

On the axiomatics of resource allocation: Interpreting the consistency principle

William Thomson

Working Paper No. 573
September 2012

University of
Rochester

On the axiomatics of resource allocation: Interpreting the consistency principle

William Thomson*

September 10, 2012

*Department of Economics, University of Rochester, Rochester, NY 14627. I outlined the main argument developed in this paper at the *7th Conference on Environmental and Resource Economics* held in Toulouse (2009), and I thank Stephan Ambec for his invitation. I am grateful for comments to Azer Abizada, Nanyang Bu, Siwei Chen, Jo Cho, Battal Doğan, Eun Jeong Heo, Paula Jaramillo, Özgür Kıbrıs, Juan Moreno-Ternero, Karol Szwagrzak, Rodrigo Velez, and Ayse Yazici.

Abstract

An allocation rule is consistent if the recommendation it makes for each problem “agrees” with the recommendation it makes for each associated reduced problem, obtained by imagining some agents leaving with their assignments. Some authors have described the consistency principle as a “fairness principle”. Others have written that it is not about fairness, that it should be seen as an “operational principle”. We dispute the particular fairness interpretations that have been offered for consistency, but develop a different and important fairness foundation for the principle, arguing that it can be seen as the result of adding “some” efficiency to a “post-application” and efficiency-free expression of solidarity in response to population changes. We also challenge the interpretations of consistency as an operational principle that have been given, but identify a sense in which such an interpretation can be supported. We review and assess the other interpretations of the principle, as “robustness”, “coherence”, and “reinforcement”.

Keywords: consistency principle; fairness principle; solidarity; punctual axiom, relational axiom, consistent extensions, converse consistency.

JEL classification number: C79; D63; D74

1. Introduction. A principle that has played an important role in the design of allocation rules in a great variety of contexts is “consistency”. Here is an informal statement. Consider a rule, some allocation problem in its domain, and an allocation the rule has chosen for the problem. Now, focus on some subgroup of the agents involved in the problem, and identify those alternatives in the problem at which the members of the complementary subgroup “receive their components” of the chosen allocation and leave. This defines the opportunities open to the remaining agents. The rule is consistent if in this “reduced problem”, it still assigns the same thing to each of these agents; no adjustments in their assignments are needed.

Given the central place that consistency holds in the axiomatics of resource allocation—several hundreds articles have been written about it¹—there is wide disagreement about what it means. We will assess here the view, which has been expressed by a number of authors—and an informal survey that we conducted over the last year met with this almost unanimous response—that consistency is not about fairness, that it has no ethical or normative justification or interpretation, We disagree: fairness is a possible rationale for it. Our main purpose here is to develop an interpretation of consistency based on fairness considerations. Fairness is a multi-faceted concept and of course, we do not pretend to have a definition that would cover all of its possible meanings and expressions, or that would be endorsed by everyone. We make the more modest claim that there is a perspective from which we can link consistency and fairness.

First of all, what does it mean to say that consistency is not about fairness? An answer could simply be that consistency cannot be interpreted as contributing to fairness, that one does not see how it helps achieve fairness: then, consistency would be “neutral” with respect to fairness. To say that consistency is not about fairness can also mean that consistency fails to prevent unfairness, that it allows unfairness, that perhaps it hinders fairness. Such a claim can be supported by the observation that certain dictatorial rules are consistent²; even worse, a number of axiom systems in which consistency play a central role have been shown to force dictatorship.

On the other hand, the restatement of the consistency principle as “Every

¹For a survey, see Thomson (2011a).

²The formal definition of these rules will have to wait until Section 11. We say “certain” because not all dictatorial rules are consistent. In a variable-population framework, for this property to be satisfied, the identities of the dictators have to be related in a manner that will be made clear at that point.

part of a fair division should be fair” (Balinski and Young, 1982) squarely places consistency as a fairness principle. Although our main goal is to establish a link between consistency and fairness, we will assert that this restatement actually obscures the nature of this link.

Other interpretations of consistency have been proposed, not referring to fairness, which we will also discuss. Some writers have called it an “operational principle”,³ usually without making it clear what should be understood by this expression, what it is that the principle is supposed to make operational, let alone showing how it applies in this case. The majority of workers in the field have simply referred to consistency as a “robustness”, “coherence”, or “stability” principle, essentially offering the definition itself as its own justification.

Which of these interpretations is right? We will argue that each of them either reads too much into the principle or too little, that each contains some of what underlies consistency, but not all.

To proceed, we need to define critical concepts and establish terminology. We believe that much of the confusion concerning consistency comes from a failure of making certain distinctions between concepts that should be differentiated, and that is our first task.

We begin by introducing a critical categorization of axioms into “punctual” and “relational” axioms (Section 3) and asserting the need for relational axioms pertaining to variations in populations (Section 4). After defining the concept of solidarity (Section 5), we emphasize the importance of teasing apart the various ingredients that may enter the formulation of an axiom, and in particular that of keeping solidarity considerations separate from efficiency considerations (Section 6). We also propose to distinguish between “pre-application” and “post-application” relational axioms (Section 7), and we define the concept of Pareto-indifference (Section 8).

At that point, we have the conceptual apparatus needed to state our position, which involves relating consistency and solidarity (Section 9). We also explain why the usual difficulty of working with choice correspondences instead of choice functions when studying solidarity objectives is not an issue for consistency (Section 10).

Having made a case for an interpretation of consistency that links it to

³Moreno-Ternero and Roemer (2012) write “Consistency has played a fundamental role in axiomatic analysis[...] even though it is mostly an operational (rather than ethical) axiom.”

fairness, we then ask whether consistency should be considered enough for fairness (Section 11). Not surprisingly, our answer is negative. We then assess a possible view of consistency as sometimes presenting an obstacle to fairness (Section 12), and the milder one that consistency is neutral with respect to fairness (Section 13). After an aside concerning various properties related to consistency, its bilateral version, its “average” version, and “flexibility” (Section 14), we introduce its converse and contrast it with consistency (Section 15).

We then evaluate four other interpretations of consistency: as a requirement of robustness to changes in perspectives when evaluating a change in a situation; as a requirement of robustness under “piecemeal” implementation of the choices made; as an operational principle; and as a reinforcement or “fractal” principle (Section 16). We believe that each of these interpretations is valid, but also that each provides an additional link to fairness.

We conclude with a short summary of our position (Section 17).

We assume the reader’s familiarity with several standard models that we use to illustrate various points. They have to do with selecting a utility vector from a set of feasible vectors (bargaining problems); selecting a utility vector when a set of feasible vectors is specified for each group of agents (coalitional games); dividing a social endowment of unproduced goods among a group of agents (classical problems of fair division); redistributing private endowments of goods among a group of agents (classical economies); and adjudicating irreconcilable claims that a group of agents have on a resource (claims problems). The first time we refer to any of these models, we summarize its main components in a footnote. Our exposition is non-technical and we avoid formal descriptions of models.

2. Defining the consistency principle. The consistency principle is applicable to a wide range of models, and in describing it, we use general language so as to cover as many of them as possible.

A **model** is a mathematical representation of a class of decision problems, or simply, **problems**. A problem involves a **group of agents** having to choose from a set of **alternatives**. An alternative can be described “abstractly”, as a payoff vector, a utility vector, or a list of welfare levels (we will use the term in an ordinal sense) that the agents may experience, no information being given on how these are achieved. Or it may be described concretely, as when the “opportunities” or “options” available to the agents are given with all of their physical detail. For instance, an alternative may

be a possible division of a vector of resources, or it may involve the specification of the levels at which public goods are produced, or the labor supplies that agents should contribute. Then, information about endowments, technologies, preferences, utilities and so on, is included in the description of a problem. We often refer to situations of this type as **economies**.

A **solution** is a mapping defined on some class of problems, its **domain**, which associates with each problem in the class a non-empty subset of its set of alternatives. Each point in this subset is a **solution outcome** for the problem. Depending upon the context, solution outcomes are interpreted as recommendations for the problem, or as predictions of the choice that the agents would make on their own. A **rule** is a single-valued solution.

The consistency principle is designed to help us evaluate solutions and rules in situations in which the population of agents involved may differ from problem to problem. This possibility is formalized by imagining a set of “potential agents” from which elements are drawn when specifying a problem. It is often assumed, for mathematical convenience, that this set is infinite. However, only problems involving finitely many agents are (usually) considered.⁴

The consistency principle. Consider a solution. To define its consistency, we select some problem in its domain of definition and apply the solution. We select an arbitrary one of the alternatives recommended by the solution for the problem. Next, we imagine that an arbitrary subset of the agents involved in the problem “collect their components of this solution outcome” and “leave”. We reevaluate the situation at this point: the options open to the remaining agents consist of all the alternatives in the initial problem at which the agents who departed indeed receive their intended assignments. The problem involving the remaining agents when their options are defined in this manner is the **reduced problem of the initial problem with respect to the set of remaining agents and the initial solution outcome**. The solution is **consistent** if, provided this reduced problem is in its domain,⁵ it selects for it the restriction of the initial recom-

⁴There are a few exceptions. For some models, it makes a difference to the analysis whether the set of potential agents is bounded or not.

⁵We need to add this proviso, because for some models, the issue arises. However, for most models of resource allocation, reduced problems are automatically in the domain.

mentation to the remaining agents.

In the above statement, two expressions are between quotation marks because they need not be taken literally; they can be understood metaphorically. Agents do not necessarily collect their material components of the solution outcome that is taken as starting point and actually leave, although in many models—we will discuss several—the principle is certainly compatible with this scenario. In general, it should simply be understood to mean that these agents’ payoffs, utilities, or welfare levels have been set, that they are “frozen”. Proceeding under that assumption, the question is then raised whether, when the situation is reassessed at that point, there is a need to make adjustments in the other agents’ assignments. In determining the opportunities available to these agents, we may take advantage of the fact that there might be a variety of concrete ways of guaranteeing the payoffs, utilities, or welfare levels, assigned to the members of the first group, and not care about how this guarantee is achieved. On the other hand, we may feel committed to a physical delivery of resources that yields these payoffs, utilities, or welfare levels. That is especially the case when we imagine agents to physically leave.

To see the difference, consider a “classical problem of fair division”⁶. An allocation decision specifies a bundle for each agent, and when an agent leaves, we can imagine that he does so with the particular bundle assigned to him. Alternatively, we may think of his leaving only with the guarantee that in the end, he will be assigned a bundle that he finds indifferent to it. We should be careful about this second interpretation however because then, the reduced economy can not be described in the same way as the initial economy, namely as a list of preference relations together with a social endowment: the domain of classical problems of fair division is not **closed under the reduction operation**, whereas it would be according to the scenario we described first.

