Common Goods: Law, Politics and Economics
Gemeinschaftsgüter: Recht, Politik und Ökonomie

edited by/herausgegeben von
Prof. Dr. Christoph Engel
Prof. Dr. h.c. Martin Hellwig, Ph.D.

Legal Series
edited by Christoph Engel

Volume 15

Christoph Engel/Lorraine Daston (eds.)

Is there Value in Inconsistency?

Nomos

Contents

Lorraine Daston / Christoph Engel:
Introduction: Can Inconsistency Be a Value? 9

Nancy Cartwright:
Against 'The System' 17

Lorraine Daston:
Comment on Nancy Cartwright: Against 'The System' 39

Michael Friedman:
Comment on Nancy Cartwright: Against 'The System' 47

Chrysostomos Mantzavinos:
Comment on Nancy Cartwright: Against 'The System' 57

Philip Pettit:
Collective Akrasia 63

Edna Ullmann-Margalit:
Comment on Philip Pettit: Collective Akrasia 83

Michael Baurmann / Geoffrey Brennan:
Majoritarian Inconsistency, Arrow Impossibility and the Comparative Interpretation: A Context Based View 93

David Bloor:
Majoritarian Inconsistency, Arrow Impossibility and the Comparative Interpretation: A Context Based View: Comment on Baurmann and Brennan 119

Christoph Engel:
The Impossibility of Turning Arrow's Impossibility Theorem Epistemic: Comment on Baurmann and Brennan 129

Michael Baurmann / Geoffrey Brennan:
A Reply to Christoph Engel's Comment 141

Russell Hardin:
Why Know? 149

Lorraine Daston:
Comment on Russell Hardin: Why Know? 171

Perri 6:
Dirty Institutions, Costly Workarounds and Fragile Settlements:
Understanding Lay Knowledge: Comment on Russell Hardin 177

William C. Wimsatt:
Inconsistencies, Optimization and Satisficing: Steps towards a
Philosophy for Limited Beings: Comment on Russell Hardin 201

Christoph Engel:
Inconsistency in the Law: In Search of a Balanced Norm 221

Edna Ullmann-Margalit:
Inconsistency and the Law: Some Observations:
Comment on Christoph Engel 283

Christopher Kutz:
Consistency's Value: Comment on Christoph Engel 291

Perri 6:
Viable Institutions and Scope for Incoherence 299

David Bloor:
Comment on Perri 6: Viable Institutions and Scope for Incoherence 355

Steve Rayner:
Back to the Roots? Comment on Perri 6 363

Alex Kacelnik, Cynthia Schuck-Paim and Lorena Pompilio:
Inconsistency in Animal and Human Choice 377

William C. Wimsatt:
Optimization, Consistency and Kluged Adaptations in Evolutionary
Theory: Can Maximization Survive? Comment on Alex Kacelnik et al. 399

Michael Baurmann
University of Dusseldorf, Institute for Social Sciences

Geoffrey Brennan
Social and Political Theory Program, RSSS, ANU
Political Science Department, Duke University
Philosophy Department, University of North Carolina, Chapel Hill

Majoritarian Inconsistency, Arrow Impossibility and the Comparative Interpretation: A Context-Based View

I. Introduction

For a social scientist of a particular kind, the first thought that springs to mind in association with 'inconsistency' is Arrow's famous impossibility theorem. The association is a natural one because, in the original formulation, inconsistency plays a double role in Arrow's theorem. At the first level, the requirement that any aggregate 'community' ranking of social states be transitive is akin to a kind of consistency requirement. At a second level, the impossibility result itself is a demonstration of logical inconsistency between various desiderata that are proposed for the aggregate social ordering – of which desiderata, transitivity (or consistency) at the substantive level is one.

Accordingly, this paper will be focused on Arrow's impossibility result in a variety of contexts in which it can arise. The issues that we shall raise here are of some interest in their own right. However, they also serve, we think, to throw some light on inconsistency as a phenomenon. Or at least, they isolate various different settings in which inconsistency might arise and suggest how, if at all, the inconsistency in question is to be resolved.

In fact, the Arrow impossibility theorem has long been a source of concern and puzzlement in the economics profession – and for that matter more broadly. The concern is reflected in the large literature that the theorem has generated, under the rubric of "social choice theory". This literature now amounts to a recognised field of specialisation within economics and related disciplines.

The puzzlement is reflected in the rather wide range of interpretations of the basic result. So, for example, the theorem has been interpreted as showing both:

- that the concept of the public interest is meaningless. As one influential commentator puts it: "it is not stating the case too strongly to say that Arrow's theorem and the research that it inspired wholly undermine the general applicability or meaning of concepts such as the 'public interest' and 'community goals'" [Ordeshook (1986) p 65]; and
- that majoritarian democracy is a totally unsatisfactory mechanism for collective decision-making. For example, according to Mackie (2003), "the main intellectual trend in American political science is the view that democracy is chaotic, arbitrary, meaningless and impossible. This trend originated with economist Kenneth Arrow's impossibility theorem". [p 2]

By contrast, some of the early assessments tended towards the dismissive. Buchanan, for example, argued that the basic result is rather obvious: on Buchanan's view, no-one should ever have expected that aggregative decision-making would exhibit the rationality of individual decision-making, so the whole exercise is really a response to a pseudo-problem [Buchanan (1954/1999)]. In a somewhat similar spirit, Tullock (1967) has argued that the theorem is largely irrelevant because the cases in which the impossibility result arises are statistically extremely unlikely.

There is, it seems, widespread professional consensus on the proposition that Arrow's theorem is important. But no consensus on what exactly the theorem amounts to. In this paper, we want to offer our own interpretation, and use it to make a general methodological point – we think an important one – about how to do ethical analysis. In the process, we shall be concerned not just with the Arrow theorem but also with its close analogue in the epistemic domain – the problem, as we shall put it, of 'majoritarian incoherence'. We think that consideration of the epistemic analogue aids understanding of what is driving the Arrow result. But perhaps more important, the wider range of applications that are thereby admitted allows us to illustrate the importance of *contextual details* in any satisfactory resolution of the Arrovian challenge, properly interpreted.

The basic methodological proposition is this: because, as a general rule, trade-offs between various normative criteria will be required, normative criteria need to be specified in 'comparative' terms – in terms, that is, of more and less – rather than in terms of stipulations that are either met or not. We believe that if Arrow's theorem had been formulated with this methodological principle in mind, much of the confusion (and one source of apparent 'inconsistency') would have been removed.

Furthermore, this alternative formulation would have served to emphasise the fact that the appropriate terms of the relevant trade-offs are themselves often heavily influenced by details of the choice context. Both the basic methodological principle and the emphasis on the importance of context are general propositions that have a life outside the Arrow application that we focus on here. However, Arrow's theorem is a good setting for illustrating the points; and, as we say, is a setting especially relevant for thinking about 'inconsistency'.

