the subject of extensive debate within legal scholarship. Rather than contributing (directly) to this debate, we identify and analyze a problem within welfarism that has received far too little attention— call this the "ex ante/ex post" problem. The problem arises from the combination of uncertainty—an inevitable feature of real policy choice—and a social

---

[†] Leon Meltzer Professor of Law, University of Pennsylvania Law School.

[††] Professor of Law, Business, and Public Policy, University of Pennsylvania Law School and Wharton School, Business and Public Policy Department. Many thanks to Jonathan Baron, Gerry Faulhaber, Alon Harel, Louis Kaplow, Yoram Marglioth, Gideon Parchomovsky, Ed Rock, Reed Shuldiner, and participants in a workshop at the American Law and Economics Association annual conference and law faculty workshops at Bar-Ilan University, Hebrew University, the University of Pennsylvania, Tel Aviv University, and the University of Virginia for their comments.

preference for equality. If the policymaker is not a utilitarian, but rather has a "social welfare function" that is equity regarding to some degree, then she faces a critical choice. Should she care about the equalization of expected well-being (the ex ante approach), or should she care about the expected equalization of actual well-being (the ex post approach)? Should she focus on the equality of prospects or the prospects for equality?

In this Article, we bring the ex ante/ex post problem to the attention of legal academics, provide novel insight into when and why the problem arises, and highlight legal applications where the problem figures prominently. We ultimately conclude that welfarism requires an ex post approach. This is a counterintuitive conclusion, because the ex post approach can conflict with ex ante Pareto superiority. Indeed, this Article demonstrates that the ex post application of every equity-regarding social welfare function—whatever its particular form—must conflict with ex ante Pareto superiority in specific situations. Among other things, then, this Article shows that legal academics who care about equity must abandon either their commitment to welfarism or their commitment to ex ante Pareto superiority.

INTRODUCTION

Imagine a social decision maker—a legislator, a regulator, a judge—who cares about inequality. This is not, let us say, her only concern—she would not be satisfied to see everyone equally wretched. But it does inform her choices and is sometimes decisive. Suppose further that she has to decide whether to encourage a certain new technology. The technology will make most individuals better off. But a few will die. Maybe the new technology is the automobile, and the few are the pedestrians who will be run down. Or maybe it is a new vaccine for a contagious disease, and the few are those who will have a fatal reaction. When she arrives at the portion of her policy analysis that concerns inequality, how should she assess the equity effects of this technology? Would it help if everyone were equally likely to be a pedestrian casualty or have an adverse reaction to the vaccine? Would her concern for inequality be assuaged, in other words, by an "equality of prospects"? Or would she reason that what matters is what actually happens, and what actually happens is in all cases disequalizing? That is, would she focus on the certainty that some will pay dearly while others benefit, and regard as irrelevant the uncertainty regarding who will end up in which position?

There are likely as many ways to approach this problem as there are philosophies of social choice. Our concern in this Article is how the problem should be addressed within the bounds of a particular approach to social choice called "welfarism." Welfarism holds that social choices should be evaluated solely according to how they affect the well-being of society's members. Importantly, the approach is broad enough to allow concern for inequality of well-being across individuals. But it does not admit criteria like "fairness," "rights," or "justice"—at least not as foundational criteria for choice.[1]

Debates about welfarism have, in recent years, consumed much energy within legal scholarship and related fields. Within legal studies, Louis Kaplow and Steven Shavell's recent book defending welfarism, *Fairness Versus Welfare*,[2] has prompted a vigorous and extensive

---

[1] A welfarist social planner might, for example, encourage *legal* institutions that involve legal "rights," but the planner's choice among the different options presented to her—for example, different institutional structures—will itself be resolved solely with reference to the distribution and amount of well-being.

[2] LOUIS KAPLOW & STEVEN SHAVELL, FAIRNESS VERSUS WELFARE (2002).

debate.[3] Similar scholarly discussions appear within normative eco-
nomics, although welfarism remains the dominant position.[4] Within
philosophy, a third arena of dispute, welfarism and related views are
surely not dominant, but they remain important contenders.[5]

We take a different tack. Rather than contributing directly to the
debate *about* welfarism, our primary object is to identify and grapple
with a large problem that arises *within* welfarism. For purposes of this
Article, we take welfarism as given—although, as we shall explain be-
low, the Article should also be of much interest for nonwelfarists, and
for those engaged in the debate between welfarism and nonwelfarism.

Call the problem illustrated by our initial hypothetical the "ex
ante/ex post" problem. Although it has received little attention, es-
pecially within legal scholarship, the ex ante/ex post problem is the
natural product of two essential specifications for any systematic ap-
proach to normative discourse about law.[6] The first is the ability to

---

[3] See, for example, the discussion of *Fairness Versus Welfare* in the *Journal of Legal
Studies*, 32 J. LEGAL STUD. 245-362 (2003), and the book reviews and responses by such
prominent figures as Jules Coleman, Richard Fallon, and Daniel Farber. Jules L.
Coleman, *The Grounds of Welfare*, 112 YALE L.J. 1511 (2003) (book review); Richard H.
Fallon, Jr., *Should We All Be Welfare Economists?* 101 MICH. L. REV. 979 (2003); Daniel A.
Farber, *What (If Anything) Can Economists Say about Equity?* 101 MICH. L. REV. 1791
(2003) (book review).

[4] *See, e.g.*, ROBIN BOADWAY & NEIL BRUCE, WELFARE ECONOMICS 143-44 (1984) (as-
suming welfarism in considering possible social welfare orderings); AMARTYA SEN,
CHOICE, WELFARE AND MEASUREMENT 18-19 (1982) (discussing the welfarist tradition
in economic theory); Walter Bossert & John A. Weymark, *Utility in Social Choice, in* 2
HANDBOOK OF UTILITY THEORY 1099 (Salvador Barbera et al. eds., 2004) (reviewing
scholarship on welfarist social choice); Marc Fleurbaey, *On the Informational Basis of So-
cial Choice*, 21 SOC. CHOICE & WELFARE 347, 374 (2003) (noting that the "most trodden
route" in social choice theory in reaction to Arrow's theorem has been to retain welfa-
rism but allow interpersonal welfare comparisons).

[5] *See, e.g.*, Richard J. Arneson, *Welfare Should Be the Currency of Justice*, 30 CAN. J.
PHIL. 497, 497 (2000) (claiming that "if we adopt an objective account of welfare and
properly accommodate concerns about individual responsibility," then the criticisms
against welfare as the "currency of justice" will lose force); Andrew Moore & Roger
Crisp, *Welfarism in Moral Theory*, 74 AUSTRALASIAN J. PHIL. 598, 613 (1996) (contending
that "[w]elfarism remains a powerful and attractive position in contemporary ethics").

[6] There is a small literature in welfare economics that addresses the choice be-
tween ex ante and ex post social welfare maximization and the closely related problem
of marrying egalitarian concerns with uncertainty. *See, e.g.*, Elchanan Ben-Porath et al.,
On the Measurement of Inequality Under Uncertainty, 75 J. ECON. THEORY 194 (1997); Larry
G. Epstein & Uzi Segal, *Quadratic Social Welfare Functions*, 100 J. POL. ECON. 691 (1992);
Thibault Gajdos & Eric Maurin, *Unequal Uncertainties and Uncertain Inequalities: An
Axiomatic Approach*, 116 J. ECON. THEORY 93 (2004); Peter Hammond, *Ex-Post Optimality
as a Dynamically Consistent Objective for Collective Choice Under Uncertainty, in* SOCIAL
CHOICE AND WELFARE 175 (Prasanta K. Pattanaik & Maurice Salles eds., 1983) [herein-
after Hammond, *Ex-post Optimality*]; Peter J. Hammond, *Ex-Ante and Ex-Post Welfare Op-*

account for the fact that policy outcomes are often uncertain—an inescapable difficulty of real-world policy analysis. The second is the ability to incorporate a social dispreference for inequality—a feature most would probably regard as important (to varying degrees) and one of the advertised attractions of welfarism.[7]

Because these specifications are so central to real-world policy discourse, the ex ante/ex post problem that they generate is ubiquitous in law and policy. Part VI of this Article reviews but a sampling of the kinds of the legal and policy issues that implicate the ex ante/ex post problem. How much should be spent on security against terrorist acts and other dangers that put many at moderate risk, but severely harm relatively few? What is the proper scale and scope for social insurance? What are we to make of the fact that residents in a particular neighborhood voluntarily agree to accept compensation in return for living near a toxic waste site with potential health risks? Is that the end of the issue? Should so-called "individual risk" tests continue to be used in environmental regulation? Should plaintiffs be allowed to recover for pain-and-suffering damages? Who should bear the burden of proof at trial? How severely should we sanction violations of the law? At what rate should we tax capital gains?

---

timality Under Uncertainty, 48 ECONOMICA 235 (1981); John M. Marshall, *Welfare Analysis Under Uncertainty*, 2 J. RISK & UNCERTAINTY 385 (1989); Ross M. Starr, *Optimal Production and Allocation Under Uncertainty*, 87 Q.J. ECON. 81 (1973). A few philosophers have also discussed these issues. *See, e.g.,* JOHN BROOME, WEIGHING GOODS: EQUALITY, UNCERTAINTY AND TIME 174-200 (1991); Wlodek Rabinowicz, *Prioritarianism and Uncertainty: On the Interpersonal Addition Theorem and the Priority View, in* EXPLORING PRACTICAL PHILOSOPHY: FROM ACTION TO VALUES 139 (Egonsson et al. eds., 2001). Our analysis is indebted to this scholarship but, we believe, also advances it. As mentioned, legal scholars have generally overlooked the ex ante/ex post problem. For two recent exceptions, see Alon Harel, Zvi Safra & Uzi Segal, *Ex-Post Egalitarianism and Legal Justice*, 21 J.L. ECON. & ORG. 57 (2005); Daniel Markovits, *Quarantines and Distributive Justice*, 33 J.L. MED. & ETHICS 323 (2005). Harel, Safra, and Segal conceptualize the ex post application of a social welfare function as we do, and show how the ex post application of an egalitarian social welfare approach can lead to recommendations for criminal law that differ from those of utilitarianism. This is an important piece. Ours advances on it in a variety of ways, particularly in distinguishing between different variants of equity regard, in clarifying how the divergence between ex post and ex ante approaches can arise for these different variants, in clarifying the conflict with ex ante Pareto superiority (as defined *infra* note 28), and in engaging the normative question of which approach is more attractive. Markovits is less formal and does not work within welfarism, but provides a valuable discussion of the difference between ex post and ex ante evaluations of the fairness of public health policies—specifically quarantines and mass vaccinations.

[7] *See, e.g.,* KAPLOW & SHAVELL, *supra* note 2, at 26-28 (stating that welfarism includes both utilitarianism and views that are concerned about the distribution of well-being).

To be more concrete about the nature and extent of the ex ante/ex post problem we must also become more abstract. Suppose then that we are in control of a society consisting of two individuals, call them JANET and JOHN. Imagine we face a choice between two alternative policies, UNCERTAIN and CERTAIN. The outcome of UNCERTAIN is, as its name indicates, uncertain. Half the time it increases JANET's well-being by 900 units with no effect on JOHN's. Half the time it increases JOHN's well-being by 900 units with no effect on JANET'S. The policy CERTAIN always increases JOHN's well-being by 400 units and JANET'S by 400 units also. Which policy should we choose?

In conformity with the limited information given in the example, assume that we are welfarists, so that we must decide solely on the basis of how the two policies affect JOHN and JANET's well-being. Let us further assume that we care not just about "efficiency"—the total amount of well-being summed over the two individuals—but also about "equity"—how evenly the total amount of well-being is distributed across the two individuals.

If both policies had but a single certain outcome, our task would be (relatively) simple. We would evaluate the policies along the two dimensions of equity and efficiency. In the event that neither policy was superior along both dimensions, our choice would depend upon the precise manner in which we traded off the two criteria.

UNCERTAIN, however, has two possible outcomes, and this raises the question of how we ought to combine them in assessing that policy's overall efficiency and equity. In fact, UNCERTAIN's effect on *total* well-being is the same across its two outcomes: total well-being always increases by 900 units. The question thereby reduces to how we should combine UNCERTAIN's effect on equity across its two outcomes. And this depends on whether we take an "ex ante" or an "ex post" approach.

Before we explain the ex ante and ex post approaches, a crucial point of terminological clarification is in order. These are *both* approaches to choice under uncertainty. They are both applicable at the very same moment in time—namely the time of choice—where the policymaker, uncertain about the outcome of each of the policies available to her, must pick one of the policies. The term ex post is potentially confusing, because it may suggest that the ex post approach is retrospective, applicable at a later point in time than the ex ante approach. But that is *not* what we mean by these terms. A better term for the ex post approach to choice under uncertainty might be the

"expected social welfare" approach; a better term for the ex ante approach might be the "expected individual well-being" approach. We use the terms ex post and ex ante instead because these have become the standard vocabulary in the literature within welfare economics about choice under uncertainty.[8] But the reader must take note that ex post and ex ante are terms of art (to be defined in a moment) and, again, that the two refer to competing methods for choice under uncertainty, both of which are applicable at the moment of choice.

What, then, do these competing approaches to policy choice under uncertainty involve? Under the ex post approach, we first judge the equity of each possible policy outcome and then combine these judgments to form an overall equity judgment for the policy. Assuming that JANET and JOHN are initially equally well off, the UNCERTAIN policy is certain to be disequalizing. In each possible outcome, one individual moves 900 units of well-being ahead of the other. It so happens that the individual moving ahead is different across the two outcomes. But no matter. Whatever the outcome of UNCERTAIN, inequality is created where there was none before.

Of course, UNCERTAIN's disequalizing effect has to be balanced against the fact that it always increases total utility by 900 units. Yet UNCERTAIN's disequalizing effect is nonetheless a strike against it. And depending on the relative importance that we place on equity, this may be enough to make us prefer its rival. Although CERTAIN increases total well-being by less—only 800 units—it is not disequalizing.

Under the ex ante approach, we (a) view the policy as offering each individual a "lottery" across possible outcomes; (b) consider the prospects for well-being that each individual's lottery provides her; and (c) judge the equity effect of a policy by considering the distribution of these individual prospects. Proceeding in this manner in evaluating UNCERTAIN we see that it offers equal prospects to JOHN and JANET: each has the same chance of benefiting by 900 units. From an ex ante perspective, therefore, UNCERTAIN is not at all disequalizing. Since that is true of CERTAIN as well, efficiency is the decisive criterion, and UNCERTAIN is thereby the clear winner.

We see, then, that determining which policy to choose depends not only on our view of the relative importance of equity and efficiency, but also on our view of whether policies with uncertain outcomes ought to be evaluated according to an ex ante approach, which

---

[8] *See supra* note 6 (citing this literature).

considers only the equity of prospects, or an ex post approach, which considers only the equity of actual outcomes. Under an ex post approach, UNCERTAIN appears disequalizing, and we may well find sufficient reason to choose CERTAIN despite its relative inefficiency. Under an ex ante approach, UNCERTAIN's disequalization within outcomes is not cognizable, and its greater efficiency makes it always the favored choice.

In fact, the difference between ex ante and ex post approaches is even more stark than this. Suppose we ask: what policy would *each individual* prefer? In judging the effect of policy UNCERTAIN on their individual well-being, JANET and JOHN would each have to combine two equally likely outcomes: that their well-being increases by 900 units and that their well-being remains the same. If JANET and JOHN simply averaged their well-being over the two outcomes (as is dictated by the standard theory of rational choice under uncertainty[9]), they would each view the welfare effects of UNCERTAIN as equivalent to a certain increase in well-being of 450 units. This being greater than 400, they would unanimously prefer UNCERTAIN to CERTAIN.

And yet, as we have seen, an ex post decision maker who places sufficient weight on equity might well prefer CERTAIN. The fact that the ex post approach does not necessarily respect the unanimous desires of the two individuals might seem to be an indictment of that approach. In fact, it just as well calls into question the value of ex ante unanimity, at least within the confines of welfarism.

One might argue against using the ex post approach in this example because it violates principles of autonomy and majoritarianism. But these values per se have no place in welfarism. A welfarist may argue on the basis of these values only to the extent that they increase individual well-being, which, in our example, they do not.

One might also argue against the ex post approach on the basis of the fact that it chooses a policy that makes both individuals worse off. This would implicitly add to welfarism the additional criterion— commonly referred to as the "Pareto principle"—that if everyone is better off under one policy alternative, that is the one we must choose.

But resort to the Pareto principle is too glib. For in the presence of uncertainty there is no single Pareto principle. An ex post analysis that chooses CERTAIN over UNCERTAIN may well violate the Pareto

---

[9] That is, expected utility maximization, as described in DAVID M. KREPS, A COURSE IN MICROECONOMIC THEORY 71-81 (1990). Note that the units in this example are units of utility, not wealth. Risk aversion is, therefore, not implicated.

principle applied ex ante to prospects. But it does not violate the Pareto principle applied ex post to actual outcomes. In the example, there are two possible outcomes, and in each possibility one individual is better off with UNCERTAIN (900 versus 400) and the other with CERTAIN (400 versus 0). It never turns out to be the case, therefore, that both individuals are better off under one of the policy alternatives. Consequently, the *ex post* Pareto principle imposes no restrictions on ex post welfare analysis in this example.

In short, a welfarist social planner who cares to some extent about equity, but nevertheless respects the ex post Pareto principle, may choose a policy (here, CERTAIN) despite the fact that every affected individual would prefer a different choice (here, UNCERTAIN).

We regard this as a dramatic result. Why? We have known since the seminal work of Amartya Sen that *non*-welfarism can lead to Pareto-inferior choices,[10] and Kaplow and Shavell's recent book, *Fairness Versus Welfare*, has advanced our understanding of this point.[11] Contemporary work in philosophy makes clear that welfarism that rejects the Pareto principle is also a conceivable view.[12] Finally, it is well recognized, and easy to see, that social choice may well deviate from unanimous individual preferences when individuals' views regarding the likelihood of various policy outcomes do not accord with those used in social decision making.[13]

What is *not* widely understood, and what we hope this Article will (among other things) help bring to light, is that even the welfarist social decision maker who respects the ex post Pareto principle, and

---

[10] Amartya Sen, *The Impossibility of a Paretian Liberal,* 78 J. POL. ECON. 152, 155-57 (1970) (discussing the conflict between liberal values and the Pareto principle).

[11] *See* KAPLOW & SHAVELL, *supra* note 2, at 52-58 (discussing how "fairness-based analysis leads to the choice of legal rules that reduce the well-being of *every* individual"); *see also* Louis Kaplow & Steven Shavell, *Any Non-welfarist Method of Policy Assessment Violates the Pareto Principle,* 109 J. POL. ECON. 281, 282 (2001) ("[F]or any non-welfarist method of policy assessment . . . there always exist circumstances in which the Pareto principle is violated.").

[12] *See* Larry Temkin, *Equality, Priority, and the Levelling Down Objection, in* THE IDEAL OF EQUALITY 126, 154-55 (Matthew Clayton & Andrew Williams eds., 2000) (concluding that equalizing welfare levels by reducing the well-being of better-off individuals constitutes a moral improvement in one respect, namely equality).

[13] *See* BROOME, *supra* note 6, at 152-54. Change the example so that in one policy outcome for UNCERTAIN both individuals enjoy an increase of 900 units in their well-being and in the other both experience no change in their well-being. If both individuals regard the first outcome, (900, 900), as very likely, they will favor UNCERTAIN over CERTAIN. If the social planner regards the first outcome as very unlikely, however, she will favor CERTAIN.

whose probability assessments are in line with each individual's, can choose a policy that each individual, maximizing her own expected utility, would reject. Ex ante Pareto superiority turns out to be a controversial principle even where, one would have thought, its foundation was firmest.

In this Article, we bring the ex ante/ex post problem to the attention of legal academics, provide insight into when and why the two approaches diverge, highlight legal applications where the distinction makes a difference, and ultimately defend—at least within the bounds of welfarism—an ex post approach, despite its implication that society should sometimes reject laws, policies, or rulings that from an ex ante perspective make everyone better off.

Part I clarifies some basic constructs: welfarism, Paretian social welfare functions, equity-regarding social welfare functions, and the ex ante and ex post approaches to social choice under uncertainty. Part II explores in depth precisely how and why the ex ante and ex post approaches can produce conflicting recommendations. Part III shows that this continues to be true even if the government has in place a highly effective system for forcing policy "winners" to compensate policy "losers," and even when individuals face no barriers to insuring privately against bad outcomes. Part IV focuses on the fact that an ex post Paretian can violate the ex ante Pareto principle. Part V argues for the ex post approach over the ex ante. And Part VI provides illustrative examples of the many and diverse legal implications of the ex ante/ex post problem.

Who should read our Article? Welfarists certainly should. Specifying welfarism under conditions of uncertainty—in particular, choosing between ex ante and ex post approaches to welfarist social choice—is a crucial problem with which law-and-economics scholars and other legal-academic welfarists have largely failed to grapple.

But nonwelfarists, too, ought to attend to our analysis. Choice under uncertainty is a *general* issue for normative theories, not unique to welfarism, and close analogues to the ex ante/ex post problem arise within various nonwelfarist views. For example, a theory concerned with the "fair" distribution of resources will need to decide whether fairness involves actual or expected resource holdings, and a theory that recognizes certain moral rights might specify those in ex post terms (a right not to be killed) or in ex ante terms (a right not to be put at high risk of death).

Finally, the debate between welfarists and nonwelfarists should itself be sensitive to the problem of choice under uncertainty. The two

approaches should be tested against each other, not in the artificial laboratory where decision makers know for certain what the outcomes of their choices will be, but under conditions faced by real, non-omniscient humans. By specifying and defending one approach to welfarist choice under uncertainty, we also contribute—albeit indirectly—to the broader debate *about* welfarism.

More specifically, on this last score, our Article shows that Kaplow and Shavell's arguments for welfarism in *Fairness Versus Welfare* are, to some extent, overstated. *Fairness Versus Welfare* takes pains to emphasize that it is possible to be welfarist, respect the unanimous wishes of the citizenry (i.e., respect Pareto superiority), *and* attend to the distribution of well-being.[14] But in the realistic context in which policies have uncertain outcomes, it is not that simple. Equity-regarding welfarists face a dilemma: either (1) use the ex ante approach to social welfare, which as we'll argue below is unsatisfactory; or (2) use the more sensible ex post approach and be prepared to ignore the unanimous objection of the citizenry. In other words, welfarists who wish, justifiably we believe, to take an ex post approach to social welfare cannot *both* be non-utilitarian (equity regarding) *and* satisfy the Pareto principle in its ex ante version.

Kaplow and Shavell correctly characterize welfarism as a moral view that is potentially sensitive to distribution.[15] But they do not explain the crucial point that this attractive feature of welfarism is purchased at the cost of either relinquishing full-blown Paretianism or adopting the questionable practice of ex ante social welfare analysis. Of course, that cost may be worth bearing.[16] Still, it is important to be clear about the features of equity-regarding welfarism, and its potential inconsistency with a Paretianism that insists not only on the Pareto principle for outcomes, but on ex ante Pareto superiority as well. Notwithstanding the argument in *Fairness Versus Welfare*,[17] the tension with the Pareto principle is not *unique* to fairness approaches; it is also characteristic of (what we take to be) the best versions of welfarism,

---

[14] *See, e.g.,* KAPLOW & SHAVELL, *supra* note 2, at 28-38.

[15] *Id.*

[16] Indeed, we argue below that the ex ante Pareto principle is not compelling, at least not within the bounds of welfarism. *See infra* Part V.

[17] *See* KAPLOW & SHAVELL, *supra* note 2, at 52-58 (stressing conflict between the Pareto principle and fairness theories).

namely welfarist theories that are equity regarding and ex post. One of the contributions of our Article is to clarify these points.[18]

I. BASIC CONCEPTS: SOCIAL WELFARE FUNCTIONS, EQUITY REGARD, AND EX ANTE VERSUS EX POST SOCIAL WELFARE ANALYSIS

In this Part we introduce some of the basic concepts that will be applied throughout the Article. For expositional purposes, we begin by assuming that policy alternatives have certain outcomes. In this artificial environment, we first describe the basics of welfarist social choice. We then discuss the rudiments of including a social dispreference for inequality in welfarist analysis. We conclude this Part by introducing uncertainty to the analysis, showing how this alters the welfarist's object of choice, and distinguishing between the ex post and ex ante approaches to this new choice problem.

A. *Social Welfare Functions Under Certainty*

What is welfarism? First, welfarism is, of course, a *normative* view— one that tells decision makers how they "ought" to act, and that evaluates choices as "better" or "worse." Second, welfarism adopts an approach to evaluating choices that is at least weakly *consequentialist* in the following sense. Whether the decision maker should choose one or another course of action depends solely on the outcomes[19] of those

---

[18] Kaplow is clear about the tension between equity-regarding welfarism and Paretianism in other work. In a 1995 article, he identifies the conflict between ex post egalitarianism and ex ante Pareto superiority and argues for utilitarianism. Louis Kaplow, *A Fundamental Objection to Tax Equity Norms: A Call for Utilitarianism*, 48 NAT'L TAX J. 497, 499-501 & n.11, 508 (1995). One may question whether this position is fully consistent with the eclecticism of *Fairness Versus Welfare*, which embraces *both* utilitarianism *and* distributionally sensitive variants of welfarism. KAPLOW & SHAVELL, *supra* note 2, at 27-28. One might well read Kaplow's earlier article to argue that the only *acceptable* version of welfarism—the only version consistent with both the sure-thing principle and ex ante Paretianism—is utilitarianism. *See also* Kaplow, *supra*, at 503-04 (accepting Harsanyi's axioms for rational social choice).

[19] What exactly is an "outcome"? Loosely speaking, an outcome is some state of affairs. In this sense, for example, "Mount Vesuvius erupts in the twenty-first century" is an outcome. But "outcome" in the loose sense is not a perspicuous way to capture the core of welfarism and the disagreements within welfarism, particularly the disagreement between ex ante and ex post approaches to social choice, which is our ultimate concern. The problem is that loose outcomes can embed uncertainty. For example, the fact that "Mount Vesuvius erupts in the twenty-first century" leaves open whether it erupts in 2005 or 2010, whether the eruption involves a large or small lava flow, hot lava or cooler lava, and so on. And it is precisely in how to cope with uncertainty that ex post and ex ante welfarists disagree. Therefore, we use "outcome" in a

choices. Finally, welfarism insists that *individual well-being* is the sole feature of outcomes that differentiates them as a normative matter.[20] Thus, if the (certain) outcomes of two policy choices are identical with respect to well-being—in the sense that each individual is equally well off in both outcomes—then the decision maker must be indifferent between the choices.

Welfare economists typically make certain additional assumptions about the structure of welfarism above and beyond those just described.[21] Because these standard assumptions facilitate our analysis and because we wish to demonstrate that the complex issues described in this Article arise within mainstream social welfare analysis, we will adopt all of these additional assumptions.

First, it is assumed that an individual's well-being in various outcomes can be represented by "utility" numbers. This means that the policy outcomes can be represented as *vectors* of utility numbers—with one coordinate for each individual in the population.

Second, welfare economists assume that society has a complete ordering over utility vectors. Each vector is better than, worse than, or exactly equal to every other vector, and the ordering is transitive.