In an abstract model, we do not have the option of thinking that agents leave with physical resources. Nevertheless, and although the consistency principle can simply be seen as a “cross-population” principle, for convenience we will often use language that evokes actual changes in populations.

⁶Such a problem is defined as a list of preference relations defined over some Euclidean commodity space, one for each agent, and a vector of resources in that space, the social endowment. An allocation is a list of bundles, one for each agent, whose sum is equal to the social endowment. *Consistency* is first studied in this context by Thomson (1988).

Consistency belongs to the category of **invariance axioms**: starting from some problem involving some group of agents and possibly given a solution outcome of that problem, a second problem is introduced that bears a certain relation to the first problem and possibly that solution outcome; the requirement is that this outcome, or its restriction to some subgroup of agents, should be a solution outcome of the second problem.

The reason why we prefer speaking of consistency as a “principle” is to emphasize the wide scope of the concept. In each model, when applying it, we are led to a specific consistency “axiom”. There is no universal, model-free, way of reducing problems. In fact, for some models—the domain of “transferable utility coalitional games” (von Neumann and Morgenstern, 1944)⁷ is a particularly striking example in this regard—the alternatives that a reduced problem may consist of can be specified in several plausible ways.⁸ However, it is also true that almost always, there is a “most natural” way of writing down an axiom that best translates into a mathematical statement the spirit of the consistency principle.

Externalities are not easily accommodated by the consistency principle if agents leave “for good”: indeed, if the mere presence of a particular agent affects the welfare of someone else, his departure has effects on this second agent over and beyond the fact that resources have been assigned to him, the first agent. Then, we would have to insist on the metaphorical interpretation of agents “leaving”.

Finally, as usually formulated and as we presented it, the consistency principle pertains to decreases in population, but we can also write it for increases in populations, although doing so may not be quite as natural. The restatement is as follows. After applying the solution to some problem, we augment the population of agents and specify a problem for them, which we solve too. We check whether the reduction of the second problem with respect to its solution outcome and the initial population of agents is the initial problem. If that is the case, we require that this second solution outcome be the result of an “augmentation” of the solution outcome of the

⁷A transferable utility game is defined as a list of numbers, one for each subset of the agent set, or “coalition”, that is interpreted as what the coalition can achieve. It is the “worth” of the coalition. A solution selects a set of payoff vectors whose coordinates add up to no more than the worth of the “grand coalition”, the coalition of all the players. *Consistency* is first studied in this context by Davis and Maschler (1965).

⁸Three definitions have played a major role, and numerous variants have been proposed too.

initial problem: this means that it can be obtained from this initial solution outcome by “adding” components interpreted as what the new agents are supposed to receive.

3. Punctual versus relational axioms. A **punctual** axiom applies to each problem separately and a **relational** axiom relates choices made across problems that are related in certain ways.

A relational axiom makes conditional statements: *if* a certain alternative has been found desirable for a particular problem, for whatever reasons, *then* for each problem that is related to that initial problem in a particular way, an alternative that is related to the initial one in a particular way should be declared desirable. A relational axiom assesses how changes in problems should be reflected in changes in the choice of alternatives. By contrast, a punctual axiom makes a recommendation for each problem on its own, without reference to how other problems should be settled.

To see how relational axioms may help us evaluate rules, in fact, why relational axioms may well be needed, consider the following “claims problem” (O’Neill, 1982).⁹ There are two agents, whose claims are 20 and 30; the endowment is 10. Does the division (3, 7) provide a good resolution of it? Here, we ask the reader to think about this particular problem *only*. It is quite difficult, almost impossible, to confidently answer this type of questions. Certainly, we may notice that agent 1 is assigned less than agent 2, which seems to be a minimal test that an awards vector should pass because his claim is smaller than agent 2’s claim.¹⁰ But such considerations do not take us very far.

Next, what about the division (4, 6)? Noting, or being told, that this division is the proportional outcome ((4, 6) is proportional to (20, 30)) makes it a little easier to form an opinion, the reason being that the expression “proportional outcome” suggests that a formula, the proportional formula, has been applied, whereas the division (3, 7) seemed to come out of thin air. When an outcome is generated by a rule, especially a familiar one, it gains credibility, so to speak. More importantly, by placing the specific problem

⁹Such a problem is defined as a list of claims over a resource, or “endowment”, one claim for each agent, and these claims add up to more than the endowment. A solution recommends for each claims problem a list of numbers, one for each claimant, whose sum is equal to the endowment. It is an “awards vector” for the problem. *Consistency* is first studied in this context by Aumann and Maschler (1985).

¹⁰The punctual axiom of “order preservation” (given two agents, the one whose claim is greater should be assigned at least as much as the other), is met.

we face in a broader context, we do not make things more complicated. The opposite is true: this context helps us understand better how to deal with the example. By thinking about rules instead of individual problems, we can more easily form a judgment because we can evaluate recommendations in relation to each other.¹¹

Indeed, once the scope of our enquiry shifts from a single claims problem to an entire domain of claims problems, and from the search for a single awards vector for a specific problem to the search for rules defined on this domain, an additional battery of tests, the relational axioms, become available to guide us. For instance, let us keep the claims vector fixed at $(20, 30)$ and vary the endowment. It is a little easier to think about how to divide an endowment of 25 *if* the choice of $(10, 23)$, say, is made for the endowment of 33, or at least to put constraints on the variations that are needed to accommodate the decrease in the endowment.

4. Consistency as an expression of coherence of choice. Thus, as is the case for all relational principles, what underlies consistency is, first of all, the desire for something like “regularity”, “coherence”, and “predictability”. This is not saying much, but it is a first step in understanding it. Moreover, properties of this type are often thought of as fundamental for a fair society. Thus, we already have here a connection to fairness, albeit minimal.

When the parameters defining a problem or economy change, we could in principle switch back and forth among several solutions in some arbitrary manner. For instance, in a classical problem of fair division, we could apply a first formula if the endowment has equal coordinates and a second formula otherwise. Such solutions do not come naturally to our minds however, and most people would consider them strange and artificial. However, they are very instructive, and when we teach the axiomatic method to our students, we encourage them to give free reign to their imagination, to “think outside of the box” and develop examples of this type. The exercise is necessary to fully realize the various ways in which each axiom constrains our options in solving problems. The example we just gave forces us to confront the issue whether the equality of the coordinates of the endowment should have any

¹¹A second reason is that economists as well as the man on the street are conditioned to thinking of proportionality as, at the very least, a reasonable or plausible criterion, a natural reference. In fact, some authors think of proportionality as the definition of fairness itself: recall Aristotle’s often quoted statement that “What is fair is what is proportional and what is unfair is what violates the proportion.”

relevance to how we allocate resources. Our intuition is that it probably should not, and that using this criterion (equality of coordinates) in defining a rule is at the cost of something. In fact, through an examination of the behavior of solutions defined in this way by piecing together several others as a “crazy quilt”, we are quickly led to (i) examples in which they seem to make unreasonable choices; then (ii) the formulation of properties that express our intuition about why such configurations of choices are unreasonable and what configurations would be reasonable, even desirable; finally (iii) the investigation of the compatibility of these properties and the study of their implications when imposed in various combinations and together with others.

Returning to our example, discontinuities in recommendations will probably occur when a sequence of social endowments with unequal coordinates converges to an endowment with equal coordinates. *Continuity* is the formal property that prevents this behavior and it is natural to want to understand when the property can be met and probe its compatibility with others.

One of the parameters entering the description of a problem is the population of agents involved, and allowing variations in populations may be quite important: an arbitrator may have to settle a dispute between management and labor in a firm, or one between provinces in a federation; it may be a personal bankruptcy involving a dozen merchants each attempting to recover what is due to him, or a commercial bankruptcy with thousands of anonymous claimants, that a bankruptcy judge may have to settle. Thus, the rule that an arbitrator or judge applies should be defined on domains of problems in which the agent set is not fixed. Let us refer to such a domain as a “variable-population” domain.

When presented with solutions defined on such a domain, including in their evaluation the manner in which they respond to changes in populations seems informative, probably necessary. Variable-population axioms in general, and *consistency* in particular, are important components of this process.

5. The solidarity principle. The concept of solidarity is a central tenet of the theory of fairness. Here is a statement of the “solidarity principle” (here too, the term “principle” seems appropriate). It is most easily formulated for rules and in the next pages, we use language that best fits *single-valuedness*:

The solidarity principle. When the circumstances in which some group of agents find themselves change—the group could be the entire population of agents present or some subset—and if none of them bears any particular responsibility for the change,

or deserves any particular credit for it, their welfares should be affected in the same direction: all members of the group should end up at least as well off as they were initially, or they should all end up at most as well as off as they were initially.¹²

We emphasize that no direction, positive or negative, is specified for the changes in the welfares experienced by the agents in the group to which it is applied: the only requirement is **uniformity in the direction of these changes**, and that is how the term solidarity should be understood from this point on. Thus, no attempt is made at measuring the magnitude of the changes, let alone at comparing these magnitudes across agents. However, the solidarity principle can certainly be complemented by engaging in finer assessments of the situations to be handled, and by imposing additional restrictions on the manner in which the chosen alternative responds to changes. If some agents do deserve particular credit or blame for the change, and if credit or blame can be meaningfully compared across agents and welfare changes can also be measured and meaningfully compared across agents, we may partition society into groups of agents to be placed on the same footing, and apply the principle to each of these groups, positively or negatively, in the manner intuition suggests in each particular application. We may additionally require some comparability of the welfare changes across the groups. Cardinal formulations of such requirements are more likely to come to mind, but ordinal formulations may be possible too.

The change to be contemplated could be in any one of the parameters of the problem, but which parameter varies will typically affect the scope of the requirement we write down. Let us distinguish between different parameter types.

1. A parameter may be **collective**. This means that it pertains to the entire agent set. In a problem of fair division, an example is the vector of resources over which agents are assumed to have equal rights—the **social endowment**. Then, in the expression of the solidarity principle, all agents would be covered: an arbitrary change in the social endowment should affect

¹²We should note there that we do not use the term exactly as in common language. There, we speak of a group of agents taking some action in “solidarity” with some other group. For instance, the first group may transfer resources to the second group in order to diminish the harm that has been inflicted upon that second group by a natural disaster. Thus, solidarity is not usually “imposed” by some outside authority. By contrast, for us, “solidarity” is embedded in a rule: the rule is such that the welfares of all agents are affected in the same direction by particular changes in their circumstances.

all of their welfares in the same direction. The axiom would be “solidarity with respect to changes in resources”, or **resource solidarity**.

