

On the terminology of economic design: a critical assessment and some proposals

William Thomson¹

Received: 28 March 2018 / Accepted: 11 May 2018 / Published online: 8 June 2018
© Springer-Verlag GmbH Germany, part of Springer Nature 2018

Abstract The purpose of this paper is to provide a critical examination of some of the terminology that is common in economic design and to suggest ways in which it can be improved.

Keywords Economic design · Terminology · Consistency

JEL Classifications D4 · D5 · D6

1 Introduction

Why should we care about good language? Because “Mal nommer un objet, c’est ajouter au malheur de ce monde” (Camus 1944). [“Misnaming an object adds to misery in this world”]

The purpose of these notes is to provide a critical examination of some of the terminology that is common in the literature on economic design and to suggest ways in which it could be improved. Mainly it is to encourage a conversation about language.

Research frequently generates new concepts but terminology develops in a haphazard way; it is rarely the result of deliberate and carefully thought-out choices authors make. Also, a term from common language that indicates with no ambiguity a concept that has to be named is often difficult to find.

Thanks to Battal Doğan, Eun Jeong Heo, Alexander Nichifor, Szilvia Pápai, and Utku Ünver for their comments.

✉ William Thomson
wth2@mail.rochester.edu

¹ University of Rochester, Rochester, NY, USA

Whether a new concept will be important in the development of a subject is not immediately clear. For a while, it will only be discussed by a small community of specialists. Unfortunately, during this period, usage of the name that designates it solidifies; as time passes, replacing it becomes increasingly difficult, even if better ones are found.

The common argument against changing a term that has been used for a while is that it is too late: people know the concept under that name and you would be confusing them. For a concept that may vanish in the near future, the change may indeed not be worth the trouble. But when a subject is alive and well, adapting and improving our language can only help it develop. If we switch to a better name for an important concept, people will not be confused for long. Besides, what about people who are new to the field, or the new generations of students to whom, year after year, we teach the subject? Challenging established terminology is worth it, and worth it at any point in time.

Examples of terminological messes abound in the literature that I discuss here. A striking one concerns the various families of rules that have been identified in the study of the assignment of indivisible resources, called “objects”. These rules are referred to as “dictatorships”, “priorities”, “queueing”, each of these terms being qualified by adjectives such as “sequential”, “serial”, “hierarchical”, and “lexicographic”. How can one possibly tell from their names how a “sequential priority rule” differs from a “lexicographic dictatorial rule”, say?

Specialized terminology is of course familiar to workers in the field, but when we address a general audience (even when we give a lecture or a seminar in an economics department), as opposed to a conference audience, it is not likely to be known. Choosing good terms will make it easier to someone with no prior knowledge of the subject to develop an understanding of it.

Another reason why we should feel free to improve terminology is that authors do not always agree on it anyway. This is true even for concepts as basic as efficiency à la Pareto. In order to distinguish between the two primary notions (when the test on an outcome is whether there is another one that everyone prefers, or simply one that everyone finds at least as desirable and at least one person prefers), most authors have picked two of the following three expressions, “strong Pareto-efficiency”, “Pareto-efficiency”, and “weak Pareto-efficiency”, and each pair has been used by someone (sometimes, expressions that do not include Pareto’s name are used). Until we read the definitions in each specific paper, we do not know which is meant.

Also, terminology eventually changes. Besides, an old term may not be applicable in some new situation, and some other term has to be invented anyway.

Unifying language is particularly important when surveying a literature because one of the goals then is to show how results are linked. In each specific study, there will only be few expressions designating related concepts but the terms used in different papers will often differ and it would be confusing to jump from one term to another when discussing them. Choices will have to be made.

If a family of related concepts have to be introduced, it is good to have a mold, a template, to create descriptive names for them. Several such templates are proposed below. The expressions that result are sometimes long or unwieldy, but shorter ones, mnemonic labels, and even transparent abbreviations can often be found.

Another guiding principle in choosing terminology is that jargon is better avoided if possible. If our work is to have an impact in the real world, we should be able to communicate our recommendations to the practitioners and non-economists in a language they can easily understand.

Thus, here are some suggestions, or rather, invitations to the reader to think about certain terminological issues.

2 General terminological and notational issues

- *Utility versus preferences* Many studies of resource allocation problems pertain to models in which agents are simply equipped with preferences, not utility functions. Authors often work with numerical representations invoking “convenience”, and calling them “utility functions”, but upon inspection, one often finds that there is actually no place where these functions are more convenient. The reference to utility is in fact a little dangerous because it may plant in the reader’s mind the thought that there is a meaningful cardinality to the agent’s evaluation of outcomes. If you use the term in a seminar, inevitably an audience member will ask a question or make a comment involving some type of utility comparison: he or she will observe that “agent 1’s utility is greater than agent 2’s utility” or that the “sum of the agents’ utilities is greater at x than at y ,” for example. Such statements are not meaningful in a model that only includes preferences.

Confusion can be prevented by only talking about preferences, and avoiding the term “utility” altogether. Also, if only preference information is included in the model, economic agents cannot be assumed to compare lotteries by means of their expected utilities. (Meaningful comparisons based only on preferences are possible however, such as ones derived from stochastic dominance considerations.)

- *Matching versus assigning versus allocating* Resource allocation theory has to do with deciding who should get what. The term “matching”, which originated in the theory concerning “marriage problems”, is frequently used to refer to physical resources being assigned to a group of people. This is an unnatural use of this term in this context.

Why object to this usage? After all, in a classical exchange economy, we are matching buyers and sellers, and in a marriage problem, we are matching men and women. Yet, in common language, we do not normally say that we are matched to an office or to task. We say that we are assigned to an office, that we are assigned a task. What gives standard two-sided “matching problems” their mathematical specificity is that a partition of the agent set is given beforehand. Each agent has to be matched to one or several agents on a prespecified “other side”.

By contrast, consider a classical exchange economy with two goods. At an allocation chosen by a rule, we can certainly partition the agent set into two subsets: some agents will give up some of good 1 in exchange for more of good 2, and the others will do the opposite. However, the partition is not known ahead of time, before the rule is applied. Besides, to achieve the trade specified for him, a given agent may have to get together with several agents trading in the opposite direction. Then, there is no specific trading

partner that he is matched with. In fact, it might be best to think of agents coming to a clearing house that delivers their net trades to them, instead of their being put in touch with specific partners to achieve these trades.

What we call a “roommate problem” is a closer counterpart to a standard economy because there is no a priori specification for each agent of some group of agents with whom that agent is required to be paired with. Still, analytically, the discreteness of roommate problems makes them close to “marriage problems”.

Returning to our earlier example, suppose now that there are more than two goods. Things are more complicated here because to implement the trade that a rule requires of an agent (his Walrasian equilibrium trade if the rule is the Walrasian rule), he may have to get together with different agents depending upon the good. The metaphor of a central clearing house to coordinate these exchanges is even more appealing here.

Because certain tools initially developed in the context of one-to-one matching have been found useful in the context of object assignment problems, the latter are sometimes covered under the umbrella of “matching theory”, but that does not seem to be enough of a justification.

- *Matching versus pairing* When matching two entities to each other, pairs are formed, and the term pairing would be more precise than matching. If the two entities are taken from two separate sets, the expression “two-sided pairing problems” could be used. A roommate problem is a pairing problem in which each person can be paired with any other person.

When the formation of larger groups is being considered, the term “grouping” may be appropriate. The expression “coalition formation” is common and when groups are formed on the basis of strategic considerations, the term “coalition”, which does have a strategic connotation, is a good one.