The paper is structured as follows. In section II, we briefly describe the Arrow theorem and its epistemic analogue. In section III, we lay out the methodological principle of 'comparativity' and the interpretation of the Arrow theorem that that principle invites. In section IV, we use the varieties of context in which the Arrow theorem and its close relatives apply, to indicate the role of contextual details in resolving the trade-offs that the theorem implies. Section V underlines some of the implications for 'inconsistency' more generally understood.

II. Arrow and Majoritarian Incoherence

Arrow's theorem states roughly that there exists no way of aggregating individual rankings of social states into a 'social ranking' that has certain apparently desirable properties. The properties in question are: domain completeness (ie no possible individual ranking can be excluded on a priori grounds); independence of irrelevant alternatives; the Pareto principle; transitivity of the social ranking; and non-dictatorship. Formally, the result is quite general. It can be interpreted as involving the aggregation of rankings based on individual *values* or on individual *preferences*. It can be understood as a proposition in *ethics* (ie the denial of a 'social welfare function' with the requisite properties) or in the theory of political institutions (ie a denial of the existence of any collective decision-making procedure that has the relevant properties).[1]

Furthermore, with a little massaging, the theorem can be reconfigured in epistemic terms, so as to refer not so much to the aggregation of individual values but to the aggregation of individual beliefs.[2] This latter fact should not in itself be surprising. Suppose that everyone has the same basic values but that people have different beliefs about how those values are best realised in the world. In this setting, all differences in individual rankings of social states can be traced to differences in beliefs rather than differences in values, but the formal impossibility result will still go through. Put another way, the phenomenon of majoritarian cycles, central to the Arrow impossibility result, has an analogue in terms of majoritarian incoherence. Just as a 'majority' can prefer a to b, and b to c, and c to a, without any individual exhibiting intransitive preferences, so a 'majority' can believe x, and believe y, and yet not believe (x *and* y) or any logical implication of (x *and* y) even though all individuals hold coherent beliefs. The similarity of the two claims is self-evident. We can just think of x as the proposition 'a is preferable to b', and of y as the proposition 'b is preferable to c'. And the reason for the apparent anomaly is the same in both cases – namely, that the composition of the 'majority' in question changes across preferences/belief-states.

The familiar cyclical preference case is illustrated in Table 1 for the case of three individuals denoted I, II and III, and three options a, b and c.

Table 1: Preferences of:

	I	II	III
Most preferred	*a*	*b*	*c*
not most nor least	*b*	*c*	*a*
Least preferred	*c*	*a*	*b*

Here, the majority consisting of I and III prefers a to b;
the majority consisting of I and II prefers b to c;
while the majority consisting of II and III prefers c to a.

We cannot rule out a preference structure like that illustrated in Table 1 on *a priori* grounds (Arrow's domain completeness requirement); and we want any 'social preference/ranking' to reflect only individual preferences/rankings (the Pareto requirement). So the only way to suppress inconsistency in general is to have some individual (at least

1 Buchanan (1954/1999) makes much of this distinction. Sen (1999) favours the latter interpretation.
2 For an attempt to do the formal analytics, see List & Pettit (2002a; 2002b).

one) who is a member of every decisive majority (violating the non-dictatorship requirement).

A simple verbal reformulation would, as we have mentioned, allow us to translate Table 1 into its epistemic analogue. But it is helpful to offer a slightly different depiction – as in Table 2. Here, x and y are independent propositions and z is some logical equivalent of $(x \wedge y)$ (ie of x and y). The individuals are denoted I, II and III as before. And each is stipulated to have coherent beliefs in the sense that those beliefs obey the laws of logical entailment in each case. So, the overall structure of belief in our strategically selected example is as illustrated in Table 2.

Table 2: Proposition

individual	x	y	$x \wedge y \Leftrightarrow z$	z
I	true	true	true	true
II	true	false	true	false
III	false	true	true	false
majority	true	true	true	false
composition	I and II	I and III	All possible	II and III

Reading across rows allows us to verify that each of I, II and III have totally coherent beliefs. That is, I believes x and y so believes z as logic requires him to. Equally, II doesn't believe y and hence is not logically required to believe z; and here does not in fact believe z. Finally, III doesn't believe x so is not logically required to believe z (and doesn't). But although each of the individuals has totally coherent beliefs, this coherence property need not translate to the beliefs of majorities – as the structure shown reveals. A majority believes x; and a majority believes y; but there is no majority that believes (x *and* y), and so no majority that believes z. (Here, the majority believes ~z.)

In section IV below, we shall develop examples of this epistemic case – known in the literature as the 'discursive paradox' or the 'doctrinal dilemma', depending on context. Our aim in developing those examples will be to contrast various of these epistemic cases both with each other and with the preference/value cases that Arrow originally focused on. The examples will, we hope, help illustrate the kinds of problems that the 'impossibility results' expose. We also hope that they will serve to relieve any suspicion that the impossibility result is of the 'quaint contrivance' type – of which economists and philosophers are sometimes excessively fond. We think there are real issues at stake here. But some of those issues are methodological in nature; and the 'bottom line' in this paper is really directed at those.

III. Comparative and Categorical Thinking

One of the differences between economists and philosophers is that the latter tend to think categorically, while the former tend to think comparatively. As with any generalisation, there are exceptions – and as we shall argue, Arrow's impossibility theorem is one such. Nevertheless, there are such things as general disciplinary inclinations and we think the comparative/categorical distinction is one such.

The difference is especially conspicuous in the context of normative analysis. So, for example, the philosopher's first instinct is to divide the set of objects of evaluation into mutually exclusive classes – good and bad; permissible and impermissible; free and unfree; just and unjust; and so on. Analysis proceeds by attempting to stipulate the defining properties of the normative category under consideration – what does true justice require? what is entailed in a proper understanding of freedom? and so on.

The economist's instincts are different. Economists are inclined to the view that the right way to view normative questions is through the lens of choice (usually choice of action in some domain); and choice always involves something forgone – often something *desirable* forgone. That is, choice is almost invariably exercised between options, each of which is less than totally ideal. This fact is one aspect of 'taking feasibility seriously' and applies whether the objects of choice are apples and oranges or ends of a more basic normative character. Achieving total justice or total freedom or total human flourishing won't normally be accessible – either because it is literally infeasible or because it 'costs too much' in terms of other desirable things forgone. Trade-offs among the various dimensions of normative desirability will be necessary. And the trade-offs that are made can be more or less satisfactory.

On this view, the failure, for example, to achieve the ideal of 'true justice' doesn't mean that all bets are off in the justice domain. It means rather that we must search for a decent metric of justice/injustice which will enable us to determine in the range of cases in which choice must be exercised which of two situations is the less just (or the less unjust). And we require a similar metric of, say, liberty, so that in trading off liberty for justice – as we will invariably have to do *at the margin,* as economists are fond of saying – we can assess just how much liberty forgone an additional measure of justice will involve.