2. A parameter may be **personal** (or “individual”), that is, “attached” to a specific agent. When such a parameter changes, the agent to whom the parameter is attached is part of the environment in which the others find themselves, and solidarity *among them* would say that their welfares should be affected in the same direction. In a “classical exchange economy”¹³, an example of such a parameter is an agent’s endowment. When an agent’s endowment changes, we would demand that the welfares of all the others should be affected in the same direction (Thomson, 1978).

A stronger expression of the principle is that, irrespective of the fact that it is the personal parameter of a particular agent that changes, the welfares of *all* agents, including that agent, should be affected in the same direction. In the application just mentioned, it would be appropriate if the agent whose endowment changes is thought of as the “accidental” recipient of these resources, because this would justify limiting his rights over them. This would not imply that individual endowments should be ignored. The requirement that he should end up at least as well off as he would be by consuming his endowment for example (the “individual endowment lower bound” requirement) could be meaningfully imposed together with this strong version of solidarity.

The types of resources under consideration may well matter when we think about ownership, in particular whether they are material resources, such as physical endowments of goods, which are external to the individual, or intellectual resources, skills or expertise for instance, which are internal to him.

3. A parameter may be **semi-collective**, that is, attached to a group of agents. For example, we can imagine particular resources, an unproduced bundle of goods say, or a technology, to be held, owned, or controlled by a group of agents. When such a parameter changes, the solidarity principle could be applied twice, first within this group, and a second time within the complementary group. Or, as before, it could be applied to the entire population.

Often, expressions of solidarity run into conflict with the recognition of

¹³Such an economy is defined as a list of preference relations defined over some Euclidean commodity space, one for each agent, and a list of vectors of resources in that space, one for each agent, their individual endowments. An allocation is a list of vectors, one for each agent, whose sum is equal to the sum of the endowments.

the power that agents may be able to derive from the control they have over resources, external or internal. We just discussed, for a classical exchange economy, changes in an agent’s private endowment. If an agent’s endowment increases and the forces of competition are left to operate, the fact that he may not need the others as much as before may result in a loss of welfare for them. For instance, he may now own a resource that initially he acquired from them; this may diminish the benefit that he derives from trading with them and as a result, they would have less of a reason to demand a reward or a compensation. The extent to which we allow the welfares of agents to be related to the control they can exercise over resources on the one hand, and the scope we give to solidarity considerations on the other hand, are decisions we need to make when choosing allocations and rules.

6. Solidarity and efficiency. We emphasize that we defined the solidarity principle as requiring that the welfares of all agents in a relevant group should be affected in the same direction by certain changes in their environment, no direction in this change being specified. However, most expressions of the principle have been written in situations in which we can tell whether the change that is contemplated permits a Pareto improvement for the group, or would have to be accompanied by a Pareto deterioration. Then, a particular direction of the impact on the welfares of its members may be specified as part of the requirement.

Recall that an alternative is (Pareto) **efficient** for a problem if there is no other alternative that all agents find at least as desirable, and at least one agent prefers. We will sometimes refer to the correspondence that associates with each economy its set of efficient allocations as the **Pareto solution**. Using this terminology, we can rephrase the observation made in the previous paragraph by saying that “some” efficiency has been incorporated in the usual formulation of the solidarity requirement, often unwittingly. For the classical problem of fair division, **resource monotonicity** (Thomson, 1978; Roemer, 1986; Chun and Thomson, 1988), which says that if the social endowment increases, all agents should end up at least as well off as they were initially, is an example. Indeed, when preferences are monotone, such an increase does permit a Pareto improvement. This is what the axiom requires.¹⁴

The same comment can be made about other axioms. It applies to **pop-**

¹⁴I say “some” *efficiency* because for the problem of fair division, the rule that throws away half of the social endowment and divides equally the other half satisfies “solidarity with respect to changes in resources” but it is not *efficient*.

ulation monotonicity (Thomson, 1983b,c), which concerns decreases in population such that the range of payoff, utility, or welfare profiles attainable by the agents who remain is the same as it would be in the initial problem if the agents whose departure is being contemplated were ignored. It says that all of the remaining agents should end up at least as well off as they were initially, any obligation to the departing agents having vanished. For the classical problem of fair division, a decrease in population is compatible with a simultaneous improvement in the welfares of the remaining agents, and *population monotonicity* has been usually stated as the requirement that such an improvement should occur. The comment also applies to **technology monotonicity** (Roemer, 1986), an axiom that concerns economies in which production possibilities are specified. It says that if a commonly owned technology improves, all agents should end up at least as well off as they were initially.

In formulating a list of axioms to be imposed jointly, it is most desirable that these axioms express conceptually distinct ideas. That is certainly the case of efficiency and solidarity ideas. Yet, efficiency considerations are so engrained in our thinking as economists that, taking once again the classical problem of fair division as an example and considering changes in resources, we feel uncomfortable simply requiring that as a result of an increase in the social endowment, *either* all agents should end up at least as well off as they were initially, *or* that they should all end up at most as well off. Indeed, we know that the change does permit the former, and that is what we demand: we impose *resource monotonicity*, not *resource solidarity*.

However, every property has its cost in terms of other properties. That is of course the case for *efficiency*, and for that reason it is not one that we should necessarily impose, or can always afford to insist on. To illustrate, for the classical problem of fair division, it is often incompatible with minimal requirements of symmetry among players and *resource monotonicity* (Moulin and Thomson, 1988), or *population monotonicity* (Kim, 2004), or *strategy-proofness*¹⁵ (Serizawa, 2002). Then, searching for rules that are *resource*

¹⁵In the “direct revelation game form” associated with a rule, each agent is given a strategy space that is the space of his possible preferences, and the outcome function (the function that maps each profile of strategies to an allocation) is the rule itself. A game form together with a preference profile for the agents involved define a game. A rule is *strategy-proof* if, for each preference profile, in the associated direct revelation game that results, no agent ever benefits from misrepresenting his preferences: telling the truth about his preferences is a dominant strategy.

monotonic or *population monotonic* or satisfy an expression of solidarity in response to changes in resources or populations that is free of efficiency, but are only required to satisfy a minimal efficiency notion, is a justified objective. It is in the same spirit that an extensive search has been conducted for rules that are *strategy-proof* but not necessarily efficient. We are referring here to the literature on the Clarke-Groves-Vickrey mechanisms, which grew out of the position that *efficiency* may justifiably be sacrificed if this allows recovering *strategy-proofness*. This literature has indeed delivered a strong incentive-compatibility property at the expense of a (rather significant) weakening of *efficiency*.¹⁶

Moreover, the solidarity principle can be applied to changes in parameters that cannot unambiguously be said to be desirable or undesirable in terms of social welfare. For the classical problem of fair division, the supply of some goods may increase and that of others decrease. Then, we do not know before preferences are specified and the rule is applied, whether the change would allow a Pareto improvement, or would have to be accompanied by a Pareto deterioration. For the quasi-linear social choice problems considered by Chun (1986), a change in population is also such that one cannot tell beforehand the direction in which the welfares of the agents who are present before and after the change can be affected, which led him to work with the general expression of the solidarity objective.

Finally, solidarity can also be applied to changes in parameters that belong to spaces that are not equipped with an order structure, an example here being preferences. A change in the preferences of an agent cannot in general be said to be unambiguously good, or unambiguously bad, for the others. However, once again, the general idea of solidarity remains applicable; it is this idea that underlies “welfare dominance under preference replacement”, the requirement that when the preferences of an agent change, the welfares of all others should be affected in the same direction. (The literature devoted to the study of this requirement is reviewed in Thomson, 1999.)

7. “Pre-application” and “post-application” relational axioms.

¹⁶This literature pertains to “quasi-linear” economies: there, an alternative includes the distribution of some infinitely divisible good and an allocation decision concerning one or several other goods, and preferences admit representations that are separable with respect to that special good and the other goods, and linear with respect to the special good. *Efficiency* can then be decomposed into two components, budget balance in the special good being the requirement that has been sacrificed in the literature.

A meaningful way of organizing relational axioms is according to whether the problems appearing in the hypotheses (i) can be specified independently of each other, or if they are related, it is in a way that does not refer to the rule; (ii) are related in a way that requires an application of the rule to one of them, the other problem having to satisfy certain conditions involving the first one and its solution outcome.

Let us illustrate the distinction in the context of “Nash’s bargaining problem” (Nash, 1950).¹⁷ A property of the first type is **domination** (Thomson and Myerson, 1980): given *two arbitrary problems*, either the solution outcome of the first one should vectorially dominate the solution outcome of the second one, or the reverse should hold. Another is **strong monotonicity** (Kalai, 1977): given *two problems that are related by inclusion*, the solution outcome of the larger one should vectorially dominate the solution outcome of the smaller one. A property of the second type is **contraction independence**¹⁸ (Nash, 1950): given *two problems that are related by inclusion, if the solution outcome of the larger one belongs to the smaller one*, then it should also be the solution outcome of the smaller one.

Because an axiom of the first type does not require that the rule be applied when specifying the hypotheses, let us refer to it as a **pre-application** relational axiom; for the opposite reason, let us call an axiom of the second type a **post-application** relational axiom.

When introducing a relational axiom, we sometimes tell a story that unfolds over time: we imagine a *first* problem, calculate its solution outcome, then hypothesize a change, calculate the solution outcome of the *second* problem, and finally, we compare the results. For a relational axiom of the first type, we can just as well start with the second problem, switch to the first one, and compare the results. In fact, we can place the two problems “on the same level”, calculate their solution outcomes separately and compare the results.

¹⁷A bargaining problem is defined as a subset of a Euclidean “utility” space of dimension equal to the number of agents involved, each point of which is interpreted as a utility vector they can achieve if they agree on it, together with a point in this set, called the “disagreement point”. The pair is supposed to satisfy certain regularity properties. A solution associates with each problem one of these vectors. When invariance of solutions with respect to choices of origin of utilities is imposed, which is the case for almost all of the theory, it is convenient to take the disagreement point to be the origin, and that is what we will do throughout. Thus, references to disagreement points are missing from our statements of axioms and definitions of rules.

¹⁸This property is commonly called “independence of irrelevant alternatives”.

For relational axioms of the second type, this is not as easily done. For instance, *contraction independence* is best described by imagining that the objects appearing in its definition are introduced in a specific order: a rule is applied to a first bargaining problem, and a second problem is considered that is related to the first one *and its solution outcome* in a certain way. Many other properties, in bargaining theory, but in other fields as well, are of this type. It is in particular the case for *consistency*.