Third, welfarists typically assume that this ordering is "continuous." The definition of continuity need not concern us here.[22] What is important for our purposes is an implication of continuity—that the social *ordering* of outcomes can be represented by a *social welfare function*.[23] The social welfare function maps each utility vector onto a number, which represents the position of that outcome in the ordering.

---

stricter sense—to mean what philosophers call a "possible world," a complete description of a possible history of the universe. *But see infra* note 20.

[20] In regard to our definition of "outcome," *supra* note 19, note that some differences between possible worlds make no difference to anyone's welfare. This calls for refinement of our notion of "outcome." An outcome, as we mean it, is a possible state of affairs that is maximally specified *with respect to human well-being.* Although the description of the outcome may not be absolutely specific—it may still embed some uncertainty—all further specifications of it are equally good with respect to everyone's welfare.

[21] *See generally* BOADWAY & BRUCE, *supra* note 4, at 137-69 (reviewing welfarist social choice theory).

[22] For a discussion of the continuity property for orderings and its relationship to functional representation, see KREPS, *supra* note 9, at 31, 36-37.

[23] *See* BOADWAY & BRUCE, *supra* note 4, at 137, 139.

Fourth, many contemporary welfare economists are willing to assume that utilities are interpersonally comparable to some degree.[24] This is a crucial assumption—Arrow's famous theorem demonstrates that the absence of comparability, coupled with a few other assumptions, implies that no social ordering exists—and one that we believe to be independently justified.[25] One of us has presented a lengthy case for interpersonal comparability elsewhere.[26]

Fifth, it is commonly assumed that the social welfare ordering respects the Pareto principle.[27] If one person is affirmatively better off in one outcome than in another, and no one is worse off, then the first outcome must be ranked higher in the ordering. Translated into the language of social welfare functions, the Pareto principle requires that the social welfare function be monotonically increasing in each of its arguments. We will use the term "Paretian social welfare function" to mean a social welfare function that respects the Pareto principle.

Whether the application of that function under conditions of uncertainty respects the ex ante Pareto principle is, of course, a separate question. By calling the social function itself "Paretian," we simply mean that the ranking of outcomes—utility vectors—generated by the function conforms to the Pareto principle.[28]

---

[24] *See id.* at 152-66. The bulk of optimal tax theory, for instance, is founded on this assumption. *See, e.g.*, Nicholas Stern, *The Theory of Optimal Commodity and Income Taxation: An Introduction, in* THE THEORY OF TAXATION FOR DEVELOPING COUNTRIES 22, 46-49 (David Newberry & Nicholas Stern eds., 1987) (discussing the use of the social welfare function in optimal tax analysis).

[25] AMARTYA K. SEN, COLLECTIVE CHOICE AND SOCIAL WELFARE 89-104 (1970) (describing the role of interpersonal comparability in Arrow's impossibility theorem and arguing for partial interpersonal comparability).

[26] *See* MATTHEW D. ADLER & ERIC A. POSNER, NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS 39-52 (2006); Matthew D. Adler & Eric A. Posner, *Rethinking Cost-Benefit Analysis*, 109 YALE L.J. 165, 204-09 (1999).

[27] *See, e.g.*, Bertil Tungodden, *The Value of Equality*, 19 ECON. & PHIL. 1, 8 (2003) ("The Pareto Principle is the core of normative economics . . . .").

[28] Another point of terminology: the Pareto principle (in either its ex post or ex ante variants) is sometimes framed in terms of preferences and sometimes in terms of well-being. For example, the ex post Pareto principle in the preference form says that an outcome that at least one person prefers and no one disprefers is a better outcome; the ex post Pareto principle in the well-being form says that an outcome which improves the well-being of at least one person, and makes no one worse off, is a better outcome. The ex ante Pareto principle in the preference form says that a set of individual lotteries that at least one person prefers, and no one disprefers, is a better set of lotteries; the ex ante Pareto principle in the well-being form says that a set of individual lotteries which is better for at least one person, and worse for no one, is a better set of lotteries. The preference-based formulations of the Pareto principle presuppose a preferentialist conception of well-being. Since we do not presuppose that conception,

Although not technically implied by welfarism—which says only that well-being is all that matters, but not how it matters—Paretianism is so automatically adopted by most welfarists that it is virtually a core commitment of the approach.[29] This accords with intuition. If the ultimate normative concern is human well-being, how can an outcome that promotes the well-being of some individuals, and harms no one, be worse? Thus, if an outcome increases the well-being of some individuals, and harms no one, we will say that it is an improvement.[30]

A sixth very plausible assumption is that the social welfare function is "symmetric" or "anonymous," meaning that the identity of individuals is irrelevant to the ranking of outcomes. If two utility vectors have the same utility numbers, but in different orders, they must be ranked the same.

Thus far, our presentation has adopted the standard premises of social choice theory. But there is one standard assumption we reject—namely, that well-being consists in the satisfaction of preferences. The nature of well-being is highly controversial, as the philosophical literature shows. Hedonic views and "objective-list" views compete with preferentialism, and for our purposes, there is no reason to choose among these views.[31]

---

see infra text accompanying note 31, we use the term "Pareto principle" throughout the Article, to mean the principle framed in terms of well-being. In other words, the "Pareto principle," for our purposes, is equivalent to the second prong of what John Broome calls the "principle of personal good." *See* BROOME, *supra* note 6, at 165 (defining the principle of personal good to mean: "(a) Two alternatives are equally good if they are equally good for each person. And (b) if one alternative is at least as good as another for everyone and definitely better for someone, it is better"). Since the first prong of Broome's principle is less central to our discussion, we use the term "Pareto principle" to mean the second prong, rather than both prongs taken together.

[29] It is conceivable that one could be so concerned about equity that one would reject an outcome that makes the most well off even more well off without affecting anyone else. *Cf.* Temkin, *supra* note 12, at 154-55 (emphasizing that a Pareto-inferior move may increase equality). We do not explore this line of argument in this Article. Whatever its plausibility, the aim of this Article is to show why, even if the argument is wrong and egalitarianism is best specified by coupling it with the Pareto principle for outcomes, a conflict with ex ante Pareto superiority can arise.

[30] Superficially, it seems possible that the unaffected individuals might envy those made better off, or might suffer a loss in status—but any "envy" or "loss of status" will tend to reduce the welfare of those who suffer it, in which case the ex post Pareto principle will no longer apply.

[31] The simple preference-satisfaction view says that an individual has greater welfare in outcome O as compared to outcome $O*$ if and only if he prefers O to $O*$. Refined variants of this view require the individual's preferences to be well-informed or otherwise idealized. Hedonic views say that an individual's well-being in O as compared to $O*$ depends on the pains and pleasures he experiences in the two outcomes.

Welfarism, utility numbers, the concept of a social welfare function, the Pareto principle, equity regard, and the choice between ex ante and ex post approaches to social choice are *generic* constructs or problems that apply regardless of the specific nature of well-being. For example, the welfarist who adopts a hedonic view will see an individual's utility number as representing her hedonic state; will see each outcome as a vector of hedonic utilities; will conceptualize the social welfare function as mapping each vector of these hedonic utilities onto a number representing the vector's place in the social ordering; and can insist that this function be both monotonically increasing in each of its arguments (Paretian) and symmetric. Similarly, if certain basic measurement axioms are satisfied, the objective-list theorist will have at her disposal utility numbers representing individual well-being in the objective-list sense; can then represent each outcome as a vector of these sorts of utilities; and so on.

To be clear: we do not wish to reject the preference-based view of welfare, but simply to remain agnostic about whether well-being consists in preference satisfaction, in the occurrence of pleasurable states and the avoidance of pain (hedonism), in the attainment of objective goods, or in something else. These issues are just orthogonal to the problem of ex post versus ex ante approaches to social planning.

To sum up: the normative framework for our analysis assumes that the appropriate choice for a social planner in some choice situation depends on the outcomes of each choice; that such outcomes can be represented by utility vectors, with utility representing individual well-being, whatever that may consist of (not necessarily preference satisfaction);[32] that the appropriate ordering over utility vectors can be

---

Objective-list views specify a list of welfare goods (such as friendship, accomplishment, the pursuit of knowledge, and so on) that are objective in the sense that an individual's life in one outcome can realize more of these goods without the individual preferring the outcome. For good discussions of these competing views and overviews of the philosophical literature on well-being, see JAMES GRIFFIN, WELL-BEING 7-72 (1986); DEREK PARFIT, REASONS AND PERSONS 493-502 (1984); T.M. SCANLON, WHAT WE OWE TO EACH OTHER 108-43 (1998); Mozaffar Qizilbash, *The Concept of Well-Being*, 14 ECON. & PHIL. 51 (1998); *see also* ADLER & POSNER, NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS, *supra* note 26, at 28-39 (discussing different views of well-being, and ultimately defending an account that looks to both actual and idealized preferences).

[32] One important technical question here involves admissible transformations of the utility vector. In effect, how much information about well-being do the utility numbers contain? Can they be transformed by a common affine transformation, a common multiplicative transformation, or perhaps by no transformation at all? Interpersonal comparability of welfare levels and differences, without more, would permit common affine transformations of the utility numbers. So would the Pigou-Dalton re-

represented by a social welfare function defined on these vectors; that this function respects the Pareto principle and is symmetric; and, finally, that it is equity regarding, a construct we will now explore.

Before we proceed, however, we should address one concern that the reader may have. Who exactly is this "social planner" who is supposed to be guided by the social welfare function? The answer, in brief, is that she is an idealized governmental official—one who is not constrained by the computational limitations of "bounded rationality."[33] A normative framework that identifies appropriate choices by deriving them from an ordering over all possible outcomes (specified in sufficient detail to capture everything that matters to well-being) cannot be implemented by real officials, since real officials do not have the mental machinery to grasp such an ordering. Still, the generic concepts of a social welfare ordering and social welfare functions—like many idealizations in normative theory—are useful in figuring out what real-world actors should do. The specifics of real-world application are examined in Part VI.

B. *Equity Regard Under Certainty*

One of the attractive, yet underappreciated, features of welfarism is that it is broader than utilitarianism. The utilitarian social welfare

---

quirement itself. (If one vector is Pigou-Dalton equalizing relative to another, then a common affine transformation of all the utilities in the two vectors preserves this relationship.) *See infra* text accompanying note 39. But particular social welfare functions may preclude common affine transformations. For example, the ordering of utility vectors generated by the sum of the square roots of individual utilities is not invariant to the addition of a constant. On these issues, see, for example, BOADWAY & BRUCE, *supra* note 4, at 144-47 (discussing invariance requirements in social choice theory, which specify "the set of transformations that can be applied to [a] . . . household utility vector without changing the [social welfare ordering]"); Bossert & Weymark, *supra* note 4, at 19-29 (discussing invariance requirements).

For purposes of this Article, we will assume that the utility numbers are absolute—that no transformations are permissible. We conjecture, but will not attempt to demonstrate, that the ex ante/ex post divergence results demonstrated in this Article would obtain even if some transformations of the utility vectors are permissible—either common affine transformations or the class of transformations permitted by the social welfare function itself, whichever is more restrictive. *Cf.* BROOME, WEIGHING GOODS, *supra* note 6, at 217 (suggesting that prioritarian and utilitarian approaches to social choice may collapse into each other, given limitations on the measurability of well-being); Rabinowicz, *supra* note 6, at 143-44 (discussing Broome's view).

[33] *See* 3 HERBERT SIMON, MODELS OF BOUNDED RATIONALITY 291 (1997) ("The term 'bounded rationality' is used to designate rational choice that takes into account the cognitive limitations of the decision maker—limitations of both knowledge and computational capacity.").

function simply adds individual utilities, ranking utility vectors in accordance with their total utility.[34] It is insensitive to distribution, in the sense that the change in social welfare produced by a given change in any individual's utility does not depend either on that individual's utility level or on the utility levels of others in society. Social welfare changes by the same amount regardless of how well off or badly off that individual is, either in absolute terms or compared to others.[35] The utilitarian social welfare function falls within the general category of Paretian symmetric social welfare functions.[36] But there are Paretian social welfare functions that, unlike utilitarianism, are equity regarding; and there are also non-utilitarian, non-equity-regarding social welfare functions. These logical relationships are illustrated in Figure 1.

Figure 1: Paretian Social Welfare Functions

Non-Equity Regarding, Non-Utilitarian	Equity Regarding
Utilitarian	

Utilitarianism's insensitivity to distribution is not only well known, but the source of much criticism—criticism that we find persuasive and need not recapitulate here.[37] Relatedly, we find it very hard to see why the appropriate social welfare function would be some other,

---

[34] This characterization of utilitarianism is accurate for the "fixed population" case (where the same individuals exist in every outcome), but not necessarily for the "variable population" case, which we do not discuss in this Article. On "variable population" problems in social choice theory, see generally CHARLES BLACKORBY ET AL., POPULATION ISSUES IN SOCIAL CHOICE THEORY, WELFARE ECONOMICS, AND ETHICS 129-208 (2005).

[35] The utilitarian may not be indifferent as to the distribution of resources, since the function from resources to utility (the measure of well-being) may not be linear. For example, it is widely assumed that money has diminishing marginal utility. Utilitarianism, however, *is* indifferent to the distribution of well-being itself.

[36] Henceforth we drop the term "symmetric," which is assumed.

[37] *See, e.g.*, Amartya Sen & Bernard Williams, *Introduction, in* UTILITARIANISM AND BEYOND 1, 4-5 (Amartya Sen & Bernard Williams eds., 1982) (noting that utilitarianism involves a "drastic obliteration of usable information," including information about the distribution of utility).

non-utilitarian variant of a non-equity-regarding social welfare function. The strongest criticisms of utilitarianism—within welfarism—involve its failure to attend to equity, and such arguments also apply to other non-equity-regarding social welfare functions. Henceforth, and without belaboring the arguments in favor of distribution as a normative consideration independent of maximizing overall welfare, our analysis will assume that the social welfare function is equity regarding in some degree.

But what exactly makes a Paretian social welfare function equity regarding? A substantial literature in economics addresses distributive justice[38] and provides a canonical answer which, we think, is quite powerful. An equity-regarding social welfare function satisfies the so-called "Pigou-Dalton" principle, also known as the "weak principle of transfers."

Pigou-Dalton/Principle of Transfers

> Given any outcome, if a fixed amount of utility is transferred from an individual with an initially higher utility to one with an initially lower utility, decreasing the distance between their utilities, then the social value of the outcome must increase.[39]

The Pigou-Dalton principle requires, in effect, that a zero sum equalizing transfer between any two individuals increase social welfare. Suppose, for example, that there are three individuals in our society, labeled 1, 2, and 3, with initial utilities of 100, 200, and 300, respectively. If we were to take 50 units of utility from person 3 and transfer them to person 1, the utility vector would become 150, 200,

---

[38] For overviews of the economic literature on equality, see FRANK A. COWELL, MEASURING INEQUALITY (2d ed. 1995); AMARTYA SEN, ON ECONOMIC INEQUALITY (expanded ed. 1997); Charles Blackorby et al., *Income Inequality Measurement: The Normative Approach, in* HANDBOOK ON INCOME INEQUALITY MEASUREMENT 133 (Jacques Silber ed., 1999); F.A. Cowell, *Measurement of Inequality, in* 1 HANDBOOK OF INCOME DISTRIBUTION 87 (Handbooks in Econ. Series No. 16, Anthony Atkinson & François Bourguignon eds., 2000); Bhaskar Dutta, *Inequality, Poverty and Welfare, in* 1 HANDBOOK OF SOCIAL CHOICE AND WELFARE 597 (Handbooks in Econ. Series No. 19, Kenneth J. Arrow et al. eds., 2002).

[39] *See, e.g.*, Dutta, *supra* note 38, at 607. As discussed in Appendix Section B, another technically weaker restriction found in the literature, also referred to as the Pigou-Dalton principle, tests only those utility vectors across which individuals do not change rank. Given symmetry, this alternative definition is equivalent to the one we provide here. We choose our definition (even though it is stronger in the absence of symmetry) because we find it more intuitive.

250. If the social welfare function satisfies the Pigou-Dalton condition, then it must assign a higher value to the second utility vector than to the first. If instead we transferred 150 units of utility from person 3 to person 1, the utility vector would become 250, 200, 150, and person 3 and person 1 would change rank. The transfer would overshoot the initial dispersion between persons 1 and 3. The distance between their utilities, however, would still decrease (from 200 to 100) and any Pigou-Dalton respecting social welfare function would have to increase under the transfer.[40]

---

[40] Several technical notes regarding the Pigou-Dalton condition are in order. First, the Pigou-Dalton principle gains additional power in combination with symmetry. If we transferred 250 units of utility from person 3 to person 1, in the example in the text, the vector would become 350, 200, 50. Persons 3 and 1 would change rank as above, but, unlike above, the distance between their utilities would increase and the Pigou-Dalton principle *by itself* would impose no requirements on an equity-regarding social welfare function in this case. Combined with symmetry, however, the Pigou-Dalton condition would require that the social welfare function prefer *the original outcome.* This is because the second outcome would, by symmetry, be regarded as equivalent to 50, 200, 350. And the original vector 100, 200, 300 could be generated from 50, 200, 350 by a zero sum transfer, bringing the utilities of persons 3 and 1 closer together.

Second, the Pigou-Dalton principle also gains some power in combination with transitivity. Compare the initial utility vector 100, 200, 300 with the utility vectors 225, 225, 150. At first glance, the Pigou-Dalton principle seems to impose no requirement regarding the ranking of these two vectors. Because all three utilities differ, the two vectors are not related by a pairwise zero sum transfer (let alone one that is equalizing). However, the utility vectors are related by a *sequence* of pairwise zero sum transfers that are equalizing. The vector 100, 200, 300 becomes 250, 200, 150 with a transfer of 150 units of utility from person 3 to person 1, and then 225, 225, 150, with a transfer of 25 units of utility from person 1 to person 2. The welfare function must prefer the second vector to the first and the third to the second, and so, by transitivity, must prefer the third to the first.

Third, the Pigou-Dalton principle is directly related to another common way of stating that one utility vector displays more inequality than another, namely "Lorenz Dominance." The "Lorenz curve," applied to a utility vector, specifies, for all x between 0 and 100, the percentage $l(x)$ of total utility possessed by the least well off x percent of the population. The endpoints of the Lorenz curve are the same for all utility vectors: 0% and 100%. In between, a greater Lorenz curve signifies that each least well off x percent possesses a greater portion of total utility. Hence, a greater Lorenz curve represents greater equality. Accordingly, utility vector A is said to "Lorenz dominate" utility vector B if A's Lorenz curve is somewhere above B's and nowhere below. One can show that this is equivalent to saying that A can be generated from B by a sequence of equalizing transfers, a la Pigou-Dalton. *See, e.g.*, Dutta, *supra* note 38, at 607-09 (discussing the connection between the Pigou-Dalton principle and Lorenz dominance).

Finally, it should be noted that the Pigou-Dalton condition (and the Lorenz dominance condition) is not a complete description of inequality. A complete measure would, for *every* pair of utility vectors, specify either that one vector is more unequal than the other, or that the two vectors are the same in terms of inequality. In contrast,

For the remainder of this Article, we shall accept the view that defines equity regard as Pigou-Dalton respecting and focus on Paretian social welfare functions that are equity regarding in this sense. It should be stressed that the Pigou-Dalton condition is not particularly demanding. There are, in fact, a large variety of Paretian equity-regarding social welfare functions. The literature discusses the forms that an equity-regarding social welfare function might take.[41] Most of our analysis of the ex ante/ex post divergence is not tied to a particular social welfare function, and so we will not discuss the possibilities in detail here. Our interest here is not in advocating a particular conception of distributive justice, but in clarifying that—whatever the conception—maximizing the expectation of equity and maximizing the equity of individuals' expected utilities can conflict, even with compensation and insurance in play.

One important distinction within the family of equity-regarding social welfare functions that should be mentioned is that between functions that are *separable* in individual utilities, and those that are not.[42] Separability means that the ranking of two outcomes cannot depend on the utility of indifferent individuals: those whose utilities are the same in both outcomes. Separability corresponds to a non-comparativist conception of equity, which says that changes to an individual's well-being have greater weight the worse off that individual is in *absolute* terms. The term used in the philosophical literature for this conception is "prioritarianism."[43] By contrast, "comparativists" *are*

---

the Pigou-Dalton condition makes no statement whatsoever about the relative inequality of two utility vectors that are not related by a sequence of mean-preserving equalizing transfers. Rather, the Pigou-Dalton condition is a minimum requirement for any complete measure (and any equity-regarding welfare function). As we will discuss, one complete measure of inequality that respects the Pigou-Dalton condition is the Gini coefficient. *Infra* text accompanying notes 45-46.

[41] *See, e.g.*, AMARTYA SEN, ON ECONOMIC INEQUALITY (1973); sources cited *supra* note 4.

[42] *See* Marc Fleurbaey, Equality Versus Priority: How Relevant is the Distinction? (Sept. 2001) (unpublished manuscript), *available at* http://cerses.shs.univ-paris5.fr/marc-fleurbaey.htm (follow "Equality vs priority: How relevant is the distinction?" hyperlink) (discussing the separability of individuals in social rankings).

[43] Much of the recent philosophical literature on equality focuses on the choice between "prioritarian" and nonprioritarian conceptions of equality. *See generally* LARRY S. TEMKIN, INEQUALITY 245-82 (1993) (discussing prioritarianism, using the term "extended humanitarianism"); THE IDEAL OF EQUALITY (Matthew Clayton & Andrew Williams eds., 2000); Tungodden, *supra* note 27. "Prioritarianism" is a bit of a fuzzy term. It is not clear whether prioritarianism means (1) any equity-regarding welfarism that generates a social ordering of outcomes that is separable in individual utilities; or, more narrowly, (2) any such approach which, in generating a ranking of *policies* (un-

essentially concerned with the pattern of well-being. They make holistic judgments: the fact that individual 1 moves up a given increment and individual 2 down a given increment might be a good thing given one background allocation of utility to other individuals, but not given another.

The choice between prioritarian/noncomparativist and comparativist views of equality has been much discussed in the recent philosophical literature.[44] We will not enter this debate, because we need not. *Both* views are represented in the family of equity-regarding Paretian social welfare functions. Prioritarians/noncomparativists accept the Pigou-Dalton condition *and* the additional requirement of separability. Comparativists accept the Pigou-Dalton condition, but reject the separability requirement.

Prioritarianism/noncomparativism is most naturally expressed via one of the standard functional forms employed by economists theorizing about equality, namely the additive social welfare function, $\varphi(u_1)+...+\varphi(u_I)$ where φ is strictly concave. (The square root social welfare function, $\varphi(u)=\sqrt{u}$ is an example.)

On the other hand, a social welfare function that first weights individual utilities according to the rank of the individual's utility in the population, with greater weight given to individuals with lower utility, and then adds the weighted utilities is an example of a comparativist social welfare function. The social welfare gain from increasing any given individual's utility depends on her own utility and the utility levels of others via that individual's rank in the population. Rank-weighted social welfare functions gain importance by association with the Gini coefficient, a fairly natural way to measure inequality and one of the most popular ways of doing so. Applied to utility vectors, the Gini coefficient is proportional to the average distance between all possible pairs of utilities divided by mean utility.[45] One can show that

---

certain sets of outcomes) consistent with the sure-thing principle, can be represented by an additive social welfare function applied in an ex post manner. *See* Fleurbaey, *supra* note 42, at 6-7, 17-18 (stressing that a ranking of outcomes that satisfies the Pigou-Dalton principle and separability can, but need not, be represented by an additive function of utilities). At a minimum, it seems correct to say that approaches to social choice that fall in the latter, narrower category are "prioritarian," and that is how we use the term in this Article.

[44] *See* sources cited *supra* note 43.

[45] The Gini coefficient is typically written as $G=\frac{1}{2}\frac{1}{\mu}\frac{1}{I^2}\sum_{i=1}^{I}\sum_{j=1}^{I}|u_i-u_j|$, where $\mu=\frac{1}{I}\sum_{i=1}^{I}u_i$ is mean utility. However, we find the following equivalent formulation easier to interpret: $G=\frac{1}{2}\frac{I-1}{I}\times\frac{1}{\mu}\times\frac{1}{I(I-1)/2}\sum_{i=1}^{I-1}\sum_{j=i+1}^{I}|u_i-u_j|$. The rightmost factor $\frac{1}{I(I-1)/2}\sum_{i=1}^{I-1}\sum_{j=i+1}^{I}|u_i-u_j|$ is the average distance between individuals' utilities. More spe-

the Gini coefficient respects the Pigou-Dalton condition: Pigou-
Dalton equalization always reduces the Gini coefficient. One can also
show that the simplest equity-regarding rank-weighted social welfare
function—one that weights the worst-off individual in a population of
I individuals by I, the second-worst-off by I-1, etc.—may be rewritten as
U-UG, where U is the simple utilitarian sum of utilities and G is the
Gini coefficient.[46] The first term, U, is an "efficiency" term, measur-
ing the size of the utility pie. The second term, UG, is an equity term,
where inequality is specifically measured by the Gini coefficient (mul-
tiplied by total utility). Thus, rearrangements of the total utility pie
that do not change its size (i.e., do not change U) are evaluated as fa-
vorable under this welfare function if and only if they reduce the Gini
coefficient. Conversely, as shown in the Appendix, any rank-weighted
social welfare function with weights that increase by a constant incre-
ment in the progression downward through the ranking, from more
well-off to less well-off individuals, can be represented as the subtrac-
tion of a utilitarian term and an equity term stated as a function of the
Gini coefficient.

We have explained and defended the Pigou-Dalton definition of
equity regard, and have demonstrated how this is a quite catholic
definition—encompassing *both* "prioritarians" *and* equality theorists
who care essentially about patterns. What does the definition leave
out?[47] Some conceptions weaken the Pigou-Dalton condition. Pigou-

---

cifically, $I(I-1)/2$ is the number of unordered pairs of utilities $\{u_i, u_j\}$ where $i \neq j$.
And $\sum_{i=1}^{I-1}\sum_{j=i+1}^{I}|u_i - u_j|$ is the sum of the distance between the elements of all such un-
ordered pairs. This average distance between individuals' utilities is then normalized
by the mean μ. Initial multiplication by $\frac{I-1}{I}$ has little effect when I is large. For more
on the Gini coefficient, see Appendix Section C.

[46] More precisely, this social welfare function may be rewritten as $\frac{1}{2}(I+1)U - \frac{1}{2}IUG$,
which for large I is roughly proportional to the formula in the text. Two social welfare
functions that are proportional to each other always give the same ranking, whether
both are applied ex ante or both are applied ex post. For more on this see Appendix
Section C.

[47] Technically, adopting the Pigou-Dalton principle leaves out two social welfare
orderings derived from John Rawls's work: maximin, and its cousin, leximin. *See, e.g.*,
Philippe Mongin & Claude d'Aspremont, *Utility Theory and Ethics, in* 1 HANDBOOK OF
UTILITY THEORY 371, 415-19 (Salvador Barberà et al. eds., 1998) (discussing leximin
and maximin). In fact, however, both are effectively included as limiting cases.

Maximin ranks outcomes by looking solely to the utility level of the worst-off per-
son in each outcome. Maximin fails Pigou-Dalton: transferring from the richest to the
second-to-worst-off makes no difference to the social ranking. Indeed, maximin fails
the Pareto principle: improving the well-being of the second-to-worst-off (or any one
other than the worst-off individual) and leaving everyone else unaffected makes no dif-
ference to the social ranking.