Broome (1999), for example, puts the point this way:

> "... economists ... think naturally in comparative terms, and when they come to think of good, they will naturally ask not what is good but what is better than what. Philosophers seem not to have this same instinct to think comparatively." [p9]

And for Broome, this is among the things that philosophers can learn from economists. As Broome puts it, the important concept in ethics is not 'goodness' but 'betterness'. (Broome also argues that goodness is 'fully reducible' to betterness. But that, whatever exactly is at stake in the claim, is a different one. Betterness could be a supe-

rior 'dimension' for ethical choice – even if we had to give something up in order to get the advantages that the betterness-perspective offers.)

Consider Arrow's formulation in this light. When Arrow formulates what is to be aggregated, he clearly reflects this kind of approach. He focuses on *rankings* of social states (or whatever the objects of evaluation are): individual rankings are what is available and a similar 'social ranking' is deemed to be required for the collectivity. However, when Arrow comes to postulate the requirements of an appropriate aggregation procedure, he reverts to the 'categorical' mode of reasoning. He lays down his desiderata, not in terms of desirable features that might be more or less realised, but in terms of requirements that are either fulfilled or not. And the fact that his result is emerges as an *impossibility theorem* is entirely an artifact of that mode of formulation!

To see this, let us simply take Arrow's theorem at face value. And let us assume with Arrow that the only information about individual preferences/values available is that contained in the rankings (ie no inter-personal measure of preference intensity is available). Then the obvious conclusion to be drawn from Arrow's result is that something has to give. That is, *some trade-off among the various specified meta-level desiderata – Paretianism, completeness, transitivity, non-dictatorship, etc – will be necessary.*

Arrow's formulation suggests that the nature of the trade-off required is that we shall just have to give up on one or other of these desiderata. But that in itself is rather odd. Extrapolating from other more prosaic cases of trade-offs, one would think that the 'optimal' response would involve trade-offs at all relevant margins. One doesn't give up consumption of oranges entirely because one cannot have all the oranges (and apples) one wants. We do not give up on liberty entirely because that would entail some sacrifice of distributive justice (supposing for argument's sake that it would). Rather we seek to trade the one object of desirability off against the other until we have an optimal combination.

Why should Arrow's case be any different? Is it meaningless to think that we might formulate domain completeness and transitivity and non-dictatorship in comparative terms? Is it the case, for example, that if we give up on domain completeness at all, we might as well give up on it completely? Is any lapse of Paretianism ethically equivalent to giving up on the idea entirely? Might we not imagine a metric of Paretianism a little like that of statistical correlation between individual and collective rankings?

One might perhaps be suspicious of the ethical status of an ethical norm that is such that, if violated in the slightest degree, loses all its ethical authority. Perhaps there are deep normative reasons why the properties of a good aggregation mechanism have to have this all-or-nothing character. But if this *is* so, then some argument would seem to be required to show why. We all think that dictatorship is undesirable in general – but is it really true that we cannot intuitively think of systems that are more and less dictatorial? Or that the degree of dictatorship does not matter? The fact that devising a satisfactory measure of 'the degree of dictatorship' (or for that matter any other of Arrow's stipulations) is difficult, is not in itself decisive.

We are not going to attempt ourselves to specify appropriate metrics for the various Arrovian criteria. To do this at all plausibly is a very serious challenge, and would take more thought and creativity than it is reasonable to expect here or would be appropriate to our ambition for this paper. Our object here is really just to emphasise that possession of such metrics is usually taken by economists to be an essential pre-requisite for proper normative theorising, and that if these were available, trade-offs between the various desiderata that did not involve implausible 'corner-solutions' would be on offer.

But more can be said here. For suppose that we did have access to some measure of 'more or less Paretianism' [P]; 'more or less completeness' [c]; 'more or less transitivity/consistency' [t]; and so on. Then overall normative evaluation could be formulated in terms of a general function $\Omega(P, c, t,\ldots)$, where the various partial derivatives Ω_P, Ω_c, Ω_t are all positive (ie more P more t and more c all add to the normative desirability of the aggregation rule). Then selecting the 'best' aggregation rule would be a matter of confronting the set of feasible options in the particular setting with the relevant Ω-function. Under this formulation, there would be no paradoxes and no impossibility theorems on offer. The functional form together with the metrics of desirability in the various dimensions would serve to isolate the 'best choice' among available aggregation rules (or decision procedures).

Seen in this light, it seems clear that the more extravagant interpretations of the original Arrow result are decidedly overblown. There is no denial of the notion of the common good. Nor of the possibility of devising a best procedure for the determination of that good, on the basis of what we can know of individual preferences/values. The resultant decision-procedure won't be 'perfect'. But, for the economist at least, that is no big deal. There are very few contexts where trade-offs don't have to be made. The reality of life in the economist's world-view is always a matter of getting along as best we can in a less than totally perfect world. The important thing in that world is not to do 'perfectly' – which is impossible – but rather to do as well as one can: and that in principle is always possible!

One point of clarification might be in order here. Arrow assumes that individual rankings do not admit of a common metric: his concern is to derive a collective ranking in a setting where inter-personal comparisons of utility are completely impossible.[3] One might think that, in this spirit, he is also logically committed to absence of trade-offs between his meta-level desiderata. But he is not. The claim that inter-personal comparisons of individual utilities are impossible does not mean that all comparisons of any sort are impossible. Arrow's meta-level norms are independently derived ethical principles, and there is nothing to suggest that those ethical principles are themselves totally incommensurable.

---

3 This is of course a further dimension in which the stringency of the Arrow formulation might be relaxed. It is one thing to argue that we cannot know everything about different individuals' utility levels – and another thing entirely to argue that we can know absolutely nothing! There seems to be considerable scope for inhabiting the middle territory here, wherever exactly the divides between the 'knowable' and the 'unknowable' and the merely suspected are to be drawn!

If this is accepted, then we must conclude that the *impossibility* character of Arrow's result is essentially an artefact of his formulation – and specifically of formulating his meta-level norms in categorical rather than comparative terms.

IV. Questions of Context

Our interpretation of the general Arrow result is that trade-offs between his desiderata will be necessary. There are two ways in which such trade-offs might be formulated. One is to retain the original formulation of the Arrow meta-level norms and ask which of those requirements it is best to give up on – noting that, when we give up on any one of them to any degree, we must give up on it entirely. The other is the method we argued for in the preceding section – namely, develop appropriate metrics of the things to be traded off, and then optimise continuously across the whole set of desiderata.