Although, for these properties, it is most convenient to have a time line in mind, one could argue that the comparisons they require can simply be based on thought experiments. For *consistency*, we may just be comparing problems with nested agent sets facing opportunity sets that are related in a particular way. To illustrate, in its expression for the classical problem of fair division, a conditional statement is listed in the hypotheses of the axiom: the social endowment of the economy with the smaller population should be equal to the sum of the assignments to its members when they are part of the larger economy. Still, the question would remain why opportunity sets should be related in this particular way. The simplest explanation would be that the rule has been applied and that some agents have left with their assignments.

Let us add here that for some models, an additional critical point can be identified on the time line along which decisions have to be made or can be made. When these decisions include production and distribution plans, we may reevaluate options before production has taken place or after. If we do so before production has taken place, we have more options than if we do so after production has taken place; then, we are only left with distributional choices. We could use the phrases **pre-production consistency** and **post-production consistency** for the corresponding requirements. (The issue is discussed in greater detail in Thomson, 2011a).

8. Pareto-indifference. A solution satisfies **Pareto-indifference** if the desirability of an outcome is transferred to any other outcome obtained by “moves along indifference curves”. Specifically, starting from some problem, if an alternative is chosen by the solution for the problem and there is some other alternative that all agents find indifferent to it, then the latter should also be chosen.

The requirement appears innocuous enough, and in many models, it is indeed satisfied by most solutions, but it is also important to realize that for some models, central solutions violate it. For the classical problem of

fair division, consider the **no-envy solution** (Foley, 1967).¹⁹ It is easy to construct two-commodity and two-agent examples revealing that this solution violates *Pareto-indifference* (Thomson, 1983a). Nevertheless, even if a solution is not *Pareto-indifferent* on some domain, interesting subdomains can often be identified on which it is. For the domain just mentioned, if preferences are strictly convex and an allocation is efficient, no other efficient allocation is Pareto-indifferent to it. Then, any subsolution of the Pareto solution satisfies *Pareto-indifference*. This obviously holds for the solution that selects, for each economy, all the envy-free and efficient allocations.

9. Consistency as an expression of solidarity. The necessary conceptual preliminaries being out of the way and language established, we now return to *consistency* and reach the core of our argument.

In a variable-population model, two main scenarios involving changes in populations, some agents leaving say, can be imagined for which we may want to require solidarity among the remaining agents, namely that their welfares should be affected in the same direction.

1. *When agents leave, they do so empty-handed.* Perhaps they have been found to have no rights on the resource. Or they have been miscounted. Or some agents have died. For most models, when *efficiency* is imposed, the solidarity principle takes us then to the axiom of *population monotonicity* as usually formulated (Section 6).

2. *When agents leave, they do so with their components of the solution outcome.* Here too, for many models, and again in the presence of *efficiency*, the same solidarity principle leads us to the standard formulation of the (invariance) axiom of *consistency*. For other models, in a reduced problem, there could be alternatives that are Pareto indifferent to the restriction of the alternative initially selected and at which the solidarity principle would be satisfied. However, if the rule satisfies *Pareto indifference*, this restriction would indeed be selected as one at which the solidarity principle is honored.

In paragraph 1 above, at the point at which we invoke solidarity, the rule has not been applied yet; in paragraph 2, it has. Because, as argued earlier, the solidarity principle is silent on *efficiency*, we can think of *population monotonicity* and *consistency* as pre-application and post-application expressions of the solidarity principle in situations in which we would be im-

¹⁹An allocation is “envy-free” if each agent finds his assignment at least as desirable as each other agent’s assignment. The “no-envy” solution associates with each economy its set of envy-free allocations.

PLICITLY interested in *efficiency*, or at least in “some” *efficiency*; in both cases, solidarity is in response to population changes. It is because of this solidarity underpinning that the consistency principle can be linked to fairness. What we have proposed here is a sort of conceptual “electrolysis” of *population monotonicity* and *consistency* that revealed both of them to be composed of *solidarity* and some *efficiency*, *solidarity* being invoked at different points in time, before the solution is applied in one case, and after it is in the other.²⁰

Let us consider some models and in the light of this discussion, revisit the *consistency* axioms that have been studied for them.

- For the classical problem of fair division, it is easy to see that if preferences are strictly convex, a selection from the *efficiency* correspondence satisfies *the post-application efficiency-free expression of solidarity in response to population changes* if, when some agents leave with their assignments, each of the remaining agents is assigned the same bundle as initially (Thomson, 1988). This is what the standard axiom of *consistency* states for this class of economies. Indeed, the restriction of an efficient allocation to a group of agents is efficient for the reduced economy associated with the group and the allocation (the Pareto solution is *consistent*), and in this reduced economy, if an allocation is efficient, no other allocation is Pareto-indifferent to it.

- The same observation applies to the problem of fully dividing a social endowment of a single good among a group of agents each of whom has single-peaked preferences over how much he consumes (Sprumont, 1991; a study of *consistency* for this model is Thomson, 1994a). Here too, if an allocation is efficient, no other allocation is Pareto-indifferent to it.

- In the axiomatic theory of bargaining, the standard formulation of the *consistency* of a rule—solutions are (almost) always required to be *single-valued* in this literature—is the following. We start from a bargaining problem and apply the rule, thereby obtaining its solution outcome—let us call it x . We define the **reduced problem with respect to a subgroup of players and x** to be the set of all payoff vectors for the members of the subgroup that are obtained as the restriction to the subgroup of some payoff vector of the initial problem at which the members of the complementary group receive their components of x . We require that, provided this reduced

²⁰The conceptual connections that we are bringing out here between these properties are also reflected, not surprisingly, in logical relations that can be established between them, *resource monotonicity* being instrumental in linking them.

problem is in the domain, its solution outcome should be the restriction of x to the subgroup. For a selection from the Pareto solution, this requirement is implied by *the post-application efficiency-free expression of solidarity in response to population changes*. The Nash rule²¹ satisfies this property; the egalitarian rule²² does not. (The first characterizations based on *consistency* for this model are due to Lensberg, 1985, 1988).

The requirement discussed by Thomson (1984), under the name of “weak consistency”, is that in the reduced problem as just defined, the payoff of each player in the subgroup should be at least as large as at x . This requirement can also be seen as the result of adding “some” efficiency to the post-application efficiency-free expression of solidarity in response to population changes, but “less” efficiency than what leads to what is usually called *consistency*. The egalitarian solution satisfies this requirement, and we wrote “less” efficiency because this solution is only **weakly** (Pareto) **efficient** (it is such that there is no outcome that all agents prefer to the one that it selects).

- There are many other models in which the post-application efficiency-free expression of solidarity in response to population changes, when imposed in the presence of *efficiency*, gives us the invariance axioms of *consistency* that have been studied for them: examples are claims problems, surplus-sharing problems, their union, strictly comprehensive bargaining problems,²³ transferable utility coalitional games, non-transferable utility coalitional games²⁴ whose feasible set for the grand coalition is strictly comprehensive, the problem of allocating a social endowment of indivisible goods, or “objects”, among a group of agents, one per agent, when each agent has strict preferences over the objects.

When *efficiency* is not imposed, this equivalence does not persist: it is usually not true that the post-application efficiency-free expression of solidarity in response to population changes is equivalent to *consistency* written

²¹It is the rule that selects, for each problem, the feasible point at which the product of the agents’ utilities is maximal.

²²It is the solution that selects, for each problem, the maximal point of equal coordinates.

²³A subset of \mathbb{R}^ℓ is “comprehensive” if whenever it contains a point x , it contains all points y such that $y \leq x$. It is “strictly comprehensive” if whenever it contains a point x , it contains in its relative interior all points y such that $y < x$.

²⁴Such a game is defined as a list of sets, one for each coalition, that are interpreted as sets of utility vectors that coalitions can achieve. A solution selects a set of payoff vectors that lie in the set associated with the grand coalition.

as an invariance axiom. The former property is weaker.

Incidentally, working with the former makes it easier to show the independence of *efficiency* from the other axioms in a characterization. However, as usual, it is always more informative to show independence of axioms by exhibiting examples of solutions satisfying stronger versions of the remaining ones, even satisfying additional ones (because in that case, we can say then that “even” if these other axioms are strengthened, or these additional axioms are added, we still cannot derive the missing property). Thus, it would then be preferable to exhibit solutions for which *consistency* is given its usual invariance form. Conversely, if a characterization of a class of solutions that are not necessarily *efficient* but satisfy the usual form of *consistency* have been derived, (for bargaining problems, Thomson, 1984, is an example), the question remains what additional solutions would become admissible when instead of the usual form of *consistency*, we would impose the post-application efficiency-free expression of solidarity in response to population changes under discussion.

We observed earlier that the standard formulation of the consistency principle does not accommodate externalities well. At least, this formulation creates difficulties when the hypothesis that agents “leave” is taken literally. This is because the welfares of the remaining agents may be affected positively or negatively by the presence of these agents. When the consistency principle is given its solidarity interpretation, this difficulty disappears.

10. Relational axioms and single-valuedness. One additional issue should be discussed when asking whether the consistency principle can be given a fairness interpretation based on solidarity considerations. It is that for most models, the *consistency* axioms that have been written down apply to solution correspondences, not just to *single-valued* mappings, what we call “rules”. By contrast, the simplest expression of solidarity is for rules, because the welfare comparisons that it requires are unambiguous then. We stated it for rules for that reason. Working with solution correspondences when discussing the consistency principle may contribute to the apparent difficulty in linking it to solidarity and fairness.

There is no real difficulty however, because even when we work with correspondences, in the statement of *consistency*, we take as point of departure in each problem a specific outcome chosen by the solution for it, and the solidarity requirement is applied *from that choice*: if that choice is made among the ones that the solution recommends for the problem, then, when

the solution is applied to any associated reduced problem, it should choose at least one outcome at which the welfares of all remaining agents should be affected in the same direction. Thus, in fact, we are not comparing sets.

In general, when solutions are interpreted as providing recommendations, *single-valuedness* is a desirable property because it leaves no room for dispute. However, mathematical reality often intrudes and makes it difficult to achieve this goal. Thus, whether we work with rules or not mainly has to do with how restrictive *single-valuedness* happens to be for the type of problems in which we are interested. For several models (bargaining problems, cost allocation problems, claims problems), a large number of *single-valued* mappings can be defined that are well-behaved from a variety of viewpoints, and the literature has been mainly written for rules. For other models, *single-valuedness* is very demanding and solution correspondences have been the principal object of study (classical exchange economies and classical problems of fair division are examples here).