- *One-sided versus no-sided problems* It makes sense to speak of a matching problem as being “two-sided”, but should we say that a “roommate problem” is “one-sided”? Roommate problems have no sides. Would one refer to a segment as a one-sided polygon and to a standard exchange economy as one-sided? As we already pointed out, in such an economy, there are no pre-specified sides. The direction in which each agent is instructed to trade depends on the preference and endowment data of the economy and on the rule that is being operated.
- *School choice problem versus priority-augmented object allocation problem* A **school choice problem** has to do with the allocation of a set of objects organized in “types” (school seats) when to each object type is attached a priority order over the possible recipients. These recipients are equipped with preference relations over the object types and each type contains multiple identical copies of some object. In the base model, preference over object types are strict and so are priorities.

The literature has expanded beyond this specification; indifference is sometimes considered and priorities are not necessarily required to be strict. However, a school choice problem in which for each school, all students belong to the same indifference class is not a school choice problem anymore that an economy in which the production set consists of the origin of commodity space is a production economy; such an economy is an exchange economy.

What justifies making production economies a separate object of study is the presence of non-trivial technologies. Technologies raise conceptual and technical questions that do not occur in exchange economies. For example, whether returns to scale are constant, decreasing or increasing matters much for the existence of competitive equilibria.

Similarly a school choice problem in which all priorities are degenerate is a plain object allocation problem. What makes a school choice problem a school choice problem is the existence of non-trivial priority relations. The challenge is to interpret priorities, to give them operational meaning, to study how the profile of priorities in a problem should be structured for allocation rules to satisfy properties of interest, and to understand for each particular allocation rule how changes in priorities may affect the welfare of the participants.¹

Calling a school choice problem a “priority-based object allocation problem” is better but not enough because priorities are of course not the only data on which allocation can be based. Assuming an understanding of what an object allocation problem is, we propose an expression indicating that a school choice problem is an object allocation problem “augmented” with priorities, a **priorities-augmented object allocation problem**.

We should also note some ambiguity in the use of the term priority. Some authors speak of a school’s priority relation over the students. Others speak of it as the student’s priority at the school; we may say for example that a student has a high priority at a school if he lives close-by. Either usage is in agreement with common language. However ambiguity may occur: saying that “a school has a high priority for a student” can be understood to mean that in the school’s priority, he is ranked highly or that in the student’s evaluation, the school is ranked highly.

- *Separable and responsive preferences.* Different types of restrictions have been imposed on preferences in the literature. An illustration is when the issue is the assignment of workers to firms, and let us use it as a running example. These restrictions have to do with how the ranking of sets of workers is related to the ranking of individual workers by firms. Here are common terms.

A relation (over sets of workers) is **separable** if the following holds: given any set S of workers, and given any worker not in S whom on his own the firm prefers to having no worker, adding this worker to S yields a set that the firm prefers to S .

A relation is **responsive** (to its ranking over individual workers) if the following holds: given any set S of workers, and given two workers neither one of whom is in S , adding to S the worker whom on his own the firm prefers to the other, yields a set that the firm prefers to the one obtained by adding the other worker.

Some writers use the adjective “separable” to cover both of these definitions.

The similarity between the properties is not reflected in their names. We could say instead that a firm’s preference relation over sets **reflects the (absolute) desirability of individual workers** in the first case and that it **reflects the relative desirability of individual workers** in the second case. We could also say that a firm’s preference

¹ It is also confusing to refer to an object allocation problem in which there is only one copy of each object as a school choice problem.

relation over sets of workers **fully respects the desirability of individual workers** if both parts are satisfied. (We could go further and require the following: given any set S of workers, and given two other sets neither one of which intersects S , adding to S the one in the pair that the firm prefers to the other, yields a set that the firm prefers to the one obtained by adding to S the other set in the pair.)

An odd feature of these definitions is that the relation over individual workers is presented as a primitive and the relation over sets a derived concept. Should that be the case? In applications, it does not make sense to consider a firm employing only one worker. To operate, a firm should have a minimal workforce. The relation over individual workers could be interpreted as representing an evaluation of the workers' particular skills, but the relevance of that evaluation to the impact of the addition of a worker to a group of workers currently employed will in general also depend on what this group is, on how well this worker will fit in. For instance, the composition of its workforce in terms of gender, racial or ethnic diversity, or national origins, may matter to a firm. Altogether, a firm could be described in terms of two objects, (i) a ranking of individual workers and (ii) a ranking of sets of workers of minimal size for it to operate, with (ii) being related to (i) in certain ways along dimensions or characteristics of workers and sets of workers that are not directly related to (i) and (ii).

- *Money versus compensation medium versus one additional, infinitely divisible, good* A **quasi-linear economy** is one in which there is a "special" good such that for each agent, each pair of consumption bundles, and each amount of the special good, if the agent is indifferent between the two bundles, then he is also indifferent between the two bundles obtained from them by adding that amount of the special good. Under appropriate topological assumptions, such preferences, if continuous, can be given continuous numerical representations that are separably additive and linear in that good. The good is often referred to as "money".

Reference to money is not desirable because money is a concept that makes sense in the context of a particular economic institution, namely when allocation is guided by prices and involves agents maximizing preferences subject to a budget constraint. By contrast, the property under discussion pertains to the psychology of the agents involved, as reflected in their preference relations.

The way in which this special good enters preferences has significant technical benefits. First, it greatly simplifies the structure of the Pareto set. If an allocation is efficient, and unless this would take some agent to the boundary of his consumption space, any allocation obtained by arbitrarily transferring the special good across agents is efficient as well.²

Second, when we have fairness objectives in mind, it broadens our options in achieving it. To explain this, suppose first that the only resources to allocate are objects, and that each agent is supposed to consume only one. We then have what we called a "plain" object allocation problem. There is in general no hope of achieving any kind

² It is important to realize that the benefit comes from the fact that it is the *same* good that plays that special role for all agents. Having good 1 as the special good for agent 1 and good 2 as the special good for agent 2 would not help. This is why it is preferable to speak of quasi-linear *economies*, not of quasi-linear *preferences*.

of fairness here, certainly not in a deterministic way. For instance, if all agents most prefer the same object, a violation of *equal treatment of equals*, a requirement that is often seen as a “minimal” fairness requirement, is unavoidable. If the commodity space is enlarged, and in particular if some amount of an infinitely divisible good is added to the social endowment, it is tempting to say that we can use it to “compensate” the agents who are not receiving their most preferred object. But what does this really mean?

When the good is introduced, consumption space expands, and preferences should be redefined over this expanded space; people do not have preferences over individual objects any more, but over pairs consisting of some amount of the divisible good and one object. If a bundle a is preferred to a bundle b , some amount t of the special good is needed to add to b to reestablish indifference. For two bundles a' and b' obtained from a and b by adding to each of them the same amount of the special good, the same amount t of the special good is needed to add to b' to reestablish indifference to a' . Without the quasi-linearity assumption, a bundle a may be preferred to a bundle b , but adding to both a and b the same amount of the divisible good may reverse preference.

If preferences are quasi-linear, the special good can be used to make compensations so that some notion of fairness is met. Calling the good a “compensatory good”, or “compensation good” makes sense then. However, if preferences are not quasi-linear, these expressions become problematic. It remains true however that the good provides an extra “instrument” in achieving fairness (Alkan et al. 1991; Tadenuma and Thomson 1993; Su 1999; Velez 2016), but that does not seem enough of a reason to call it a “compensation” good.

- *Templates for definitions* To show relationships between definitions, it is convenient to have molds and templates for them. Here are several proposals.

1. *The “sequential” template* This template is discussed in a later section where the use of the “conditioning” prefix is also brought up.

2. *The “wise” template* A number of rules are defined in two steps, by first considering a simple version of a model, a version with only one good, or only one type of objects, or only one candidate—let us refer to these characteristics as “dimensions”—and then in generalizing the definition to what can then be called “multiple dimensions” by applying the original definition dimension by dimension. For this to be possible, preferences may have to satisfy certain separability properties, namely it should be possible to extract from each preference relation enough information in each dimension for the one-dimensional version of the rule to be applicable. It is also possible for an entire family of rules to emerge. We may obtain in the first case a family of rules by specifying some parameter in some set, and in the second case a family of rules by specifying for each dimension a parameter in the set. The list of these sets is then called a “profile”. They can be chosen independently.