Whichever of these approaches is adopted, the appropriate trade-offs are likely to be sensitive to context. In the next section, we want to illustrate this claim by appeal to a variety of contexts in which the Arrow theorem might be understood and applied. It is, however, worth noting that the relevance of context is a matter that the original Arrow formulation tends to suppress. An impossibility theorem is, after all, just that – a *theorem*. It applies in *all* contexts where its axioms are applicable. Formulating the theorem in terms of the Ω-function not only underlines the necessity of trade-offs but is also suggestive of the idea of context-dependence. This is so, in part because the economist will naturally tend to assume that the general form that the Ω function will take will be that usually ascribed to utility functions – namely, that it will exhibit diminishing marginal goodness in each argument. In other words, the trade-offs between ethical desiderata will be pretty much like trade-offs between apples and oranges. The standard assumption is that one will be prepared to pay a higher price in apples forgone to get another orange, the more apples and the fewer oranges one has. The analogous assumption in the ethical case is that one will be prepared to forgo more liberty in order to reduce injustice, the more liberty one has and/or the greater the injustice. This seems generally plausible to us, but we do not think anything much of what we say in the following sections hinges on that assumption. And as it happens, these are not the contextual differences that we intend to focus on.

Contextual effects are not limited to issues of 'convexity' of the Ω-function. Ethical choices are exercised within a specific social environment. To implement ethical principles we need a fabric of tools and instruments – social, political and legal institutions that embody those principles or serve to give them efficacy. And the social institutionalisation of the underlying ethical principles does not necessarily leave the principles themselves untouched. Significant changes may be required. A utilitarian principle which could well be justified as a guiding maxim in an 'ideal' world with no frictional resistance, may not be suitable in a real world in which the principle's application is entrusted to political institutions based on force and hierarchy. It might well be, for example, that the whole idea of 'rights' is best understood against the background of problems of institutionalisation. According to this view, rights are to be thought of as second-best solutions for an imperfect world – optimal formulations against a naturally interventionist political system. In this sense, contextual matters can bite pretty deeply. They can operate not just to shape the terms on which independently justified ethical principles are to be traded off: they may be interwoven with the principles themselves, change the ethical maxims that operate in the arena of action, and/or constitute an intrinsic part of justifying argument.

More important for what follows is the lesson the social scientist can teach the economist when pointing to contextual factors. Because, as important as the economist's focus on trade off and metrification in ethical questions, is the insight that there may well be no 'context-free' ethical indifference curves – or at least, none that are defined over the objects of ethical choice that the world serves up. For example, the relative value of transitivity/consistency and non-dictatorship may be quite different in the context of voluntarily exchange institutions from that prevailing in the context of institutions that are intrinsically coercive. The terms of the ethical trade-offs, and the practical ethical principles themselves may differ considerably across different arenas of choice and different institutional settings. Much like his philosopher-colleague, the economist can easily underestimate the impact of different kinds of institutions for the comparative value of competing ethical principles.

V. The Specific Cases

In what follows, we want to point up the variety of contexts in which the Arrow problem and its close cousins can arise. As indicated, our object in emphasising that variety is to suggest that the terms of trade between the various desiderata are likely to vary between contexts. As indicated in the previous section, we believe that the most plausible way of dealing with such trade-offs is by developing a metric for each of the meta-level criteria that Arrow postulates and comparing situations that have a bit more of one and a bit less of others. However, because such metrics have yet to be devised, the 'trade-offs' in question take a much less nuanced form: we simply ask in each case which of the desiderata seems to be the lowest priority. This procedure gives us a way of distinguishing between contexts, but as should be totally clear it falls well short of what we consider is ultimately required.

On the basis of this more modest ambition, we shall develop a catalogue of specific cases, indicating in each, the apparent value of consistency/transitivity and the significance of the non-dictatorship condition. The cases we shall consider are:

- the Arrow theorem in its original setting, with political preferences conceived along *homo economicus* lines;
- the original setting, with 'values' rather than preferences as the rankings to be aggregated;

- the Arrow theorem in its individualised version suggested by Kavka (1991);
- three different variants of the Arrow theorem in the belief context.

The latter examples take up most of the discussion – partly because they are the more novel versions of the Arrow family. We shall indicate the shape of our results at the end of this section via a simple summary 'score-sheet' focusing on just two of the meta-level criteria: non-dicatorship and transitivity/consistency.

1. The Relevance of Transferable Property

Imagine that we are operating in a domain where there is some good (income) that is a universal object of desire and that can be transferred among individuals. This is the kind of domain that economists normally have in mind when they consider the operation of collective decision mechanisms. And it certainly seems to be the kind of domain envisaged by Arrow in his original formulation.

Note that in such a setting, there is always an incentive for any agent to use any power he exercises to the disadvantage of those over whom that power is exercised: the fact of transferable goods lends all social relations a certain intrinsically rivalrous character. A decisive majority, for example, can always gain by maximally exploiting the corresponding minority. The outcome that is best for me is likely to be the worst outcome for all others. Of course, I may be constrained by my own internal morality or dispositions of benevolence not to reduce you to penury. But the fact that maximising my own income involves minimising yours is a significant contextual fact.

Now, in standard public choice theory, individual agents are normally ascribed what we might call *homo economicus* motivations. That is, they are assumed to exercise any power they possess in their own narrow interests. In other words, in such a setting, individuals not only have the capacity to redistribute virtually everything in their own favour; but they also have the presumed desire to do so. The individual rankings of social states therefore have an entirely predictable character: each agent values most highly the outcome where she gets almost everything and others get almost nothing. It is entirely clear that in such a setting, non-dictatorship gets a very high value. We would be prepared to give up on consistency/transitivity of the social ranking in order to avoid dictatorship[4].

Consider now the political setting with a different account of political motivations – an account that lends greater weight to the role of moral and other symbolic elements in voter choices.[5] In this setting, it is still possible that different outcomes will involve

4 This seems to be the preference setting assumed by Buchanan (1954). And this is his conclusion: allow majority coalitions to vary even if it means intransitive choices, because maintaining the same coalition as dominant just leads to minority exploitation.

5 Perhaps along the lines developed in Brennan & Lomasky (1993) or in a very different way by discursive democrats, presumed to operate under the constraints of something like 'ideal speech' conditions.

very different distributions of transferable goods; but there is no necessary presumption that the individuals (electorally revealed) rankings of such outcomes will be dominated by unconstrained self-interest. Each citizen is *vulnerable* to exploitation. But making one person dictator will not predictably lead to exploitation. In this setting, non-dicatorship remains highly valued vis-à-vis transitivity/consistency but not as highly as in the *homo economicus* case.

As a next step, consider what we might think of as a 'values' setting in which the domain of collective choice does not include some interpersonally transferable good.[6] And let the values issue be one of significance and considerable controversy – for example, whether abortion is to be legally permissible or not. Of course, each will rationally want her own values to prevail: there is natural competition among persons in that sense. But there is no intrinsic advantage to me from having everyone else's values frustrated: there is no respect in which I can benefit just by virtue of the fact that others lose. In the values context, it is at least conceptually possible that all could agree on the 'best' outcome. The world of interests is, by contrast, an intrinsically rivalrous world, a world of necessary conflict. In the values setting so described, non-dictatorship remains of considerable value; but that value is less than in the more naturally rivalrous settings. And so transitivity/consistency could be expected to make relatively stronger claims.