11. Consistency is not enough for fairness. In a fixed-population model, a **dictatorial rule** is one that always selects an alternative that a specific agent, chosen beforehand and once and for all, most prefers. This agent is called a **dictator**. For simplicity let us assume this alternative to be unique. (In a classical problem of fair division, an alternative is an allocation, and if preferences are strictly monotone, there is a unique alternative that any agent most prefers, namely the allocation at which he is assigned the entire social endowment.²⁵)

In a variable-population model, the definition can be generalized in two ways. (i) We apply the fixed-population definition population by population: for each population, we select a dictator. (ii) We proceed as in (i) but in addition, we require the identities of the dictators to be related across populations in the following manner. We select a “reference order” on the set of potential agents. Once a population is chosen and a problem involving it is specified, in that population, we identify the agent who is first in the reference order, and we select the outcome that he most prefers. Let us maintain our assumption that this outcome is unique so as to get a well-defined rule, and

²⁵If preferences are not strict, we proceed as follows. We specify an order on the agent set. Among the outcomes that the first agent most prefers, we identify the one that the agent who is second most prefers; if this outcome is not unique, we turn to the agent who is third, and so on . . . , thereby obtaining a “sequential”, “recursive”, or “lexicographic” extension of dictatorship, but let us ignore this possibility here.

let us still call any rule defined in this manner a **dictatorial rule**.

Because such a rule is *consistent*, as is easily verified, it may be somewhat puzzling that *consistency* could be seen as contributing to fairness. The puzzle is easily resolved however when we remember the distinction between punctual and relational axioms. *Consistency* does not say, nor does it imply, that dictatorship is fair; it simply implies that *if*, for whatever reasons, the choice has been made to maximally favor a particular agent in a given problem, *then* “similar” choices should be made in other problems. Thus, whatever fairness interpretation we may be able to give to *consistency*, it surely cannot be enough for fairness. However, this observation does not deprive *consistency* of its *relational* fairness content.

The same point can be made about other properties that we have discussed, such as *resource monotonicity* and *population monotonicity*. Indeed, dictatorial rules are *resource monotonic*, and in a variable-population framework, their extensions obtained, as explained above, by choosing for each population the dictator to be the one among the agents who are present to be first in some fixed reference order are *population monotonic*. Yet the fairness content of the two axioms is hardly questionable, but once again, they are *relational* fairness axioms.

Besides, fairness may require an asymmetric treatment of agents if they enter asymmetrically in the description of a problem. In a claims problem, agents differ in how much they are entitled to, and differences in claims are a good reason why they should not be assigned equal amounts. A decision to award all of the resource available to a single claimant, or to a particular group of claimants, may well be judged punctually fair. *Consistency* would transfer any asymmetric treatment across problems, which is quite natural: two agents whose characteristics justified that they be treated asymmetrically in some initial problem still have the same characteristics in a reduced problem to which they both belong, and one can argue that they should still be treated asymmetrically. How asymmetrically? *Consistency* says that they should be treated “as asymmetrically as” they were initially.

Agents may also differ in characteristics that are not included in the description of the formal model that we study, and fairness may then call for two agents to be treated asymmetrically *even though* they enter symmetrically in the model. How asymmetrically a rule may want to treat them can be uncovered by applying it to some well chosen problem(s) (Thomson, 2001, refers to the step in a characterization of a family of rules where a particular member is distinguished from the others as “calibration”), and *consistency*

will force this asymmetric treatment to be carried out systematically across all problems.

Fairness is not a single, indivisible notion. It has “components” and in each application, to be fully satisfied that fairness has been achieved, we may want to insist on several of these components. However, we should also recognize the constraints placed on us, constraints that the purpose of axiomatic work is to help uncover: not all combinations of axioms are compatible, and choices may have to be made. To emphasize the generality of the point, let us take another example that has nothing to do with *consistency* or even variable-population considerations.

For the classical problem of fair division, many writers have proposed *no-envy* as a fairness requirement. Here is not the place to debate the concept; for our purposes, it is enough that the reader accepts that this may be a legitimate position to take. Now, consider two problems of fair division that differ only in their social endowments, one dominating the other. The *no-envy* requirement by itself allows us to select allocations for these two economies such that some agent is better off when the social endowment is larger and some other agent is worse off. Then *resource monotonicity* is violated. On that basis, one can argue that *no-envy* is not sufficient for fairness.

On the other hand, dictatorial rules (their fixed-population versions as well as their variable-population versions) are *resource monotonic*. So, one can say that *resource monotonicity* is not sufficient for fairness either. Some punctual fairness notion, whether it is *no-envy* or some other notion, but let us focus on *no-envy* for the sake of argument, seems necessary to prevent the extremely skewed outcomes, in terms of welfare, that these rules select.

Moreover, we will probably be interested in *efficiency* too, and therefore, we may well look for rules satisfying at least three properties, *efficiency*, the punctual fairness requirement of *no-envy*, and the relational (solidarity) requirement of *resource monotonicity*. But it turns out that no such rule exists (Moulin and Thomson, 1988). In the face of such an impossibility, it would be pointless to declare either *no-envy* or *resource monotonicity* “necessary” for fairness, or say that *efficiency* itself is “necessary” for social choice.

The same observation applies to *consistency*. We had already argued that it may not be sufficient for fairness, and we are now claiming that *consistency* is not necessary either. Also, the decision to impose punctual fairness requirements on a solution can be made separately from the decision to impose on it relational fairness requirements.

12. Consistency an obstacle to fairness? For some models, it seems fundamentally impossible to obtain any kind of punctual fairness. Consider for instance the problem of assigning two objects to two agents, one for each, and suppose that both agents prefer the same object. Here, we have no choice but to favor one of the two agents. Now, turn to the profile in which they both have the reverse preferences. Wouldn't it be fair if for that profile we favored the other agent? To compensate the agent who has been assigned his worse object in the first profile, we would assign him his most preferred object this time. In fact, since for the two remaining profiles, the ones in which they have distinct most preferred objects, it is possible to assign to each agent what he most prefers, the four choices that we have now listed for the four possible profiles, because they distribute symmetrically across profiles which one of the two agents is treated more favorably, seem to have achieved a sort of "overall fairness", "fairness across profiles", or "fairness behind the veil of ignorance" (Rawls, 1971).²⁶

Reducing an economy from two agents to only one agent is (usually) trivial (in our example, in a one-object and one-agent economy, there is only one feasible allocation), so let us allow for more than two objects and two agents, but for simplicity, let us still assume that there are as many objects as agents. Let us once again start from a situation in which all agents have the same preferences. Here too, we have no choice but to treat agents differently even though nothing distinguishes them. One agent will get his most preferred object, another will get his second most preferred object, and so on. Our choice of allocations can be interpreted as reflecting our preference over the agents, who we want to favor at the expense of the others, but it can also simply be interpreted as the logical consequence of our being unable to treat equally agents who have the same characteristics: in fact, we may have no reason to want to favor anybody. Now, observe that *consistency* forces us to serve agents in the same order (technically, the "induced" order) in many other economies, namely all the reduced economies associated with subgroups of agents and the choice we have made in the initial economy. Thus, it "propagates" the way in which we have violated punctual fairness, violations that were forced on us initially.

Of course, we could have started from some other profile of like preferences, and for that profile, as we suggested earlier in the two-agent example, we would have the opportunity to serve agents in a different order, thereby

²⁶That is, before preferences are assigned to agents.

partially “making up” for our inability in the first economy to treat equally agents whom we would have wished to treat equally. *Consistency* will have additional implications, starting from that profile.

This is not the end of the story however. Indeed, we may actually not have the option of doing what we want for all profiles of like preferences, because these profiles can be obtained by reducing economies with more agents in which not all preferences are the same. For *consistency* to be met overall, further relationships between choices for different economies will have to be met. Altogether then, what is needed for *consistency*?

The answer is provided by Ergin (2000) and it is not very good news. It turns out that the only *efficient* rules that pass the test are the **sequential priority** rules, which can be seen as generalizations of the sequential dictatorial rules that we encountered earlier (Section 11). They are defined as follows. A reference order on the agent set—a **priority relation**—is specified once and for all. For each problem, among the agents who are involved, the one who is first in that order is assigned his most preferred object. The agent who is second is assigned his most preferred object among the remaining ones, and so on, until each agent has been assigned an object.²⁷ Thus, *consistency* gives us no choice but to systematically favor agents according to a fixed reference order on the set of potential agents. Shouldn’t we say then that *consistency* presents an obstacle to fairness?

To understand the issue, it is useful to distinguish between the punctual fairness of an outcome for a specific problem and the overall fairness of a collection of outcomes for a collection of problems. As the latter property has to do with the manner in which (in this case, unavoidable) “favors” are distributed across problems, an expression such as “distributional fairness” might be appropriate here, but we will avoid it because the adjective “distributional” is commonly used in the literature to refer to the fact that we are distributing resources *across people*. What we are discussing is what could be called instead the “*personal* distribution of an agent’s assignments *across problems*”.²⁸ Using this language, we can indeed conclude that *consistency* is at odds with fairness; it prevents fairness in the personal distributions of assignments across problems.

In some models, one can meaningfully calculate for each problem the

²⁷We use the expression sequential priority rule here as it suggests a less drastic asymmetric distribution of resources than the term dictatorial rule. Indeed, even if an agent is served first, there are constraints on how much he can receive.

²⁸We could also speak of “statistical fairness”.

various outcomes produced by possibly unfair but “symmetrically unfair” rules, even maximally unfair, and recover fairness by an averaging operation. This requires that convex operations be possible, and if we are also interested in *efficiency*, that *efficiency* also be so preserved. This is how the Shapley value (Shapley, 1953) is defined for transferable utility coalitional games, and how the so-called “random arrival rule” (O’Neill, 1982) is defined for claims problems. It turns out that *consistency* is not preserved by averaging operations (Thomson, 2011a). Even though the dictatorial rules (Section 11) are *consistent*, their average is not.²⁹ Thus, we have here an (admittedly narrow) sense in which *consistency* gets in the way of fairness: it does not allow us to take advantage of averaging operations that could help in this regard.

13. Consistency neutral to fairness? The argument presented in the previous section is somewhat limited, and it may be more accurate to say that *consistency* is neutral with respect to punctual fairness: it simply transmits across economies of different sizes whatever punctual fairness or unfairness the decisions made for them may embody. For the classical problem of fair division, the equal division rule satisfies many punctual (and relational) fairness tests and it is *consistent*. The dictatorial rules satisfy no interesting punctual fairness notion, and they are *consistent* too.