When preferences over a k -dimensional space for $k > 1$ are separable, rules can be constructed from rules defined for $k = 1$ by applying them to each dimension separately. Using the construction “dimension-wise” seems to be a good way of naming them. Here are examples:

- (a) For the problem of fully allocating a social endowment of a single commodity among agents with single-peaked preferences, a central rule is the “uniform

- rule” (Bénassy 1982). Let us allow for multiple commodities, and assume that agents have preferences that are single-peaked commodity by commodity, and such that for each commodity, the “marginal” preferences for that commodity is independent of the consumptions of the other commodities. Then, applying the uniform rule commodity by commodity gives us the **commodity-wise uniform rule** (Amorós 2002; Adachi 2010; Morimoto et al. 2013; Cho and Thomson 2011).
- (b) For the problem of reallocating objects among agents each of whom is endowed with one object, the rule of arguably greatest interest is the core. When the model is enlarged so as to accommodate goods organized in types, each agent being endowed with one object of each type, and preferences have the separability property described in the previous paragraph, applying the core object-type by object-type gives us the **object-type-wise core** (Miyagawa 2004; Klaus 2008). The slightly shorter expression **type-wise core** might be sufficient here.
- (c) For the problem of deciding whether to accept a person who has applied for a position (to be an emissary, say), and each agent has preferences over this 0-1 choice, a rule may be based on a “tournament”. More generally, consider the problem of making such a decision about each candidate in some set of candidates when any number of candidates can be approved of (for instance, the issue might be to form a delegation). If preferences are separable and the decision is made candidate by candidate on the basis of a tournament for each candidate, to define the rule, we need a list of tournaments indexed by candidates: this list is a “tournament profile” and the rules are **candidate-wise tournament-based rules**, or **candidate-wise tournament rules**. If there are restrictions on the number of emissaries, additional constraints should be placed on the tournaments; other restrictions may relate the tournaments across candidates: we then have a **structured** tournament profile (Ju 2005).

3. *The “augmented” template* A rule may be defined on a class of economies the specification of which includes certain parameters, and may depend on this information in a particular way. A sparser description may not include such information, but rules can be defined by adding to it, by “augmenting” the model with it.

To illustrate, for the problem of allocating objects, a quota may be specified for each agent, that is, a number of objects that he should receive, rules being required to respect the quotas. An example of such a rule is a rule that assigns to each agent in turn, and in a predetermined order, the set of objects of cardinality equal to his quota that he most prefers among all sets of that size. As before, such a rule can be called a **sequential priority rule**. Alternatively, a rule can be defined on a sparser class of economies, economies whose description does not include quotas, by first specifying a quota profile—this profile is a parameter of the rule as opposed to being a parameter of the economy—and for each economy, applying the sequential priority rule to the economy obtained by adding the quota profile. We speak then of a **quota-augmented sequential priority rule**.

4. *The “lower bound” or “upper bound” template* In many situations, some lower bound is imposed on the welfare that an agent should experience at an allocation. These lower bounds are defined by reference to some right that the agent has, or some choice that he may be given. Examples of such rights or choices are the following:

- (a) “consuming his endowment”, in a model in which to each agent is attached a bundle of resources that he “owns” or is “endowed with”;
- (b) “consuming an equal share of society’s resources”, in a model in which a bundle of resources is specified that is interpreted as a social endowment; in such an economy, an equal share of the social endowment is not necessarily a “right” that an agent has, but it may still be a meaningful reference; (the notion of an envy-free allocation can also be understood as making operational the notion of collective ownership; yet, at such an allocation, an agent may be worse off than he would be at equal division);
- (c) “consuming the bundle he most prefers among those that he can attain using his endowment optimally and some technology to which he is given access”, in a model whose description includes a technology;
- (d) “being unmatched”, in a matching model in which agents have the option of not being matched, not being married (for a man or woman), or not working (for a worker), or not operating (for a firm), and so on.

In general one could say that a rule satisfies a “lower bound”, adding as a prefix the data that are used to define this bound. Thus, one can speak of a rule satisfying the **individual-endowments lower bound**, or the **equal-division lower bound**, or the **autarchy lower bound**, or the **being-single lower bound**, or the **being-unmatched lower bound**, and so on.

When upper bounds are discussed, the same terminological template can be used: we can speak of a rule satisfying the **identical-preferences upper bound**, although the adjective “identical” seems to place unneeded emphasis on the fact that agents have the same preferences. An expression such as the **like-preferences lower bound** might suffice.

In none of these expressions is it really needed to indicate that the bound is a bound on welfares.

- *Markets versus problems versus economies* The term “market” is often used to refer to a class of problems in which agents are individually endowed with resources and the question is how to best redistribute these resources among them. This is the situation considered by Shapley and Scarf (1974) when resources are objects. It does suggest individual ownership, and the suggestion is good, but it also brings to mind exchanges mediated through prices, thereby implicitly assuming that a particular economic institution is already in place, namely the institution that is commonly referred to as the market (or some variant of it). Expressions such as “market economies” and “market forces” strengthen the implied connection. This is not desirable because our goal is to identify economic institutions to guide exchange, with no a priori conception of what is best. In particular, there is no reason why prices should be part of their definition.³

Thus, the expression “housing market” does not seem appropriate to describe the problem of reallocating houses (more generally, objects) among house owners, espe-

³ The fact that the allocation obtained by applying Gale’s algorithm can also be obtained as an equilibrium of something that looks like a market is a mathematical coincidence that has no bearing to the issue at hand, especially because the price-taking behavior would hardly be justified given the unique character of each of the objects owned by the various agents.

cially when the expression “house allocation problem” is used to refer to the problem of allocating a social endowment of houses among agents who start out with no house of their own. Pairings such as house-allocation versus house-reallocation, or more abstractly and generally, **object-allocation** versus **object-reallocation** appear preferable. In one case, ownership is public; in the other, it is private.⁴

Similarly, using the expression “market design” to refer to the research program devoted to the identification of the most desirable ways of organizing production and distribution suggests that these decisions are mediated by prices. This is too narrow a definition and it obscures the ambition of this program, which is to ignore all constraints implied by the operation of price-based rules. In fact, in several central classes of allocation problems in which the issue of design—one should say, the need for design—arises precisely because prices cannot be used—they would be impractical; sometimes, they are forbidden by law—and speakers presenting their work on “market design” routinely begin by saying that they are concerned with situations in which there is no market and in which there cannot be one, so that some alternative must be found. What their work is about is “alternative-to-markets” design. This is the case in priority-augmented object allocation problems such as school choice, and in organ allocation problems (in which using prices is indeed forbidden). An expression such as “mechanism design”, used in Hurwicz’s pioneering work on the subject, conveys better the scope of the enterprise (the appropriateness of the term “mechanism” is discussed below.)

Concerning the ownership of resources, it may be split, some resource being held privately and the remainder constituting a collective endowment. The expression **mixed-ownership economies** is well-adapted to designate a class of economies that includes standard “private ownership economies” as well as “collective ownership economies”, a situation that has been considered when the resources are indivisible. This application has been discussed under the name of “house allocation with existing tenants” (Abdulkadiroğlu and Sönmez 1999): some housing units on a campus are “privately owned” by the returning students, in the sense that they have a right to them, which for a rule means that it should assign to each such student a unit that he finds at least as desirable as the one he owns, and others are not owned; these are the units that have been freed by the students who graduated, together with newly built units. They can be thought of as a social endowment.⁵

- *Quotas versus capacities* In the context of object allocation problems, the number of copies of a particular object is sometimes called its “quota”. Is that a good term? In common language, the term “quota” has multiple meanings. One is “proportion”, which is not what is meant here. A second one is “upper bound”. The term is commonly given that meaning in international trade: an “import quota” is an upper bound on trades: at the prevailing prices, importers would want to import more

⁴ The fact that the rule that is arguably most central to handle object-reallocation problem, the so-called top-trading-cycle rule, can be given a market interpretation is irrelevant. There is no reason why a class of problems should be named after a rule that is particular relevance to it.