In the purely ethical setting, non-dictatorship is desirable but perhaps not overwhelmingly so. In the political setting, the risk of losing virtually everything makes non-dictatorship a predominant concern. But the force of that concern is influenced by what one believes the content of political preferences is likely to be.[7]

2. The Individual Arrow Problem

Kavka (1991) has argued that the Arrow problem can plausibly arise in the individual case when the individual's values are multi-dimensional. Kavka's main conclusion here seems to be that the distinction, routine among economists, between collective and individual decision-making processes is over-played. However, since collective decision-making involves all the complications of the individual case as well, Kavka thinks that "collective choice *is* more complicated than individual choice". But if our general line is right, Kavka's approach is ultimately rather misleading. Kavka's emphasis is on the question as to whether we can imagine that Arrow-like problems might arise in the indi-

6 Arrow himself seems to reject quite explicitly the possibility of any real distinction between preferences and values. We think that is a mistake. But in any event, we can surely distinguish between the Arrow theorem as a proposition about collective-decision-making mechanisms and as a proposition in a kind of individualistic ethical theory.

7 One way to think about the claims here is in terms of domain restrictions – not restrictions imposed as an independent normative exercise, but restrictions that are likely to arise 'naturally' by virtue of the way the world actually is. Domain restrictions of this latter kind seem to be entailed – and precisely not ruled out – by the requirement that we 'take feasibility seriously'. Universal domain (across the set of individual rankings) might well turn out to be an unnecessary (and highly profligate) constraint.

vidual setting. He therefore focuses on questions like the intra-personal comparability of rival values. But even if such problems might arise, there remains the question – in our view the crucial question – as to whether the priorities between the various desiderata are similar to those prevailing in the multi-person case. And we think that they are not. Dictatorship seems to involve a problem of a quite different order when the conflict is between different considerations within the individual's calculus from when the individual stands to have all his property confiscated by some other. For the within-person case, on the other hand, intransitivity may well be a larger problem because of the possibility of exploitation by others via Dutch-book/money-pump devices. If the real force of the Arrow theorem is that trade-offs between desiderata are necessary, the real distinction between individual and collective settings is whether the trade-offs are similar between the settings. If, as we believe, they are not, then the collective/individual distinction continues to bite – though the distinction is to be understood in different terms from those the economist normally applies.

3. Beliefs

When we come to examine the generic differences between the Arrow result in its original setting and the epistemic analogue, it is tempting to think that the belief case, like the values case, is one where dictatorship considerations are less likely to be crucial. And in lots of cases we think that is so. But much depends not just on the fact that beliefs are in play, but also on what the beliefs are about. We made the observation earlier that the Arrow problem could be transformed into a belief problem by the simple expedient of treating the assessments of alternative outcomes as propositions about the relative desirability of those outcomes. When the stakes are high – when for example the choice concerns who is the best person to be sent on the dangerous mission – it seems unlikely that one would give up one's right to have a say lightly.[8] In other words, much depends on what the beliefs are *about:* and in settings where others may actually gain from outcomes that are bad for you, your own participation in the decision-making process may be important.

In what follows, we examine a variety of cases in which majoritarian incoherence may arise, directing attention in each case to the relative significance of consistency/transitivity and non-dictatorship considerations.

8 Of course, a right to 'have a say' is not a right of veto. In large number settings, the absence of a dictator may not actually buy you much in security against collectively approved outcomes that are very bad for you.

a) The Medical Case

You are seeking advice from a body of medical experts. There are three of them. They are offering counsel; but the final decision on action is your own. There are three issues to be addressed:

- do you have the disease?
- is the (standard) operation appropriate to your case?
- what is the best timing for the operation?

The structure of beliefs is as indicated in Table 3. Doctor 1 believes that you have the disease and that the standard operation is the most appropriate treatment. But he thinks that the operation should be in six months time.

Doctor 2 believes that you have the disease. She does not believe that the standard operation is appropriate to your case. But if you are to have the operation the right time to have it is now.

Doctor 3 thinks that the standard operation is the best treatment in cases of this disease, and thinks that the earlier the operation is performed the better. But she doesn't think that you have the disease.

Everyone agrees that: if you have the disease; and if the operation is the best treatment; and if the best time to have the operation is now; you should have the operation now. But majority opinion does not follow that logic. Majority opinion is: that you have the disease, that the standard operation is the best treatment, and that the earlier the operation is performed the better. But the doctors are *unanimous* that you shouldn't have the operation now!

Your task is to use the expert advice you have at your disposal to give you the best available handle on the truth of the situation. The issue is whether you should follow the 'premise-based' mode of reasoning, using the best judgement on each of the premises taken separately and then allow simple logic to carry you to the conclusion; or to take the doctors' overall judgement on your best course of action now. If you do the former, you will follow majority judgement in each case and have the operation. If you do the latter, you will not have the operation. Which is it to be?

Table 3: The medical 'paradox'

	Disease?	**Operation Appropriate?**	**Now or later?**	**Operation now**
Doctor 1	yes	yes	later	no
Doctor 2	yes	no	now	no
Doctor 3	no	yes	now	no
majority	**yes**	**yes**	**now**	**no**

Clearly, yours is a somewhat unclear case. Either that, or expert medical opinion is more divided than is re-assuring! Still, given that you assign authority to the opinion, what is your best course? It seems plausible to us that your best course is indeed to have the operation now. That is what the majority of medical opinion is urging. The fact that all the medicos are against this course is totally incidental. The expert judgement on each of the separate premises is clear; and given that the issues are independent[9] you seem well advised to let logic take its course. After all, the fact that all the individual medicos are against that course is a mere artifact of how the negative reasons are distributed across the trio. It could just as easily have been the case that those negative reasons were concentrated in a single doctor (as illustrated in Table 3')

Table 3': The medical non-paradox

	Disease?	Operation appropriate?	Now or Later?	Operation now?
Doctor 1	yes	yes	yes	yes
Doctor 2	yes	yes	yes	yes
Doctor 3	no	no	no	no
majority	**yes**	**yes**	**yes**	**yes**

Here, the majority view on the premises is exactly as in Table 3, but the negative reasons are all concentrated in Doctor 3: no 'paradox' emerges. Would we have any greater reason to have the operation now if the structure of belief had been that in table 3' rather than that in table 3? We cannot see why. The role of the doctors is to give their best judgement on each of the three relevant propositions. How they happen to line up overall is irrelevant. Suppose for example that the three doctors involved in giving each piece of advice were different doctors: one set of diagnosticians; and two separate sets of clinicians, with somewhat different skills. We would then have had no reason to ask any one of them what he would choose overall: they would just do what you are doing – namely, go to the best counsellors for that aspect of the decision.