Can we say more though? Can we describe the sense in which punctual fairness, or lack thereof, of a choice is transferred across economies in a coherent way by *consistency*? Can we give a reason why we should want this to be the case? The answer is yes and it is given by referring to solidarity. In the presence of *efficiency*, whether a solution outcome is punctually fair or not for a particular problem, that quality or lack thereof, is transferred to reduced problems by insisting that the welfares of all agents should be affected in the same direction when the others leave with their components of it.

14. Variants of consistency. A weaker version of *consistency* is **bi-lateral consistency**. It differs from *consistency* in that reduced problems involve only two agents. The two-agent case is particularly interesting because our intuition is clearer then. We can abstract from the conceptual

²⁹Probabilistic rules can be thought of as achieving fairness by proceeding in the opposite direction. Such a rule may be fair ex ante, before the lottery it recommends has been drawn, but each realization of the lottery is unfair.

complications caused by having to assess options open to non-trivial subgroups, in particular overlapping subgroups.

Another version is **average consistency** (Maschler and Owen, 1989). Here, the insistence that for each problem, each agent, and (i) each reduced problem involving him, his assignment should be the same as in the initial problem, is replaced by the requirement that (ii) the average of his assignments in the various reduced problems involving him should be the same as his assignment in the initial problem. (This is meaningful only if, as discussed earlier, assignments can indeed be averaged, which is not possible in all models.)

Both *bilateral consistency* and *average consistency* are weakenings of *consistency* that are worth considering when *consistency* itself is too demanding. They do not affect the possible interpretations of this principle in any significant ways, so we will not discuss them further.

A solution is **flexible** (Balinski and Young, 1982)³⁰ if the following holds: consider a problem and identify an alternative chosen by the solution for it. Now focus on a subgroup of the agents involved in this initial problem and identify a solution outcome of the problem obtained by reducing the initial problem with respect to the subgroup and the initial solution outcome. The requirement is that the “concatenation” of this second solution outcome with the components of the initial solution outcome pertaining to the complementary group should be a solution outcome of the initial problem. For a *single-valued* solution, *flexibility* and *consistency* are equivalent. Thus, in that case, the same interpretations would apply to *flexibility*. As for solution correspondences, the property amounts to guaranteeing that the choices we make are robust to giving some freedom to subgroups—that is why we use the term “flexibility”—in possibly reassigning the resources they have received. This freedom is constrained by the requirement that the same rule be used in calculating reassignments. It seems unrelated to fairness considerations.

Finally, consider a model in which resources cannot be disposed of, or a model such that when *efficiency* is imposed, resources are not disposed of. Then, it makes no difference whether, in a reduced problem, the resources that are made available to the remaining agents are calculated as either (i) whatever is left after the departing agents have collected their bun-

³⁰In their study of apportionment, these authors present the properties that I call *consistency* and *flexibility* as the two components of a single property, which they name “uniformity”. Balinski (2005) also uses the term “coherence” for this double property. I borrow the term “flexibility” from Shimomura (1993).

dles, or (ii) the sum of their assignments. Otherwise, the implications of what could be called **pre-disposal consistency** and **post-disposal consistency** (Thomson, 2011a, proposes these expressions to label the resulting definitions.) may differ. This can be seen in the context of the classical problem of fair division: when preferences may be satiated, the equal division Walrasian rule is or is not *consistent* depending upon the formulation (Thomson, 1988). That it may make a difference has also been brought out in the context of object allocation problems (Ehlers and Klaus, 2006, 2007). However, both formulations can be given each of the interpretations discussed in this essay.

15. Converse consistency. Consider a problem and a feasible outcome for it. Suppose that its restriction to each pair of agents is chosen by a solution for the reduced problem associated with this population and the outcome. Then, the solution is **conversely consistent** if it chooses this outcome for the initial problem.

Can this property be given a fairness interpretation as well? First of all, let us address a name issue. The expression *converse consistency* makes it sound as if the concept is derivative of, or secondary to, *consistency*. There is no reason why we should think about the two properties in this way. The two notions are obviously related. Both are based on reducing some initial problem with respect to subpopulations and some outcome that is feasible for it. If we called *converse consistency* “augmentation invariance” (in some earlier text, we explored expressions of this type in order to emphasize that it should not be thought as derivative to *consistency*), *consistency* could be presented as the auxiliary property, something like “converse augmentation”.

We do not see any fairness underpinning to *converse consistency* and we will present later (Section 16c) what we perceive to be the most natural interpretation of this requirement. Both *consistency* and *converse consistency* do say that some information can be ignored in making a decision. In reducing a problem, information is lost.

This is a good place to recognize that the solidarity principle also implies that certain types of information should be ignored, specifically that when a change occurs in the problem to be solved, certain shifts in the configuration of feasible payoff, utility, or welfare profiles towards certain agents and against others should not be taken into account. An example is when a technological improvement occurs that causes opportunities to shift towards

some categories of workers and against others. Then, solidarity—again, *efficiency* makes an appearance—would say that transfers should be made so that everyone benefit. Consider the development of computer-assisted machinery. Workers who have computer skills, or whose opportunity costs of acquiring these skills is low would be in the first category; workers without such skills or whose age or family commitments prevent them from acquiring them would be in the second category. One need not go all the way to advocating the compensations necessary for everyone to end up benefiting, but such compensations may be possible, and the extent to which they should be carried out may nevertheless be an important element of the debate that citizens and politicians engage in.

16. Other underpinnings for consistency. Other justifications for the consistency principle have been proposed in the literature, which we assess in this section.

(a) **Consistency as a robustness requirement with respect to choices of perspectives when evaluating a change in a situation.** After a rule has been applied to a problem, imagine some agents leaving with their components of the outcome. Two positions or perspectives can be taken in assigning payoffs to the remaining agents. (i) Their initial assignments are confirmed; this amounts to declaring that the departure of the agents who left is irrelevant. (ii) The opportunities open to the remaining agents are recalculated taking into account the fact that the agents who left have done so with some resources. The rule is applied to the resulting reduced problem. Someone who think that the two positions are equally legitimate, a natural requirement on the rule is that it should not make any difference which of them is taken.

For abstract or unstructured “Arrovian” choice problems, “path independence” (Plott, 1973), for bargaining problems, “step-by-step negotiations” (Kalai, 1977), for claims problems, “composition up” (Young, 1987), for the problem of allocating a social endowment of indivisible goods, “composition up” (Abizada and Chen, 2011), are also expressions of this principle of robustness to choices of perspectives.

Of course, there could be protocols that specify how changes in a situation should be handled, but in practice, none may have been written. It is indeed rarely the case that all contingencies are explicitly covered in a constitution or contract. More importantly, it is often not clear on what basis such protocols would be written, and much arbitrariness would have to occur in laying them

down. The robustness principle we stated says that this arbitrariness can be avoided altogether, because the application of the solution is independent of the choice of perspectives. One possible interpretation of the term “fairness” is “not letting irrelevant or unimportant features of the situation we face matter in choosing an alternative”. Thus, this robustness argument lends additional support to the view of the consistency principle as contributing to fairness.

(b) **Consistency as a robustness requirement under piecemeal implementation.** Our next interpretation is closely related to the previous one. The story developing over time that we used earlier to explain the principle is taken seriously here, as opposed to being an expository device. Now, payoffs are indeed delivered in a particular temporal sequence. Time passes between dates at which different agents collect their payoffs and leave. For sure, if after some agents have done so, we make an adjustment in the payoffs to the remaining agents, any one of the agents who ends up worse off will claim that the adjustment is unfair. Note that no *interpersonal* statement is involved here. A promise has been made to an agent and the promise has not been kept. Many would say that renegeing on a promise is a violation of *intrapersonal* fairness, at least in one of its commonly accepted meanings. Thus, we have here yet another interpretation of *consistency* in terms of fairness.

The disappointment experienced by an agent who is receiving less than he was initially promised may be compounded by the fact that for some other agent, the change in welfare may go in the opposite direction: this second agent now receives more than he had been promised. Requiring that the assignments to the remaining agents should not be affected by the delivery to the others of their own assignment, what could be called “robustness under piecemeal implementation”,³¹ would eliminate the problem.

This motivation is developed by Lensberg (1985) who imagines a group of workers who have agreed to take on a job, but the nature of the job requires that they move in a particular temporal sequence. To build a house for instance, the bricklayer goes first, the carpenter second, and the painter last. After the bricklayer is done, he picks up his check and leaves. Some time goes by until the carpenter is done, picks up his own check, and leaves. The painter leaves last. In the context of intergenerational allocation, this interpretation

³¹We do not give to the term “implementation” the technical meaning it has in the modern theory of mechanism design, but instead we use it as it is used in common language.

of *consistency* is particularly compelling because time is explicitly modeled and consumption takes place over time. Some agents leave because they have consumed. It is then natural to look at the situation anew at that date.

One difficulty with this interpretation however is that the physical or technological reasons why time matters usually call for a specific order in which agents leave, whereas *consistency* is written for arbitrary groups of agents leaving. Of course, we can imagine weaker versions of the requirement, in which restrictions are placed on the groups of agents with respect to which the reduction operation is performed. Graph restrictions of this type are discussed in Thomson (2011b).

(c) **Consistency as an operational principle.** To evaluate the assertion that the consistency principle is an “operational principle”, we need to first clarify the possible meanings of this expression. Let us distinguish between two possibilities.

First, we can speak of giving “operational meaning” to an abstract idea. Fairness itself, the concept of “equal rights”, are elusive concepts, but we can say that they are “made operational” by concrete and formal mathematical definitions, such as that of an envy-free allocation, the property of *resource monotonicity*, and so on. We can work with these definitions: we can investigate their logical relations, check whether certain rules satisfy them, study their mutual compatibility and explore their mathematical implications.

Thus, one cannot simply say that the consistency principle is an operational principle. One has to explain what it is that it makes operational. One possibility was already discussed at the beginning of this essay, when we talked about “crazy quilt solutions”. Our intuition is that there is some undesirable arbitrariness to such solutions. *Consistency* makes operational our desire for “order” and “coherence”. This observation does not free us from having to attempt to understand the specific way that it does so however. We offered solidarity as a possible explanation.

The second meaning is related to the first one, but somewhat narrower, more focused on the fact that *consistency* is a cross-population principle, and it has to do with certain computational issues. To explain the point, it will be useful to return to *converse consistency*, because this principle can more clearly be interpreted in this way: recall that a solution is *conversely consistent* if whether it chooses a particular outcome for a particular problem can be deduced from the fact that it chooses the restriction of the outcome to each pair of agents for the associated reduced problem this pair faces.