⁵ In “Consistency in generalized economies” (Thomson 1992, 2012), I used the expression “generalized economies” to refer to what I am here proposing to call “mixed-ownership economies”. The term “generalized” is not enough, as it provides no clue as to the type of generalization that is involved.

but they are prevented from doing so. In common language, the term quota is also used to designate a lower bound: a quota on female candidates is imposed in some European elections. A third meaning of “quota” is “desired amount”, a goal to be achieved: a “work quota” imposed on a worker may refer to the number of units that he is required to produce in a workday.

If a school is required by a school board to enroll a minimal number of students of a certain type, for instance minority students, or a maximal number of students of some other type, for instance foreign students or out-of-state students, then the term quota makes sense. In the literature on “school choice” or “college admission” however, it has been used to simply mean the capacity of the school or of the college, and that does not seem to be a good use of the term.

“Quorum” has only one meaning: it designate a lower bound. The term is correctly used in this sense by Monte and Tumennasan (2013) in their study of a certain class of object allocation problems as a lower bound on the number of people consuming a particular object. It is a plural genitive of a Latin noun, so it has no plural, but we need a plural form when there is one such number for each object. Given that the term is well implanted in English, speaking of quorums might be acceptable, as Monte and Tumennasan do. (For a discussion of how much of a violation of the grammatical rules of the languages from which we have imported certain terms we should indulge in when we need to make room for them, perhaps modify them, within the structure of the English language, Pinker (2014)’s nuanced position is compelling.) Certainly, one should not write “quora”, using the distinction “datum-data” as template, “data” being the Latin plural of “datum”. Perhaps one of the expressions “exact quota”, or “target”, or “numerical target”, could be used. Quorums are attached to objects, not to recipients.

“Quorum” unambiguously designates a lower bound, but “quota” does not always designate an upper bound. The expressions “lower bound” and “upper bound” are unambiguous, and perhaps preferable to the pair “quorum-quota”.

- *Assignment versus allotments versus allocations* In a branch of the literature under discussion, the expression “assignment problem” is applied to a specific class of pairing problems in which matching entities on two pre-specified “sides” produces value, this value can be divided among everyone, and parties only care about their share and not about who they are matched with (Shapley and Shubik 1972). An example is when workers and firms are matched, each match produces output, the dividend being the revenue that is generated by selling this output. (I use the term “dividend” in the sense of the “quantity to be divided”.) “Value” simply means “money” then.

There is no reason why the expression should be understood in this way. In common language, these connotations are absent: the term “assignment” simply refers to whatever an agent is given by a rule; the resources could all be divisible; none of them could be; there could be mixtures of resources of two types, some resources being divisible and the others not. In a school choice problem, a student can be assigned to a school and a school seat can be assigned to a student.

Similarly, the term “allotment” has appeared in several theoretical studies to designate what an agent is assigned by a rule when the resource is an infinitely divisible

commodity and agents' preferences are single-peaked over how much they consume (Barberà et al. 1997). Again, there is nothing in common language that makes this term particularly appropriate for that application.

- *Phantom voters versus calibration points versus signature* Consider the problem of choosing a location in an interval when agents have single-peaked preferences over the interval. A *strategy-proof* rule can be described in terms of a number of points in the interval, indexed by coalitions and satisfying certain relations. If the rule is *anonymous* and *efficient*, the number of parameters is equal to the number of agents minus one. Given a typical profile of preferences, to find out what the rule selects, we take the median of the agents' most preferred locations and these points. The points are called "phantom voters".

The expression originates in Border and Jordan (1983),⁶ although Moulin (1980) attributes the concept to Murakami (1968), himself calling "fixed ballots" the ballots supposedly cast by these phantom voters: "Adding fixed ballots to the voters' ballots is a technical device already used by Murakami to describe the so-called representative system of social functions between two alternatives".

Speaking of "phantom voters" is confusing. First, knowing a voter's most preferred level is of course not in general sufficient to deduce his whole preferences. To the extent that a rule only depends on the profile of most preferred levels, the so-called "peak-only" property, this may not matter much, and the initial literature on the subject was written under that restriction. One could also say that "phantom voters may not have all the characteristics of actual voters but we could specify preferences for them whose preferred levels would be these parameters". But what would be the purpose of introducing these additional data since they would not be used in calculating the choice? Besides, when the peak-only restriction is not imposed, the expression "phantom voter" would not be well adapted.

The term "phantom" suggests that the parameters attached to a rule come out of nowhere, and it almost absolves us from attempting to interpret them. But they do have a clear interpretation: to the extent that they reflect preferences, it is actually the preferences of the rule (as implicitly would the parameter attached to any rule that is a member of a parameterized family), and not the preferences of fictitious voters.

The term "phantom" may suggest that the parameters are exogenous. And indeed, it is for the authority in charge of solving the problems under consideration to specify them ahead of time, before preferences are known. But if it is this exogenous character that the term "phantom" is meant to suggest, this usage is not very compelling. When we solve object allocation problems by means of a sequential priority rule, by letting agents choose according to a particular order the object they most prefer among those that remain, we certainly would not want to refer to this order as a "phantom order".

More importantly, the proof of the characterization of the class of *strategy-proof* rules makes it clear (again, see Moulin's paper), that the parameters are simply the choices that such a rule makes for certain kinds of societies, societies of "extremists", that is people whose most preferred alternative are either endpoint of the interval of possible choices. The structure of the characterization proof is as follows. You first

⁶ This is the earliest reference that I could find.

apply the rule to a society composed only of “leftists” (agents whose most preferred point is the leftmost point of the interval), and make a note of the choice the rule makes. If the rule is efficient, we can actually skip that step since only one point is efficient then. Then you apply it to a society in which all agents but one are leftists and the remaining one is a “rightist” (an agent whose most preferred point is the rightmost point of the interval); this gives you another point; if the rule is *anonymous*, this point is independent of the identity of the rightist. Then you apply it to a society in which all agents but two are leftists and the remaining ones are rightists; this gives you a third point; again if the rule is *anonymous*, this point is independent of the identities of the two rightists ... You continue in this way until all agents have become rightists. Again, you need not apply the rule then if efficiency is required. Then, you show that *strategy-proofness* implies that the rule enjoys certain independence and monotonicity properties (as happens on any domain). This allows you to establish a number of relationships between the points. Finally, you determine how the rule handles any economy, once again using the independence and monotonicity properties it satisfies.

In the search for good terminology here, physics may help. To calibrate a thermometer, you dip it in freezing water, which gives you 0C, and then in boiling water, which gives you 100C. Equipped with a theory of how density of mercury varies with temperature, you then know how to measure any temperature (in practice, only within a certain range of course). To calibrate a *strategy-proof* rule, you dip it in societies of extremists. Whether you bother with societies composed only of leftists or only of rightists depends on whether you have imposed *efficiency* (because in either of these cases, there is a single efficient choice.) Whether you care about the identities of these extremists depends on whether you have imposed *anonymity*. Together with what you know of the behavior of *strategy-proof* rules, you then deduce what the rule does for any society.

Going beyond this example, proceeding in the manner just described is in fact a standard way to identify the members of a family of rules satisfying a list of axioms that do not force uniqueness. For instance, if you drop *symmetry* from Nash’s (1950) characterization of the so-called Nash bargaining solution—I will refer to it as Nash’s rule—and add *strong individual rationality*, the requirement that each player be assigned a payoff that is larger than his disagreement payoff (the *strict disagreement point lower bound* would be a better expression), you obtain a one-parameter family. Consider a rule satisfying *Pareto-optimality*, *scale invariance*, *strong individual rationality*, and *contraction independence* (Nash’s list of axioms from which *symmetry* has been deleted and to which *strong individual rationality* has been added). In order to find out the value of its parameter, apply the rule to the simplex (rather, the comprehensive hull of the simplex in order to get a well-defined bargaining problem). The calibration of a rule consists of a single test.