Now, one might think – along lines canvassed in Brennan (2003) or otherwise – that the fact that there is no consensus on any of the premises is a relevant fact. One might think specifically that one ought to be 67% persuaded as to the majority view on each of

9 One might be concerned that the issues may not be independent – in the sense that a doctor who thinks that surgery is undesirable will tend to think that surgery, if it is to be undertaken, is better taken as late as possible – say, twenty years hence! One might also insist that one cannot take any proper decision on the medical issue unless one specifies what the likely outcomes of the various courses of action are and how much those outcomes are valued. All this is true; and in its own way important. But our object here is to make a specific point about majoritarian consistency and contrast this case with others in which Arrow-like considerations are in play. For such purposes, the details of the case can be finessed.

the premises. On that basis, the probability that all of the premises holds is $[2/3]^3$ and the odds favour rejecting the conclusion. After all, there is a one in three chance that you don't have the disease; and a one in three chance that the operation will do more harm than good; and so on. Those chances combine to give you mutually supportive reasons for not having the operation. As suggested by Brennan (2003) however, it is not obvious how to ground this kind of 'mandate theory of truth' in any plausible statistical foundation. In any event, the conclusion still follows that the pattern of final judgements across the three doctors is irrelevant: however one interprets the epistemic weight of the majoritarian judgement on the premises, the particular distribution of individual misgivings across the individual doctors is irrelevant. Table 3 and Table 3' ought to give you the *same* reasons for having the operation – or not, as the case may be.

The implication of these remarks is that in this example the question of transitivity/consistency is really a total non-issue. On the other hand, it does not seem as if non-dictatorship is of great significance either. If it turns out that one of the experts is a dictator in the sense that you are following that doctor's advice throughout, this does not seem on its face to be an issue of great moment. Certainly, you would not be prepared to pay much to avoid that possibility. So although the logic of the Arrow theorem transfers to the epistemic case, it is not by any means clear that the meta-level desiderata for good decision-making that Arrow isolates are appropriate for that case.

The important thing here seems to be to ask the right question(s) – the questions, that is, on which the epistemic authority of the experts is greatest. So for example it would be a mistake to ask each of the doctors: would you recommend my having the operation now? The right questions are those on which the doctors' expertise is focused: namely, the diagnostic and clinical premises on which the medical conclusion is based. Perhaps greater epistemic authority could be secured by a yet finer division of intellectual labour. Perhaps diagnosticians differ in their abilities to read the results of different tests; and yet others have special expertise in aggregating the results of those tests in an overall judgement as to whether the condition is present or not. And even where the same doctors are used as the sources of advice on these various aspects, the division of the decision seems likely to be desirable.

The question arises as to whether there are any contexts where the doctors' decisions on the overall conclusion – whether to have the operation now or not – would be of any significance. One possibility in this connection is where the doctors who give the advice are also responsible for conducting the operation. It may well be important for the success of the operation that the doctor who performs it believes it to be necessary. In that event, there *will* be a significant difference between Tables 3 and 3'. And there will be some reason in a case like Table 3 to follow the unanimous decision of the experts – not because that decision is truth-revealing, but because of psychological effects on doctor competence in cases where doctors perform operations they believe unnecessary or inappropriate.

b) The Judicial Case

Our second example of majoritarian inconsistency is one that has already had some exposure in the literature [Kornhauser & Sager (1993); Chapman (1998)]. There is a bench of three judges who have to decide a breach of contract case. There are three issues at stake in the case:

- Was there a contract?
- Was the contract conscionable?
- Was there a breach?

The pattern of individual judgements is as shown in Table 4. Judge I thinks there was a contract; and that there was a breach; but that the contract was unconscionable – so he finds for the defendant. Judge II thinks that there was a contract; and that the contract was totally conscionable; but that there was no breach. Judge III believes that if there had been a contract, it would have been perfectly conscionable and there would have been a breach; but in fact she thinks there was no contract actually formed. So each judge finds for the defendant and against the plaintiff. Yet each finds for the defendant on quite different grounds. When the bench comes to write up the reasons for the decision, they will have a problem. because though there is unanimity on the verdict, they will not be able to agree on the reasons for that verdict. In fact, a majority of the bench believes that there was a contract, that it was conscionable and that there was a breach. If the bench were to act 'consistently' on the basis of the majority judgements on each aspect of the case taken separately, they would find for the plaintiff.

As before, we could reconstruct the example in such a way that the bench's majority position on each aspect of the case was the same as in table 4, but no dilemma arose. This would be the case when the negative reasons all happened to lie with the same judge. But in the absence of any particular reason for thinking that the minority judge was systematically less competent than his colleagues on the interpretation of the law, it

Table 4: The judges' dilemma

	contract	**Conscion-ability**	**breach**	**c + C + b ⇔ plaintiff**	**find for**
Judge I	yes	no	yes	yes	defendant
Judge II	yes	yes	no	yes	defendant
Judge III	no	yes	yes	yes	defendant
Bench majority	yes	yes	yes	yes	Defendant?

seems that we have to regard the particular distribution of views across the bench on the various aspects of the case as irrelevant.

Does the dilemma matter? Indeed it does – because the judges have two tasks before them. They must decide the overall verdict of the case before them: they must do justice to the parties. But they must also provide *reasons* for their finding – and this, because the case becomes part of the stock of precedents on the basis of which similar cases will be settled in future. Those reasons provide information not just for future judges and/or lawyers, but also for future possible contractors. These future contractors want to know what a contract *is* in the law's eyes, and what the conscionability constraints *are*, and what exactly *would* constitute a breach. In this sense, Table 4 reveals a potential serious conflict between the requirements of justice to the parties and the requirements of doctrine. And that conflict matters, because the bench would fail in its duty if its public reasons and its verdict were at odds.

So unlike the medical case examined earlier, here 'consistency' *does* have a value: we want justice and doctrine to cohere. There are of course institutional mechanisms that might support such coherence. One such mechanism involves compartmentalising the various elements in any case, so that the court has to decide on each aspect separately. The overall verdict would then just be driven by the logic of the separate decisions. In the case in Table 4, the final verdict would be for the plaintiff – notwithstanding the misgivings that all three judges would have on the justice of that outcome. To some extent, the requirement that judges write a 'majority report' justifying their verdict performs a similar function to *separation* of reasons. This requirement means that the judges not only have to arrive at a common verdict but also have to find common *reasons* for the verdict they arrive at.