Dalton counts *any* equalizing transfer as increasing social welfare. It is possible to stipulate that only a subset of equalizing transfers are equity improving, and that any Paretian social welfare function that respects the thus diluted Pigou-Dalton condition is equity regarding.

In particular, scholarship on equality suggests that the Pigou-Dalton principle might be weakened to regard as equity improving only the following types of transfers:

- Extreme: Any fixed transfer from the best-off person to the worst-off person.[48]
- Poverty line: Any fixed transfer to someone below a poverty line (most simply, some absolute utility level).[49]
- Across the mean: Any fixed transfer from someone above the mean to someone below the mean.[50]
- Below the mean: Any fixed transfer from someone below the mean to someone worse off below the mean.[51]

Leximin, which compares two outcomes by comparing the worst-off in both, then (if they are equal) the second-to-worst-off, and so on, does satisfy the Pigou-Dalton principle. But we still exclude it from our analysis—first, for the technical reason that leximin is a discontinuous ordering that cannot be represented by a social welfare function; second, and more substantively, because leximin preserves one of the troubling features of maximin, namely the lexical priority to the worst off. A tiny improvement in the well-being of the worst-off individual guarantees a social improvement, even if that tiny improvement is secured at large costs to the well-being of other poor individuals.

Nonetheless, both leximin and maximin can be approximated within the set of rank-weighted social welfare functions that assign all individuals strictly positive weight and strictly greater weight to those with strictly lower utility. All such social welfare functions are Paretian and satisfy the Pigou-Dalton principle. Thus, for all practical purposes, we *are* including maximin and leximin: on any finite set of possible policies, however large and comprehensive, there is a Paretian, Pigou-Dalton social welfare function that ranks these choices in precisely the same way.

[48] *Cf.* Tungodden, *supra* note 27, at 11-13 (discussing a "conditional contracting extremes" requirement, a stronger principle that stipulates that an outcome is better if the best-off individuals are made worse off and the worst off are made better off by any amount, leaving everyone else unaffected).

[49] *See* Roger Crisp, *Equality, Priority, and Compassion*, 113 ETHICS 745, 755-63 (2003) (discussing the moral significance of "compassion" thresholds).

[50] The "relative mean deviation" measure of inequality—the sum of the absolute value of the differences between each individual's well-being and mean well-being—is sensitive to transfers across the mean, but not to transfers on one side of the mean. SEN, ON ECONOMIC INEQUALITY, *supra* note 38, at 25-27.

[51] *Cf.* Peter Vallentyne, *Equality, Efficiency, and the Priority of the Worse-Off*, 16 ECON. & PHIL. 1, 10 (2000) ("The intuitions of most egalitarians will hold . . . that equality is increased . . . by slightly inefficient transfers from one person below the mean to a worse-off person."). Vallentyne himself does not advocate this particular principle.

The extreme transfers principle seems entirely too weak, and the mean-based tests give implausible relevance to the mean. Poverty line views are more plausible, but fail to account for equalizing transfers from the wealthiest to those who are much less well off, and yet not impoverished.

However, rather than flesh out these substantive objections to diluted Pigou-Dalton conceptions of equity, we will point out that these conceptions are not immune from the ex ante/ex post divergence. What drives the divergence (as demonstrated below) is that a transfer may be disequalizing *within* an outcome, but equalizing (or indifferent) in terms of expected utilities *across* outcomes. This point is true whether the equalizing transfers are all Pigou-Dalton transfers, or only a subset. For ease of exposition, and because the undiluted Pigou-Dalton condition is both the most widespread definition of an equity-regarding social welfare function and intuitively quite supportable, our detailed analysis will focus on equity regard in the Pigou-Dalton sense—with the conjecture that the analysis generalizes to social welfare functions that are sensitive to a proper subset of Pigou-Dalton transfers.

C. *Adding Uncertainty: Ex Post vs. Ex Ante Social Welfare Analysis*

In the foregoing analysis, we assumed that policies resulted in certain outcomes and were thus represented in our social welfare function as utility vectors. When we add the possibility that policies have uncertain outcomes, their representation becomes more complicated. To make the complication manageable, let us assume that the planner knows the set of possible outcomes for each social choice, that this set is finite, and that the planner knows the probability of each outcome.[52]

Instead of utility vectors the welfarist now associates policies with (1) a probability distribution over outcomes, plus (2) a utility matrix. The latter is depicted for three individuals and four states of the world in Table 1. Each column of the utility matrix corresponds to a different possible "state of the world." Each row of the matrix corresponds to a different individual in the population. Thus the column of individual utilities corresponding to a given state represents the policy outcome in that state. Each column is associated with a probability.

---

[52] *See supra* notes 19-20 and accompanying text (discussing the meaning of "outcome").

Table 1: Policy as Utility Matrix

		States of the World			
		A	B	C	D
Individuals	1	u_1^A	u_1^B	u_1^C	u_1^D
	2	u_2^A	u_2^B	u_2^C	u_2^D
	3	u_3^A	u_3^B	u_3^C	u_3^D

The difference between ex post and ex ante is now easy to state. The ex post planner assigns a value to the policy matrix by first applying his social welfare function to each column (state) regarded as a utility vector, and then calculating the expected value of these state-specific social welfare function values. The ex ante planner reverses the order: she first aggregates across columns, determining an expected utility for each individual, and then applies the social welfare function to these expected utilities. In short, the ex post planner first applies the social welfare function and then applies the expectation operator; the ex ante planner first applies the expectation operator and then the social welfare function.[53]

For readers who prefer formulas,[54] the difference between the ex ante and ex post approaches in evaluating the policy in Table 1 may be expressed as follows. The ex ante planner first calculates expected utility for individual 1:

$$\mathrm{E}\,u_1 = p^A u_1^A + \ldots + p^D u_1^D,$$

where p^A, p^B, p^C, p^D are the probabilities of the states. She then does the same for individuals 2 and 3. Lastly, she calculates the social welfare of the vector of expected utilities:

---

[53] A note on terminology: we reserve the term "social welfare function" for the function from utility *vectors* to numbers. Ultimately, however, the social planner's preferences concern not utility vectors, but *lotteries* over utility vectors (i.e., utility matrices with probabilities assigned to each column), and one might use the terminology "social welfare function" to describe the corresponding function on the latter set that represents these preferences. The "social welfare function" in this larger sense is, in the ex post case, the expected value of what we are calling the "social welfare function," and, in the ex ante case, the value under what we are calling the "social welfare function" of the expected value of the utility vectors.

[54] More general and precise formulas are included in the Appendix.

$$w\left(\mathrm{E}\,u_1, \mathrm{E}\,u_2, \mathrm{E}\,u_3\right). \tag{1}$$

This is the number that she uses to rank the policy at hand relative to alternatives. The ex post planner first calculates social welfare in state A based on the utility vector in state A, $w^A = w\left(u_1^A, u_2^A, u_3^A\right)$. He then does the same for states B, C, and D. Lastly, he calculates the expected value of these welfare values:

$$\mathrm{E}\,w = p^A w^A + p^B w^B + p^C w^C + p^D w^D. \tag{2}$$

He then uses *this* number to compare the policy at hand to alternative policies. Thus, while the ex ante planner ranks policies according to the social welfare of expected utilities, (1), the ex post planner ranks policies according to the expected social welfare, (2).

A simple example further illustrates the difference between the two approaches. Imagine that the social welfare function w is equal to the sum of the square root of individual utilities. (This is an equity-regarding function, specifically a prioritarian one.) Imagine, moreover, that there are two states that are equally likely, and two individuals, Jim and June. In the status quo, both individuals receive a utility of 3.7 in both states. By contrast, as depicted in Table 2, the policy's outcomes are uncertain: in one state, Jim has utility 4 and June utility 0; in the other state, Jim has utility 9 and June utility 4.

The ex post application of w assigns a value of 3.5 to the policy. The ex ante application of w assigns to the policy a value of 3.96. Because the status quo is certain—the same utility vector results regardless of the state—the ex post and ex ante application of w to the status quo assigns it the same value, namely $3.85 \cong \sqrt{3.7} + \sqrt{3.7}$. So in this case the differing methods for applying w lead to divergent recommendations for policy choice. The ex post method says that the status quo is a better choice than the policy (because 3.85 is greater than 3.5), while the ex ante method says that the policy is a better choice than the status quo (because 3.85 is less than 3.96).

Table 2

	State A $p=.5$	State B $p=.5$	Expected Utility
Jim	4	9	6.5
June	0	4	2
w (Jim, June)	$\sqrt{4}+\sqrt{0}=2$	$\sqrt{9}+\sqrt{4}=5$	w applied ex post: $2\times.5+5\times.5=3.5$ w applied ex ante: $\sqrt{6.5}+\sqrt{2}\cong 3.96$

As we have already explained, the terms ex ante and ex post may be a little misleading. Both approaches take account of probabilities. And both approaches are guides to choice, not just tools for retrospective assessment. The difference—subtle but vital—has to do with the specific way in which the choiceworthiness of a probabilistic choice is evaluated: the order of integration across states and persons. Perhaps a better terminology might distinguish between "expected individual well-being" (rather than ex ante), and "expected social welfare" (rather than ex post) approaches to choice under uncertainty. However, for the sake of terminological consistency with the existing literature on social choice under uncertainty, we will stick with ex ante and ex post.

The ex ante and ex post valuations of a policy might coincide. In fact, if the social welfare function is utilitarian (or more generally linear with weights that depend only on individual identity), the two valuations *must* coincide. It makes no difference whether one takes the expected value of the (weighted) sums of the columns (the ex post approach) or the (weighted) sum of the expected values of the rows (the ex ante approach). Indeed, the converse also holds.[55] If the social welfare function is such that ex post and ex ante valuations always coincide, then the social welfare function must be some weighted sum of individual utilities (plus some constant term, which is irrelevant for social choice[56]), with weights that depend only on individual

---

[55] *See infra* Appendix Section A.

[56] Adding a constant k to the social welfare function w changes the social welfare of all policies in all states by precisely the same amount, namely k. It, therefore, in-

identity.[57] If such weights are equal, the social welfare function is utilitarian. But if the weights are unequal, then the social welfare function violates the symmetry condition. All told then, a Paretian, symmetric social welfare function is guaranteed to make ex post and ex ante valuations coincide if and only if that social welfare function is utilitarian.

The proposition that ex post and ex ante are guaranteed to converge only if the social welfare function is a weighted sum of utilities is, mathematically, quite straightforward. That proposition is also a direct implication of a substantially stronger theorem that John Harsanyi established fifty years ago.[58] But it is not obvious what the exact

---

creases the *expected* social welfare of every policy by *k*. It also increases the social welfare of *expected* utilities under every policy by *k*. It is thus irrelevant to our analysis.

[57] This "converse" proposition is a straightforward consequence of the definition of linearity. One way to see this is to notice that this converse is merely an extension to a multidimensional domain of the basic point from individual risk analysis that risk neutrality implies a linear (actually "affine," to be precise, *see infra* note 58) utility function. Risk neutrality means that the expected utility of a lottery over dollar amounts equals the utility of the expected value of the numbers. The statement that ex post social welfare equals ex ante social welfare is directly analogous to this: the social welfare function plays the role of individual utility; utility vectors play the role of dollar amounts; ex post social welfare is like expected utility; and ex ante social welfare is like the utility of the expected value.

Another way to think of it—one that takes account of the fact that we are dealing with lotteries over (utility) vectors and not over (dollar-valued) scalars—is this: First, imagine a function of two variables onto one. Think of the function in a three-dimensional graph with function values on a third axis, a "z axis," coming up out of the page. Then take any finite group of points on the x-y plane and connect the dots to make a polygon in the x-y plane. Finally, look at the function values over the region bounded by this polygon. The statement that ex post social welfare equals ex ante social welfare means that the function values form a plane over this polygon (though, of course, the plane will generally have a slope). Since this works for any finite group of points, the function values must form a plane globally.

[58] *See* John C. Harsanyi, *Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility*, 63 J. POL. ECON. 309 (1955). For a detailed discussion, see John A. Weymark, *A Reconsideration of the Harsanyi-Sen Debate on Utilitarianism, in* INTERPERSONAL COMPARISONS OF WELL-BEING 255 (Jon Elster & John E. Roemer eds., 1991). Harsanyi is famous for identifying the tension between ex ante Pareto superiority and non-utilitarian approaches to social choice. More precisely, Harsanyi proves that, if individuals and the social planner both obey the expected utility axioms, and the social planner is indifferent between options when all individuals are also indifferent, then the social planner's decisions are representable as the maximization of a social welfare function equaling a weighted function of individual utilities (plus a constant). Within the bounds of our simplified structure, which takes utility functions and social welfare functions as given, Harsanyi's theorem implies the following: if, whenever the expected utility of each individual is the same across two policies the ex post social welfare is also the same, then the social welfare function is affine—meaning it is a linear function of individual utilities plus a constant. To see why the proposition that ex

connection is between ex ante/ex post divergence and equity regard, nor whether the phenomenon of divergence is a pervasive one. Moreover, it is crucial for lawyers to consider the role of compensation and insurance in mitigating divergence. Finally, the extent of the conflict between the ex post application of an equity-regarding social welfare function and ex ante Pareto superiority needs to be explored. These are the issues that we now examine in turn.

II. UNDERSTANDING THE DIVERGENCE BETWEEN EX POST AND EX ANTE APPROACHES TO SOCIAL CHOICE

The object of this Part is to illuminate how ex post versus ex ante approaches to implementing an equity-regarding social welfare function can produce divergent assessments. Section A explicates the basic phenomenon by demonstrating that for all equity-regarding social welfare functions, there are some choice situations that are evaluated differently ex ante and ex post. The remaining Sections establish that such divergence is ubiquitous. Differing evaluation of choice situations is particularly easy to produce for separable social welfare functions.[59] Nonseparable social welfare functions are also vulnerable to ex ante/ex post divergence when the utility rankings of individuals differ across policy outcomes.[60]

---

post/ex ante convergence implies a linear social welfare function is a straightforward corollary of Harsanyi's theorem, consider the following: (1) if individual expected utilities are equal as between two policies, then the ex ante approach must rank them as equal; and (2) if, in turn, ex post and ex ante valuations always converge, then (in particular) the ex post approach must be indifferent between two policies where each individual's expected utilities are the same.

To be sure, using Harsanyi's theorem to prove that ex post/ex ante convergence implies a linear social welfare function is a bit like the proverbial elimination of a mosquito with a bazooka. The theorem is more directly relevant to our analysis, below, of the divergence between the ex post approach and ex ante Pareto superiority. *See infra* note 93. We mention Harsanyi, here, because the body of scholarship to which this Article aims to contribute—examining the interaction between equality, uncertainty, and Pareto superiority—was spawned by Harsanyi's work.

[59] For the definition of separable social welfare functions, see *supra* text accompanying notes 42-43.

[60] As a preliminary matter, we should assuage one general worry that a technically sophisticated reader might have about our approach. The analysis in this and subsequent Parts takes a particular social welfare function and applies it ex post and ex ante to a given policy, or a given choice among policies, showing how the two approaches can diverge. The reader might worry that social welfare functions are simply mathematical devices for representing orderings, and that we really should be comparing the ex post application of the social welfare function and all admissible transformations of

A. *The Basic Phenomenon*

In this Section we explore the basic connection between ex
ante/ex post divergence and welfare judgments about inequality.[61]
We show that, starting from any uncertain status quo, one can always
find a mean-preserving transfer that is Pigou-Dalton equalizing from
an ex ante perspective and yet Pigou-Dalton disequalizing from an ex
post perspective. This transfer, viewed as a policy, would therefore be
favored by every equity-regarding social welfare function applied ex
ante, and disfavored by every equity-regarding social welfare function
applied ex post. To make the point even more dramatically, we also
show that a policy may be Pigou-Dalton disequalizing in *all* states si-
multaneously, while at the same time Pigou-Dalton equalizing ex ante.

---

that function, on the one hand, to the ex ante application of the social welfare func-
tion and all admissible transformations of that function, on the other hand.

More precisely, the reader might worry that we should be comparing the ex post
application of the social welfare function and all affine transformations of that func-
tion to the ex ante application of that function and all increasing transformations of it.
The ex post approach ranks lotteries by taking expectations over the value of the social
welfare function in each state. Thus, the ex post application of any affine transforma-
tion of a social welfare function will preserve the ordering of lotteries generated by the
initial function. The ex ante approach ranks lotteries by applying the social welfare
function to the vector of expected utilities associated with the lottery. Thus, the ex
ante application of any increasing transformation (affine or not) of a social welfare
function will preserve the ordering of lotteries generated by the initial function.

The answer to this subtle but important worry is straightforward. If a particular
social welfare function generates ex ante/ex post divergence in some choice situation,
then any affine transformation of that function, applied ex post, will diverge from any
increasing transformation of that function, applied ex ante. Thus, the apparently na-
ive approach our analysis takes—showing the conditions under which divergence over
choices can arise when a given social welfare function is applied ex post as opposed to
ex ante—reaches exactly the same results as a more complicated analysis that explicitly
considers certain admissible transformations of the function.

[61] Equity regard is not the only source of ex ante/ex post divergence. Suppose,
for example, that individual utilities in each of the two equally likely states are (100,
100) and (100, 100) in the status quo and would be (0, 0) and (200, 200) after imple-
menting the policy—meaning the individuals have equal utility in all circumstances.
The policy leaves the individuals' expected utilities at their status quo level of (100,
100). But the policy "spreads" both individuals' utilities away from 100 across the two
states, making them 0 in one state and 200 in the other. The policy is perfectly neutral
from an ex ante perspective (and a slight change of the numbers would make it strictly
improving). But this is not necessarily the case from an ex post perspective. If mar-
ginal social welfare diminishes in individual utilities, then the policy is ex post social
welfare *reducing* (and would remain so were the policy numbers slightly changed). Be-
cause marginal social welfare diminishes, the 100 unit reduction in utilities in the first
state takes more away from social welfare than the 100 unit increase in utilities in the
other state gives back.

Consider first a hypothetical policy that affects only one state. Suppose that the vector $t_1^s, ..., t_I^s$ represents the change in utility effected by the policy in state s, with individuals indexed by subscripts 1 through I. Each t_i^s may be positive or negative. To focus on equity, let this "transfer" be mean preserving, so that $t_1^s + ... + t_I^s = 0$. Viewed from the perspective of state s taken alone, this transfer might be Pigou-Dalton equalizing, Pigou-Dalton disequalizing,[62] or neither.[63] Which of these it is depends on the initial level of utilities in that state $u_1^s, ..., u_I^s$. Roughly speaking, if the transfer tends to go toward those with high (low) initial utility in state s, it is disequalizing (equalizing).

The key insight is that this ex post transfer in state s also produces a "transfer of expected utility" ex ante. In particular, where p^s is the probability of state s, expected utilities change according to $p^s t_1^s, ..., p^s t_I^s$. To wit, when we transfer t_1^s units of utility to individual 1 in state s, we also transfer $p^s t_1^s$ units of *expected* utility to individual 1 from an ex ante perspective. Whether this transfer of expected utility is equalizing, disequalizing, or neither from an ex ante perspective depends, analogously, on the initial level of *expected* utilities.

The question is whether the transfer might be Pigou-Dalton disequalizing in state s and yet Pigou-Dalton equalizing ex ante. It is easy to see that this is possible. The simplest Pigou-Dalton *disequalizing* transfer within state s is a single pairwise transfer from an individual with lower utility to an individual with higher utility. (In this case, only two t_i^s's are nonzero.) This transfer can, nevertheless, be *equalizing* ex ante if the individual with the lower utility in state s is the individual with the initially higher *expected* utility. In this case, the transfer of t units of utility from "poor" to "rich" within state s becomes a transfer of $p^s t$ units of utility from *rich to poor* ex ante.[64]

---

[62] As discussed in *supra* note 40 and at *infra* Appendix Section B, a multiperson transfer is Pigou-Dalton equalizing if and only if it can be generated by a sequence of pairwise transfers that are zero sum and reduce the distance between the utilities of the two individuals. A multiperson transfer is Pigou-Dalton disequalizing if and only if its additive inverse is Pigou-Dalton equalizing. The Pigou-Dalton *principle*, as defined in *supra* Part I.B, plus transitivity implies that a multiperson transfer that is Pigou-Dalton equalizing (disequalizing) must increase (decrease) social welfare.

[63] Given a *particular* inequality measure respecting the Pigou-Dalton condition, it will be equalizing, disequalizing, or perfectly *neutral.*

[64] Of course, in order for the transfer to be equalizing ex ante, it must also not overshoot the ex ante inequality between these two individuals. For example, if two individuals start out with utilities (100, 90), and we transfer 3 units of utility from the first to the second, we close the gap by 6 units. If we transfer 5, we close the gap entirely. If we transfer 7, we open a gap in the other direction, but not one greater than the initial gap. However, if we transfer more than 10, we open a gap in the other di-

Suppose, as shown in Table 3, that status quo utilities in two equally likely states are (10, 20) and (200, 100), so that expected utilities are (105, 60). Consider a policy that transfers 10 units of utility from the first individual to the second in the first state, thus changing the first state's utility vector from (10, 20) to (0, 30). This is disequalizing in the first state. But the corresponding transfer of expected utility—namely, $\frac{1}{2}10 = 5$ from the first individual to the second—is an equalizing transfer from an ex ante perspective, wherein the first individual is better off.

Table 3

		State A $p = .5$	State B $p = .5$	Expected Utility
Status quo	**Individual 1**	10	200	105
	Individual 2	20	100	60
Policy	**Individual 1**	0	200	100
	Individual 2	30	100	65

It is vital to see that this example does not trade on a particular kind of equity-regarding social welfare function. In this case, because the policy is disequalizing in the first state in terms of utilities, but equalizing in terms of expected utilities, *any* social welfare function that is equity regarding in terms of satisfying the Pigou-Dalton principle—whether the function is comparativist or noncomparativist and (more technically) whether it is continuous or discontinuous—will

---

rection that is larger than the initial gap. Thus, equalization is a matter not just of the direction of the transfer, but also of its magnitude. But, returning to the example in the text, since a transfer in this direction of any magnitude is disequalizing in state *s*, it is still easy to see that the transfer can be made small enough to avoid overshooting ex ante inequality.

This discussion points to another way in which the transfer can be disequalizing in state *s* and yet equalizing ex ante. If an equalizing pairwise transfer must be both in an initially equalizing direction and also not so large that it overshoots and causes greater inequality in the other direction, then there are two ways for a transfer to be disequalizing, in particular, in state *s*. It can be initially disequalizing, as above. Or it can be initially equalizing, but go too far. In the latter case, so long as it does not also go too far ex ante, it will be disequalizing in state *s* and equalizing ex ante. Notice that this arises when the same individual is initially better off in both state *s* and ex ante. *See infra* Appendix Section D.

count the policy as worse than the status quo if applied ex post, and better than the status quo if applied ex ante.

More generally, given any status quo with at least one *rank-contrary state*—that is, one state in which the utility ranking of individuals ex post is not precisely the same as the expected utility ranking of individuals ex ante—we can always find some mean-preserving policy change to this status quo that is Pigou-Dalton equalizing ex ante and Pigou-Dalton disequalizing ex post. The significance of using a mean-preserving transfer is that the discrepancy between ex ante and ex post evaluations arises solely due to different views regarding what is equalizing and what is disequalizing. The significance of the application of the general Pigou-Dalton definition of equalization is that the existence of this discrepancy is independent of the particular equity-regarding social welfare function employed.

Less obvious from the foregoing discussion in this Section is that the same result holds even if there is no rank-contrary state in the status quo.[65] We can still find a mean-preserving transfer that is Pigou-Dalton disequalizing ex post and Pigou-Dalton equalizing ex ante. How so? Roughly speaking, it is always possible to find a transfer from an initially better-off individual to an initially worse-off individual that is great enough to "overshoot" inequality within the state wherein it is made, but not so great that its (probability-discounted) ex ante impact overshoots inequality in expected utilities.[66]

Of course, a policy will, in general, affect more than one state, and so the next question that arises is whether such a general policy can be ex post disequalizing—taking account of its effects across all states—and ex ante equalizing. Difficulties arise if the policy has differing equity effects across states: that is, if it is Pigou-Dalton disequalizing in some, Pigou-Dalton equalizing in others, and neither in still others. Which effects dominate in ex post social welfare will depend on the probability distribution across states and the particular manner in which the social welfare function incorporates equity regard. No general statements can be made for the full class of equity-regarding social welfare functions.

---

[65] The result holds so long as there is some inequality in some state. Note that this assumption was unnecessary in the previous paragraph because it was implied by the existence of a rank-contrary state. If there is *no* inequality in any state, then there is no inequality ex ante, and all mean-preserving transfers in any state are disequalizing both ex post and ex ante.

[66] This logic is more completely spelled out in the second paragraph of note 64 *supra.*

However, we can answer a more specific, but still informative question: is it possible for a policy to be disequalizing in *every* state taken individually, and yet equalizing ex ante? In this case, we can be assured that the policy is disequalizing in ex post social welfare when all states are aggregated, and so to this extent ex post social welfare reducing (and at the same time ex ante social welfare increasing for the fact that it is ex ante equalizing).

It is not difficult to see that even this is possible. Focus again on policies that, now in every state, effect a simple pairwise transfer between a fixed pair of individuals, say individuals 1 and 2. Recall that one way for such a policy to be disequalizing in state *s* is for it to transfer from the individual with low utility in that state to the individual with high utility. Suppose that it does this in every state. Now in some states the relative ranking of individuals 1 and 2 differs from their ranking ex ante. This is similar to the case discussed above. In each such "rank-contrary" state, therefore, the transfer, disequalizing intrastate, is equalizing ex ante, just as above. In other states, the relative ranking of the two individuals is the same in that state as ex ante. For these "rank-conforming" states, the within-state disequalizing transfer is disequalizing ex ante. It is clear then that ex ante equalization requires that the expected value of the transfers in rank-contrary states must exceed the expected value of the transfers in rank-conforming states.[67]

A simple example will clarify. Suppose that initial utilities for two individuals in each of the two equally likely states are (40, 60) and (80, 20), so that expected utilities are (60, 40). Therefore, the first state is rank contrary, and the second rank conforming. Suppose that the policy transfers 20 units of utility to the higher utility individual in the first state, and only 10 units to the higher utility individual in the second state, making the utility vectors (20, 80) and (90, 10). In this case, expected utilities become (55, 45). Thus, the policy is equalizing ex ante and yet disequalizing in every state. As a result, the policy is ex ante social welfare increasing, but ex post social welfare decreasing. What is happening here? The expected value of the disequalizing transfer in the rank-contrary state, $\frac{1}{2}20 = 10$, exceeds the expected

---

[67] Moreover, so as not to overshoot ex ante to such an extent as to cause ex ante disequalization, the excess of the first state over the second state can not exceed the initial difference in expected utilities.

value of the disequalizing transfer in the rank-conforming state, $\frac{1}{2}10 = 5$.[68]

A conceptually important special case of this phenomenon arises when there is initially perfect equality ex ante. In this case, the disequalizing transfers in each state must perfectly cancel in expected value in order to disequalize ex post and yet not disequalize ex ante. For example, imagine that individuals in both of the states have utilities of 100, and the first state has a probability of 1/3. If the policy changes the first state to (130, 70), then the second must be changed to (85, 115) to preserve ex ante equality. The expected transfer in the first state (30 × 1/3 =10) and in the second (15 × 2/3 =10) must be equal. In the even simpler case where the two states have equal probability, the transfers in both states must be the same. So if the policy changes the first state to (120, 80), the second must be changed to (80, 120). We now see that this very simple case, which itself includes the JANET and JOHN example in the introduction, is a special case of a more general phenomenon.[69]

B. *Rank Uncertainty, Rank-Weighted Social Welfare, and the Gini Coefficient*

The prior section identified the kind of equity-implicating policy changes that produce ex ante/ex post divergence for every equity-regarding social welfare function. A different approach is to pick out a subclass of equity-regarding social welfare functions and give a fuller description for that subclass of the kinds of policy choices that lead to ex ante/ex post divergence. This Section, along with the next three, takes this alternative approach. The present Section shows that ex ante/ex post divergence is *not* an issue for certain kinds of policy choices and a certain set of equity-regarding social welfare functions.