Depending upon the context, what you obtain from a calibration test is a point in a vector space, a probability distribution, an order, and so on, and each such object is a “calibration point”, or a “calibration probability”, or a “calibration order”, and so on. Put together, these data constitute the **signature of the rule** (a term used by Moulin 1980).

Of course, there is usually more than one way to calibrate a rule. In the case of a thermometer, we could use other materials whose melting or boiling points we know.

In the context of Nash's problem, we could apply the rule to the restriction to the non-negative quadrant of a circle centered at the origin; or to any other problem whose Pareto boundary is smooth, and whose endpoints belong to the axes. However, we should not use a problem whose Pareto boundary has a kink in its relative interior because then, there would be a whole range of weighted Nash rules that would pick the point. Thus, applying a rule to this problem would not identify it uniquely; it would simply restrict it to belong to a certain class.

As we have seen, more than one calibration test may be needed. Part of the analysis of a class of problems involves identifying a family of calibration tests that allows us to tell each rule in the family apart from the others. I refer to this collection of tests as a **calibration protocol**. Ideally, the family should be minimal (there should be no redundant test). Even better, it should be simple and the tests should have a transparent interpretation. For Nash's bargaining problem, a one-problem protocol suffices, and the simplex certainly passes both tests. For the problem of selecting a point from an interval when agents have single-peaked preferences over the interval, societies of extremists are simple too and they have a clear meaning. Sometimes, the tests composing the calibration protocol can be administered in any order. Sometimes, the protocol has to be structured in particular ways.

The larger the family of rules to be characterized, the larger the number of calibration tests that may be needed, and the more complex the calibration protocol may have to be. To illustrate, when characterizing the class of monotone path rules in bargaining theory as the only rules to be *weakly Pareto optimal* and *strongly monotonic* (Thomson and Myerson 1980), we need an infinite family of calibration tests. The simplices in \mathbb{R}_+^N of equation $\sum_N x_i = k$ for $k \in \mathbb{R}_+$, can serve this purpose. Indeed, to each unbounded monotone path in utility space emanating from the origin is associated a monotone path rule. (Actually, considering only those simplices associated with a dense subset of the non-negative reals would suffice because these rules are continuous.)

I have already pointed out that if we drop the *peak-only* assumption, it would become unclear whether a rule would have to be described in terms of a list of complete preference relations, the preferences of phantom voters. In fact, rules are still parameterized by a list of points. Here are additional arguments against "phantom voters" terminology.

1. Suppose that we generalize the class of admissible preferences by allowing plateaus (see above for a discussion of this term) instead of a peak. Should we then give the phantom voters plateaus instead of peaks? The answer is no (Moulin 1984) The acceptable rules are still parameterized by a number of points in the interval of possible choices and these points obtained by applying the rule to societies of extremists.
2. Next, let us turn to probabilistic rules (rules that choose probability distributions on the line instead of points). The answer here is that if a rule is *strategy-proof*, there is a list of probability distributions on the interval that are ordered by stochastic dominance (the counterpart of monotonicity properties that the parameters indexing a rule in the deterministic case have to satisfy), such that for each preference profile, the choice it makes is the median of the agents' preferred locations and additional locations chosen according to these probabilities (Ehlers et al. 2002).

We have the habit of representing probability distributions as bell-shaped, that is, single-peaked, so once again, it almost looks like we are adding phantom voters! In fact, we are not, for two reasons. The single-peaked functions that we use to represent what we call single-peaked preference relations are defined only up to monotone transformations, but probability distributions are uniquely defined. Second, the bell shape of a probability distribution is only a special case for the result under discussion. The distributions are not subject to any restrictions; they may have more than one peak; they may have atoms. Clearly, these distributions cannot be interpreted as preference relations. To identify them, proceed as usual and apply the rule to societies of extremists. Here, the calibration protocol delivers “calibration probabilities”.

Of course, for each of these extensions, simply speaking of calibration and signatures does not immediately tell us what to do. It is clear from the various papers just cited that much work remains. But it might help a little in guessing results and proving them. And for sure, it will help a lot in explaining the material to newcomers to the literature on *strategy-proofness* as well as other literatures, the process being common in many other areas.

3 Rules

- *Mechanism versus rule* The term “rule” is used in these notes to designate a *single-valued* mapping from a class of economies to a set of outcomes. Having a separate term to refer to the manner in which outcomes are calculated is useful, one could argue, necessary. This calculation may involve a geometric construction, an algorithm, a system of differential equations, a tâtonnement process, and so on. The term “mechanism” could be used for the general category. When we open a grandfather clock, what we see is a mechanism (cogwheels, pendulum, weights, pulley,...). If we are not particularly interested in the mechanism but only in the mapping defined by the mechanism, a second term is needed. “Rule” seems to be appropriate to that purpose.⁷ A rule can even be defined without a specific mechanism being specified that would help calculate the outcomes it prescribes.⁸

Thus, one can speak of the **deferred acceptance rules** as the mappings from the space of matching problems to the space of matches defined by operating the (men-proposing and women-responding, or women-proposing and men-responding) **deferred acceptance algorithms**. These algorithms are defined by lists of instructions that you feed a computer. One can similarly distinguish between the **Walrasian rule**

⁷ Admittedly, the term “rule” has other meanings. For instance, in the description of an algorithm, it can legitimately be used to designate what is to be done in each of several possible cases that are enumerated. This usage is also common in implementation theory; different cases are distinguished in describing strategy profiles and in each case, a “rule” is applied to determine the outcome.

⁸ We noted earlier that the research program associated with Hurwicz’s name has to do with more than the identification of processes that lead to good outcomes, but also with the identification of desirable mappings from economic environments to outcome spaces. However a primary concern of Hurwicz’s as well as of the writers in the so-called socialist controversy indeed was processes. This concern was reflected in the focus of that literature on planning procedures.

(to designate the mapping that associates with each economy the allocation at which, at some prices called “equilibrium prices”, demand equals supply, on domains on which there is indeed only one such allocation; otherwise the expression **Walrasian correspondence** can be used) and the **Walrasian mechanism**, an expression that seems appropriate to designate the price adjustment process that sometimes (if preferences satisfy certain properties) helps us find equilibrium prices (“tâtonnement”).

In some fields (notably matching and object-allocation problems, more generally, when discrete resources have to be allocated), most rules are defined through algorithms, and they are given the names of these algorithms. Using the prefix “deferred acceptance” when naming the rules induced by the deferred acceptance algorithms is fine, but the rules themselves are not algorithms.

In fact, I can well imagine an objection to the expression “deferred acceptance rule” too because the rules can also be defined in other ways, a fixed point theorem for example (Adachi 2000). To emphasize their striking welfare implications, we could alternatively call them the “men-optimal rule” and the “women-optimal rule”. In these expressions, the sense in which this “optimality” is achieved is unclear but it would be too much of a mouthful to be explicit and say “optimal-within-the-set-of-stable-matches”. As is almost always the case, an expression of reasonable length that would leave no ambiguity is a terminological impossibility.

It may be true that studying the properties of rules that are defined through algorithms might be facilitated by having access to these algorithms, but that does not seem enough of a reason not to call them algorithms. Besides, for other properties, alternative definitions may be more convenient.

- *Dictatorship and range-dictatorship* The Gibbard–Satterthwaite theorem asserts that a rule satisfying certain properties is such that there is an agent, specified beforehand and once and for all, such that for each preference profile, what the rule chooses is an alternative in the range that this agent finds at least as desirable as any other alternative in the range. It does not assert that the outcome is an alternative that this agent finds at least as desirable as any other feasible alternative.