In a predicament like this, it seems likely that the judges will themselves decide that something has to give. And it seems very unlikely that the thing to give up on will be consistency itself. Rather, once confronted with the inconsistency between reasoning and verdict, members of the bench are likely to adjust their views so that consistency will be preserved. There are of course a variety of ways this might be done. If judge III can persuade one or other of his colleagues that the requirements for a contract have not been met, then the verdict for the defendant can stand on that ground. Equally, if judge II can persuade one or other of her colleagues that the contract was unconscionable then the verdict can stand on those grounds. And so on. If no judge can do this persuasive work on one of the separate aspects of the case, then all may have to change their verdicts to preserve the coherence of the law.

Some such accommodation by one or other of the judges seems extremely likely. After all, these judges are likely to form a bench on other occasions. There is likely to be a certain amount of horse-trading across cases when the judges' dilemma occurs. If as a judge you never give way, you are unlikely to be able to persuade your colleagues to give way in future cases when you feel even more strongly. And the judges are appointed for, and adopt a culture of, treating the law with respect. Among other things,

this means accepting previous benches' findings, even where those findings are not entirely congruent with personal judgements. And it would seem a trifle odd to treat other benches' determinations with respect, while counting as irrelevant the judgements of one's immediate colleagues.

If in cases of unresolvable conflict, the judges decided to determine somehow[10] that one of them should act as 'dictator' and write the report on the basis of her own individual judgements, it seems unlikely that we would find this particularly objectionable. Or at least, it would not be *intrinsically* objectionable. Perhaps we might be worried that the bench was throwing away valuable expertise. Or perhaps on other grounds we might be committed to premise-based procedures. But 'non-dictatorship' doesn't seem to deserve great weight as an independent normative criterion in cases like this. Dictatorship would be just one means to secure the required consistency between doctrine and justice; and in this case, consistency is trumps.

c) The Academic Case:

In the medical case, the issue was one of pure counsel. In the judicial case, the issue was one of using the expertise of the judges in the best way to reconcile the possibly competing claims of justice and doctrine. In neither case however was there a collective action problem of the kind that often arises when multiple decision-makers are involved. Political decision-making is often conceptualised as a matter of collective action, but as we have already noted, politics involves the capacity for coercion and does so in a context where there may well be both scope for and incentives in relation to transfers of well-being within the collective.

Here, we wish to take an example of decision-making for collective action when transfers as such are ruled out. So consider the following example. A small academic centre consists of three scholars who do collaborative work. The centre has been approached by some external body to do a particular job. The three academics involved are totally agreed as to the criteria for deciding whether to accept or not:

- Is the project likely to prove intellectually interesting?
- Is the work likely to be professionally rewarding (eg by leading to publications in the better journals)?
- Will the money offered cover the full cost of the project?

And they are unanimous that they will do projects if and only if these criteria are met. There is no conflict of interest in the sense that all their work is fully collaborative, and they are assumed to have identical preferences for intellectual and professional aspects. And they have consented to operate by majority rule.

10 In determining who will be decisive, the judges might choose randomly, or take it in turns, or defer to the most senior – or assign the chore of writing the report to the most junior. What they will *not* do, it seems, is to offer reasons for a verdict other than that delivered.

Each scholar has to make a judgement on each of the three criteria, and they all take themselves and each other to be equally competent to make such judgements. The only difference between this case and the other cases already examined is that if the project goes ahead, the same individuals who pool their expertise to make the decision will have to do the work.

In a manner that should by now be familiar, majoritarian inconsistency problems can arise; and the relevant case is depicted in Table 5.

Table 5: The academics' dilemma

	Academic interest	Professional reward	Full costing	A+ P + F ⇔ accept	accept
Scholar A	yes	no	yes	yes	no
Scholar B	yes	yes	no	yes	no
Scholar C	no	yes	yes	yes	no
majority	yes	yes	yes	yes	yes

Accepting the epistemic authority of the majority view on premises means that the group should indeed proceed with the project. Each of the criteria is satisfied, according to the majoritiy's judgement. But there is a unanimous verdict *against* proceeding. This example has the same structure of individual beliefs as the medical and judicial examples: there is the same weight of logic behind the judgement that one ought to accept the project as there is that the patient ought to have the operation, for example. But there is an additional complication here. In the medical example, all the expert advisors acted strictly as experts. They gave their best judgement; and the patient was left to decide. Here, however, the same scholars who deliver the advice have to act on the advice given. The separation of action and advice characteristic of the medical case does not apply. There is therefore a tension between the 'external point of view' that each scholar brings to bear in assessing the group's best judgement, and the 'internal point of view' associated with the beliefs that the individual himself holds. I recognise that I do not hold the majority view on one of the aspects: but that fact in itself does not seem likely to be sufficient to induce me to change my belief. There is something odd about a scholar whose 'best judgement' can be unseated simply by virtue of the fact that it is a minority judgement. You might, for practical purposes, go along with your colleagues in such cases: but 'going along' is not going to alter your judgement that the choice of action is not the best one available. And of course, in this example, all the scholars are in the same predicament. They all believe that proceeding with the project is a mistake – though of course they believe this for different reasons.

The scholars could, of course, adopt a premise-based procedure, examining the collective wisdom on each aspect of the decision separately and then allowing logic to

drive the final action choice. But unlike the judicial case where the reasons have independent (doctrinal) status the only reason for following the premises here is that it gives rise to a final decision that all the scholars believe is superior. So when it comes to the implementation of the premise-based decision, it seems likely that the group will face a commitment problem. Someone is surely likely to observe that none of them 'really believes' that this project is worth pursuing. And that fact seems bound to carry some weight. Going along with one's colleagues in a case where none of one's colleagues actually believes in the action at issue seems like obeying a phantom majority.

In short, the scholars are likely to feel that they are faced with a genuine dilemma. Abstract institutionalised reasoning suggests one course of action: but following that reasoning requires one to suspend your best judgement as to the proper course of action. In such a case, the decision to just suspend 'consistency' – to decide instead on the basis of individual final judgements – doesn't seem either psychologically implausible or necessarily mistaken. On the other hand, seeing the academic example through the lens of the medical one, it is not clear that suspending the demands of consistency in this way is necessarily right either.

The simple point is that the group's epistemic authority seems to depend on whether you yourself are part of the expert advisory team or not. And this difference is not just a matter of protecting one's interests, as might be the case in more conventional 'political' settings. The difference is rather a matter of a divergence between judgements made from the 'external' and the 'internal' points of view.

* * *

Other examples of Arrow's epistemic cousin could be devised. The characteristic feature of all such cases is that there is no conflict of interest among the parties. All differences between persons are differences in their beliefs. The value to be placed on 'consistency' in such settings seems to us to vary according to context. All the cases examined are constructed in such a way that there is inconsistency between reason-based and final-action-based majority decisions.

In the medical case, where premise-based decision-making seems fundamental, consistency drives the final action decision – whether or not that decision conflicts with what the experts would choose. *In*consistency is scarcely an issue here because the experts' *overall* judgement is irrelevant.