Let us say that a policy (be it the status quo or a policy that alters the status quo) is "rank homogenous" if each individual has the same rank in each possible state. A "rank-heterogeneous" policy is one that is not rank homogenous. For example, the policy in Table 4 below is rank-homogenous because individual 2 is the worst off in all three

---

[68] Moreover, in conformity with note 67 *supra*, the excess 10 – 5 = 5 does not exceed the initial difference in expected utilities 60 - 40 = 20.

[69] In accord with our discussion of the single state case in note 64 *supra*, there are other ways for the same phenomenon to arise. In some states, the disequalization can come from overshooting the inequality. For such states, the transfers are counted for ex ante equalization if the state is rank conforming and against if the state is rank contrary. Finally, note that by the same kind of reasoning described herein, a policy can be disequalizing ex ante and equalizing ex post.

states, individual 3 is always the best off, and individual 1 is always in the middle.

Table 4: A Rank-Homogenous Policy

| | | States of the world | | |
		A	B	C
Individuals	1	50	41	85
	2	20	40	60
	3	51	60	100

Rank-homogenous policies are policies that involve no uncertainty about the relative standing of different individuals. If a policy is rank homogenous, there can be uncertainty about how much well-being various individuals will have and what the differences between their well-being levels will be, but not about whether any given individual will be better or worse off than any other.

In an important class of equity-regarding social welfare functions, ex ante social welfare equals ex post social welfare on the set of rank-homogenous policies. In other words, the social welfare number that one would obtain applying such a social welfare function to a rank-homogenous policy in ex post fashion would be the same as the number one would obtain applying it in ex ante fashion.

The class of equity-regarding social welfare functions referred to here are the rank-weighted social welfare functions discussed above in Part I.B. Recall the simplest social welfare function of this form: that ranking the most well-off person as 1, the second most well-off person as 2, the third as 3, etc., all the way down to the least well-off person, whose utility would be multiplied by *I* (the size of the population) before being added to social welfare. To be sure, not all rank-weighted social welfare functions are normatively plausible. But some are. In particular, as discussed in Part I.B, the Gini coefficient—a very common measure of inequality, indeed perhaps the measure most frequently used by social scientists—corresponds to a rank-weighted social welfare function.

Why do rank-weighted social welfare functions yield the same welfare number when applied ex ante or ex post to rank-homogenous policies? The answer piggybacks on our earlier discussion of utilitari-

anism and generalizes Harsanyi's famous result.[70] We noted in Part I.C that if we are just taking the simple sum of utilities, it makes no difference whether we sum expected utilities over individuals or take the expected value of the sum of utilities in each state. More generally, there is no difference between ex ante and ex post social welfare when the social welfare function is a *weighted* sum of individual utilities with weights that do not change across states. Because individual ranks do not change across states for a rank-homogenous policy, rank is effectively fixed across states for purposes of evaluating such a policy. In this case, it does not matter whether we take the expected value of a weighted sum of utilities or a weighted sum of expected utilities.

Because the ex post and ex ante value of a rank-homogenous policy are necessarily the same under any rank-weighted social welfare function, it follows that, for any choice situation between *two* rank-homogenous policies, even if the rankings differ *across* the two policies, a rank-weighted social welfare function will pick the same option whether that function is applied ex ante or ex post. If, under an ex post approach, the social welfare function assigns the first policy a value of 1000 and the second a value of 2500, then under an ex ante approach it must make the same value assignments. Thus, it must pick the same policy. Of course, the same logic applies to the choice among more than two individually rank-homogenous policies.

C. *Strictly Concave Social Welfare Functions*

Rank heterogeneity is not a necessary condition for ex ante/ex post divergence across all equity-regarding social welfare functions. Consider, for example, the class of social welfare functions that are "strictly concave." Roughly speaking, a social welfare function is strictly concave if each individual's utility makes a diminishing marginal contribution to social welfare.[71] That is, giving one more unit of utility to a less well-off individual increases social welfare more than giving one more unit of utility to a more well-off individual.

A rank-weighted social welfare function is not strictly concave. In particular, the marginal impact on social welfare of changes in an in-

---

[70] *See supra* notes 57-59 and accompanying text.

[71] More generally, increasing all individuals' utilities in any fixed proportion has a diminishing impact on social welfare.

dividual's utility that do not upset utility *rankings* across individuals is constant rather than diminishing.

A strictly concave social welfare function always assigns a greater ex ante value to a policy[72] than its ex post value. Indeed, the converse is also true. In other words, *the ex ante value of a social welfare function exceeds its ex post value for every uncertain policy if and only if that social welfare function is strictly concave.*

An explication of the reasoning behind this result is beyond the scope of our present discussion. Suffice it to say that this reasoning has a close analogue in the theory of individual choice, where strict concavity of *individual utility* for money translates into individual risk aversion with respect to money—that is, an individual preference for the expected value of a money gamble as opposed to the gamble itself.[73]

Because strictly concave social welfare functions always value a policy greater when applied ex ante than when applied ex post, it follows that: ex post versus ex ante social planners, implementing strictly concave social welfare functions, may evaluate policy *choices* differently. Again, unlike the case of rank-weighted social welfare, this is so even if all the policies at issue are rank homogeneous. Indeed, the policy choice may be differently evaluated under a strictly concave social welfare function even if, additionally, all such rank-homogeneous policies share the same ranking. Take a rank-homogenous status quo (also a "policy" in our framework), whose ex ante value is s and whose ex post value is $r < s$. Consider a policy that is certain to result in some given utility vector. Assume that each individual has the same rank in this utility vector as in each state in the (uncertain but rank-homogenous) status quo. Because the certain-outcome policy leads to the same result in each state, the social welfare function's ex ante and ex post evaluation of the policy must be the same—call it q. Certainly, we can imagine the case wherein q lies between r and s.[74] In that case, ex ante application of the social welfare function prefers the status quo, but ex post application prefers the policy.

---

[72] That is, if the policy has some uncertainty—in other words, some divergence in utility vectors across states. This is Jensen's inequality. *See, e.g.,* Tristan Needham, *A Visual Explanation of Jensen's Inequality,* 100 AM. MATHEMATICAL MONTHLY 768 (1993).

[73] *E.g.,* KREPS, *supra* note 9, at 81-84.

[74] By the intermediate value theorem, such a q will exist, no matter how close r and s are, so long as the social welfare function is continuous on $\mathbb{R}^I$ and we can find a certain policy with a value less than r and a certain policy with a value greater than s.

D. *Permuting Policies*

In a sense, the "opposite" of a rank-homogenous policy is a permuting policy. A permuting policy has the following characteristic: each state's utility vector is just a reordering of the other states' vectors. A given ordering can occur more than once and not every ordering need be represented. Table 5 gives an example. Recall that rank-homogenous policies have no uncertainty regarding rank: we know how individuals are ranked; we just don't know the levels of their utilities. By contrast, rank is the only uncertainty in a permuting policy: we know the levels of utilities; we just don't how they are assigned.

Table 5: A Permuting Policy

		States of the world		
		A	B	C
Individuals	1	50	20	20
	2	20	50	51
	3	51	51	50

As we prove in the Appendix, all equity-regarding social welfare functions, including not only strictly concave social welfare functions but also rank-weighted ones, must give a greater ex ante than ex post valuation to a permuting policy.[75] Intuitively, the ex ante planner sees less inequality in a permuting policy than does her ex post counterpart, because the switching of rank across states turns the expected utility calculation into an averaging across individuals.[76] Permuting policies therefore provide another ready recipe for generating ex ante/ex post divergent choices—one directly parallel to the recipe we used above for generating ex ante/ex post divergence from strictly

---

[75] More precisely, the permuting policy must exhibit some uncertainty and some inequality. *See infra* Appendix Section F.

[76] Specifically, one can show that taking weighted averages over permuted outcomes, as we do in taking expected utilities, corresponds to making a (multiperson) mean-preserving equalizing transfer relative to any one of the outcomes. *See supra* note 62. It follows that expected utilities under any probability distribution over outcomes must generate a higher social welfare value than any one of the permuted outcomes. Therefore, under any probability distribution the social welfare of expected utility will exceed the expected value of social welfare over such outcomes.

concave social welfare functions. Given *any* equity-regarding social welfare function and any permuting policy, produce an alternative policy whose outcome is certain, and whose social welfare value (the same for ex ante and ex post) lies between the divergent ex post and ex ante values of the permuting policy.[77] Then ex ante application of the social welfare function will prefer the original, permuting policy, while ex post application will prefer the alternative policy.

E. *Prioritarianism and Separable Social Welfare Functions*

The rank-weighted social welfare function is the leading example of a nonseparable social welfare function. It is nonseparable because the welfare contribution of increases in each individual's utility depends on how the individual's utility is weighted, which depends on her rank, which in turn depends on the utility of other individuals.

What if we turn to separable social welfare functions, where the social preference between two utility vectors cannot depend on the utility of indifferent individuals? As discussed in Part I.B, separable social welfare functions correspond to a prioritarian/non-comparativist conception of social welfare. Our discussion there indicated that, at least in the context of evaluating policies with certain outcomes, separable social welfare orderings could be represented by an additive social welfare function—that is, one of the form $w(u_1,...,u_I) = \varphi(u_1) + ... + \varphi(u_I)$.[78]

The representation of prioritarianism/noncomparativism in the context of uncertainty remains unsettled.[79] Let us assume that such views imply an additive social welfare function in this context, so that ex post social welfare is the expected value of $\varphi(u_1) + ... + \varphi(u_I)$, and ex ante social welfare is the same sum of functions applied to expected utilities.

---

[77] *See supra* note 74 (describing the assumptions that guarantee the existence of this alternative policy).

[78] This assumes that the ordering is continuous. *See supra* note 22 and accompanying text. And the fact that w is the same for all individuals follows from anonymity/symmetry.

[79] *See supra* note 43. Although a separable ordering over certain outcomes can be represented (assuming continuity) by an additive social welfare function, it can also be represented by a nonadditive function. *See* Fleurbaey, *supra* note 42, at 6-7, 16-17. The "correct" function, under conditions of uncertainty, is the one that accurately represents the social planner's preferences over lotteries. If, for example, she is an ex post planner who is risk averse in social welfare, then her "real" social welfare function may be a strictly concave transformation of the sum of some function of individual utilities, rather than simply the sum of some function of individual utilities.

Here, it emerges—just as in the case of strict concavity—that rank heterogeneity is *not* a necessary condition for ex ante/ex post divergence. This is no coincidence. All equity-regarding additive social welfare functions must be strictly concave. Why is this? The intuition is simple. Suppose that individuals 1 and 2 have the same utility and start repeatedly transferring small amounts from 1 to 2. After the first transfer, the transfers are disequalizing: 1's utility is now below 2's. Thus, they must reduce social welfare. But this can happen only if w increases at a decreasing rate, so that the additions to 2's utility are worth less than the subtractions from 1's.

Because additive social welfare functions are strictly concave, it follows that every such function gives a greater ex ante than ex post value to every policy. Moreover, ex ante versus ex post application of the function can lead to divergent assessment even where all the policy choices involved are rank homogenous.

F. *A Summary*

What have we established in this Part? To begin, *every* equity-regarding social welfare function is susceptible to ex ante/ex post divergence. For every such function, there will be some choice situations that are evaluated differently depending on whether the function is applied ex ante or ex post. Permuting policies provide a ready recipe for generating ex ante/ex post divergence, since every equity-regarding social welfare function assigns a greater ex ante value to a permuting policy than its ex post value.

However, some social welfare functions are especially vulnerable to ex ante/ex post divergence. Here, philosophical debates about prioritarianism come into play. Prioritarianism corresponds (plausibly) to an additive equity-regarding social welfare function, which in turn is one kind of strictly concave function. The popular Gini coefficient corresponds to a rank-weighted social welfare function, which is not strictly concave. Rank-weighted functions are immune from ex ante/ex post divergence if all the policies at stake are rank homogenous—that is, if there is no uncertainty about what any given individual's rank in the utility distribution will be. By contrast, strictly concave social welfare functions, including additive, prioritarian functions, assign a greater ex ante than ex post value to every policy,

and therefore can generate divergent choice assessments even if all of the policies at stake are rank-homogenous.[80]

III. COMPENSATION AND INSURANCE

This Part examines the extent to which ex post compensation by the government and insurance by private individuals alleviate ex ante/ex post divergence.

Section A considers ex post compensation by the government. It shows that ex ante/ex post divergence is alleviated only if utility is "freely transferable," not just across individuals (the usual result), but also across states. When utility is not "doubly freely transferable," the divergence remains an issue in two respects. First, the ex post planner may regard as unjustifiably expensive the compensation that would be required to assuage her concerns with a policy's ex post disequalizing effects, concerns that her ex ante counterpart may not share. Second, even if the ex post planner finds such compensation cost justified, so that she comes into agreement with the ex ante planner regarding the policy itself, she and her ex ante counterpart may disagree about the propriety of the compensation, for which the ex ante planner may see little purpose, but which is essential for the ex post planner to see the policy as beneficial.

Building on these results, Section B discusses compensation from a different angle quite relevant to current debates about relegating equity considerations to tax law. What if we disaggregated social welfare analysis across levels of government? We hypothesize a special

---

[80] A superficial reading of John Broome's work may suggest a contrary result: namely, that prioritarian social welfare functions are actually immune from ex ante/ex post divergence and, in particular, are immune from a conflict between the ex post approach and ex ante Pareto superiority. But, in fact, the difference between our results and Broome's work arises because—in the case of prioritarianism—the ex ante approach can be conceptualized in different ways. In the case of a prioritarian function—one with the form $\varphi(u_1) + ... + \varphi(u_I)$—it can be applied either ex ante (as we have defined that) or "ex ante+." Ex ante+ means determining each individual's $\varphi(u_i)$ value in each state, determining her expectations of these values, and summing these expectations. Clearly, the ex ante+ application of a prioritarian function and the ex post application will always converge. On these issues, see Rabinowicz, *supra* note 6, at 147-54. We do not find the notion of an ex ante+ application of a social welfare function particularly compelling. The notion is simply inapplicable for comparativist social welfare functions, and the more natural expression of ex ante egalitarianism and ex ante Pareto superiority, we believe, looks to the vector of expected well-being associated with each policy, not the vector of expectations of some concave transformation of well-being. For these reasons, we do not further discuss the ex ante+ approach to equity regard in this Article.

compensation agency (perhaps within the IRS, perhaps the legislature itself). All other agencies are instructed to ignore equity and choose the policy that maximizes the simple sum of utility across individuals. The compensation agency, however, is given an equity-regarding social welfare function by the legislature and is told to implement equity concerns as follows: for any chosen policy, the compensation agency should adjust the policy using a compensation scheme that is optimal in light of the equity-regarding function. Would this two-step disaggregated approach reliably produce policies that are at least as good as those produced by a holistic, one-step application of the social welfare function? Absent free transferability of utility across individuals and states, the answer is no. Without these special conditions, the two-step approach can yield lower social welfare.

Section C discusses private insurance, which may not only fail to eliminate ex ante/ex post divergence, but may even frustrate governmental efforts to compensate policy losers.

A. *Will Compensating Policy Losers Eliminate the Ex Ante/Ex Post Divergence?*

The ability to compensate policy losers from the winnings of policy winners may in certain special cases eliminate the divergence between ex post and ex ante social welfare. To see this, consider a simple case where there are two individuals and two equally likely states. (The example is similar to the JANET and JOHN example in the Introduction.)[81] Under the status quo, each individual receives 300 in both states. The policy changes the possible outcomes to (200, 500) and (500, 200). In other words, in the first state, individual 1 suffers a loss (of 100) and individual 2 enjoys a greater gain (of 200). In the second state, the individuals' roles are precisely reversed. Both individuals' expected utility is 350. The policy is ex ante Pareto superior to the status quo and will therefore be favored by all ex ante planners whatever their particular social welfare function. An equity-regarding ex post planner, however, recognizing the certainty of ex post disequalization, might well disfavor the policy.

Now suppose that the winner in each state could compensate the loser, leaving both better off. Imagine, for example, that the winner could transfer a sufficient number of "utils" to the loser (150) to eliminate the discrepancy in their utilities. With such transfers in

---

[81] *See supra* pp. 285-88.

place, the ex post policy outcome in each state would be Pareto improving relative to the status quo (providing each individual an additional 50 utils). Hence, every Paretian ex post planner would now favor the policy (*cum* transfer). Since the same can still be said of every Paretian ex ante planner, the divergence between ex ante and ex post would be eliminated.

The question then is how far can we push this example. Under what conditions does compensation eliminate the divergence between ex ante and ex post social welfare? We highlight two necessary conditions that by themselves seem exceptional enough to tilt the balance toward concluding that compensation cannot eliminate the issues we have highlighted thus far in the paper.

1. Free Transferability of Utility Across Individuals

One implicit assumption in the example above is that a unit of utility can be taken from the winner and given to the loser with no loss of utility along the way. In general, compensation is capable of eliminating ex ante/ex post divergence only if utility is "freely transferable" across individuals. (We discuss free transferability across *states* below.) Utility is freely transferable across individuals (within a given state) if we may "transfer" units of utility from any individual to any other individual without changing the total utility in that state in the process.

To see why freely transferable utility matters, let us return to the initial example. Suppose now that the transfer "bucket" is leaky, so that for every util leaving one person, only one-third of a util arrives at the other person. (Below we describe why the bucket might leak in this way.) Thus, within each state, every util taken from the winner and transferred (in part) to the loser reduces the utility distance between them by one (the winner's utility reduction) plus one-third (the loser's utility increase). But it does this at a cost of two-thirds in total utility across the two individuals, since only one-third of the transferor's utility loss makes it to the transferee. This trade-off between inequality reduction and total utility might not be acceptable to an equity-regarding social planner. In other words, the policy, even when combined with any transfer under such terms, might still be ex post welfare reducing.

But even if transfers of utility between the two individuals on those terms are capable of bringing the ex post planner around to the policy (*cum* transfer), imperfect transfers still may not eliminate the divergence between ex post and ex ante social welfare analysis. The dis-

agreement will merely be relocated. The ex ante social planner may well prefer not to make the kind of costly transfer necessary to make the policy acceptable to the ex post planner.

This is easiest to see in the special case where the transfer is restricted to be the same across the two states. In that case, the transfer reduces both individuals' expected utilities and so reduces (Paretian) ex ante social welfare. To wit, in return for every util taken from individual 1 when he is the winner (in the second state), the same individual 1 receives only 1/3 of a util when he is the loser (in the first state). The states being equally likely, individual 1's expected utility changes by $(1/2)(1/3) + (1/2)(-1) = -1/3$ for each unit taken from the winner (simultaneously in both states). By the symmetry of the example, the same is true for individual 2. The same logic can be generalized beyond the special case of equal transfers across states.[82]

Why might the bucket leak? The simplest answer is that there are administrative costs. Every practical means of transferring across individuals—such as unemployment insurance, disability insurance, litigation, welfare policy, and taxation—involves the expenditure of substantial social resources in administration.

Another reason that the bucket might leak concerns differences in marginal utilities for the transfer medium. Utility per se is never physically transferred. Rather, some resource—a medium of transfer—is what passes between the individuals. The marginal utility of this medium may vary across individuals.

Suppose, in particular, that the loser in each state not only suffers a drop in the level of his utility, but also a decline in his marginal utility of money. If losses are purely economic, a drop in marginal utility is admittedly unlikely, given the usual assumption of declining marginal utility of income. On the other hand, it is also unlikely that all losses are economic. We might well imagine, for example, that the losses caused by the policy are health related and that the suffered condition reduces the marginal utility of money, as when the loser dies or becomes vegetative. It might well be, then, that the loser's marginal utility of money is 1/3 that of the winner in each state.

In general, a positive correlation between those who lose under the policy (more precisely, those whom we would like to compensate)[83] and those whose marginal utility of money decreases under the

---

[82] See *infra* Part III.A.3 for more general comments regarding this point.

[83] These individuals may differ when the status quo exhibits inequality, as when the loser under the policy is still better off than the winner.

policy—a special problem for policies whose losses are predominantly nonmonetary—has the potential to render monetary compensation cost ineffective at mitigating the ex post disequalizing effects of a policy, and thus incapable of closing the gap between ex ante and ex post social welfare analysis.

A third reason why the bucket might leak concerns the information problems associated with knowing whom to tax and whom to compensate, a problem that has spawned whole literatures in economics, including a broad discourse on "optimal taxation."[84] When taxes and transfers can be made contingent only on individuals' observable choices, as opposed to the unobservable characteristics that drive those choices, such taxes and transfers will distort economic choices and thus cause efficiency losses.

2. Free Transferability of Utility Across States

Even if utility is freely transferable across individuals, there still may be divergence between ex ante and ex post social welfare if utility is not also freely transferable *across states*. Utility is freely transferable across states (for a given individual) if it is possible to transfer that individual's utility in one state into utility for that individual in another state without changing that individual's expected utility across states. For example, if two states have the same probability, then free transferability across states implies that we can transfer utility for this individual, one for one, between these states. If the first state is twice as likely to occur, then free transferability across states implies that we can transfer two units of utility into the second state for every one we take from the first.

The government's means of transfer across states is its ability to force its citizens to "insure" against utility losses, where insurance is broadly defined to include trading across states of the world.[85] The government might force payment of an "insurance premium" in the form of a tax imposed upon the individual in states where her utility is high, paying out the "insurance proceeds" to the individual in the form of a government transfer in the states where her utility is low.[86]

---

[84] *See, e.g.,* Stern, *supra* note 24, at 27-46 (reviewing the theory of optimal taxation).

[85] Because insurance is a trade, we can reciprocally view the government as forcing citizens to sell insurance to it.

[86] We consider *private* insurance and how it affects the government's ability to do this *infra* Part III.C.

A more specific example is government-provided health care financed by progressive taxes. Note that these are, of course, examples of *some* transferability across states, not necessarily *free* transferability. Free transferability requires that the expected value of premiums equal the expected value of payouts.

Free transferability across states is necessary for the convergence of ex ante and ex post social choice in the same way that it is necessary for a risk-averse individual to act as if she is risk neutral. When insurance policies are available that permit free transferability across states, the individual judges uncertainties solely by their expected value, anticipating the purchase of insurance. With free transferability across states, the state judges policies by their effects on expected utility, anticipating the imposition of insurance on individuals.

To make this concrete, consider a two-state/two-individual example wherein the individuals' utilities are equal to each other both in a certain status quo *and* in each of two equally likely states. Suppose that the individuals' utility falls by 100 in the first state and increases by 200 in the second. The policy increases both individuals' expected utility by 50 and so is again ex ante Pareto improving. But the level of individual welfare and total welfare is risky—these vary across states—and an equity-regarding ex post policymaker may be sensitive to this kind of riskiness. In particular, if the social welfare function is such that the marginal impact of individual utility is diminishing, the ex post social planner may find the policy to be social welfare reducing. This would occur if the loss of 100, acting on lower levels of welfare where each utility unit is relatively dear, counts for more than the gain of 200, acting on higher levels of welfare where each utility unit has a relatively low impact.

If, however, the government could force individuals to insure against the bad outcome in the first state, then the ex post planner would also favor the policy. Suppose, for instance, that the government could purchase insurance from Lloyd's of London that enables it to raise each individual's welfare by 150 in the first state in return for lowering each individual's welfare by 150 in the second state. (Because the individuals' fates are correlated, the government would need to access an external insurer of some kind.) In that case, the policy is Pareto improving in each state, increasing individual utilities each by 50. Hence, it would be approved by an ex post planner.

However, just as in the case of intrastate transfers across individuals, disagreements will remain so long as interstate transfers are at all costly—which they might be for similar reasons. While the ex post

planner will not, by hypothesis, favor the policy without interstate transfers, the ex ante planner will disfavor adding any leaky transfer to the policy. Even if only one util is lost in transit across states, so that for each individual 150 utils are transferred away from the good outcome and 149 toward the bad, expected utility for both individuals decreases by 1/2, and the move is ex ante Pareto inferior. The general case is again worked out in the Appendix.

3. General Case

In the Appendix, we show that if utility is freely transferable across individuals and across states, and if policies are ranked by their maximal attainable social welfare when combined with feasible transfers, then ex ante social welfare analysis generates the same ranking of social policies as ex post social welfare.

The proof proceeds in two steps. First, we show that "double free transferability" renders any social welfare function, whether applied ex post or ex ante, ordinally equivalent to utilitarianism, similarly applied. This generalizes, to the case of uncertainty, a well-known result regarding the free transferability of utility across individuals in a *deterministic* setting. Next, we rehearse the well-known result from Harsanyi that ex post utilitarianism is fully equivalent to ex ante.[87]

Examples such as those above suffice to establish the converse proposition that no such ordinal equivalence obtains when either form of free transferability does not hold. More general comments regarding the converse are also possible. First, as shown in Appendix Section I, a Paretian equity-regarding ex ante planner would never agree to a costly transfer, however small the cost, starting from an initial position of ex ante equality. This is because the ex ante effect of every costly transfer scheme can be decomposed into two ex ante welfare-reducing changes in the expected utility vector: (1) a (Pareto inferior) reduction in some individuals' expected utility, and (2) a zero sum transfer from these individuals to the rest. The latter zero sum transfer will be ex ante disequalizing if the individuals begin with identical expected utilities.

On the other hand, equal expected utility does not mean equal ex post utility in any given state. And, depending on the relationship between equity regard and transfer costs, the ex ante planner's ex post

---

[87] *See supra* note 58.

counterpart might well consider a costly transfer worthwhile if it equalized ex post utilities.

More generally, even if we do not start from a position of perfect equality ex ante, the ex ante planner will generally be less willing than her ex post counterpart to pay for costly ex post transfers to the extent that the equity effects of these transfers wash out ex ante. Thus, suppose that individual utilities in two equally probable states are (150, 200) and (250, 100), so that the first individual is better off ex ante. The ex post planner might be willing to institute a costly transfer that in each state takes, for example, 25 utils from the individual who is better off in that state and gives 24 to the other individual. But the ex ante planner will see no point in making this transfer, for its positive ex post equity effects wash out ex ante, leaving only an equal reduction in each individual's ex ante utility of $1/2$ util.