To make it clear that the power given to this privileged agent can be specified in several ways, let us refer to a rule of the first type as a **range dictatorship**, and to the latter as a **full dictatorship**. If the rule is required to satisfy certain additional properties, the most prominent one of which is *efficiency*, then full dictatorship will be the outcome. To illustrate, for a wide class of models, if a rule is *efficient*, then any alternative is in its range, so the requirement that the rule be *onto* yields full dictatorship. For others, only a few alternatives can be in the range of a dictatorial rules: examples are the classical problem of fair division, the partitioning of a non-homogeneous linear or circular, continuum, and so on.

- *Sequential dictatorships versus sequential priority rules* In discrete allocation models, a standard way of obtaining *strategy-proofness* is to assign to each agent in a sequence—depending on the model, certain constraints may have to be respected—what he most prefers. For example, in abstract Arrovian social

choice,⁹ if agents are assumed never to be indifferent between any two alternatives, the choice of the “dictator” determines the alternative. If indifference is allowed, it does not, and a second agent may have to be called upon to break ties; if needed, a third agent and so on. These successive choices limit the opportunities of the agents who come later in the order. The order may be given exogenously, or it may be generated endogenously: the identity of the agent whose preferences are maximized at each step may depend on the assignments to the agents who have come before him, or even on their preferences.

I have already discussed this example in the introduction and noted that rules of this type have been variously referred to as “dictatorial rules”, “priority rules”, or “queueing rules”, these expressions being often prefixed by the qualifiers “sequential”, “serial”, “lexicographic”, or “hierarchical”. The adjective “dictatorial” has also been applied to domains (specifically, to domains on which *strategy-proofness* implies that power is distributed in the most skewed way). Unfortunately, common language does not allow us to relate phrases such as “sequential priority rule” or “lexicographic dictatorial rule” to their formal definitions, and virtually each of the other combinations has appeared somewhere.

What complicates matters is that other rules are defined in a sequential manner, the sequencing not reflecting or implying an asymmetric treatment of agents, and in fact being fully compatible with their symmetric treatment. The terms “serial” and “recursive” have been used in that connection.

Here is a proposed resolution. In situations in which agents are treated asymmetrically, as described two paragraphs above, two expressions could be used, **sequential dictatorial rules** and **sequential priority rules**, depending upon the “degree” of this asymmetry. Given the very negative connotation of the term “dictator”, the term could be reserved for rules for which whoever comes first generally leaves the other agents with their most undesirable outcomes, the term priority being applied to situations in which not being first does not necessarily mean being left with no choice, or very undesirable choices.

To illustrate the difference, in a classical problem of fair division, if an agent’s preference relation is maximized first over the entire feasible set, and his relation is strictly monotonic, he ends up being assigned the entire social endowment. Nothing is left for the others, the worse outcome for them. If his preferences are not strictly monotonic but only weakly monotone, there may be an entire set of allocations at which his welfare is maximized, and some other agent may be given the chance to maximize his own welfare within that set. However, this will be the rare case; at least, there does not seem to be enough of a departure from the situation when one agent always consumes everything to justify searching for a milder term than dictatorship: adding the adjective “sequential” seems quite enough.

By contrast, when what is assigned to whomever is listed first, and then to each of his successors in turn, is constrained by certain quantity requirements, or so as to

⁹ By this expression, I mean the problem of either ranking the alternatives in a set or selecting one of them as a function of the list of the preferences of a group of agents, this set not being endowed with any particular structure.

guarantee certain welfare levels or more generally, opportunities, to these successors, the milder expression “sequential priority rule” might be appropriate.

For instance, if the issue is how to assign “objects”, (say rooms to students in the house they are renting together, offices to professors in an academic department, cubicles to co-workers in a company), a meaningful constraint is that each agent should receive only one of these objects. Because there may in fact be little difference between them (admittedly, cardinality notions have to be brought in for this statement to be meaningful), the term “dictatorship” does not describe well this process of going from agent to agent to assign objects.

For the adjudication of conflicting claims, it is natural to require that each agent’s assignment be bounded above by his claim. In such cases, giving priority to a particular agent (honoring his claim first) does not necessarily imply leaving all the others with their worst outcome (receiving nothing). A sequential priority rule here ends up partitioning, for each problem, the agent set into three groups, a group of agents who get full satisfaction, a “group” consisting of one agent who is partially compensated, and a group of agents who indeed get nothing. Still, in this context, the term dictatorship does not seem best.

Some rules have been characterized that allow limited departure from strict priorities. They are “diluted” sequential priority rules. The extent of the dilution may be very minimal, sometimes occurring only at the top two levels; one could then use the expression “diluted-at-the-top-two-levels” sequential priority rule although the expression is somewhat awkward. Other axioms systems have led to rules for which, slightly more generally, the dilution occurs for pairs of successive agents; then one could speak of “pairwise-diluted” sequential priority rules. Any departure from the systematic preferential treatment of particular agents of the kind mandated by sequential priority rules should be seen as a good thing from the viewpoint of punctual fairness. The departure is suggested by the term “dilution”, but admittedly the desirability of the departure is not conveyed well by the term.

It was mentioned earlier that as one moves down the sequence, some conditioning may also take place: at each step, the identity of the agent whose preferences are maximized (subject to some constraint) may depend on what has happened earlier. An expression that includes a reference to this conditioning could be **conditional sequential priority rule**, and a prefix could be added to indicate the extent and the nature of this conditioning. For instance, the adjectival phrase **previous-assignments-conditional** can be added to designate a conditional sequential priority rule for which the conditioning at each step is based on the list of the maximizers of the relations of the agents who have come before.

The use of priorities is often an implication of axioms that we impose on rules; it prevents treating like agents alike, contrarily to what fairness usually requires. However, fairness sometimes requires the opposite, namely that two agents be treated differently even though they have the same characteristics *in the model*: there could be differences in seniority, marital status, record of service and so on, that demand that agents who are otherwise alike in the model be treated differently. Of course, one could argue that if this information is relevant, it should be incorporated in the model, but that is not always an option, especially when the relative merits of agents are not easily quantifiable. How more deserving is a war veteran than a civilian and

how more deserving is a war veteran who lost an arm as compared to one who lost a leg? Instead, we drop the symmetry requirements on rules and leave it up to the user of the theory to calibrate rules on his perceived need to treat certain categories of agents better than some other category. More often than not, when no such external reason exists to justify a differential treatment of agents, priority orders are chosen on the basis of the outcome of a random device or on “orthogonal” information such as date of birth, alphabetical order of the recipients, order of registration, and so on. Dates of birth are in a sense the outcome of a lottery, and in the face of the difficulty in achieving fairness, people accept that they be used to break ties. The device does not make sequential priority rules fundamentally fairer, but it makes them acceptable.

The advantage of the language proposed here is two-fold. One is its greater suggestiveness and precision. The other is that, because the expressions come from the same mold, relations between concepts and results are more visible. The disadvantage is the length of some expressions. However, in the context of the discussion of each model, there is rarely any need to use them more than a few times.

The adjective “serial” has been used in several contexts. It appears in the list above, but also in two other literatures. One is cost sharing. The other is the probabilistic allocation of indivisible good. It too suggests that allocations are selected by means of a step-by-step process. However, as already noted, in neither of the last two contexts does it imply an asymmetric treatment of agents. Since it is convenient to have a term for these occasions, that is, occasions when rules are defined by a step-by-step process that does allow for an asymmetric treatment of agents, the term could be kept (although the two contexts are too different for one to be able to say that the concepts are related in a meaningful way). Thus, a rule can be said to be **serial** if it is the result of a step-by-step process. Generalizations of the notion can be defined that do treat agents asymmetrically. The term “serial” can then be qualified with the adjectives “weighted” or “asymmetric”.

On other occasions when the sequential nature of a process to define allocations is not intended to bring about an asymmetric treatment of agents, the adjective “recursive” could be used.