Handling problems with many agents may be computationally difficult, but if a rule is *conversely consistent*, we have a procedure to determine the desirability of an outcome, in the form of a “checklist”. It is not as much as one could hope for, because it certainly would be more useful to have a way of “finding” such an outcome instead of simply being able to “verify” whether a pre-specified outcome is a desirable one. Nevertheless, this is a well-defined sense in which *converse consistency* can be said to be an operational principle. Besides, one can go further. Algorithmic procedures can indeed be developed in some contexts that do help us achieve the more ambitious goal of finding the desired outcomes. (Such procedures are proposed in Thomson, 2011b.)

Can the same thing be said about *consistency*? A claim that it does help us find the solution outcomes of reduced problems is hard to support because the reduced problems are not given beforehand. It would be different if, starting from any problem that we may be interested in, we could embed it in a larger one whose solution outcome we could easily find, and such that reducing it with respect to the initial population and this solution outcome would give us the problem we care about. But that is not what *consistency* says.

However, we would like to propose another sense in which the term “operational” may apply to *consistency* itself. It is that, when the choice has been made of a two-agent rule, *consistency* may provide a way of settling problems involving arbitrarily many agents: indeed, there may be a rule defined over the entire domain of problems that is *consistent* and in the two-agent case, agrees with the given two-agent rule. Such a rule is a **consistent extension** of the two-agent rule.³²

Not all two-agent rules have *consistent* extensions however, and an interesting question is determining when that is the case. In the context of claims problems, answers have been provided for a number of two-claimant rules and general procedures have been developed that are often applicable. (Thomson, 2003, 2012, surveys some of these results.)

Altogether then, in situations in which the two-agent case has been settled, *consistency* may provide the basis for the construction of operational procedures to handle all problems. There does remain the question of choos-

³²This should not be confused with the concept of a minimal “consistent enlargement” of a solution (Thomson, 1994), which is the smallest solution that is *consistent* and contains the solution.

ing a two-agent rule, and of course we should still subject to a critical examination the manner in which *consistency* allows us to extend the way in which we solve two-agent problems to the way in which we solve problems involving more than two agents.³³

(d) **Consistency as “reinforced fairness”**. In the introduction, we quoted Balinski and Young (1982)’s expression of the consistency principle, “Every part of a fair division should be fair”. It is remarkable that in spite of its lapidary quality, translating this sentence into a mathematical statement actually takes us quite close to our formal and somewhat complex definition.³⁴ “Every part of a fair division” seems to refer to the result of a projection operation relative to a subpopulation of agents, this operation being performed on an outcome that is thought fair for some initial problem, which has to mean an outcome that satisfies some punctual fairness notion. The predicate “should be fair” is what the axiom requires: the restricted outcome should be fair for some second problem, presumably it should satisfy the same punctual fairness notion; that second problem should of course be one in which this outcome is feasible, and it is natural to choose as its social endowment the sum of what the members of the subpopulation had been assigned initially. Altogether then, “Every part of a fair division should be fair” can be restated, in our words, as the requirement that “the reduction operation should preserve punctual fairness”. Balinski (2005) refers to “Every part of a fair division should be fair” as a “slogan”.³⁵ The term is not too strong: phrased in this seductive way, how could anyone object to consistency?

There are two problems with it however. The first one is its more limited scope than that of the general idea under discussion in this essay: it seems

³³Aumann and Maschler (1985) distinguish between a rule being (i) “consistent with a particular two-claimant rule”, in our language, being a *consistent* extension of a two-claimant rule, and (ii) being “self-consistent”, which we refer to as being *consistent*. Young (1994) also makes the distinction, calling (i), which he attributes to Huntington (1921), as a “primitive” form of *consistency*, later on, as a “rather pale form” of *consistency*. We do not endorse these epithets.

³⁴Balinski and Young intend the sentence to capture *flexibility* as well. It is harder to read this second property in the definition. However, as we have already noted, for a *single-valued* mapping, *consistency* and *flexibility* are equivalent. Balinski and Young write that consistency as they formulate it is an “inherent part of any fair division”. They add “For example, one property of a fair division of an inheritance should be that no subset of heirs would want to make trades after the division is made.”

³⁵Balinski rephrases it slightly as “Any part of a fair division must be fair”.

mainly to apply to the problem of dividing privately appropriable resources. We do not perceive this in itself to be too much of an issue and in fact we have mainly considered this problem in this essay. Like an overstretched empire, it is partly because of the considerable breadth of its dominion, the fact that it has been applied to such a great variety of problems, that the consistency principle may present a challenge to interpret. By focusing on a narrower class of applications, some of these difficulties may in fact disappear.³⁶

The second problem is more serious: it is that in the usual formulation of the consistency principle as well as in Balinski and Young's own mathematical expression of it—it is the one that we have been exploring—as distinct from their maxim, there is no assumption that the outcome that is the point of departure should satisfy any punctual fairness notion. The only assumption is that it should be selected by the solution. The technical content of the consistency principle is simply a requirement of invariance and robustness: what has been chosen for some problem, *for whatever reasons*, when appropriately restricted, should still be chosen in each associated reduced problem. Balinski and Young's slogan does capture the idea of preservation (by the reduction operation) and its compact elegance makes the desirability of preservation virtually self-evident, to the point that we may feel absolved of finding any additional reason for it. However, it unnecessarily and unjustifiably adds, again using our language, that it is punctual fairness that should be preserved. Once again, this is not part of the consistency principle (recall the *consistency* of the dictatorial rules). We have argued that if the principle has something to do with fairness, it is mainly with relational fairness, in particular solidarity, not punctual fairness.

Balinski and Young's maxim does raise an interesting question though, namely whether *consistency* is particularly justified, or only justified, when applied to a solution that is also required to satisfy some punctual fairness notion. If we care about fairness, shouldn't we be "systematic" and insist on its various expressions. Can we ask for one form of fairness and ignore other possible forms?

We do not see why this should be the case, any more that we see any of the monotonicity properties that have discussed (with respect to resource, or population, or technology) to be only meaningful, or even particularly meaningful, for a rule that already satisfies some punctual fairness notion, or

³⁶Young (1994) does apply the principle broadly, to abstract problems such as bargaining as well as to concrete resource allocation problems.

that we cannot impose one without imposing the other. Clearly, as already noted (Section 11), fairness has “components”. Each relational fairness axiom pertains to possible changes in one of the parameters that enter into the description of a problem. In each particular type of situations that we face, only some of these parameters are likely to change and the others remain fixed. It is important to understand how the implications of general requirements depend on the scope of their expression, and if possible, to explore the implications of each relational axiom separately. Besides, as we have seen all too often in the development of the axiomatics of resource allocation, insisting on overall fairness would be pointless: in many models, in particular when *efficiency* is imposed as well, the conclusion would be an impossibility.

The same comment applies to axioms of robustness to strategic behavior. Should we say that a rule is “strategy-proof” *only* if it is immune to *any kind* of strategic moves that agents may make: misrepresenting their preferences, transferring endowments among one another, lying about the expert knowledge they have on the probability of some future event, and so on? Probably not. Instead, researchers in this field have found it meaningful to break down the quest for rules that are robust to strategic behavior into components, and we believe that this is a sound approach. What it delivers is an understanding of the tradeoffs between the various ways in which this robustness can be expressed and other properties, which in turn allows us to better tailor rules to applications.

(e) **Consistency simply as “reinforcement”**. Consider a solution and an economy in its domain. Specify a list of assignments, one for each of the agents involved in this economy. For each group of agents, including the entire group, focus on the problem of allocating among them the sum of the assignments intended for them. Let us require of the initial list that for each such group, the restriction of the list to this group should be one of the choices the solution makes.

Here, the larger economy is not given a particularly prominent role, by contrast to our previous story involving agents leaving with assignments calculated for them in that economy. All groups of agents are placed on the same footing. I should say “almost” placed on the same footing, the qualification being needed because the data of the problem include a social endowment for the largest group only—this endowment is given exogenously—not for the subgroups. For these subgroups, preferences are given beforehand, but there is no social endowment; this parameter is calculated from the allocation that

is evaluated.

The “multi-part” test just described is on the allocation and the solution together. However, it naturally leads to a requirement on the solution itself: for each economy, there should be at least one allocation passing the test. To better understand the requirement, let us rewrite Balinski and Young’s maxim so as to remove from it the reference to punctual fairness. We obtain “Every part of a socially desirable allocation (that is, one chosen by the solution) should be socially desirable”, but to be explicit about the quantification over the allocation that is the point of departure, let us restate it again as “Every part of *every* socially desirable allocation should be socially desirable”. By contrast here, we are requiring of the solution that “There should be *at least one* allocation that it deems desirable and whose every part it also deems desirable”.

When applied to rules, this last statement is in fact *consistency*, but for correspondences, there is a difference: returning to our previous language and including all variables, our formal definition of *consistency* for a solution was that for each economy, *for each allocation* chosen by the solution for that economy, and for each subgroup of the agents involved, the restriction of the allocation to the subgroup should be chosen by the solution for the associated reduced economy. The statement under scrutiny here is that for each economy, *there should be at least one allocation* chosen by the solution such that, for each subgroup Let us refer to this property as **reinforcement**.

Reinforcement, which is obviously weaker than *consistency*, is not equivalent to the requirement that the solution should contain a *single-valued* subsolution, a **selection**, that is *consistent*. Indeed, for the problem of fair division, the equal-division Walrasian solution satisfies *reinforcement*, (in fact it is *consistent*.) but it has no *consistent* selection.

Reinforcement should simply be seen as a strengthened version of the one-economy test on an allocation, as we said, a “multi-level” application of it. The desirability of an allocation for an economy would be “confirmed” by the fact that its restrictions to subgroups are also judged desirable for the economies whose endowments are derived from it by summations. Would an analogy to “fractals” be too far-fetched here (Mandelbrot, 1982), and to the property, called “self-similarity”, of a geometric object that its characteristics are reproduced on a smaller scale by each part of it?

Under that interpretation, there would be no need to invoke extraneous considerations to justify the idea except to say that a requirement imposed on society as a whole should be satisfied in its subsocieties. The question of

existence of allocations passing this multi-level test would then be settled if the solution on which it is based is *consistent* because what could also be called this “confirmation test” would be passed by *each* of the allocations the solution chooses in the economy involving everyone. This would certainly be another sense in which *consistency* could be seen as an operational requirement too, but it is *reinforcement* that it would make operational. It would be a very minimal sense however because only existence would be obtained.