B. *The Suboptimality of Compartmentalizing Equity Concerns*

This Section asks a different question about compensation—one that relates to the division of institutional responsibility for equity. Imagine that regulatory bodies and other governmental agencies generally ignore equity, and simply maximize total utility.[88] A set of transfers is then arranged by the legislature or a specialized body. This body applies an equity-regarding social welfare function, using the ex post approach, and chooses a set of transfers that is optimal, given this function and this approach. Will those transfers overcome any possible divergence between total utility and equity that might arise? More precisely: if the regulatory body is faced with a choice between P and P', such that utilitarianism favors P, but ex post application of an equity-regarding social welfare function favors P', is it necessarily the case that ex post application of that same social welfare function would rank P optimally modified by the compensation body as better (or at least no worse) than P' similarly modified?

The answer to this question is straightforward, given the connection between utilitarianism, double free transferability, and equity re-

---

[88] The point we make in this Section applies as well under conditions of certainty, and so is not directly related to the ex ante/ex post problem. We include it because it is of interest to legal scholars and because it is an immediate implication of the analysis in the previous Section. A standard claim in law and economics is that distributive considerations are best handled by the tax system. This Section undermines that claim. *See generally* Chris William Sanchirico, *Deconstructing the New Efficiency Rationale,* 86 CORNELL L. REV. 1003 (2001) (describing, and criticizing, the view that legal rules should ignore distributive considerations).

gard discussed in Part III.A above and demonstrated in the Appendix. For any equity-regarding social welfare function, the following two propositions are true. First, given free transferability of utility within and across states, the ex post ranking of P and P' conjoined with optimal transfers is the same as the utilitarian ranking of P and P', for all P and P'. Thus, double free transferability enables the institutional separation of efficiency and equity. Conversely, without double free transferability, there will be some pairs of policies P and P' such that the utilitarian criterion prefers P, but the ex post application of the social welfare function prefers P' optimally modified by the compensation body to P optimally modified.

In other words, absent the special conditions of double free transferability, the following scenario can occur: The utilitarian regulatory agency chooses one policy over an alternative; the compensation body then arranges compensation to maximize society's equity-regarding social welfare function, but produces a worse result than if the regulatory agency had been sensitive to equity and chosen the alternative (perhaps subsequently optimized by the compensation body). Absent double free transferability, the institutional separation of efficiency and equity may lead to some choices that reduce social welfare.

C. *Private Insurance*

One might wonder whether individuals' purchase of private insurance helps to mitigate the divergence between ex post and ex ante evaluation. Returning to our earlier example with two individuals and two possible states, each individual might purchase private insurance against the state in which she is the loser. Suppose, for example, that fair insurance is available. Therefore, individual 1 might insure against her 100 utility loss in the first state by buying an insurance policy that pays her in that state in return for a "premium" paid in the second state. The hope might even be that private insurance of this kind would make up for the likely shortcomings of government transfers, both between individuals and between states, as discussed above.

But this story relies on individuals purchasing insurance in a way that pleases the ex post social planner. In fact, there is no reason to think that they would. Indeed, the most important point to make about private insurance is that it may work against the government's attempts to alleviate a policy's ex post disequalizing effects.

In the canonical example, for instance, the ex post planner is concerned with individuals' low utility states and would like to see

them transfer resources to those states through private insurance markets. But why would individuals make their insurance decisions on the basis of the *level* of their utility? In particular, why would they strive to transfer utility from high to low utility states, as the planner would prefer? If individuals are expected utility maximizers, they would rather strive to transfer resources away from states with low marginal utility for the transfer medium (say, money) and toward states with high marginal utility. Once again, the correlation between utility levels and the marginal utility of the transfer medium is a crucial factor. If the low marginal utility states are also low utility level states—as with some health-related losses—then expected utility-maximizing individuals will use private insurance markets to shift resources toward high utility states, contrary to the desire of the ex post planner.

Indeed, private insurance might even undo the government's efforts at reallocating across states. We have already noted the importance of the government's ability to transfer utility across states. With private insurance readily available, the government's transfer of utility to a given state can be undone by individuals. Individuals will rearrange this transfer by moving it to where they prefer, as opposed to where the ex post planner prefers. If the level of utility and the marginal utility of money are correlated, then (under certain plausible conditions) [89] the individual would use government compensation to purchase insurance in such a configuration that the net effect of the compensation would be to increase utility in high utility states more than in low utility states. That is, the individual would rearrange the government compensation in such a way that the compensation makes matters worse for the ex post social planner. In the canonical example, individual 1 will shift any government compensation paid in the first state, where she is the loser, toward the second state, where she is the winner.

IV. EX ANTE PARETO SUPERIORITY VERSUS EX POST SOCIAL WELFARE

We have focused to this point on the possible divergence of ex post and ex ante assessments. Let us now return to the possible conflict between ex post assessments and Pareto superiority.

---

[89] The assumption is that the individuals' absolute risk aversion is not increasing, which is an implication of constant relative risk aversion.

First, it should be underscored that the conflict is with ex ante Pareto superiority. Our focus, throughout this Article, is on the class of equity-regarding social welfare functions that respect the Pareto principle. These functions generate orderings of outcomes such that a Pareto superior outcome is necessarily ranked above a Pareto inferior outcome.

Second, because the equity-regarding social welfare functions that we are interested in respect the Pareto principle for outcomes, both ex ante and ex post applications of such functions under conditions of uncertainty guarantee the selection of policies that are ex post Pareto superior. One policy—that is, one lottery over outcomes—is ex post Pareto superior to another policy if the outcome of the first policy in each possible state of the world is Pareto superior to the outcome of the second policy in that state. Pareto superiority in this sense is respected by the ex post approach to social choice, as well as by the ex ante approach.

Further, it should be stressed that cases of ex ante/ex post divergence in policy choice are not necessarily cases in which the ex ante preferred policy is Pareto superior. In some cases of divergence, the policy preferred by ex ante application of the social welfare function will be ex ante Pareto incomparable to the other policy rather than ex ante Pareto superior. In other words, there will be some individuals who have greater expected utility with the ex ante preferred policy, but others who have lower expected utility.

On the other hand, it is obvious that—because the social welfare function is Paretian—a policy that is ex ante Pareto superior relative to some alternative must be preferred by ex ante application of the social welfare function. The upshot, then, is that *the conflict between ex ante Pareto superiority and ex post valuation is a subset of the general ex ante/ex post divergence.* Of all possible choice situations in which ex ante and ex post applications of a given social welfare function diverge—recommending different policies—a proper subset are choice situations in which ex post application of the social welfare function recommends a policy that is ex ante Pareto inferior to some alternative.

How general is the ex ante/ex post Pareto superior conflict? Bracketing the possibility of insurance and compensation, it is easy to show that *every* equity-regarding social welfare function is susceptible

to this conflict.[90] We can demonstrate this using the notion of a permuting policy, discussed above in Part II.D. As mentioned there, *every* equity-regarding social welfare function gives a greater ex ante than ex post valuation to any permuting policy.[91] Take any permuting policy with the following particular characteristics: each possible permutation of the utility vector corresponds to one and only one possible outcome, and the probability of each outcome is the same. Assume that the ex ante social welfare valuation of the policy is s and the ex post is $r < s$. Because this is a permuting policy, mean utility across individuals within every state is the same (k). Because the possible outcome vectors have equal probability, and all permutations are accounted for, each individual's expected utility is also k. Next, construct a policy where every individual is certain to get k. The ex ante valuation of this alternative is s, as it is for the original permuting policy. The ex post valuation must be greater than r, because the policy equalizes with mean preservation in every state. Indeed, the ex post valuation must also be s, since the policy can lead to only one utility vector (k for each individual) and thus must have the same ex post and ex ante valuations. Now alter this certain-outcome policy by subtracting a small amount (d) from each individual's utility, so that each individual receives $k - d$. If d is small enough, the ex post and ex ante valuation of this new policy will be greater than r, but less than s.[92] Thus, the new policy will be ex post preferred to the original permuting policy, but ex ante dispreferred. Further, because individuals are perfectly symmetric, the original permuting policy must be not only ex ante preferred, but also ex ante Pareto superior.[93]

---

[90] More specifically, this result holds for all continuous equity-regarding social welfare functions.

[91] *See supra* note 75 and accompanying text.

[92] This assumes that the social welfare function is continuous in the vector of individual utilities, so that a small change in that vector produces a small change in the social welfare function value.

[93] This result is an implication of Harsanyi's theorem, *see supra* note 58 and accompanying text, combined with the assumption of continuity. Stated in contraposition formulation, Harsanyi proved that, if the social welfare function is not an affine function of individual utilities (the case with an equity-regarding social welfare function), then there exist two policies to which all expected-utility-maximizing individuals are indifferent, but one of which the social planner strictly prefers to the other. Adding the continuity of the social welfare function, the conclusion here implies the existence of two policies, one that is strictly preferred by all individuals, and another that is strictly preferred by the social planner. Take the indifferent pair of policies given to us by Harsanyi and reduce each individual's utility in the socially preferred policy by an amount small enough not to upset the social ranking. All individuals will now strictly disprefer the modified policy that the planner strictly prefers.

For reasons discussed in Part III above, government compensation and private insurance may well be ineffective at closing the gap between ex ante Pareto superiority and ex post welfare valuation. Compensation, for example, would affect the argument above only if it were capable of raising the ex post valuation of the uncertain policy to equal that of the certain policy. But if the "bucket leaks" sufficiently, compensation might well be ineffective at this task. Compensation might be capable of increasing the ex post valuation only somewhat— or perhaps not at all.

In sum, the ex post application of any equity-regarding social welfare function will produce a conflict with ex ante Pareto superiority, in some choice situations, even if social and individual probability assessments are identical (which we are assuming throughout the Article). This is true even with insurance and government compensation in play (unless utility is fully and freely transferable across individuals and states).

This conclusion may seem a *reductio ad absurdum* of the ex post approach. Many scholars—certainly mainly legal scholars—reflexively assume that ex ante Pareto superior rules, institutions, or other policies must be normatively attractive.[94] The ex post approach entails the possibility of conflict with ex ante Pareto superiority, while the ex ante approach guarantees that a conflict will not arise. Is that not a decisive argument for the ex ante approach? Now that we have a better appreciation of the nature of ex ante and ex post evaluation, let us turn to this crucial question.

V. WHAT IS THE CORRECT APPROACH TO WELFARIST SOCIAL CHOICE UNDER UNCERTAINTY: EX POST OR EX ANTE?

The utilitarian social welfare function necessarily produces convergence of ex ante and ex post assessments. Assuming that the social planner and all individuals in the population assign the same probabilities to outcomes, there can never arise a choice situation in which ex post and ex ante application of the utilitarian social welfare function select different policies. But the egalitarian case against utilitarianism—the case for a shift to some kind of equity-regarding social welfare function—is powerful. It is this shift that requires a choice between ex ante and ex post approaches to social decision making

---

[94] *See* Matthew D. Adler, *The Puzzle of "Ex Ante Efficiency": Does Rational Approvability Have Moral Weight?*, 151 U. PA. L. REV. 1255, 1255 n.1 (2003) (citing sources of law and economics work that "explicitly invoke[] the criterion of 'ex ante efficiency'").

under uncertainty. As we have fleshed out in detail in preceding Parts, any equity-regarding social welfare function (as defined by the Pigou-Dalton condition)[95] leads to divergence of ex ante and ex post evaluation in some choice situations. Further, in a subset of these choice situations, a conflict between ex post evaluation and ex ante Pareto superiority will arise.

So which approach is correct? The case for the ex post approach relies on the "sure-thing principle." More precisely, it rests on the applicability of the sure-thing principle at the level of social choice—as a constraint on the choices of the social planner. In Section A, we describe the sure-thing principle and explain why it demands the ex post approach. In Sections B and C, we defend the principle itself. We show, in particular, that a social planner who violates the sure-thing principle is apt to be disturbingly, irrationally inconsistent over time.

A. *The Connection Between the Sure-Thing Principle and the Ex Post Approach*

The sure-thing principle can be formulated in different ways, all of which are logically equivalent—or at least imply each other given further basic axioms.[96] Our formulation of the sure-thing principle (at the level of social choice) will be as follows: Given two policies, *P* and *Q*, which have the same probability of producing two outcomes that the social welfare function ranks as equally good, changing that probability (making equiproportional compensating changes in the probabilities of other states, but without changing anything else about the policies) or substituting two other equally good outcomes for the original outcomes (without changing anything else about the policies) should not change the ranking of the policies.

Table 6: The Sure-Thing Principle

	State *A*	State *B*
Policy *P*	*O*	*O*+
Policy *Q*	*O*′	*O*++

---

[95] *See supra* note 39 and accompanying text.

[96] *See, e.g.*, BROOME, *supra* note 6, at 90-120 (discussing the sure-thing principle in the context of expected utility theory); KREPS, *supra* note 9, at 103-04 (discussing treatment of the sure-thing principle by Savage, Anscombe, and Aumann).

Referring to Table 6 above, imagine that the social planner prefers policy P to policy Q; that states A and B have probabilities z and $(1 - z)$ respectively; and that the social welfare function assigns equal value to outcomes $O+$ and $O++$. $O+$ and $O++$ might be identical utility vectors; or they might be reorderings of each other; or they might be wholly distinct vectors that the social welfare function ranks as equally good. Then if z is changed to y, the social planner must prefer the altered P to the altered Q. Similarly, if $O*$ and $O**$ are substituted for $O+$ and $O++$, respectively, and $O*$ and $O**$ are equally good outcomes, then the planner must prefer the altered P to the altered Q. If instead the social planner is initially indifferent between P and Q, then after a substitution of y for z or $O*$ and $O**$ for $O+$ and $O++$, the planner must remain indifferent between the altered policies.

It is immediately obvious that the ex post approach to social choice satisfies the sure-thing principle. This is because the ex post formula is additively separable across outcomes. It takes the form $p^A w(O^A(P)) + p^B w(O^B(P)) + \ldots$, where p^A is the probability of state A and $O^A(P)$ is the outcome of state A given policy P. If this additive formula ranks one policy over another, increasing the probability of some state with indifferent outcomes across the policies (with an equiproportional compensating change in the probabilities of other states), or changing the w values of the indifferent outcomes, will not alter the ranking.

A somewhat less obvious, but crucial, point is that ex ante application of an equity-regarding social welfare function necessarily violates the sure-thing principle. The reason implicates our discussion in Part II.A, above, of precisely why it is that ex post and ex ante welfare values can diverge. Suppose we take a given policy and change the utility vector in some state A, transferring 1 unit of utility from individual 1 to individual 2. Suppose further that we change the utility vector in no other state. The sure-thing principle demands that the social welfare status of this state A change be independent of the utility vectors in the unaltered states $B, C, D, \ldots$. But this will not be the case in an ex ante calculation of social welfare. The transfer in state A effects, as we have previously noted, a transfer of $p^A 1$ units of *expected* utility from individual 1 to individual 2. If this equalizes expected utilities, then the transfer improves ex ante social welfare, given equity regard. If it disequalizes expected utilities, it lowers ex ante social welfare. Whether it is equalizing or disequalizing depends on whether individual 1's expected utility was initially larger or smaller than individual 2's. This, in turn, depends—and this is the point—not only on utili-

ties in state A but also on utilities in states B, C, D, ... If initial utilities in A, B, C, D, . . . are such that individual 1 starts out with greater expected utility than individual 2, then the transfer in state A is ex ante equalizing. Replace the utility vectors in B, C, D, . . . with vectors that are different enough to make individual 2's expected utility initially greater, and the transfer in state A is now ex ante disequalizing.

A different way to see why the ex ante approach violates the sure-thing principle draws on our above analysis of ex ante/ex post divergence in the ranking of permuting policies.[97] Take any permuting policy. As noted, any equity-regarding social welfare function must assign the permuting policy a greater ex ante than ex post valuation. Reorder the vectors in each state so that individuals are ranked the same across states in terms of their utility levels. The sure-thing principle (coupled with symmetry) requires that the planner be indifferent between the permuting policy and the reordering. This condition is met for the ex post social planner but not for the ex ante planner. The ex post valuation of the reordering *is* the same as the ex post valuation of the original permuting policy, because the reordering does not change anything for the ex post calculation. Thus, the ex post planner is indeed indifferent between the two policies. However, because the reordering has the same utility vector in each state, the new policy has no uncertainty, and its ex ante value must equal its ex post value. Therefore, the ex ante value of the permuting policy must exceed the ex ante value of the reordering. So a social planner following the ex ante approach will prefer the permuting policy to the reordering, in violation of the sure-thing principle.[98]

---

[97] For that analysis, see *supra* Part II.D.

[98] The sure-thing principle represents a *strongly* consequentialist approach to social choice. As we explained earlier, *see supra* note 19 and accompanying text, both the ex ante and ex post application of a social welfare function are weakly consequentialist, in the sense that the ranking of a policy is wholly determined by its possible outcomes, and that two policies yielding identical utility (i.e., well-being) vectors in any given state must be ranked as equal. The sure-thing principle is more robustly consequentialist than this. To understand the point, it is crucial to see that outcomes can be nonidentical—in the sense that the well-being of some particular individuals varies—but equally good. The utility vector $(1, 0, 0, 1)$ is identical to itself; it is equally good as, but not identical to, the vector $(0, 1, 1, 0)$, because the identity of the individuals who end up at utility levels 0 and 1 is varied. In short, reorderings of a utility vector, given the symmetry principle, always yield equally good but nonidentical vectors. And, if the social welfare function is the sum of the square root of utility, the vector $(1, 0, 0, 1)$ is equally good as—but clearly not identical to—the vector $(4, 0, 0, 0)$. The sure-thing principle says, in effect, that all that matters about a state-conditional outcome is its goodness. Swapping one possible outcome of an uncertain choice for an equally good outcome (albeit one in which particular individuals may achieve different levels of

B. *Is the Sure-Thing Principle Attractive?*

We have not yet defended the sure-thing principle at the level of social choice. We have simply clarified that the ex post approach satisfies this principle, while the ex ante approach violates it.

Should we accept the sure-thing principle at the level of social choice? It should be stressed that the issue here is normative, not descriptive. Much empirical work in political science and psychology has shown that the principle has tenuous descriptive credentials in tracking the actual behavior of governments or individuals.[99] But here we are asking: how *should* governmental officials behave? Is the most attractive normative account of governmental choice under conditions of uncertainty one that embraces the sure-thing principle?

As already explained, the proponent of the ex ante approach must answer "no" to this question. That puts the ex ante proponent in an uncomfortable position. At least as typically formulated, the ex ante approach rejects the sure-thing principle at the level of social choice, but embraces it at the level of individual well-being. Why? Remember that the ex ante approach applies the social welfare function to individual *expected utilities*. In other words, it assumes that the ranking of different policies for a given individual can be determined by using the standard expected utility formula, multiplying the probability of each possible outcome by the utility of the outcome for that individual, and then adding up the results. Given two policies, which in some state will produce outcomes that are equally good for a given individual, the sure-thing principle stipulates that changing the probability of the state (with equiproportional compensating changes in the probabilities of other states), or substituting another pair of outcomes that are equally good for that individual, must not alter the

---

well-being) should leave the decision maker indifferent. And altering a pair of policies by exchanging a pair of equally good state-conditional outcomes in some state with another such pair should leave the decision maker with the same ranking of the policies. The ex post approach is consistent with these strongly consequentialist axioms, while the ex ante approach is not.

[99] A basic result in expected utility theory is that, if an individual's preferences satisfy the sure-thing principle plus a few other very basic axioms, her choices can be represented as maximizing the expectation of some utility function. But much work in psychology and empirical economics, going back to Allais, shows that individuals often do not behave as if they are maximizing the expectation of some utility function. *See, e.g.*, REID HASTIE & ROBYN M. DAWES, RATIONAL CHOICE IN AN UNCERTAIN WORLD: THE PSYCHOLOGY OF JUDGMENT AND DECISION MAKING 289 (2001) (surveying scholarship on human decision making and concluding that "subjective expected utility theory is not a valid descriptive theory of human behavior").

ranking of the policies for that individual. Indeed, the expected util-
ity formula for individual well-being satisfies this requirement.

The question at hand, then, is *not* whether we should wholly reject
the sure-thing principle. Rather, the question is whether we should
adopt a hybrid position, as the ex ante approach does, accepting the
sure-thing principle at the level of individual but not social choice; or
instead wholly embrace the principle, accepting it at both levels—as
the ex post approach does. Arguments that flatly deny the rational
force of the sure-thing principle do not, in fact, provide ammunition
for the defender of the ex ante approach.[100] What she needs, instead,
are more nuanced arguments, which would show that the principle is
a rational constraint on individuals but not on society as a whole.

1. Diamond's Example

Is it possible to bite the bullet and defend the ex ante approach by
rejecting the sure-thing principle at the level of social choice, but em-
bracing it at the level of individual well-being? Yes. Indeed, this is just
the strategy that the economist Peter Diamond suggests in a well-
known article criticizing Harsanyi's argument for utilitarianism.[101]
Diamond provides the following as a counterexample to the sure-
thing principle at the level of social choice:

---

[100] In fact, we believe that the case for the sure-thing principle at the level of indi-
vidual choice—again, as a normative, not a descriptive principle—is powerful. The
argument against the individual-level principle sometimes invokes a preference for
certainty. Imagine that (x, y, z) is a lottery, which yields the individual utility x in the
first possible state, y in the second, and z in the third. It is alleged that an individual
might rationally prefer (5, 3, 1) to (4, 4, 1) but (4, 4, 4) to (5, 3, 4) in violation of the
sure-thing principle, giving extra weight to the certainty that (4, 4, 4) provides. But
this certainty will either be reflected in the individual's mental states in the different
outcomes, in which case the correct vector would be $(4+c, 4+c, 4+c)$, where c is the in-
dividual's benefit from knowing what the benefit would be; or it will not be, in which
the preferences just described *are* irrational. *See* BROOME, *supra* note 6, at 97-98. Note
also that violating the sure-thing principle at the level of individual choice would lead
to time inconsistency, for the same reason that doing so at the level of social choice
does. *See infra* Part V.B.2-3 (introducing and discussing the concept of time inconsis-
tency).

However, because the ex ante approach and the ex post approach disagree about
the sure-thing principle at the level of social choice, not individual choice, our defense
of the sure-thing principle, below, will focus on the social level.

[101] *See* Peter A. Diamond, *Cardinal Welfare, Individualistic Ethics, and Interpersonal
Comparison of Utility: Comment*, 75 J. POL. ECON. 765 (1967).

Table 7: Diamond's Example[102]

	State A $p=.5$	State B $p=.5$	Expected Utilities
Policy P	$(1, 0)$	$(1, 0)$	$(1, 0)$
Policy Q	$(1, 0)$	$(0, 1)$	$(1/2, 1/2)$

Diamond points out that the sure-thing principle plus welfarism and anonymity/symmetry require that the social planner be indifferent between the two policies. Welfarism with anonymity/symmetry requires that the planner be indifferent between outcomes (i.e., utility vectors) $(1, 0)$ and $(0, 1)$. Therefore, the sure-thing principle requires that substituting one for the other in state B, without changing the outcome in state A, not change the planner's assessment of social welfare. The social welfare impact of making individual 2, rather than individual 1, the beneficiary in state B (as happens in switching from policy P to policy Q) should be nil. In particular, it should not turn on which individual is the beneficiary in state A.

Diamond claims, quite tersely, that this conclusion is counterintuitive. Policy Q "seems strictly preferable to me," he says, "since it gives [the second individual] a fair shake while [policy P] does not"—as captured by the equality of expected utilities under Q but not P.[103] Diamond concludes: "I am willing to accept the sure-thing principle for individual choice but not for social choice, since it seems reasonable for the individual to be concerned solely with final states while society is also interested in the process of choice."[104]

We find Diamond's argument, such as it is, unpersuasive, for reasons elaborated by John Broome. Diamond appeals to "fairness"— specifically, a concern for a "fair process"—as a grounds for preferring the ex ante approach to the ex post. But in fact, "fair process" concerns, if persuasive, move us beyond welfarism and argue against *both* the determination of social choice by ex post application of a so-

---

[102] This table does not conform to the matrix representation of policies that we introduced earlier in the Article. Nevertheless, its meaning should be clear.

[103] Diamond, *supra* note 101, at 766.

[104] *Id.* A similar sentiment is expressed in Roger B. Myerson, *Utilitarianism, Egalitarianism, and the Timing Effect in Social Choice Problems*, 49 ECONOMETRICA 883, 884 (1981).

cial welfare function *and* the determination of social choice by ex ante application of a social welfare function. The sort of information incorporated by a social welfare function—namely, the different utility (well-being) vectors that would occur in different possible states— leaves out the sort of nonwelfarist procedural facts to which a fairness assessment would be sensitive. The fact that ex ante utilities are equal, and the fact that individuals have been treated fairly, coincide only haphazardly.

Broome provides the following sort of example.[105] Imagine that Diamond's Table represents the allocation of an indivisible item between two individuals, who are equally "deserving" of it. The individual who gets the item will receive utility 1; the other individual will receive utility 0. Imagine that a fair coin is flipped and the first individual has won the flip. We now consider two "policies": giving the valuable item to the first individual, or flipping the coin again. Giving the valuable item to the first individual corresponds to policy *P* in the table; flipping the coin again corresponds to policy *Q*. Policy *Q* equalizes expected utilities; *P* does not. But, intuitively, *Q* is no fairer than *P*. By what theory of "fair process" is flipping a coin twice fairer than flipping it once? To engage in "fair process" means to follow some stipulated procedure, perhaps a randomizing one, for allocating benefits and burdens. Once the procedure has been run, the "policy" of conforming to the allocation chosen by the procedure (policy *P* in this case) may well fail to equalize expected utilities, relative to some other policy (in particular, the policy of rerunning the procedure). Another way to put the point is that anyone who believes in "fair procedure" must imagine a procedure that ends at some juncture, producing an allocation—and at that juncture, the "fair" policy and Diamond's policy of equalizing expected utilities will diverge.

Alternatively, imagine that the social planner has not done anything yet to allocate the item. The first individual is evidently shorter than the second. Neither is evidently less Semitic.[106] Moreover, both are equally likely to turn out to be the least Semitic upon further investigation. The social planner is choosing between policy *P* (which says to allocate the item to the shorter individual) and policy *Q* (which says to undertake genealogical research and allocate the item to the individual who turns out to have less Jewish ancestry). Could one still then argue that *P* is less fair than *Q* because *Q* equalizes expected utili-

---

[105] John Broome, *Uncertainty and Fairness*, 94 ECON. J. 624, 628 (1984).

[106] Again, this example is based on Broome's article. *See id.*

ties? At best, both policies are equally arbitrary and "unfair." Indeed, one might argue from the comparative history of discrimination against tall people and people of Semitic origin, that *Q* is less fair than *P*. The point is that the decision criteria themselves drive fairness. The fact that application of one of the criteria in this case is initially uncertain seems irrelevant.

In short, the defender of the ex ante approach needs an argument that is *welfarist*—one that accepts that outcomes can be represented by utility vectors, thus summarizing only welfare information and ignoring nonwelfare facts—and that further shows why the sure-thing principle applies at the level of individual well-being, but not at the level of social choice. Diamond's argument, by appealing to fairness concerns that (if cogent) move us beyond welfarism, does not fit the bill.