- *Serial rule versus sequential object-wise equal division rule* For the problem of allocating objects when each agent consumes one (versions have also been defined for situations in which each agent consumes several) and assignments are probability distributions over objects, the rule that is known as the “probabilistic serial rule”, and to which we referred to earlier as the BM rule (Bogomolnaia and Moulin 2001), is based on an implementation of the goal of equal division. Each agent starts consuming probability shares of his most preferred object. All agents consume at the same rate. When an object reaches exhaustion, each of the agents who were consuming it turns to the object, among those whose supply is not exhausted yet, that he most prefers. The process continues until all objects are exhausted. The underlying scenario bears some similarity to that underlying a rule defined for the so-called “airport problem”: arranging airlines in the order of lengths of the runways they need, the cost of each successive addition to the runway, or “segment”, is shared equally by all airlines using it. In a review of the literature on the airport problem (Thomson 2006), I call this rule the **sequential equal division**

rule. Because the division of each segmental cost can be determined in any order, the expression “segment-wise equal-division rule” might be preferable.

Returning to the problem of assigning objects, a good way to refer to the BM rule could be to add the adjective “sequential” to the expression “object-wise equal division rule”, yielding the expression **sequential object-wise equal division rule**. Here, by contrast to several applications mentioned earlier, the adjective “sequential” is not intended to reflect an asymmetric treatment of agents, but there is no reason why it should. (By contrast, the “sequential priority rules” treat agents asymmetrically.) In fact, the rule is *anonymous*. The expression is simply a reference to the fact that one proceeds from object to object in a sequential manner. Objects are treated symmetrically too; what determines the order in which they are handled are preferences. (Of course, generalizations of the idea can easily be defined that treat neither agents nor objects symmetrically.)

As to the fact that the phrase “equal division” is used in a model in which the resources to be assigned are objects, a term that is used to designate entities that are indivisible, we have to remember that in the context of probabilistic assignment, divisibility is recovered by thinking of people being assigned chances of receiving objects: thus the supply of each object is a number between 0 and 1, each agent may be assigned a share of an object that lies between 0 and 1, and feasibility constraints are written as in a classical economy: the sum of the shares of an object assigned to various agents should not exceed its supply.

The expression “sequential object-wise equal division rule” is quite awkward however, but here is something that may help. The scenario that underlies the rule has each person going for whatever object he most prefers whose supply is not exhausted. From a strategic viewpoint, it may be smarter to consume something else. One could say that it is shortsighted for someone to always go for the object that is most attractive at each point, in other words, to seek immediate gratification. Along the algorithm defining Bogomolnaia and Moulin’s (2001) “immediate gratification rule”, and mimicking the cricket in “The Cricket and the Ant”, everyone is looking for short-term pleasure instead of planning for the future.

- *Boston mechanism versus immediate acceptance rule* I have already discussed the importance of keeping separate the notion of a mechanism from the notion of a rule; here, I focus on the “Boston” qualifier. The expression “Boston mechanism” refers to the algorithm (that was) used in the Boston school district to assign students to schools. It is widely used to handle similar situations, when entities in two sets have to be paired, each agent on one side being equipped with a preference relation over the entities on the other side and each entity on this other side being equipped with a priority order over the agents on the first side.¹⁰

¹⁰ I have had a number of conversations about the school choice problem and a very frequent response when the Boston mechanism was brought up has been “I have attended many presentations where it was discussed but I don’t remember the definition”. All of these conversations were with people who would state the definition of the deferred acceptance rule with no hesitation and I have no doubt that any of them would immediately guess what the expression immediate acceptance could refer to, or at least would easily remember it if presented with the definition.

An algorithm that, until recently, had been mostly studied in the context of two-sided matching problems is the so-called “deferred-acceptance algorithm”. When adapted to school choice, a student’s enrollment in a school at any stage is only tentative; each student has to wait for the algorithm to run its course to know whether he is accepted for good. (In fact, a student who is tentatively accepted in a particular school at the first stage may be turned down by that school at the penultimate stage.¹¹)

By contrast, acceptance in a school at any stage of the Boston algorithm is never rescinded. It occurs if a student’s rank in the school’s priority order among those students who are applying to it at that stage is higher than the number of remaining seats then: that school is the one to which he is assigned. For that reason, I suggest that it be referred to as the **immediate acceptance** algorithm. For sure, acceptance is not as immediate as one could imagine; it does not occur as soon as the student applies, in a single step. There is indeed an algorithm that is operated to produce the assignment: a student who is rejected has to keep making offers until he is accepted somewhere. Still, the pair of expressions “deferred acceptance” and “immediate acceptance” reflects well the conceptual distinction between the two rules.

A class of rules parameterized by how long a student is kept waiting until his acceptance is confirmed is introduced by Chen and Kesten (2017). A name for them is a little harder to come by.

Incidentally, the common expression “deferred acceptance” itself could be improved upon. When describing the algorithm, people say that acceptance at each step of the algorithm is “tentative”, sometimes “provisional” or “temporary”, but almost never say that it is “deferred” and indeed, it is not: it is acceptance at some school that is deferred. You certainly cannot say that, when a student is accepted at a school at some step, his *acceptance* there is deferred. Nor could you say that his *admission* is deferred. In common language, the expression “deferred admission” is used when a student is admitted to a program, but for some reason, typically a personal reason (illness for example), the student requests to enroll the following year, that is, to have his admission *in that particular school* deferred. Altogether, expressions such as **tentative acceptance** (perhaps “provisional acceptance”, or “temporary acceptance”) would better convey what is important about the algorithm and the rule.

Returning to the immediate acceptance rule, there are actually two main versions of it. In the version that is commonly discussed, a student who is rejected by a school at some step of the algorithm applies next to the school that is next in his preference list, and does so even if the school is full. What appears to be a more natural way to proceed is to allow the student to skip any such school. This is what Alcalde (1996) proposes in the context of two-sided matching, calling the resulting rule “now-or-never”: if someone applies to a school at a step and is not accepted, it is because the school is filled with students with higher priority, and since admissions are never rescinded, that applicant will never gain admission there. This is a feature of both versions of the immediate acceptance rules however. The expression **immediate acceptance rule with skips**, proposed by Harless (2015), helps distinguish the two versions. (Other terms have been used such as “modified Boston mechanism”, Dur 2015, and “adap-

¹¹ By definition, the last stage is defined to be the one at which all students are accepted.

tive Boston mechanism”, Mennle and Seuken 2014, for closely related rules.) This second version does seem more appealing, but whether it is requires of course that the properties of the two versions be analyzed. Dur (2015) and Harless (2015) compare their strategic properties, and there is indeed a sense in which the version with skips is more satisfactory, but only marginally so. As for the normative properties, the version without skips generally performs better (Harless 2015).

- *Acyclicity conditions on priority structures versus no-reversal conditions* Conditions on the data of priority-augmented object allocation problems (the primary application is to school choice) have been shown to be necessary and sufficient for the student-proposing deferred acceptance rule to satisfy certain properties. They are referred to as “acyclicity properties”. Several points should be made about this expression.

First, we usually speak of the acyclicity of a single binary relation; it is important to realize that the property pertains to more than one priority relation.

Second, and because the capacities of the schools enter the definition, the property should not be referred to as pertaining to the priority profile only but as pertaining to the pair of the priority profile and the capacity profile, the priority-capacity profile.

Third, because the definition actually involves pairs of schools, the absence of a cycle could simply be described as the absence of a reversal.

At this point, it already appears preferable to speak of a **pair of priority-capacity pairs for two schools exhibiting no reversal**, and to say that a **profile of priority-capacity pairs exhibits no reversal** if it exhibits no reversal for any pair of schools.