In the judicial case, consistency between reasons and verdict is critical; but the verdict does indeed matter. So cases where inconsistency would arise constitute a problem and have to be resolved by some modification of the judges' expressed views – whether on verdict or doctrine is an open matter.

In the academic case, where the scholars act both as sources of expert judgement on premises and as final actors, there is a tension between internal and external points of view – a tension that is embedded in the requirements of consistency that each scholar imposes on his own views. In this case, however, it seems that the majority judgement on premises is likely to be less compelling. The fact that there is total consensus on the best course of action seems likely to be decisive.

We can bring this section to a close by a summary table, showing the relative values of 'consistency' and 'non-dictatorship' in the whole set of cases we have considered. We do this in Table 6 – and we trust that in the light of the foregoing discussion, its message will be self-evident.

Table 6: Score-sheet for non-dictatorship and consistency across cases

CASE	Non-dictatorship weight	Consistency weight
Politics with egoistic preferences	Very high	moderate
Values aggregation	moderate	moderate
Pure epistemic (medical) case	low	negligible
Judicial case	low	Very high
Academic case	low	low

VI. Inconsistency

What does all this tell us more generally about inconsistency?

The form of inconsistency we have been concerned with here is of a special kind. It relates to the peculiarities involved in aggregating individual values/preferences/political judgements to form a 'collective' analogue. Does the fact that the individuals all have consistent rankings of social states or consistent beliefs imply that the consistency carries over to the collective ranking/beliefs. And if not does it matter?

In the majoritarian coherence analogue to the Arrow theorem, consistency plays a double role – both as a desideratum for any decision-making process (analogous to tran-

sitivity) and in terms of consistency among the different meta-level desiderata. As we might put it, if we insist on aggregative consistency, is that consistency 'consistent' with other properties we would like the aggregate to exhibit?

Focussing on the latter aspect first, what we have tried to argue is that inconsistency between desiderata is more a matter of construction than a logical necessity. Recognition that conflicts can arise between normatively desirable properties does not necessarily invoke fundamental inconsistency issues of the kind implied by 'impossibility' results. Such conflicts are, on the economist's world-view, just a standard fact about the world: they are always present in every genuine choice situation. Conflicts only imply radical inconsistency when trade-offs are ruled out – either by some form of incommensurability, or by imposing an all-or-nothing structure on some elements of desirability. That incommensurability can be either an explicit piece of the normative structure or just an implicit feature of the way in which normative desiderata are formulated. In the original Arrow theorem, for example, the 'impossibility of inter-personal comparisons' – the fact that the only information available about individual preference/value is that contained in rankings – is an explicit assumption underlying the whole construction. By contrast, the fact that the meta-level norms applying to alternative aggregation procedures are formulated in categorical terms seems arbitrary. If those norms were specified in terms of more or less (more or less dictatorship, more or less transitivity, etc.) then the impossibility result as such would disappear. We would be left with the need to make trade-offs; but there does not seem to be anything conceptually troubling about that kind of exercise. In fact, such trade-offs are what normative judgements will normally require of us.

However, the precise terms of such trade-offs are, we have argued, likely to differ from context to context. Different Arrow-like problems suggest different priorities among the different desiderata he specifies. Within the 'majoritarian coherence' version of the Arrow theorem, 'consistency' re-appears at the substantive level as one of the desiderata – playing a role more or less analogous to the role that transitivity plays as a requirement of the aggregated value-judgement or 'social preference' in Arrow's original formulation. With inconsistency so understood, there are several conclusions we seek to draw about the cases of inconsistency we have addressed.

First, inconsistency is not an intrinsic value: in no case is it desired for its own sake. However, inconsistency can be derivatively desirable, in the sense that it can buy you things that are more preferred. Whether giving up on consistency is too high a price to pay depends on details of the setting. Sometimes it is. In some cases, consistency is an absolute requirement and demands other adjustments so that it can be secured (as in the judicial case). Sometimes, consistency is irrelevant, because the judgement on the final question is irrelevant (as in the medical case). Some cases are in between.

But this variety is exactly as we ought to expect. It is a basic message of economics that relative prices matter: and that relative prices can differ as between cases. It is perhaps re-assuring, if somewhat uninformative, to discover that this message is the bottom line of our discussion here. Sometimes the 'goods' to be traded off against one another

will be a little unconventional. And sometimes genuine incommensurability will obtrude to block trade-off possibilities. But incommensurability has to be argued for. To impose it merely as a matter of convenient formulation, as philosophers often do in formulating normative criteria, is improper. This is one context where the economist's instincts in favour of comparative formulations can do useful work.

References

Arrow, K. (1951) *Social Choice and Individual Values* Wiley, New York

Baurmann, Michael and Geoffrey Brennan (2006). Majoritarian Inconsistency, Arrow Impossibility and the Continuous Interpretation. A Context-Based View. *Is There Value in Inconsistency?* Lorraine Daston and Christoph Engel. Baden-Baden, Nomos: 93-118

Brennan, G (2001) "Collective Coherence?" *International Review of Law and Economics* **21,** 197-211

Broome, J. (1993) "Goodness is Reducible to Betterness: The Evil of Death is the Value of a Life" in Koslowski, P. & Shionoya, Y. *The Good and the Economical: Ethical Choices in Economics and Management* Springer-Verlag, Berlin 70-84 [reprinted in Broome, J. (1999) *Ethics out of Economics* Cambridge University Press, Cambridge 162-173]

Buchanan, J. (1954) "Social Choice, Democracy, and Free Markets" *Journal of Political Economy* **62,** 114-123 [reprinted in Buchanan (1999) *The Logical Foundations of Constitutional Liberty: The Collected Works of James M. Buchanan* Vol 1 Liberty Fund, Indianapolis, 89-102]

Chapman, B. (1998) "More Easily Done than Said: Rules, Reason and Rational Social Choice" *Oxford Journal of Legal Studies* **18,** 290 -330

Kavka, G. (1991) "Is Individual Choice Less Problematic than Collective Choice?" *Economics and Philosophy* **7**, 143-166

Kornhauser, L. & Sager, L. (1993) "The One and the Many" *California Law Review* **81,** 1-59

List, C & Pettit, P. (2002) "Aggregating Sets of Judgements: An Impossibility Result" *Economics and Philosophy* **18,** 89-110

List, C. & Pettit, P. (2004) "Aggregating Sets of Judgments: Two Impossibility Results Compared" *Synthese 140*, 207-235

Mackie, G. (2003) *Democracy Defended* Cambridge University Press, Cambridge

Ordershook, P. (1986) *Game Theory and Political Theory* Cambridge University Press, New York

Sen, A. (1999) "The Possibility of Social Choice" *American Economic Review* **89,** 349-378

Tullock, G. (1967) "The General Irrelevance of the General Impossibility Theorem" *Quarterly Journal of Economics* **81,** 256-270