2. Welfare Contractarianism

Is there any argument that would fit the bill? Perhaps. We can imagine a kind of welfarist and "contractarian" defense of the ex ante approach that would satisfy the conditions just articulated. It would say the following: (1) Social choices should be evaluated by asking what the affected individuals would hypothetically agree to. (2) Each individual is assumed to care only about her own well-being. (3) The determination of individual well-being under conditions of uncertainty satisfies the sure-thing principle (i.e., equals expected individual utility). (4) Social choice should respect ex ante Pareto superiority, because individuals would contract to ex ante Pareto superior policies. (5) Social choice among policies that are ex ante Pareto incomparable should be determined using a social welfare function applied to expected utilities.

Contractarian approaches to social choice are, of course, worth serious consideration. They have their roots in Locke and Hobbes and were revitalized in modern political theory by John Rawls. But it is not clear whether this philosophical tradition in fact provides the resources to defend the kind of *welfare* contractarianism that the ex ante approach represents. Rawls himself is not a welfarist,[107] nor is

---

[107] Rawls focuses on the distribution of "primary goods," such as income, rather than on the distribution of well-being itself. JOHN RAWLS, A THEORY OF JUSTICE 92 (1971). This may seem like a small difference, but, in fact, the divide between theorists who see well-being as the "currency" for distributive justice and theorists who see "primary goods," "resources," "capabilities," or other such constructs as the correct "cur-

Tim Scanlon, the most prominent living contractarian philosopher, who stresses that individuals in actual or hypothetical contracting scenarios might well have non-self-interested aims.[108] In short, the proponent of the ex ante approach who defends it as a variant of contractarianism must show that—*pace* the leading contemporary contractarians—contractarianism is best specified along welfarist lines.

Perhaps this showing can be made. But the proponent of the ex ante approach must also confront a second, perhaps even more important difficulty: time inconsistency. Rawls and Scanlon avoid time inconsistency by using the construct of a hypothetical contract, not as a criterion for governmental or individual choice, but as a higher-order criterion that is used at a single hypothetical moment in time to identify the moral rules and principles that will be applied by individual and government actors over time.[109] By contrast, the ex ante approach is time inconsistent, as we shall now discuss at length.

---

rency" is a major theme in contemporary philosophical scholarship. *See generally* Arneson, *supra* note 5.

[108] *See* T.M. SCANLON, WHAT WE OWE TO EACH OTHER 213-18 (1998); *see also* Claire Finkelstein, *A Contractarian Approach to Punishment, in* THE BLACKWELL GUIDE TO THE PHILOSOPHY OF LAW AND LEGAL THEORY 207, 214-15 (Martin P. Golding & William A. Edmundson eds., 2005) (arguing that contractarians would deny that basic social institutions are justifiable merely because they increase each individual's expected well-being).

[109] On this feature of Rawls's and Scanlon's views, see Matthew D. Adler, *Against "Individual Risk": A Sympathetic Critique of Risk Assessment,* 153 U. PA. L. REV. 1121, 1232-37 (2005) (distinguishing between act and rule contractualism and discussing Rawls and Scanlon). In effect, Rawls's and Scanlon's contractarian theories have a two-tiered structure, similar to the two-tiered structure of rule utilitarianism. The hypothetical contract picks out the correct set of moral rules; those rules then guide choices over time. This approach avoids time inconsistency, at least potentially, because the rules under consideration by the hypothetical contractors might be limited to time-consistent rules.

Why not salvage the ex ante approach by giving it a similar two-tiered structure? Throughout this Article, we have understood both the ex ante and ex post approaches as choice criteria that will be applied, directly, to governmental policy choices over time. But why not reframe the approaches as higher-order criteria for choosing moral rules? The short answer is that a two-tiered structure is, at most, weakly welfarist; if the chosen moral rule does not itself focus solely on well-being, then the theory is hardly welfarist in the robust sense meant by Kaplow and Shavell and other proponents of welfarism. Kaplow and Shavell and other contemporary welfarists believe, we take it, that the morally relevant features of particular governmental choices include only welfare information. Within this normative tradition, the social welfare function is understood as a device that focuses solely on welfare impacts, and that is directly applicable to governmental choices—not merely to some initial, and hypothetical, choice of moral rules.

3. Time Inconsistency of the Ex Ante Approach

Perhaps the most compelling affirmative argument for the sure-thing principle, and against the ex ante approach, is that violating the sure-thing principle makes the decision maker unsettlingly inconsistent over time.[110]

A time-inconsistent decision maker, facing a foreseeable and fully specified contingency, prefers one course of action for that contingency when she views the contingency prospectively, and prefers another when and if the contingency actually arises. She plans, for instance, to take her umbrella to work if it rains the next morning, but when it does rain, she dons her rain poncho instead. Furthermore, if the time-inconsistent decision maker has taken certain actions up front in anticipation of a contingency, she may try to undo those actions when the contingency arises. If she purchases an umbrella on the way home from work the evening before, she may try to return it on her way to work the next morning, having decided that her rain poncho is the best option. Indeed, the time-inconsistent decision maker may well realize upfront, at the time she formulates a contingent plan, that she will later resist that plan and try to undo it should the contingency for which it is designed materialize. That is, even while shopping for an umbrella the evening before, she may well realize that she will be returning it on her way to work the next morning. In fact, realizing that her future contingent self will have different ideas, the time-inconsistent decision maker may attempt to "tie herself to the mast." She may destroy her rain poncho the evening before, or she may set her alarm fifteen minutes later, so that she will have no time to return the umbrella.

To be sure, the rational decision maker may "change plans" based on new information. But if she anticipated that such information might arise, then she is "changing plans" only in the sense that her original plan called for changing actions in the event that she received this new information. That is, her original plan was a contingent plan, and that contingent plan has not changed. Furthermore, a rational decision maker may not anticipate a contingency in the first place. When that contingency arrives, she may revise her contingent plan by adding yet one more contingency to the array of possible contingen-

---

[110] On time inconsistency and the choice between ex post and ex ante social welfare, see Hammond, *Ex-Post Optimality, supra* note 6, at 177-78 (advocating the ex post approach over the ex ante approach on the basis of the finding that only the former is time consistent for general social welfare functions).

cies, and then deciding what to do therein. And, of course, a rational decision maker may change her contingent plan if her preferences over outcomes change between the time that she adopts her contigent plan and the time that contingencies arise. The time-inconsistent decision maker, on the other hand, rejects a contingent plan upon occurrence of a fully anticipated contingency even when her preferences over outcomes remain the same. A time-inconsistent decision maker is, in a particular but very pointed sense, irrational.

It is easy to see that departures from the sure-thing principle—in particular, the ex ante application of a social welfare function—can be time inconsistent. We will first consider the issue in general terms and then present two specific examples.

When we first encounter the ex ante social planner, before the actual social state is revealed, we may find her, as in many of our examples, unbothered by the fact that a policy produces ex post inequalities, because all individuals have equal prospects under the policy. In making this judgment, she is violating the sure-thing principle because she is evaluating the outcome in each state not independently, but with reference to outcomes in other states. When we meet her again, however, after she has implemented the policy, and after the state has been revealed, she may be singing a different tune. The other state(s) that did not occur, and that were informing her choice ex ante, are now irrelevant to her view regarding the current state of affairs. States that did not occur are now, in effect, zero probability states and, consequently, do not figure in her calculation of individuals' expected utilities. In this after-the-fact world, then, inequality in the policy outcome is not mitigated by the possibility that the reverse inequality might also obtain. The possibility of the reverse inequality has passed.

In essence, the sure-thing principle requires that the contribution of each possible outcome to social welfare be separately evaluated—that it not turn on the nature of other possible outcomes. The resolution of uncertainty over time operates to separate the outcomes that actually occur from what else was possible. If the planner was not originally treating the outcomes separately, she will be eventually, and this may cause odd shifts in her attitudes and behaviors.

For a simple example, imagine that the social planner at time T_0 faces a status quo that, like policy Q in the Diamond example,[111] involves two possible and equally probable states, leading to outcomes

---

[111] *See supra* Part V.B.1 tbl.7.

with unequal utilities, but producing equal expected utilities across individuals. At time T_1, one of the states will be revealed to be the actual state. Redistribution of welfare within a given state is possible, at some cost. The ex ante social planner at T_0 may choose to take a costless "hand-tying" measure that will prevent her from redistributing utility at T_1. If her hands are not tied at T_0, then at T_1, whichever state is revealed, she will want to redistribute (assuming that the costs of redistributing are sufficiently low). Forecasting her own unfettered behavior at T_1, however, the ex ante social planner will undertake the hand-tying measure at T_0. From her perspective at T_0, the future redistribution "solves" a nonexistent equity problem at positive cost; the positive equity effects of such compensation wash out across states from her perspective, leaving cost without benefit. But having thus tied her hands at T_0, the planner will regret having done so when she reaches T_1 and the state is revealed.

A modification and extension of Broome's example of repeated randomization[112] provides an even more striking example of the irrationality of the ex ante approach. Suppose there are two individuals and an indefinite number of periods. In the first period, the social planner flips a coin. After seeing the outcome of the toss, the planner may choose to stop or to flip again. If the planner chooses to stop, and the coin lands on heads, then the first individual gets 100 utils and the second gets nothing. If the coin lands on tails, then the second individual gets 100 utils and the first gets nothing. If the planner chooses to flip again, she "pays" 1 util for the right to do so, thus leaving 99 utils left to allocate. She then flips the coin again. Seeing the outcome, she decides whether to stop or flip again. If she decides to stop, and the coin lands heads, the first individual gets all 99 remaining utils. If it lands tails, the second individual gets the 99. If the planner decides to flip again, then the total is reduced to 98, and so on.

When will the ex ante planner stop flipping? The answer depends in part on her degree of equity regard relative to the cost of flipping again. Under very plausible assumptions, she will continue to flip until almost nothing of the 100 utils is left. Consider, for example, the basic social welfare function, $\left(u_1 + u_2\right) - \frac{1}{2}\left|u_1 - u_2\right|$, which is Paretian and equity regarding. As shown in detail in the Appendix, an ex ante planner employing this social welfare function will continue to toss the coin period after period until only 2 utils remain to be allocated.

---

[112] *See supra* text accompanying note 105.

Until this point, the social welfare loss from reducing total utility by 1 util (and so reducing social welfare by 1 unit) is less than the social welfare gain from equalizing prospects (which gain equals half of the total that could currently be allocated).

C. *Violating Ex Ante Pareto Superiority*

And what of the fact that the ex post approach conflicts with ex ante Pareto superiority? As we demonstrated above in Part IV, every equity-regarding social welfare function, if applied ex post, diverges from ex ante Pareto superiority in some choice situations. This means that, in some choice situations, the ex post planner will stand alone against the unanimous wishes of her constituency. Isn't that a decisive objection to the ex post approach?

We think not. Imagine that our normative theory tells us that social choice should be welfarist and sensitive to equity considerations. Imagine, too, that we are further inclined to think that welfarist social choice, like individual choice, should respect the sure-thing principle—that it should be consequentialist rather than "contractarian." At this point, we learn that equity-regarding social choice that respects the sure-thing principle will conflict with ex ante Pareto superiority. The social planner, who cares about the distribution of well-being in each possible outcome, might pick a policy different from that which each individual, maximizing her expected well-being, would select. But this occurs because the individuals are each focused on their own well-being. None of them cares about the distribution of well-being in different possible outcomes (indeed, none of them cares about equity at all), while the social planner does. Because this is how the conflict between the ex post approach and ex ante Pareto superiority arises— as a conflict between the aggregate of individual self-interested perspectives and a social perspective that is impartial and equity regarding—why would that conflict persuade us to abandon the ex post approach?

It might be objected that "fair process" concerns cut in favor of ex ante Pareto superiority. Doesn't a commitment to democratic processes require the social planner to choose ex ante Pareto superior policies? But this is to make the same mistake that Diamond does.[113] To plump for "fair democratic process" as a determinant of social

---

[113] *See supra* text accompanying notes 105-106 (criticizing Diamond's argument that "fair process" considerations argue for an ex ante approach).

choice is to move beyond welfarism. The outcome of a unanimous democratic process may not be ex ante Pareto superior (because the individuals involved may be motivated by considerations other than their own well-being), and a policy can be ex ante Pareto superior, by virtue of maximizing each individual's expected utility, without any actual voting or other democratic choice at all.

It might also be objected that ex ante Pareto inferior options will be "unstable." If the social planner chooses the ex ante Pareto inferior option, the affected individuals will overturn it, or at least will make efforts to do so, with social costs. But no actual choice can ever be "overturned"; at most, the affected individuals can force the planner to make a later choice that is similar in some respects to the earlier, Pareto superior one. Moreover, many Pareto inferior choices will not be "overturned" even in this weaker sense, given collective action problems among the individuals. Thus, it is a contingent question how costly or unstable ex ante Pareto inferior choices will be. Those costs will show up in the possible outcomes of the Pareto inferior choice. They may or may not be large enough to sway the ex post planner toward a different option. In any event, ex post social planning will be sensitive to these "instability" costs; their possibility does not amount to an argument against that approach.

A final objection to the ex post approach in terms of its disregard for ex ante Pareto superiority concerns the hypothetical example of a "Robinson Crusoe" economy, consisting of a single individual. In this special case, ex ante Pareto superiority collapses to the maximization of the individual's expected well-being. One policy is ex ante Pareto superior to another if and only if Robinson Crusoe's expected utility under the first policy is greater than the second. Ex post application of some equity-regarding social welfare functions, in the Robinson Crusoe case, may reject the policy that maximizes Robinson's expected well-being.

How could this occur? If the welfare function is, for example, strictly concave—so that individual utility has diminishing social welfare returns—then the social planner will, in effect, be more risk averse than the individual. Because utility has a diminishing positive impact on social welfare, the planner will be less impressed than the individual with the potential utility upside of any risk. The planner may, therefore, disagree with the individual regarding which policy risks are worth taking.

Isn't this perverse? One arguable response is that social welfare functions are inapplicable in a Robinson Crusoe economy. Social wel-

fare functions represent different moral views and therefore, arguably, should come into play only when there are multiple individuals. This answer may or may not be persuasive in the Robinson Crusoe case. Even if it *is* persuasive in that hypothetical setting, it does not answer a related case—one in which the policies have identical welfare impacts on all individuals, meaning that, for a given policy, each individual's utility in any state is the same as every other individual's utility in that state. The ex post application of an equity-regarding social welfare function, even in the case of identical individuals, may reject a policy that is ex ante Pareto superior.

What are we to make of this objection to the ex post approach? To begin, some readers may find the choice of an ex ante Pareto inferior policy less counterintuitive in the many-person, identical-individual case than in the Robinson Crusoe case. Consider the prioritarian view that individual welfare has diminishing marginal moral weight. If one is wedded to that view, the government's imposition of the Pareto inferior result in the many-person case might well seem like an instance of morally justified paternalism. Morality requires more risk aversion than the individuals display.

But for those who are troubled by these cases—and that reaction is plausible—we stress that the cases are *not* an objection to all equity-regarding social welfare functions. They involve a particular *kind* of equity-regarding social welfare function: one that is strictly concave everywhere—that is, as between any two utility vectors.[114] Such a social welfare function (no matter how small or large the population) imbues the social planner not only with equity regard, but also with an extra degree of risk aversion—one that operates over and above the potential risk aversion of individuals and that exists independent of the planner's preference for equity. Not all equity-regarding social welfare criteria are everywhere strictly concave. For instance, rank-weighted social welfare functions—which, as discussed above in Part II.B, take a variably weighted sum of individual utilities with larger weights for individuals who are less well off—are not everywhere strictly concave, and yet *are* equity regarding. These social welfare functions add no extra societal risk aversion to social choice—at least not apart from their preference for equity. In the special case of a single individual, or of multiple individuals who have the same utility

---

[114] It is thus an objection to prioritarian equity-regarding social welfare functions, but it extends as well to social welfare functions that are strictly concave but not additively separable in individual utilities.

as each other in any given state, they evidently collapse to simple utilitarianism. So even applied ex post, they never disagree with the individual(s) in these cases.

There are many equity-regarding social welfare functions. Some iron-rice-bowl functions might seem to give extreme weight to equity. Some let-them-eat-cake functions might seem to give extreme weight to efficiency. Some, in addition to displaying equity regard, display a societal level of risk aversion, which may also independently cause such functions to violate ex ante Pareto superiority. We do not here argue for the acceptability of the entire set of equity-regarding social welfare functions. At most, we argue only that the social criterion should be located somewhere in that set. And, in fact, the main thrust of our argument is conditional: *if* an equity-regarding social welfare function is adopted, it should be applied ex post rather than ex ante.

*   *   *

To summarize: welfarism is best specified as accepting the sure-thing principle at the level of social as well as individual choice, and thus as accepting the ex post approach to social choice under uncertainty—notwithstanding the conflict with ex ante Pareto superiority. This conclusion, like virtually any normative position, is contestable, but we find it very plausible and hope the reader does as well. Even if she does not, it is critical to understand the logical connections between the sure-thing principle, welfarism, nonwelfarism, individual choice, social choice, equity regard, ex post Paretianism, and ex ante Paretianism. We hope that this Part, along with the preceding ones, has helped illuminate those connections.

VI. LEGAL APPLICATIONS

This Article has sought to contribute to an important current debate in legal theory: that regarding the justifiability and meaning of "welfarist" approaches to law. We have argued for the following, very general position: *if* social choice should be welfarist and equity regarding, then (ideally) decision makers should handle uncertainty by applying the social welfare function in an ex post rather than ex ante manner. Put another way, policymakers should focus on the prospects for equity, not on the equity of prospects.

We do not attempt to provide concrete policy recommendations. To do so would mean specifying a particular social welfare function. Some would argue that this is a quintessentially legislative task. Whether or not that is true, the problem of specifying the form of an

equity-regarding social welfare function certainly requires much analysis and debate, which a substantial number of scholars in philosophy and economics have undertaken, but which we lack space to pursue here.

Still, the reader might wonder whether the theoretical work undertaken in this Article—the argument for an ex post rather than ex ante approach to welfarist social choice—has relevance to law and policy. This Part aims to allay that concern. Ex ante/ex post divergence can arise in a host of areas. Across a multitude of legal and policy domains, ex post planners may end up choosing courses of action that their ex ante counterparts would reject—which, indeed, may be ex ante Pareto inferior. This Part provides a sampling of illustrative examples drawn from some of the most salient legal and policy issues of our time.[115]

A. *Risk-Reduction Policies for Widely Borne Risks*
(Herein of Terrorists, Asian Birds, and British Cows)

One important difference between the ex post and ex ante approaches to welfarism is this: the ex post policymaker is willing to go farther than her ex ante counterpart in alleviating hazards that expose many to risk, but few to actual injury. Both policymakers may find intervention worthwhile for the usual reasons—externalities, market failures, paternalism, etc. But in the narrow incidence of actual harm,

---

[115] One worry that may persist is that our specifications of the ex ante and ex post approaches to welfarist social choice under uncertainty—throughout the Article, including this Part—assume a highly idealized decision maker. Like many other scholars in contemporary decision theory, social choice theory, and game theory, we have developed a normative theory which is suitable for an "unboundedly rational" chooser—someone who is not omniscient, but who has the cognitive and computational resources to handle uncertainty by considering each choice as a lottery across numerous, highly specified states of the world to which she can attach probabilities. But what about bounded rationality? We do not have a satisfying answer to this question, because no one has yet developed a satisfying account of rational choice for uncertain and boundedly rational decision makers. In particular, the intellectual tools that would take us from an account of social choice for unboundedly rational social planners to a decision procedure for boundedly rational planners—administrators, legislators, judges—have yet to be developed. All we can say here is that basic differences in ideal theory among welfarists—different specifications of the social welfare function, or the difference between ex ante and ex post approaches to choice under uncertainty—would surely be reflected in nonideal theory, in recommendations for boundedly rational planners. But this is just a promissory note. Our analysis provides a general template for how unboundedly rational governmental decision makers should handle uncertainty, assuming welfarism. It does not provide a general template for how real-world officials, given their cognitive limitations, should do so.

the ex post policymaker sees an additional equity problem deserving of additional attention, which her ex ante counterpart, focused as she is on the relative equality of prospects, does not recognize. This will induce the ex post planner to push harder for costly prevention measures that avoid or negate such solely ex post inequality.

The September 11 attacks on New York and the Pentagon, the London subway bombings, the Madrid commuter train bombings, and other terrorist attacks around the globe bring to the fore the question of how much a society should spend on the security of its citizens. If the society cares about equity, the answer depends on whether it takes an ex ante or an ex post view. In particular, an ex post social planner is willing to spend more on security than the ex ante planner, because the ex post planner sees an additional equity problem with a terrorist attack that the ex ante planner does not.

To understand why, imagine a city whose citizens start out identical in every respect, and imagine that every citizen has an equal chance of being injured. The ex ante city planner spends on security until the value of additionally reducing the sum of individuals' expected loss from the attack (a function of the harm from the attack and the probability of an attack) is not worth the additional spending required to do so. The ex post planner spends more than this because the unequal ex post impact of the attack adds insult to injury, as it were, and alleviating this inequality by preventing an attack in the first place is an additional reason to spend on prevention.[116]

The same analysis applies to public spending on vaccinations or food inspections. If policymakers care about equity, the proper amount to spend will depend again on whether they take an ex ante or ex post approach. As with the terrorist attack, the ex ante planner is not troubled by the fact that not everyone shares in the social cost of the virus or the tainted meat. The ex ante planner cares only about ex ante equity, and the fact that a virus or tainted meat may affect everyone with roughly equal likelihood (supposing this to be the case) means that equity plays almost no role in her spending calculations. The ex post planner, however, sees an additional equity problem in the narrow incidence of disease and is thereby willing to spend more on vaccinating and inspecting because, in addition to reducing suffering from the disease, these measures reduce the inequity caused by the disease's random incidence.

---

[116] See Appendix Section K for a more formal discussion of this point.

Daniel Markovits's scholarship on the distributive implications of quarantines and vaccinations makes a related point about risk prevention: just as the amount of spending for risk prevention may depend on the choice between ex ante and ex post approaches, so too may the particular kind of risk prevention program selected.[117] Consider the policymaker deciding whether to prevent a future epidemic through mass vaccination (which will impose time and discomfort costs on everyone, but avoids almost all cases of the disease) and quarantine (which avoids vaccination costs, but results in those infected incurring large costs down the road, including both the direct cost of the illness and the cost of being quarantined). If the quarantine option increases expected overall welfare, then—if everyone is equally exposed to the disease—that option is presumably ex ante Pareto superior and will be selected by the policymaker over the vaccination option. But the ex post policymaker may insist on vaccinations instead, given the inequality in each state between those particular individuals who end up being infected and quarantined, and the rest of the population.

B. *Risk Compensation (Herein of Workers' Compensation,*
Unemployment Insurance, and the Compensated Siting of "LULUs")

Prevention or mitigation is not the only governmental policy option for responding to risks to health, safety, property, or other aspects of well-being. Compensation is another option—either after-the-fact compensation for those who end up suffering harms, or before-the-fact compensation for all who are exposed to the harm. In both instances, the ex ante and ex post approaches may select different policies as optimal.

Consider, first, after-the-fact compensation. Many of these programs exist even when the ex ante prospects of harm are roughly equal. This creates a puzzle for the ex ante approach, but no puzzle at all for the ex post approach. Workers' compensation is a good example. Any one of the workers on the factory floor may be injured on the job. And each may face the same prospect of injury. Why, then, incur the administrative costs of a workers' compensation program, which both reduces overall well-being and (from the ex ante perspective) has no equity benefit?

---

[117] *See generally* Markovits, *supra* note 6 (discussing the distributive implications of quarantines and vaccinations from ex ante and ex post perspectives).

One answer is that workers' compensation substitutes for failures in the insurance market—for example, the failure of actual workers to purchase the insurance that rationally they should, given their risk aversion. But, at best, this argument explains workers' compensation schemes that seek to mimic hypothetical insurance markets. By contrast, ex post planners might support schemes that provide compensation for injured workers above and beyond what they would have insured for. This is because the ex post planner sees workers' injuries as, inter alia, an equity problem within the population of workers, while the ex ante planner does not. From the ex ante planner's viewpoint, there is no equity issue to the extent that each worker is equally likely to be injured. From the ex post planner's perspective, however, an equity problem is certain to arise, as the costs of inevitable injury are actually borne not by all workers evenly, but instead by an unlucky few.

A similar analysis applies to health insurance, to unemployment insurance, and to any government program that compensates harms to which many are potentially exposed but few are actually exposed. Unforeseeable changes in the economy may cause layoffs in certain sectors, for example. The risk of sectoral layoffs of this sort is widespread, but the actual cost is more narrowly focused. Both the ex ante planner and the ex post planner may wish to put in place a system of unemployment insurance. But the ex post planner is likely to fund the program to a greater degree. On the plus side, the ex post planner, like the ex ante planner, may be concerned about the cost of unemployment risk and the lack of affordable private insurance against job loss. On the negative side, the ex post planner, like the ex ante planner, may be concerned about the administrative costs and potential perverse incentive effects of unemployment insurance. But unlike the ex ante planner, the ex post planner places an additional item on the plus side of the ledger: compensating unemployed workers, although costly, solves an ex post equity problem that the ex ante planner does not see.

Consider, now, policies that provide before-the-fact compensation for being exposed to the risk of harm, property damage, or other welfare setbacks. A good example is "compensated siting" schemes.[118] One commonly proposed mechanism for overcoming local resistance

---

[118] *See* Vicki Been, *Compensated Siting Proposals: Is It Time To Pay Attention?*, 21 FORDHAM URB. L.J. 787, 788-96 (1994) (surveying proposals for the compensated siting of LULUs, including proposals that involve ex ante compensation).

to toxic waste dumps and other "LULUs" (locally unattractive land uses) is to couple the siting of the LULU with payments to all of the residents, sufficient to induce them to accept it. The idea here is that the LULU plus payments to locals from remote individuals (those who benefit in lower product costs and so forth from the LULU, but do not bear its costs) will be ex ante Pareto superior for the entire population. That argument for compensated siting is, necessarily, compelling for the ex ante theorist. But it is not compelling for the ex post theorist. Some of those exposed to the LULU will die, or suffer injury or substantial economic or hedonic damage, perhaps leading to an increase in ex post inequality. This is not to say that the ex post theorist will necessarily favor the no-LULU option over the option of siting the LULU with compensation. (After all, the costs to society of having no waste dumps at all would be overwhelming.) Rather, the point is that the ex post theorist will need to further analyze an ex ante Pareto superior compensated siting scheme, bringing to bear her particular social welfare function. For her, in contrast with the ex ante theorist, the fact that the scheme is ex ante Pareto superior, relative to other policies, does not guarantee its optimality.