17. Conclusion. We close with a summary of what we have attempted here. It was mainly to propose an interpretation of the consistency principle based on a post-application and efficiency-free expression of the solidarity ideal in response to population changes. “Post-application” means that the rule under consideration has to be applied first, before the changes for which we require solidarity are contemplated. “Efficiency-free” means that solidarity in its “pure” form is invoked, namely the requirement that the welfares of all agents should be affected in the same direction, the direction itself not being specified; thereby, “contamination” from what should be independent considerations of *efficiency* is avoided. The specific *consistency* axioms that have been studied in the literature for various models can all be seen as the result of adding “some” *efficiency*, and in some situations, requiring invariance with respect to Pareto-indifferent moves, *Pareto-indifference*. The end-results are invariance requirements.

Thus, because solidarity is a normative requirement having to do with fairness and ethics, the consistency principle can be seen as contributing to fairness and it is concerned with ethics. We have emphasized however—and the observation applies to all of the properties that have been discussed in the fairness literature—that *consistency* should neither be considered as sufficient or necessary for overall fairness. On the one hand, in each application, other punctual and relational fairness requirements may well have to be imposed on rules for them to be judged fully satisfactory from this viewpoint. On the other hand, in situations in which one may feel that assessing how rules respond to population changes is unimportant, or not important enough, *consistency* considerations may safely be disregarded.

We have also argued that much of the confusion about *consistency* comes from the failure to distinguish between punctual axioms and relational axioms. To the extent that the consistency principle is about fairness, it is about relational fairness, not punctual fairness. No reference to punctual fairness should be made in describing it and justifying it. Conversely, its

compatibility with lack of punctual fairness—extreme violations of which are exemplified by certain types of dictatorial rules—should not be understood as meaning that *consistency* is neutral to fairness, let alone hinders fairness, although we did uncover a limited sense in which this may be true.

We have also attempted to define and develop an interpretation of the consistency principle as an “operational” principle. We have proposed objectives that the principle can be understood as making operational, and pointed out that the manner in which it does so is best justified when, once again, a reference to solidarity is made. Finally, we have evaluated other interpretations of the consistency principle, as a requirement of robustness under piecemeal implementation of allocation rules, and as a requirement of reinforcement or self-similarity.

References

- Abizada, A. and S. Chen, "Allocating tasks when some may get canceled or additional ones may arrive", mimeo, 2011.
- Aumann, R. and M. Maschler, "Game theoretic analysis of a bankruptcy problem from the Talmud", *Journal of Economic Theory* 36 (1985), 195-213.
- Balinski, M., "What is just?", *American Mathematical Monthly* 112 (2005), 502-511.
- Balinski, M. and P. Young, *Fair Representation*, Yale University Press, New Haven, 1982.
- Chun, Y., "The solidarity axiom for quasi-linear social choice problems," *Social Choice and Welfare* 3 (1986), 297-310.
- Chun, Y. and W. Thomson, "Monotonicity properties of bargaining solutions when applied to economics," *Mathematical Social Sciences* 15 (1988), 11-27.
- Davis, M. and M. Maschler, "The kernel of a cooperative game", *Naval Research Logistics Quarterly* 12 (1965), 223-259.
- Ehlers, L. and B. Klaus, "Efficient priority rules", *Games and Economic Behavior* 55 (2006), 372-384.
- and —, "Consistent house allocation", *Economic Theory* 30 (2007), 561-574.
- Ergin, H., "Consistency in house allocation problems", *Journal of Mathematical Economics* 34 (2000), 77-97.
- Foley, D., "Resource allocation and the public sector", *Yale Economic Essays* 7 (1967), 45-98.
- Kalai, E., "Proportional solutions to bargaining situations: interpersonal utility comparisons," *Econometrica* 45 (1977), 1623-1630.
- Kim, H., "Population monotonicity for fair allocation problems", *Social Choice and Welfare* 23 (2004), 59-70.
- Lensberg, T., "Stability, collective choice and separable welfare", Ph.D. Dissertation, Norwegian School of Economics and Business Administration, Bergen, Norway, 1985.
- Lensberg, T., "Stability, collective choice and separable welfare", Ph.D. Dissertation, Norwegian School of Economics and Business Administration, Bergen, Norway, 1985.
- , "Stability and the Nash solution", *Journal of Economic Theory* 45 (1988), 330-341.

- Mandelbrot, B., *The Fractal Geometry of Nature*, W. H. Freeman and Company; 1982.
- Maschler, M. and G. Owen, "The consistent Shapley-value for hyperplane games", *International Journal of Game Theory* 18 (1989), 390-407.
- Moreno-Ternero, J. and J. Roemer, "A common ground for resource and welfare egalitarianism", *Games and Economic Behavior* 75 (2012), 832-841.
- Moulin, H. and W. Thomson, "Can everyone benefit from growth? Two difficulties", *Journal of Mathematical Economics* 17 (1988), 339-345.
- Nash, J. F. (1950), "The bargaining problem," *Econometrica* 28, 155-162.
- O'Neill, B., "A problem of rights arbitration from the Talmud", *Mathematical Social Sciences* 2 (1982), 345-371.
- Plott, C., "Path independence, rationality and social choice", *Econometrica* 41 (1973), 1075-1091.
- Rawls, J., *A Theory of Justice*, Harvard University Press, Cambridge, MA, 1971.
- Roemer, J., "The mismatch of bargaining theory and distributive justice", *Ethics* 97 (1986), 88-110.
- Serizawa, S., "Inefficiency of strategy-proof rules for pure exchange economies", *Journal of Economic Theory* 106 (2002), 219-241.
- Shapley, L. S., "A value for n -person games", in H. Kuhn and A. W. Tucker (eds.), *Contributions to the Theory of Games* 2 (1953), 307-317.
- Shimomura, K.-I., "Flexible solutions to the fair division problems with single-plateaued preferences", University of Rochester mimeo, 1993.
- Sprumont, Y., "The division problem with single-peaked preferences: a characterization of the uniform allocation rule", *Econometrica* 59 (1991), 509-519.
- Thomson, W., "Monotonic allocation mechanisms: preliminary results", University of Minnesota mimeo, 1978.
- , "Equity in exchange economies," *Journal of Economic Theory* 29 (1983a), 217-244.
- , "The fair division of a fixed supply among a growing population," *Mathematics of Operations Research* 8 (1983b), 319-326.
- , "Problems of fair division and the egalitarian principle," *Journal of Economic Theory* 31 (1983c), 211-226.
- , "Monotonicity, stability and egalitarianism," *Mathematical Social Sciences* 8 (1984), 15-18.
- , "A study of choice correspondences in economies with a variable num-

- ber of agents”, *Journal of Economic Theory* 46 (1988), 247-259.
- , “Consistent solutions to the problem of fair division when preferences are single-peaked”, *Journal of Economic Theory* 63 (1994a), 219-245.
- , “Consistent extensions”, *Mathematical Social Sciences* 28 (1994b), 219-245.
- , “Welfare-domination under preference-replacement: a survey and open questions”, *Social Choice and Welfare* 16 (1999), 373-394.
- , “On the axiomatic method and its recent applications to game theory and resource allocation”, *Social Choice and Welfare* 18 (2001), 327-387.
- , “Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: a survey”, *Mathematical Social Sciences* 45 (2003), 249-297.
- , “Consistent allocation rules”, mimeo, 2011a.
- , “On the computational implications of converse consistency”, mimeo, 2011b.
- , “Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: an update”, mimeo, 2012.
- Thomson, W. and R. B. Myerson (1980), “Monotonicity and independence axioms”, *International Journal of Game Theory* 9, 37-49.
- von Neumann, J. and O. Morgenstern (1944), *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
- Young, E., “Progressive taxation and the equal sacrifice principle”, *Journal of Public Economics* 32 (1987), 203-214.
- , *Equity in Theory and Practice*, Princeton University Press, Princeton, NJ, 1994.

Index

- Abizada, 29
- alternative, 3
- Aumann, 7, 33
- axiom, 7
 - invariance, 6
 - punctual, 7
 - relational, 7
 - post-application, 15
 - pre-application, 15
- Balinski, 27, 33
- bargaining problem, 15
- Chen, 29
- Chun, 12, 14
- claims problem, 7
- classical exchange economy, 11
- classical problem of fair division, 11
- closedness under the reduction operation, 5
- coalitional game, 6
 - non-transferable utility, 19
 - transferable utility, 6, 19
- comprehensiveness, 19
 - strict, 19
- consistency*, 5
 - average*, 27
 - bilateral*, 27
 - post-disposal*, 28
 - pre-disposal*, 28
 - as an expression of coherence of choice, 8
 - as an expression of solidarity, 17
 - as operational principle, 31
 - as reinforcement, 35
 - as robustness under piecemeal implementation, 30
 - as robustness with respect to choices of perspectives, 29
- consistent* enlargement, 32
- consistent* extension, 32
- continuity*, 9
- contraction independence*, 16
- converse consistency*, 28
- Davis, 6
- dictator, 21
- dictatorial rule, 21
 - sequential, 22
- domination*, 15
- economy, 4
- efficiency*, 12
 - weak*, 19
- egalitarian rule, 19
- Ehlers, 28
- Ergin, 25
- flexibility*, 27
- Foley, 17
- individual endowment lower bound, 11
- Kalai, 15, 29
- Kim, 13
- Klaus, 28
- Lensberg, 19, 30
- Mandelbrot, 36
- Maschler, 6, 7, 27, 33
- model, 3
- Moreno-Tertero, 2
- Morgenstern, 6
- Moulin, 13, 23
- Myerson, 15

Nash, 15
 Nash rule, 19
 no-envy, 17

 O'Neill, 7
 object, 19
 Owen, 27

 parameter, 11
 collective, 11
 personal, 11
 semi-collective, 11
 Pareto solution, 12
 weak, 19
Pareto-indifference, 17
 Plott, 29
population monotonicity, 13
 post-application, 18
 post-production, 16
 pre-application, 18
 priority, 25
 problem, 3

 random arrival rule, 26
 Rawls, 24
 reduced problem, 5
reinforcement, 36
resource monotonicity, 13
resource solidarity, 11
 Roemer, 2, 12, 13
 rule, 4

 selection, 36
 sequential priority, 25
 Serizawa, 13
 Shapley, 26
 Shapley value, 26
 Shimomura, 27
 single-peaked preferences, 18
 solidarity, 9
 solution, 4

 Sprumont, 18
strategy-proofness, 13
strong monotonicity, 15

technology monotonicity, 13
 Thomson, 1, 5, 11–20, 22, 23, 26, 28, 31,
 32

 von Neumann, 6

 Walrasian solution, 36

 Young, 27, 29, 33