C. *"Population Risk" Versus "Individual Risk" and the Problem of Statistical Versus Identified Lives*

Scholarship on environmental, health, and safety regulation regularly takes the position that "population risk" (total fatalities) and "individual risk" (the risk of fatality borne by various individuals in the population) are both relevant in regulating hazards.[119] A paradigmatic example would be a choice between expending a limited regulatory budget on eliminating toxic chemicals that are present in high concentrations in a waste dump in a small town with a population of 1000, and that create an incremental fatality risk for each resident of 1 in 1000; or expending the funds instead on cleaning up much lower concentrations of a toxin that are present in an urban waste dump and that expose each individual in a population of 100,000 to an incremental fatality risk of 2 in 100,000. In this sort of case, it is standardly suggested, "population risk" considerations weigh in favor of cleaning up the urban dump, while "individual risk" considerations

---

[119] *See* Matthew D. Adler, *Risk, Death and Harm: The Normative Foundations of Risk Regulation*, 87 MINN. L. REV. 1293, 1413, 1427-28 (2003) (citing this body of scholarship).

weigh in favor of cleaning up the town dump.[120] Note that the urban dump can be expected to cause two deaths, while the town dump can be expected to cause only one death. On the other hand, because the toxins in the town dump are present at much higher concentrations, it imposes a large incremental "individual risk" of fatality on each exposed individual (1 in 1000); while the urban dump imposes much lower levels of "individual risk" (1 in 50,000 on each urban dweller).

Scholarship on the valuation of life, reflecting a widely held lay view, also sometimes takes the position that remedying the deaths of "identified" individuals has greater social value than remedying "statistical" deaths.[121] This view is a special case of the position that focuses on both "individual risk" and "population risk." It says, in effect, that hazards which pose an "individual risk" of 1 to some individuals take special regulatory priority. (To say that someone is the "identified" rather than "statistical" prospective victim of a hazard is just to say that her risk is unity, or at least is very high.) As between a hazard that poses an "individual risk" of 1 to N persons and causes N expected deaths, and a hazard that poses low "individual risk" and causes N expected deaths, regulators should choose to avert the first hazard.

A concern for "individual risk" is reflected not just in scholarship, but also in actual regulatory practice. Many environmental, health, and safety schemes incorporate "individual risk" cutoffs, seeking to ensure that the incremental "individual risk" to the maximally exposed individual, or to some other specified individual in the exposure distribution, not exceed some numerical threshold (typically, 1 in 1,000,000; 1 in 100,000; 1 in 10,000; or 1 in 1000).[122] One illustrative example (there are many): the FDA regulates food additives so that the risk to a "highly exposed" individual, one in the ninetieth percentile of food consumption, does not exceed 1 in 1,000,000.[123]

---

[120] *See, e.g.,* Adam M. Finkel, *Comparing Risks Thoughtfully,* 7 RISK 325, 343 (1996).

[121] *See* Richard Brook, *Statistical and Identifiable Deaths, in* PHILOSOPHY AND ITS PUBLIC ROLE 167, 168 (William Aitken & John Haldane eds., 2004) (noting that "[m]any individuals feel that the value of an identifiable life should be greater than a statistical life" (quotation marks omitted)); *cf.* Charles Fried, *The Value of Life,* 82 HARV. L. REV. 1415 (1969) (critically surveying arguments for distinguishing between statistical and identified deaths).

[122] *See* Adler, *supra* note 109, at 1147-79 (surveying the role of "individual risk" considerations in regulatory practice).

[123] *See id.* at 1165-67. More precisely, this is the FDA's approach for carcinogens exempt from the Delaney Clause. The FDA also focuses on "individual risk" to the ninetieth percentile consumer in regulating noncarcinogenic additives. *See id.* at 1165.

We can view the regulatory practice of considering "individual risk" along with "population risk" as an attempt to remedy a lack of equity regard in the "population risk" criterion taken alone. There are many ways to produce a given "population risk." A "population risk" of 1 death in a group of 1000 might be created "equitably," as when each individual faces a risk of 1 in 1000, or "inequitably," as when a particular person in the group is certain to die. Adding "individual risk" criteria to the analysis—whether through a prohibition on identified victims or, more generally, through a threshold that specifies the acceptable "individual risk" of death—is a way to favor a seemingly more equitable distribution of harms.

But "individual risk" tests are problematic on a number of counts. For welfarists, they are problematic because they are insensitive to the nonfatal impacts of hazards. A hazard may increase some person's risk of death but also increase her expected utility, all things considered, if the hazard has countervailing benefits. "Individual risk" tests, at least as employed in regulatory practice, are also problematic because they focus on the individual's risk of death from a particular source, not from all sources.[124] A given individual at high risk from some source might have a high life expectancy, all risks considered.

Our analysis reveals yet another difficulty with "individual risk" tests: their ex ante focus. To see this, imagine a very simple scenario in which premature death always yields utility 0, and a full life span, utility 1. In each state, the utility vector is a string of zeros and ones; each individual either lives a full life span or dies prematurely. The "individual risk" that a given policy imposes on a given individual—in this scenario—is simply 1 minus her expected utility. It is her chance of dying prematurely, given all the risks to which she is exposed, if the policy is implemented.

We are dealing, here, with a refined version of the "individual risk" approach that cannot be criticized for ignoring some sources of welfare, or some risks. Still, the approach is problematic. Imagine any methodology for social decision making that is a function not just of the *number* of people who die in each state (i.e., the welfare value of the utility vector in each state) but also the *probability* that each individual will die (i.e., each individual's "individual risk"). Such meth-

---

[124] *See* W. Kip Viscusi, *Risk Equity, in* COST-BENEFIT ANALYIS: LEGAL, ECONOMIC, AND PHILOSOPHICAL PERSPECTIVES 7, 26 (Matthew D. Adler & Eric A. Posner eds., 2001) (noting that policy focus with respect to risk equity is on incremental risk, while "[a]ny meaningful notion of risk equity . . . should be grounded in the absolute risk level of the individual").

odologies include: (1) the approach, reflected in current regulatory policy, which seeks to reduce "individual risk" below a threshold, such as 1 in 1,000,000; and (2) the limiting case of this approach, which seeks to avoid policies with identified victims. Given the analysis earlier in the Article,[125] it is relatively straightforward to show that, given symmetry, any such methodology will violate the sure-thing principle. An easy way to see the inconsistency of these "individual risk" tests with the sure-thing principle is to consider policies that are permutations of each other—that amend some status quo by changing the *identity* of the individuals who die in a given state, but not their numbers. Consider a status quo in which every individual is below a proposed "individual risk" threshold. By an appropriate reordering of the vectors in some states, we can raise some individuals above the threshold. By reordering the vectors so that one individual dies in every state, we trigger the concern with "identified" victims. But, as a matter of ex post equity, the policies are all identical. In each state, some individuals are well off (they live a full life) and others are badly off (they die prematurely). Merely shifting who dies in a given state does nothing to redress the inequity in that state; and a combination of such shifts, whatever effect it might have on various persons' "individual risk" levels, does nothing to change the expected degree of inequity.[126]

D. *Pain-and-Suffering Damages*

The ex post approach provides a new perspective on the topic of damages for pain and suffering, a controversial topic that has generated much legal scholarship.[127] At least in the case of products liabil-

---

[125] *See supra* Part V.A.

[126] The argument that "individual risk" violates the sure-thing principle generalizes beyond our simple example with ones and zeros. Even in more complicated settings, "individual risk," suitably defined, can always be changed with the kind of permutations described in the text. Such permutations will, therefore, change social preferences when "individual risks" are separately accounted for. However, given symmetry, such permutations merely replace the outcome in one state with another socially equivalent outcome. Under the sure-thing principle, this should not affect social choice.

[127] *See* Steven P. Croley & Jon D. Hanson, *The Nonpecuniary Costs of Accidents: Pain-and-Suffering Damages in Tort Law*, 108 HARV. L. REV. 1785, 1793, 1797-1834 (1995) (summarizing "the conventional wisdom" that tort compensation for pain and suffering is undesirable on insurance grounds, and arguing that individuals might rationally purchase pain-and-suffering insurance). Our analysis is different from Croley and Hanson's, because we argue that the equity-regarding social planner might favor pain-

ity, pain-and-suffering damages amount to a kind of compulsory insurance. All consumers pay a higher price for products if pain-and-suffering damages are part of liability, and those who are injured receive a payout. But we do not observe pain-and-suffering insurance in the free market—perhaps because of market failures, but perhaps, instead, because rational individuals might well refrain from purchasing pain-and-suffering insurance. Why? As discussed in Part III.C above, if the occurrence of pain and suffering lowers the marginal utility of money, then it is irrational to insure against it. Why, then, should tort law impose compulsory insurance on the consumers of potentially hazardous products? One possible answer is that pain and suffering generally increase the marginal utility of money, and there are insurance market failures. Another is that pain-and-suffering damages have a deterrence function in addition to their insurance function.

The ex post approach provides a third answer, one not available to the ex ante theorist. Even if pain-and-suffering insurance has no deterrence benefits (or benefits too small to outweigh administrative costs), and even if a population of consumers would not rationally purchase such insurance, requiring this population to purchase insurance can nevertheless increase ex post social welfare. Compulsory pain-and-suffering insurance shifts money from an injury-free state, in which the consumer is at a higher welfare level, to an injured state, in which the consumer is at a lower welfare level. Imagine that money has lower marginal utility in the injured state. In that case, shifting money to the injured state lowers the consumer's expected well-being, but may still have a sufficient benefit in terms of ex post equity to justify the shift.

E. *Criminal Law*

Harel, Safra, and Segal provide two formal examples of how ex post egalitarianism can justify non-utilitarian policies in the area of criminal law.[128] They assume, in the first example, that society is choosing a policy consisting of both a burden of proof for criminal conviction and a probability that crimes will be detected.[129] Raising the probability of detection requires increased enforcement effort and is therefore costly. A given policy affects the number of individuals

---

and-suffering damages even in cases where the injury lowers the marginal utility of money and rational individuals therefore would not purchase insurance.

[128] *See* Harel, Safra & Segal, *supra* note 6, at 63-69.

[129] *See id.* at 64-66.

and their utility levels in five different groups: convicted innocents, convicted criminals, innocent victims of crime, ordinary innocents, and unconvicted criminals. The authors show that (under certain assumptions) the ex post egalitarian will favor increasing the burden of proof for conviction, together with a compensating increase in the probability of detection (financed by taxes on innocents and nonconvicted criminals), as compared with the utilitarian policy. This shift decreases the number of individuals in the worst-off group—convicted innocents—and decreases the spread between their welfare and that of the best-off groups (ordinary innocents and unconvicted criminals).

In a second formal example, the authors consider the severity of the criminal sanction.[130] Gary Becker famously argued that the optimal sanction policy, given the costs of detection, is to threaten very high sanctions—torture or death for a parking violation!—and to use very few resources on detection.[131] In their analysis, Harel, Safra, and Segal consider policies consisting of a level of sanction severity and a probability of detection. The ex post egalitarian will choose to shift away from the optimal utilitarian policy by increasing the probability of detection and reducing the severity of the sanction. This shift reduces overall well-being—reflected again in increased taxes for crime detection—but increases the well-being of those in the worst-off group, namely those who are punished.

Harel, Safra, and Segal, in their formal analyses, compare ex post egalitarianism with utilitarianism—not with ex ante egalitarianism. The analyses are nonetheless quite helpful for our purposes here, in suggesting how the ex ante and ex post approaches to equity regard can deviate. What drives their results, of course, is the fact that a utilitarian policy leads to a substantial spread in actual well-being. That spread may not be mirrored in ex ante utilities, depending on when the policy is chosen. At the extreme (if society commits from day one to a policy), individuals may all face the same probabilities of falling into each category—convicted criminals, unconvicted criminals, innocents, victims, and convicted innocents—and thus the utilitarian and ex ante egalitarian policies will converge. Or, to examine a less extreme case, individuals may know whether they are innocents or criminals, but not know whether they will end up at the high or low end of

---

[130] *See id.* at 66-68.

[131] *See id.* at 58-59 (reviewing Becker's argument).

each group: as convicted or victimized innocents versus ordinary in-
nocents, or convicted versus ordinary criminals.

F. *Capital Gains and Lotteries*

A certain amount of stock market gains may be due to luck. If we
are to believe random-walk theorists, most are. Ex post and ex ante
equity analysts view these windfall gains differently. The ex ante ana-
lyst sees no equity problem with the kind of windfall gain that any in-
dividual might have enjoyed, had she been lucky enough to have her
particular investment take off. The fact that one person invested in
Microsoft and enjoyed a gain, while another invested in Sunbeam and
did not, is of no consequence to ex ante equity analysis because, be-
forehand, each had an equal chance of profit or loss. The ex post
analyst, on the other hand, is troubled by the fact that one person
makes out well and another does not. (To the extent that this gain is
derived from skill in investing, this concern is, of course, balanced
against the social benefits of rewarding good decisions.)

Why does this difference in attitude matter? One arena where
such differences concretely manifest themselves is the taxation of
capital gains. Both ex ante and ex post planners may have reasons to
tax or avoid taxing capital gains. The ex post planner, however, places
an additional weight on the taxation side of the balance. She is con-
cerned with the inequity produced by the actual pattern of capital
gains and losses, whereas her ex ante counterpart is not.

Similarly, the ex post equity analyst may be more hesitant to en-
dorse public lotteries. Both planners may see pluses and minuses to
raising state revenues in this manner. But the ex post planner sees an
additional minus where the ex ante planner does not. The fact that
all lottery participants have an equal chance of winning assuages the
equity concerns of the ex ante—but not the ex post—policy analyst.

G. *Tax Policy Norms*

In an influential article, Thomas Griffith argues against the tradi-
tional approach to tax policy analysis:

> Traditional tax policy analysis has focused on whether the particular tax
> provision under examination is consistent with basic 'tax norms' such as
> horizontal equity, vertical equity, ability to pay and the ideal tax base,
> typically Haig-Simons income. These norms are not grounded, however,
> in more general ethical principles. This Article will argue that this is a
> fundamental flaw and, thus, special tax principles should be discarded as

a method of evaluating tax policy. Instead, this Article recommends that the likely consequences of the policies under consideration should be determined and then judged under explicitly stated general normative principles. This approach can lead to tax policy recommendations quite different from those generated by traditional methods.[132]

Griffith illustrates his proposed analytic approach by focusing on the taxation of personal injury recoveries. The normative criterion that he uses is ex ante Pareto superiority:

> A tax policy is ex ante Pareto superior if, prior to the time any taxpayer knows his particular circumstances, each taxpayer would prefer that policy. In the context of personal injury recoveries, the ex ante perspective means that each individual must choose the tax treatment of such recoveries without knowing whether or not he will be injured.[133]

Griffith argues that—in contrast to current practice[134]—only recoveries for medical expenses, not for lost wages or pain and suffering, should be exempt from income tax, since only in the first case does an exemption shift income between states in a way that increases the taxpayer's expected utility.[135]

We fully agree with Griffith that—bracketing bounded rationality—tax policy makers should bring to bear fundamental moral principles rather than traditional tax policy norms that are only imperfectly associated with any plausible moral theory. Within welfarism, this means that—again, bracketing bounded rationality—tax policy makers should bring to bear some social welfare function. But we obviously disagree with Griffith's assumption that debates about the form of the social welfare function—or more generally, about the choice between welfarist and nonwelfarist views—can be circumvented by taking ex ante Pareto superiority as a normative principle that all plausible moral theories support. The Pareto principle for outcomes may have that status; ex ante Pareto superiority does not, as we have

---

[132] Thomas D. Griffith, *Should "Tax Norms" Be Abandoned? Rethinking Tax Policy Analysis and the Taxation of Personal Injury Recoveries*, 1993 WIS. L. REV. 1115, 1116. Thanks to Yoram Margalioth for suggesting that we engage Griffith's article.

[133] *Id.* at 1121-22 (italics omitted).

[134] Subsequent to publication of Griffith's article, Congress amended the relevant code section, I.R.C. § 104(a), to allow an exclusion only for damages awarded on account of personal *physical* injury. Small Business Job Protection Act of 1996, Pub. L. No. 104-188, § 1605(b), 110 Stat. 1735, 1838 (1996). However, in *Murphy v. I.R.S.*, 460 F.3d 79, 81 (2006), the D.C. Circuit ruled that section 104(a) as amended is unconstitutional for failing to exclude damages awarded on account of nonphysical injury, such damages not being "income" within the meaning of the Sixteenth Amendment.

[135] *Id.* at 1118-19, 1131-34.

argued at length. Tax analysts sympathetic to welfarism and equity regard should follow Griffith's advice to abandon traditional norms such as "horizontal equity," "vertical equity," or "ability to pay," but not his further premise that egalitarians should embrace tax policies that make everyone better off ex ante.[136]

CONCLUSION

This Article has trained its attention on a vital problem within welfarism, one insufficiently appreciated by legal scholars: the problem of welfarist policy choice under uncertainty. We have explicated two different possible approaches: the ex ante approach, which applies the social welfare function to the vector of expected utilities associated with each policy under consideration; and the ex post approach, which calculates the expected value of the social welfare function for each policy (by applying the function to each possible outcome, discounting by the outcome's probability, and summing).

The Article has explained how these two approaches can diverge for any "equity-regarding" social welfare function. The family of equity-regarding social welfare functions are those that satisfy the Pigou-Dalton principle—in contrast to a utilitarian function—and include both prioritarian and comparativist social welfare functions. We have shown that, for *any* equity-regarding social welfare function, there are policy choice situations in which the function applied ex post prefers one policy but applied ex ante prefers another. Yet more strikingly, we have shown that, for any equity-regarding social welfare function, there will be choice situations in which the function applied ex post prefers a policy which is ex ante Pareto inferior to another option.[137]

Further, the Article has clarified the conditions under which ex ante and ex post approaches diverge, for prioritarian and comparativist functions. We have demonstrated that governmental compensation mechanisms cure the divergence only if the mechanisms are implausibly powerful—only if utility is freely transferable between persons and states. And we have shown that private insurance exacerbates, rather than cures, the ex ante/ex post divergence.

So which approach to welfarist policy choice under uncertainty is more attractive as a normative matter? We have argued that the ex post approach is the better specification of welfarism. "Fairness" ar-

---

[136] *See id.* at 1121.

[137] This assumes continuity, as referred to *supra* note 22 and accompanying text.

guments for the ex ante approach are misconceived on this score, both because they depart from welfarism, and because concerns for fair process only coincide haphazardly with the ex ante application of a social welfare function. Within welfarism, there are two powerful considerations favoring the ex post approach: the "sure-thing principle," which many find a compelling axiom of rationality, and the principle of "time consistency." On the other hand, the potential conflict with ex ante Pareto superiority does raise important worries about the ex post approach, but we have argued that these worries can be defused. It bears emphasis that all of the equity-regarding social welfare functions we have explored are Paretian in their ranking of outcomes, and that the ex post approach therefore *will* favor policy P over policy Q if P's outcome, in each possible state, is Pareto superior to Q's outcome in that state.

Finally, the Article has sought to show that the ex post approach to welfarist policy choice under uncertainty has wide legal implications. We have surveyed some of these implications: optimal expenditures on the prevention of terrorism, epidemics, and other widely borne risks; the design of mechanisms that seek to compensate individuals who suffer physical injury or financial loss, or who are put at risk of such harm, such as unemployment insurance, workers' compensation, or compensated siting; the use of "individual risk" tests in risk regulation and the distinction between statistical and identified victims; whether tort victims should recover pain-and-suffering damages; evidentiary and punishment rules in criminal law; policies regarding financial windfalls, such as the taxation of capital gains and the use of state lotteries; and the role of ex ante Pareto superiority as a tax policy norm.

MATHEMATICAL APPENDIX

This Appendix supplements the text with formal definitions of key concepts and formal proofs of key results. Unless otherwise stated, the results herein are, to the best of the authors' knowledge, original.

A *policy* is a finite set of *utility vectors* $\left(u_1^s,...,u_I^s\right) \in \mathbb{R}^I$, one for each *state* $s = 1,...,S$, and a probability distribution over states $\left(p^1,...,p^S\right) \in \mathbb{R}^S$, with all $p^s \geq 0$ and $\sum p^s = 1$.[138] Given a policy, write $u^s = \left(u_1^s,...,u_I^s\right)$ for the *utility vector in state s* and $u_i = \left(u_i^1,...,u_i^S\right)$ for individual *i*'s utility profile across states. A *social welfare function* w maps utility vectors $\left(u_1,...,u_I\right) \in \mathbb{R}^I$ onto the real numbers. The utility vector $u^s \in \mathbb{R}^I$ generated by any policy may be regarded as a random vector, and the social welfare value of that vector, $w\left(u^s\right)$, a random variable. The expected values of these random entities are, respectively, $E\left[u^s\right] = \sum p^s u^s \in \mathbb{R}^I$ and $E\left[w\left(u^s\right)\right] = \sum p^s w\left(u^s\right) \in \mathbb{R}$. Given social welfare function w, the *ex ante social welfare* of a given policy is $w\left(E\left[u^s\right]\right)$, while *ex post social welfare* is $E\left[w\left(u^s\right)\right]$.

A. *Utilitarianism and the Equivalence of Ex Ante and Ex Post Social Welfare Calculations*

In this Section, we prove that the only social welfare functions for which ex ante welfare equals ex post are affine functions of individual utilities. That is, social welfare functions that are a weighted sum of individual utilities plus a constant. As remarked in the Article's text, this result plus symmetry, and the irrelevance of the constant for social choice, take us to utilitarianism.[139]

As discussed in note 58, *supra*, the following result is often ascribed to Harsanyi, but it is actually more elementary than what Harsanyi proved. It is a generalization to multidimensional domains of the well-known fact that risk-neutral utility functions are affine.

A function f is *linear* if $f\left(ax\right) = af\left(x\right)$ and $f\left(x + y\right) = f\left(x\right) + f\left(y\right)$. Equivalently, a linear function may be expressed as $f\left(x\right) = a_1 x_1 +...+ a_I x_I$, for some a_i, $i = 1,...,I$. Just set $a_i = f\left(\left(0,...,1,...,0\right)\right)$. A function f is *affine* if it may be expressed as $f\left(x\right) = a_1 x_1 +...+ a_I x_I + k$ for some a_i, $i = 1,...,I$ and k. In this case

---

[138] Equivalently, we could define a policy as a probability measure on the set of all utility vectors. Notice, however, that we are restricting attention to the subclass of probability measures with finite support.

[139] *See supra* Part I.C.

$f(0) = k$ and the function $\hat{f}(x) = f(x) - f(0)$ is linear. Conversely, if $\hat{f}(x)$ is linear, then $f(x) = \hat{f}(x) + k$ is affine for any k.

PROPOSITION 1: *Consider* $w : \mathbb{R}^I \to \mathbb{R}$. *If for all policies,*

$$E\left[w\left(u^i\right)\right] = w\left(E\left[u^i\right]\right) \tag{3}$$

then w *is affine.*

Proof. First, suppose that $w(0) = 0$. We will show that w must be linear. First, we show that for all $x \in \mathbb{R}^I$ and all $a \in \mathbb{R}$,

$$w(ax) = aw(x). \tag{4}$$

Take any $x \in \mathbb{R}^I$ and any $a \in \mathbb{R}$. If $a \in [0,1]$, then ax is a convex combination of x and 0. Thus, (4) is immediate from (3) and the assumption that $w(0) = 0$. If $a > 1$, then x itself is a convex combination of ax and 0:

$\frac{1}{a}(ax) + \left(1 - \frac{1}{a}\right)0 = x$.

Therefore, by (3)

$$w(x) = \frac{1}{a}w(ax) + \left(1 - \frac{1}{a}\right)w(0) \;\Rightarrow\; aw(x) = w(ax). \tag{5}$$

And if $a < 0$, then 0 is a convex combination of ax and $-ax$:

$\frac{1}{2}ax + \frac{1}{2}(-ax) = 0$,

so that by (3) and the results above for positive a,

$w(0) = \frac{1}{2}w(ax) + \frac{1}{2}w(-ax) \Rightarrow w(ax) = -w(-ax) = --aw(x) = aw(x)$.

Next, we will show that for all $x, y \in \mathbb{R}^I$, $w(x+y) = w(x) + w(y)$. Thus, consider any $x, y \in \mathbb{R}^I$. We have $x + y = 2\left(\frac{1}{2}x + \frac{1}{2}y\right)$ and so by (4) and (3),

$w(x+y) = 2w\left(\frac{1}{2}x + \frac{1}{2}y\right) = 2\frac{1}{2}w(x) + 2\frac{1}{2}w(y) = w(x) + w(y)$.

We have thus shown that w is linear if $w(0) = 0$. If $w(0) \neq 0$, then consider $\hat{w}(x) = w(x) - w(0)$. Equation (3) holds also for $\hat{w}(x)$. Thus, $\hat{w}(x)$ is linear, implying that $w(x)$ is affine. *Q.E.D.*

B. *Equity Regard*

We say that utility vector $\left(u_1', \ldots, u_I'\right)$ may be derived from utility vector $\left(u_1, \ldots, u_I\right)$ by means of a (pairwise) *mean-preserving* (strictly) *equalizing transfer*, if and only if for some pair of individuals i, j with $i \neq j$ and some transfer $t_i, t_j \neq 0$ with $t_i = -t_j$,

$\left(u_1', \ldots, u_I'\right) = \left(u_1, \ldots, u_i + t_i, \ldots, u_j + t_j, \ldots, u_I\right)$ and

$$\left| u_i + t_i - \left(u_j + t_j \right) \right| < \left| u_i - u_j \right|. \tag{6}$$

A social welfare function w is *equity regarding* (i.e., satisfies the *Pigou-Dalton principle*) if and only if, whenever $(u_1',...,u_I')$ may be derived from $(u_1,...,u_I)$ by a mean-preserving equalizing transfer, w assigns a strictly higher value to $(u_1',...,u_I')$.

A technically weaker restriction found in the literature, and also referred to as the "Pigou-Dalton principle," requires only that the social welfare function increase in the special case that the equalizing transfer does not alter the utility rank of the two individuals i and j. In essence, this weaker definition adds the condition $|t_i| = |t_j| \le \frac{1}{2}|u_i - u_j|$ to the definition of "mean-preserving equalizing transfer." However, given anonymity/symmetry, this alternative definition is equivalent to the one we provide here. As is evident from Figure 2, an equalizing transfer that reverses rank is equivalent to an equalizing transfer that does not reverse rank plus an exchange of utilities between the two individuals. The non-rank-reversing transfer must increase welfare under the weaker Pigou-Dalton condition, and the exchange of utilities cannot change welfare given the anonymity/symmetry assumption.

Figure 2
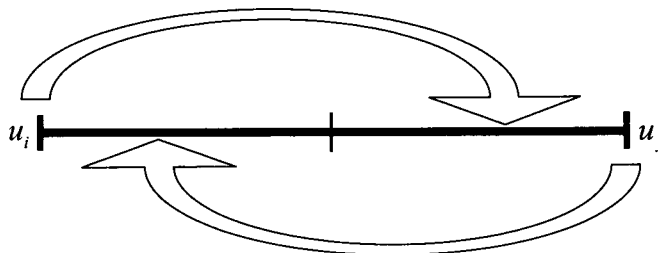
Alternatively, one can ostensibly strengthen the Pigou-Dalton principle by specifying that the social welfare function must increase if the second utility vector can be obtained from the first via any *sequence* of mean-preserving equalizing transfers. Extending the principle to sequences has no effect, however, given the assumed transitivity of the social welfare ordering (which is reflected in the fact that the ordering can be represented by a social welfare function).

C. *Rank-Weighted Social Welfare and the Gini Coefficient*

This Section supplements the discussion in Part I.B, above, regarding the interpretation of the Gini coefficient and its relationship to rank-weighted social welfare functions. The results in this Section are not original. As discussed in note 45, *supra*, the *Gini coefficient* is typically written as

$$G = \tfrac{1}{2}\tfrac{1}{\mu}\tfrac{1}{I^2}\sum_{i=1}^{I}\sum_{j=1}^{I}\left|u_i - u_j\right|,$$

where $\mu = \tfrac{1}{I}U = \tfrac{1}{I}\sum u_i$ is mean utility. In what follows it is more convenient to work with an equivalent definition, in which each pairwise difference is added only once:

$$G = \tfrac{1}{\mu}\tfrac{1}{I^2}\sum_{i=1}^{I}\sum_{j=1}^{i}\left|u_i - u_j\right|.$$

Consider the double sum in this equivalent statement: $\sum_{i=1}^{I}\sum_{j=1}^{i}\left|u_i - u_j\right|$. The number of times a given u_i is added or subtracted in this double sum depends on how often it is the larger utility in a pair. Each time it is the larger utility, it is added. Each time it is the smaller, it is subtracted. Specifically, the i^{th} smallest utility is added i times and subtracted $I-(i-1)$ times (including in both cases the pair (i,i) where it is repeated). Therefore, the double sum is, in essence, a rank-weighted sum with higher weights for more highly ranked utilities. Specifically, if we rearrange the utility vector so that $u_1 \le u_2 \le ... \le u_I$, then

$$\mu I^2 G = \sum_{i=1}^{I}\sum_{j=1}^{i}\left|u_i - u_j\right| = \sum_{i=1}^{I}\left(i - \left(I - (i-1)\right)\right)u_i$$

$$= \sum_{i=1}^{I}\left(2i - (I+1)\right)u_i.$$

From this fact follows the close relationship between rank-weighted social welfare and the Gini coefficient. Thus, take any rank-weighted social welfare function where the ranking increment is constant. Such a welfare function may always be written as

$$w = \sum_{i=1}^{I}\left((k + \alpha I) - \alpha i\right)u_i,$$

where the utility vector is arranged so that $u_1 \le u_2 \le ... \le u_i \le ...u_I$, where $k > 0$ summed with αI is a ranking constant, and where α is the constant ranking increment. Then,

$$w = \sum_{i=1}^{I}\left((k + \alpha I - \alpha i) + \tfrac{\alpha}{2}\left(2i - (I+1)\right) - \tfrac{\alpha}{2}\left(2i - (I+1)\right)\right)u_i$$

$$= \sum_{i=1}^{I}\left(\left(k + \alpha I - \cancel{\alpha i}\right) + \tfrac{\alpha}{2}\left(\cancel{2i} - (I+1)\right)\right)u_i - \tfrac{\alpha}{2}\sum_{i=1}^{I}\left(2i - (I+1)\right)u_i$$

$$= \left(k + \alpha I - \tfrac{\alpha}{2}(I+1)\right)U - \tfrac{\alpha}{2}UIG$$

$$= I\left(\tfrac{k + \alpha I - \tfrac{\alpha}{2}(I+1)}{I}U - \tfrac{\alpha}{2}UG\right).$$

For every I, this welfare function w ranks utility vectors the same as $\frac{w}{I}$. In the limit as $I \to \infty$, $\frac{w}{I}$ goes to $\left(\frac{\alpha}{2}U - \frac{\alpha}{2}UG\right)$. The case discussed in the text is that in which $\alpha = 1$ and $k = 0$.[140]

D. *Transfers That Are Disequalizing Ex Post and Equalizing Ex Ante*

This Section proves the result discussed in Part II.A. (The proposition is stated without symbolic notation, but its symbolic formulation is clear from the proof to follow. A similar statement applies to several of the propositions to follow.)

PROPOSITION 2: *If the status quo is uncertain, and there is some inequality in some state, then we can always find a mean-preserving transfer in some state that is disequalizing with respect to the utility vector in that state, yet equalizing with respect to the vector of expected utilities.*

Proof: The text essentially proves the result in the case that there is at least one rank-contrary state (as there defined).

Suppose, then, that there is no rank-contrary state and suppose that individual j's utility is at least as great as individual i's in all states. By definition of the expected value operation, $\exists s$ such that $p^s > 0$ and $u_j^s - u_i^s \leq E\left[u_j^{\tilde{s}} - u_i^{\tilde{s}}\right]$.

Define $t = u_j^s - u_i^s + \varepsilon$, for any $\varepsilon \in \left(0, \frac{1-p^s}{p^s}E\left[u_j^{\tilde{s}} - u_i^{\tilde{s}}\right]\right)$. The interval from which ε is drawn is nonempty: rank conformity plus inequality in some state implies $E\left[u_j^{\tilde{s}} - u_i^{\tilde{s}}\right] > 0$, and uncertainty in the status quo implies $1 - p^s > 0$. Because $t > u_j^s - u_i^s$, transferring t from j to i is strictly disequalizing in state s. The transfer, though initially in an equalizing direction, "overshoots." On the other hand, using the fact that $\varepsilon < \frac{1-p^s}{p^s}E\left[u_j^{\tilde{s}} - u_i^{\tilde{s}}\right]$, the transfer of expected utility is

$$p^s t = p^s\left(u_j^s - u_i^s\right) + p^s\varepsilon \leq p^s E\left[u_j^{\tilde{s}} - u_i^{\tilde{s}}\right] + p^s\varepsilon < E\left[u_j^{\tilde{s}} - u_i^{\tilde{s}}\right]. \tag{7}$$

Further, by choice of s, $p^s > 0$, and so $p^s t > 0$. Combined with (7), this means that the transfer of t from j to i in state s positively equalizes ex ante. *Q.E.D.*

E. *Strict Concavity*

This Section proves the result discussed in Part II.C above. A real-valued function $f : \mathbb{R}^n \to \mathbb{R}$ is *strictly concave* if and only if for all $x, y \in \mathbb{R}^n$ with $x \neq y$, and all $t \in (0,1)$, $f\left(tx + (1-t)y\right) > tf(x) + (1-t)f(y)$.

---

[140] *See supra* note 46 and accompanying text.

The reasoning behind the following proposition is familiar from the analysis of concave functions. The proof, which is not original, is recited here for the convenience of the reader.

PROPOSITION 3: *A social welfare function is strictly concave if and only if for all policies ex ante social welfare strictly exceeds ex post.*

Proof: Suppose, first, that the social welfare function is strictly concave. Take any policy. The proof that ex ante social welfare is strictly greater than ex post is by induction. If $S = 2$, then

$$w\left(E\left[u^s\right]\right) = w\left(p^1 u^1 + \left(1 - p^1\right)u^2\right)$$
$$> p^1 w\left(u^1\right) + \left(1 - p^1\right)w\left(u^2\right) = E\left[w\left(u^s\right)\right].$$

Now suppose that the same holds for policies with $S - 1$ states and consider a policy with S states. Consider the probability measure $p^{2|1},...,p^{S|1} \equiv \frac{p^2}{1-p^1},...,\frac{p^S}{1-p^1}$ on states $2,...,S$. Then,

$$w\left(E_p\left[u^s\right]\right) = w\left(p^1 u^1 + \left(1 - p^1\right)E_{p^{\cdot|1}}\left[u^s\right]\right)$$
$$> p^1 w\left(u^1\right) + \left(1 - p^1\right)w\left(E_{p^{\cdot|1}}\left[u^s\right]\right)$$
$$> p^1 w\left(u^1\right) + \left(1 - p^1\right)E_{p^{\cdot|1}}\left[w\left(u^s\right)\right]$$
$$= E_p\left[w\left(u^s\right)\right].$$

Conversely, suppose the function is not strictly concave. Then, by the negation of the definition of strict concavity, we can find some policy (with two outcomes) for which ex ante social welfare is not strictly greater than ex post. *Q.E.D.*

F. *Permuting Policies*

This Section proves the result discussed in Part II.D above: every uncertain permuting policy (as there defined) has a greater valuation ex ante than ex post, independent of the choice of equity-regarding social welfare function. A *bistochastic matrix* is a square matrix of non-negative elements, each of whose rows and columns add to one. A *permutation matrix* is a bistochastic matrix all of whose elements are either 1 or 0.

The definition of equity regard provided above in Appendix Section B tests whether the social welfare function increases under every mean-preserving transfer between any two individuals that equalizes their utilities. The literature provides an equivalent definition of equity regard in terms of linear transformations of utility vectors via

premultiplication by bistochastic matrices.[141] In essence, transforming a utility vector by a bistochastic matrix replaces each individual's utility with a weighted average of all individuals' utilities. The weighting may be different across individuals, but the sum (and so the mean) of all individuals' utilities must remain the same. This averaging process is equivalent to a series of mean-preserving equalizing (pairwise) transfers. From this follows the equivalence of the two definitions of equity regard.

The portion of this equivalence that is important for our purposes is as follows. We use bold capital letters to represent matrices and bold lowercase letters to represent column vectors.

LEMMA: *If* $\mathbf{Q}$ *is a bistochastic matrix and* $\mathbf{u}$ *a utility vector, then an equity-regarding social welfare function assigns no lower value to* $\mathbf{Qu}$ *than to* $\mathbf{u}$. *If* $\mathbf{Q}$ *contains at least one element that is neither zero nor one, then the social welfare of* $\mathbf{Qu}$ *will be strictly greater than* $\mathbf{u}$.

The following Lemma links the foregoing feature of equity regard to permutation matrices, and so to permuting policies.

LEMMA: *The weighted average of permutation matrices is a bistochastic matrix.*

Proof. It is clear that each element of the weighted average of permutation matrices is nonnegative. Consider row 1 of the weighted average matrix. This is the weighted average of first rows of the permutation matrices. The weighted contribution of the first permutation matrix is a row with all zeros except for one entry, which is p_1. A similar statement holds for each of the other permutation matrices. Therefore, the sum of the elements of the first row of the weighted average matrix must be 1. Similar reasoning applies to every other row and column. *Q.E.D.*

PROPOSITION 4: *Suppose each outcome of some policy is the permutation of a given vector of utilities. Suppose also that the social welfare function is equity regarding. Then, whatever the probability of each state, ex ante social welfare is no less than ex post. Moreover, ex ante welfare is strictly greater than ex post if the probabilities do not put all weight on one state.*

Proof. By hypothesis, there exists $\mathbf{u}_{I \times 1} = \left(u_1, ..., u_I \right)'$, such that for all s, $\mathbf{u}_{I \times 1}^s = \mathbf{P}_{I \times I}^s \mathbf{u}_{I \times 1}$, where each $\mathbf{P}_{I \times I}^s$ is a permutation matrix. For all probability distributions $\mathbf{p}_{S \times 1}$ over states, the vector of expected utilities may be written

---

[141] *See, e.g.,* Partha Dasgupta et al., *Notes on the Measurement of Inequality,* 6 J. ECON. THEORY 180, 182-83 (1973).

$$\left[\mathbf{P}^1_{I \times I} \mathbf{u}_{I \times 1} \quad \cdots \quad \mathbf{P}^s_{I \times I} \mathbf{u}_{I \times 1} \quad \cdots \quad \mathbf{P}^S_{I \times I} \mathbf{u}_{I \times 1} \right] \mathbf{p}_{S \times 1}$$

$$= \mathbf{P}^1_{I \times I} \mathbf{u}_{I \times 1} p_1 + \ldots + \mathbf{P}^s_{I \times I} \mathbf{u}_{I \times 1} p_s + \ldots + \mathbf{P}^S_{I \times I} \mathbf{u}_{I \times 1} p_S$$

$$= \left[\mathbf{P}^1_{I \times I} p_1 + \ldots + \mathbf{P}^s_{I \times I} p_s + \ldots + \mathbf{P}^S_{I \times I} p_S \right] \mathbf{u}_{I \times 1}$$

Therefore, by the immediately preceding Lemma, expected utility is a bistochastic transformation of the permuted utility vector. And, by the first Lemma in this Section, the expected utility vector attains at least as high or strictly higher social welfare (depending on the condition in the last sentence of the Proposition), assuming regard for equity. *Q.E.D.*

G. *Separable Social Welfare*

This Section proves the result discussed in Part II.E: a continuous, symmetric, separable social welfare function is equity regarding if and only if it is strictly concave. The results are most likely not original. A symmetric social welfare function is (*individual*) *separable* if and only if it may be written as $w(u_1, \ldots, u_I) = \varphi(u_1) + \ldots + \varphi(u_I)$ for some uniform *transformation function* $\varphi: \mathbb{R} \to \mathbb{R}$. We will need the following equivalent definition of strict concavity for continuous functions:

LEMMA: *A continuous function f is strictly concave if and only if for all* x, y, $f\left(\frac{1}{2}x + \frac{1}{2}y\right) > \frac{1}{2}f(x) + \frac{1}{2}f(y)$.

Proof: Well known.[142]

LEMMA: *A continuous, symmetric, and separable social welfare function w is equity regarding if and only if its transformation function φ is strictly concave.*

Proof: For the "if" implication: if the transformation function is strictly concave, then so is the social welfare function as a whole, implying that the social welfare function is equity regarding.

For the "only if" implication: if the social welfare function is separable, symmetric, and continuous, then it is the sum of uniformly transformed utilities, and the uniform transformation function must be continuous. If the transformation function is not, in addition, strictly concave, then by the equivalent definition of concavity for continuous functions, as stated in the immediately preceding Lemma, there exists u, u', with $u > u'$, s.t. $\phi\left(\frac{1}{2}u + \frac{1}{2}u'\right) \leq \frac{1}{2}\phi(u) + \frac{1}{2}\phi(u')$, or equivalently $2\phi\left(\frac{1}{2}u + \frac{1}{2}u'\right) \leq \phi(u) + \phi(u')$. This implies that transferring $\frac{1}{2}(u - u')$ from an individual with u to an individual with $u' < u$ does

---

[142] *See, e.g.,* PlanetMath, Convex Function, http://planetmath.org/convexfunction.html (last visited Dec. 3, 2006).

not strictly increase social welfare. Therefore, the social welfare function is not equity regarding. *Q.E.D.*

PROPOSITION 5: *A continuous, symmetric, and separable social welfare function w is equity regarding if and only if it is strictly concave.*

Proof: The transformation φ is strictly concave if and only if $\sum \varphi(u_i)$ is strictly concave. Given this, apply the immediately preceding Lemma. *Q.E.D.*

H. *Compensation*

This Section proves the result regarding "double free transferability," discussed in Part III.A.3 above.

PROPOSITION 6: [ordinal degeneration to utilitarianism with freely transferable utility across individuals and states]: *Suppose that the social welfare function w is continuous and Paretian. Consider any two utility matrices* $u^1, u^2 \in \mathbb{R}^{I \times S}$. *Then* $\forall \tau^1 \in \mathbb{R}^{I \times S}$ *with* $\sum_{is} p_s \tau_{is}^1 \leq 0$ *and* $u^1 + \tau^1 \geq 0$, $\exists \tau^2 \in \mathbb{R}^{I \times S}$ *with* $\sum_{is} p_s \tau_{is}^2 \leq 0$ *and* $u^2 + \tau^2 \geq 0$ *s.t.*

$$\sum_s p_s w\left(u_s^2 + \tau_s^2\right) \geq (>) \sum_s p_s w\left(u_s^1 + \tau_s^1\right), \tag{8}$$

if and only if

$$\sum_s p_s \sum_i u_{is}^2 = \sum_i \sum_s p_s u_{is}^2 \geq (>) \sum_i \sum_s p_s u_{is}^1 = \sum_s p_s \sum_i u_{is}^1 . \tag{9}$$

REMARK: The Lemma should be read either with weak inequality in both (8) and (9) or with strict inequality in both (8) and (9).

REMARK: The proposition implies that double free transferability "equates" utilitarianism (which itself is the same ex ante and ex post) with both ex post and ex ante application of any continuous Paretian social welfare function w. Ex post application of w is "equated" with ex post utilitarianism in the sense that utility matrix u^2 is preferred, *as a starting point for transfers,* to utility matrix u^1 (as per (8)) if and only if u^2's expected utility sum over individuals is higher (as per (9)). *Ex ante* application is similarly equated with ex ante utilitarianism *a fortiori,* since the proposition holds when $S = 1$, wherein u^2 and u^1 are single column matrices whose entries can be interpreted as expected utilities.[143] Therefore, because there is no divergence between ex ante

---

[143] Note that there is no advantage to differentiating transfers by state when the welfare function is applied ex ante; all that matters is the effective transfer of expected utility.

and ex post utilitarianism, there is no divergence between ex ante and ex post application of w (ignoring the potentially differing transfer schemes that each approach would institute). In other words, if policy A is preferred to policy B (anticipating adjustment with free transferability across individuals and states) under w applied ex ante, A is also preferred to B under w applied ex post (anticipating a potentially different adjustment with free transferability across individuals and states).

Free transferability of utility across individuals is inherent in the effect of transfers on individual utilities in (8) and in the fact that the constraint that transfers be nonpositive sum in expected value, the "budget balance constraint," applied to the constructed transfer τ^2, allows for transfer schemes that are precisely zero sum (in expectation). This is easiest to see in the simple case of within-state transfers. The budget balance constraint allows us to increase individual 1's utility in state 1 by one unit in return for decreasing individual 2's utility in state 1 by one unit.

Free transferability of utility *across states* is inherent in the fact that the budget balance constraint requires that the constructed transfer τ^2 be zero sum only *in expectation* across states. (That it may be *negative* sum is superfluous.) Free transferability across states would not hold were the constructed transfer required to be zero sum *within each state.*

Proof. Suppose $\sum_{is} p_s u_{is}^2 \geq (>) \sum_{is} p_s u_{is}^1$. Take any transfer matrix $\tau^1 \in \mathbb{R}^{I \times S}$ s.t. $\sum_{is} p_s \tau_{is}^1 \leq 0$ and $u^1 + \tau^1 \geq 0$. Define the transfer matrix $\tau^2 = u^1 + \tau^1 - u^2 (+\varepsilon 1)$, where $\varepsilon > 0$ and 1 is an $I \times S$-matrix all of whose elements are 1. (Note: as indicated by parentheses, $\varepsilon 1$ is added only in the strict inequality case.)

Notice that $\sum_{is} p_s \tau_{is}^2 = \sum_{is} p_s u_{is}^1 + \sum_{is} p_s \tau_{is}^1 - \sum_{is} p_s u_{is}^2 \leq (<) 0$ (if $\varepsilon > 0$ is small enough). Further, $u^2 + \tau^2 = (>) u^1 + \tau^1 \geq 0$, from which we may conclude not only that $u^2 + \tau^2 \geq 0$, but also that, for every s, $w\left(u_s^2 + \tau_s^2\right) \geq (>) w\left(u_s^1 + \tau_s^1\right)$, given that w is Paretian. *A fortiori,* $\sum_s p_s w\left(u_s^2 + \tau_s^2\right) \geq (>) \sum_s p_s w\left(u_s^1 + \tau_s^1\right)$.

Conversely, suppose $\sum_{is} p_s u_{is}^2 < (\leq) \sum_{is} p_s u_{is}^1$. Let $\hat{\tau}^1$ maximize $\sum_s p_s w\left(u_s^1 + \tau_s^1\right)$ subject to $\sum_{is} p_s \tau_{is}^1 \leq 0$ and $u^1 + \tau^1 \geq 0$. (This maximum exists because w is continuous, and the constraint set is compact.) The claim is that there is no τ^2 with $\sum_{is} p_s \tau_{is}^2 \leq 0$ and $u^2 + \tau^2 \geq 0$ s.t. $\sum_s p_s w\left(u_s^1 + \hat{\tau}_s^1\right) \leq (<) \sum_s p_s w\left(u_s^2 + \tau_s^2\right)$. If there were, we could set $\tau^1 = u^2 + \tau^2 - u^1 + \varepsilon 1$ with $\varepsilon > (\geq) 0$ and have a) $\sum p_s \tau_{is}^1 = \underbrace{\left(\sum_{is} p_s u_{is}^2 - \sum_{is} p_s u_{is}^1\right)}_{<(\leq)0} + \sum p_s \tau_{is}^2 \leq 0$ for small enough $\varepsilon > (\geq) 0$;

b) $u^1 + \tau^1 > (\geq) u^2 + \tau^2 \geq 0$; and c) given that w is Paretian, for all s, $w\left(u_s^1 + \tau_s^1\right) > (\geq) w\left(u_s^2 + \tau_s^2\right)$, implying $\sum_s p_s w\left(u_s^1 + \tau_s^1\right) > (\geq) \sum_s p_s w\left(u_s^2 + \tau_s^2\right)$. Given the supposition $\sum_s p_s w\left(u_s^1 + \hat{\tau}_s^1\right) \leq (<) \sum_s p_s w\left(u_s^2 + \tau_s^2\right)$, this contradicts, by transitivity, the definition of $\hat{\tau}^1$. Q.E.D.

I. *Nonfreely Transferable Utility*

This Section formalizes results discussed in Part III.A.3 above.

LEMMA: [ex ante decomposition of transfers]: *Let* $\tau_i^j \in \mathbb{R}$ *denote the transfer of utility to individual* i *in state* j. *Let* W *denote the subset of ex ante winners under the transfer, that is, individuals for whom:* $E\left[\tau_i^j\right] \geq 0$. *Let* L *denote the subset of ex ante losers, namely, individuals for whom* $E\left[\tau_i^j\right] < 0$. *Then,*

$$\underbrace{\sum_{i \in L} E\left[\tau_i^j\right]}_{-} = -\underbrace{\sum_{i \in W} E\left[\tau_i^j\right]}_{+} - \underbrace{E\left[\sum_{i=1}^I \tau_i^j\right]}_{\text{cost of transfer}}.$$

Proof. Immediate.

REMARK: This Lemma says that the aggregate transfer *to* the losers (negative) is the additive inverse of the aggregate transfer to the winners (positive or zero) less the expected cost of the transfers, $-E\left[\sum_{i=1}^I \tau_i^j\right]$. Notice that the expected cost of the transfers combines the loss, if any, within each state with the loss, if any, across states.

PROPOSITION 7: *Suppose that expected utilities are equal across individuals in utility matrix* $u \in \mathbb{R}^{I \times S}$. *Then every transfer scheme* $\left\{\tau_i^j\right\}$ *with positive expected costs (*$E\left[\sum_{i=1}^I \tau_i^j\right] < 0$ *) is ex ante welfare reducing (given Paretianism and equity regard).*

Proof. Refer to the decomposition in the above Lemma and imagine that the transfer is accomplished in two steps. First, the cost of the transfer is taken from some losers with no compensating increase in the utility of any winner. Second, the remainder of the transfer is taken from losers and given to winners. The first portion of this decomposition, which adds $E\left[\sum_{i=1}^I \tau_i^j\right] < 0$ to the aggregate expected utility for the losers, is a Pareto inferior move from an ex ante perspective and so reduces ex ante social welfare. Note that, given initial equality, every loser has no greater expected utility than every winner after this change. The second portion of the decomposition, the transfer from losers to winners of $\sum_{i \in W} E\left[\tau_i^j\right] > 0$, is then (given the initial portion of the decomposition) disequalizing. Specifically, one can represent it as the composition of a series of pairwise disequalizing transfers. Q.E.D.

REMARK: This result speaks to the persistence of ex ante/ex post disagreement when transfers are possible, but costly. It says that any

equity-regarding Paretian ex ante planner disprefers *any* costly trans-
fer from a position of equal expected utilities. Examples of utility ma-
trices are easy to find wherein expected utilities are equal ex ante, but
not ex post. Indeed, many of the examples in the text conform to this
pattern. In this case, assuming continuous welfare, whatever the ex
post planner's degree of equity regard, we can always find a transfer
cost small enough such that she prefers to make the transfer, contrary
to the preferences of her ex ante counterpart. Conversely, whatever
the cost of transferring, we can always find an ex post planner with
sufficient equity regard to prefer the transfer. Neither of these two
dimensions of variation—degrees of positive equity regard or degrees
of positive transfer cost—will affect the fact that the ex ante planner
will oppose the transfer.

REMARK: As noted by numerical example in the text, the general
point—that ex post and ex ante planners may disagree over costly
compensation schemes that mitigate ex post inequality—is not de-
pendent on perfect ex ante inequality.

J. *Time Inconsistency of the Ex Ante Approach*

This Section discusses in more detail the coin flipping example in
Part V.B.3. Recall that the planner, whose welfare function equals the
sum of utilities less half their difference, is deciding which of two indi-
viduals should receive 100 utils. The planner cannot split that total
between them. In each period she may choose to flip a fair coin, but
each coin toss reduces the amount to be allocated by 1 util. Suppose
that x_t utils are left to be allocated after the t^{th} toss, where $t = 0,1,2,\dots$.
Will the ex ante planner want to toss the coin again? If she chooses
not to toss again, social welfare is $x_t + 0 - \frac{1}{2}|x_t - 0| = \frac{1}{2}x_t$. If she chooses
to toss, social welfare is $\frac{1}{2}(x_t - 1) + \frac{1}{2}(x_t - 1) + \frac{1}{2}|(x_t - 1) - (x_t - 1)| = x_t - 1$.
Thus, the planner chooses to toss the coin so long as
$x_t - 1 > \frac{1}{2}x_t \Rightarrow x_t > 2$. Therefore, the ex ante planner spends 98 utils
deciding how to allocate the 100, finally giving the remaining 2 utils
to the individual who wins the ninety-eighth toss.

K. *Application: Spending on Prevention*

This Section formalizes the discussion surrounding note 116, *su-
pra*. Suppose that (a) there are two individuals, each starting with 100
utils; (b) no more than one individual will be attacked; (c) each indi-
vidual is equally likely to be attacked; (d) the probability of an attack is
$p(s)$, where s, the policy variable, is the (equal) cost imposed on each

individual for security spending stated in terms of utils, and $p' < 0$, $p'' > 0$, $p'(0) = -\infty$; (e) an attack causes the harmed individual to lose all 100 utils; and (f) the individual and time-separable equity-regarding social welfare function is the sum of utilities less half their absolute difference: $u_1 + u_2 - \frac{1}{2}|u_1 - u_2|$. Then, each individual's expected utility from an attack is

$$-s + \tfrac{1}{2}(100) + \tfrac{1}{2}(1 - p(s))(100)$$

$$= -s + 100 - \tfrac{1}{2}p(s)100.$$

Ex ante social welfare is

$$2\left(100 - s - \tfrac{1}{2}p(s)100\right)$$

$$= 200 - 2s - p(s)100,$$

and ex ante social welfare is maximized where the first order condition is fulfilled:

$$-2 - p'100 = 0 \text{, or}$$

$$p'(s_A) = -\tfrac{1}{50}.$$

Ex *post* social welfare, however, is

$$\tfrac{1}{2}\left((1 - p(s))2(100 - s) + p(s)(100 - 2s - \tfrac{1}{2}|100 - 0|)\right)$$

$$+ \tfrac{1}{2}\left((1 - p(s))2(100 - s) + p(s)(100 - 2s - \tfrac{1}{2}|100 - 0|)\right)$$

$$= (1 - p(s))(200 - 2s) + p(s)(50 - 2s)$$

$$= -2s + (1 - p(s))200 + p(s)50$$

$$= 200 - 2s - p(s)150.$$

The first order condition for ex post welfare is:

$$-2 - p'150 = 0 \Leftrightarrow p'(s_P) = -\tfrac{1}{75}.$$

Therefore, $p'(s_P) > p'(s_A)$. Given $p'' > 0$, this means that $s_P > s_A$.

*   *   *   *   *   *