Although these pairwise statements are what come out of the analysis, from the viewpoint of welfare distribution, the implications for the entire profile of these pairwise statements, that is, how far one can get from a list of identical priorities, is what matters. The focus should be on these implications. The issue of existence of priorities should be settled first, but it is obviously resolved positively when the priorities are all the same. Then, the DA rule that results is the sequential priority rule associated with that common priority. That is not a very desirable outcome. It has been shown that unfortunately, one cannot move far away from these rules.

- *Voting-by-committees rules versus critical-support-family-based rules* For the problem of making a binary decision when each agent is for or against, the following rules have been important in the study of *strategy-proofness*. There is a family of sets of agents such that, if the set of agents who are in favor belongs to the family, the decision is implemented, and if there is no such set, it is not. The family should satisfy certain intuitive requirements: (i) it should be non-empty; (ii) it should not contain the empty set, and (iii) it should be closed under enlargements. Because of (iii), it suffices to specify its minimal elements for it to be known. If *anonymity* is desired, only the size of a set matters: then, there is a minimal number of agents who should be in favor of the decision to guarantee that it is implemented.

These types of rules have come up in different contexts. Examples are for the problem of selecting a subset of some object set (Barberà et al. 1991); for the problem of choosing a point along a one-dimensional continuum when agents have single-dipped preferences (Manjunath 2014). How should they be named?

Barberà et al. (1991) use the expression “schemes of voting by committees” but the phrase “voting by committees” does not provide much of a clue as to what the definition is about. In fact, it may be understood to mean that committees are set up to decide which alternatives are selected. How are these committees formed? How do they decide?

A “simple game” in cooperative game theory is associated with each family of the type described above. It assigns worth 1 to each member of the family. Thus, it bears some similarity to the concept under discussion here. The expression is not very transparent, but it is not misleading. The expression “decisive groups” is sometimes used to designate the members of a family of groups that have the ability to enforce some outcome. Reference to the concept seems relevant too and it could be part of the name designating it.

We need an expression that could also be adapted to the problem of making decisions on several issues, when for each issue, decision is made in this way, that is when, for each issue, there is a family of sets satisfying properties (i)–(iii) listed above. The term “profile” could be used for the list of these objects indexed by the issues. There could be restrictions across these objects, in which case we could speak of a “structured profile” or simply a “structure”.

The properties have a clear normative content even though they emerged of strategic analysis (but that is frequently the case).

The *family* of groups of voters whose *support* for an alternative is *critical* for the alternative to be selected could be referred to as a “critical-support family” or simply as a **critical family**, and a rule that is based on the specification of such a family as “critical-support-family–based” or **critical-family based**. When a decision has to be made about several alternatives, and for each alternative, the rule operates in this way, we need a “profile of critical-support families” or a **profile of critical families**.

The term “critical” is a little close to suggesting “necessary”, but speaking of a “sufficient” family would convey the wrong message. Using the longer expression “necessary-and-sufficient” might be a little confusing: it will be understood to mean that it is necessary and sufficient that the set of agents who support a particular issue contains one of these critical groups for the issue to pass, but because the family is closed under enlargements, any group that is a superset of a critical group should also be included.

For the generalization of this family defined by Ju (2005) to cover situations in which indifference is allowed, a family of pairs of sets of agents needs to be specified for each issue. The family can be referred to as what is called a “tournament” in voting theory, a term discussed above.

4 Properties of rules

- *Resolute versus single-valued* A social choice correspondences that selects for each economy a single outcome is sometimes called “resolute”. It certainly suggests “definiteness” and lack of wishy-washiness, but it may not immediately bring to mind “single-valuedness”. As it is unambiguous, the term *single-valued* may be preferable.

- *Individual rationality versus the individual-endowment lower bound* When a model includes the specification of individual endowments, instead of saying that a rule is *individually rational* or induces *participation*, we could say that it “meets the individual-endowments lower bound”, a proposal discussed earlier in connection to the lower bound template. When a model only includes the specification of a social endowment and resources are infinitely divisible, instead of saying that it is *individually rational from equal division* or satisfies *egalitarian rationality*, we could say that it **meets the equal-division lower bound**.

The reason for the phrase “individual rationality” is presumably the following. An agent being “endowed” with resources is usually understood to mean that he has rights on these resources, in particular that he has the right to consume them on his own. Thus, it would be “irrational” for him to accept a bundle that he finds inferior to his endowment. However, the bundle to which we attach his name and that we call his endowment need not be interpreted in this way. It may simply be a “reference” bundle that the rule is supposed to take into account in some fashion. Similarly, when the problem is fair division, equal division is certainly an interesting reference point and the requirement that a rule assign to each agent a bundle that he finds at least as desirable as equal division is meaningful. However, referring to the requirement as “individual rationality from equal division” seems to be justified only when each agent is given rights over an equal share of the social endowment. We need not think that punctual fairness requires that such rights be given. In fact, an envy-free allocation does not necessarily satisfy the requirement, as is easy to see in an Edgeworth box economy; yet, many would agree that the no-envy definition is a meaningful attempt at making operational the idea of common ownership.

Returning to the problem of reallocating a profile of endowments, the counterpart of the no-envy requirement, “no-envy for trades”, is satisfied by allocations at which some agents are worse off than at their endowments. There is nothing paradoxical about this; here too, one need not think of the bundle “attached” to an agent in an economy in its initial state as one on which he has rights, although that may be the most natural interpretation.

As mentioned above, “individual rationality” suggests that agents are given the option of approving the bundle assigned to them and that it would be irrational for someone to accept a bundle that is not as desirable for him as his private endowment. To be fully justified, this interpretation would require that an approval stage be added to the model, but no such stage is included in the models in the context of which this expression is used.

- *Strategy-proofness* To apply an allocation rule, we typically need the preferences of the agents involved. A rule is **strategy-proof** if no agent ever benefits from misrepresenting his. “Ever” stands for multiple quantifications: the true preferences of the agent who is contemplating misrepresenting them, the manner in which he chooses to do so, and the other players’ announced preferences. (Whether these are the truth for them is immaterial.) Other names have been used for this property: “straight-forwardness” (Farquharson 1956; Dasgupta et al. 1979; Border and Jordan 1983), “stability” (Gärdenfors 1977) “incentive-compatibility” (Hurwicz 1972), “cheat-

proofness” (Pazner and Wesley 1977; Feldman 1979); “non-manipulability” (I could not find a primary source for this expression).

None of these expressions is particularly transparent. Certainly, “straightforwardness” does not say much about what the property is about. “Stability” is worse because the term has been used in multiple different ways in other literatures, and because it is not explicit about the test that is performed to evaluate the stability of a rule: in what sense is a rule stable? One certainly expects non-manipulability of a social arrangement to contribute to its stability, but are there terms that are better tailored to the interpretation?

“Incentive-compatibility”, an expression that has also been given a variety of other meanings in other literatures, “cheat-proofness”, “non-manipulability”, as well as “strategy-proofness” itself, do indicate concerns with agents behaving strategically, but they are not explicit about the manner in which an agent or a group of agents can do so, namely through individual or coordinated misrepresentation of their preferences. Other possible strategic moves are the destruction by an agent of part of the resources he controls (immunity to such behavior is called *destruction-proofness*, an expression that seems appropriate), the withholding by an agent of part of the resources he controls for his own private consumption (immunity to such behavior is called *withholding-proofness*), and the transfer of resources from one agent to another (before the rule is applied, or after it is, or both; robustness to such schemes could be called *pre-application transfer-proofness* or *post-application transfer-proofness*). “Incentive-compatibility” sounds too general for our purposes: it could also be understood as the property of a rule that it gives agents the incentive to exert themselves on the job, for instance.

A more informative alternative to any of the expressions listed above is “preference-misrepresentation-proofness”, which is too much of a mouthful, or the shorter “misrepresentation-proofness”.

Would **lying-proofness** or **lie-proofness** be better? Not necessarily. For one thing, in some situations, an agent could lie about things other than his preferences. For instance, an agent could misrepresent what he knows about the outside world (for an engineer, his expertise about a technology; for a weatherman, the likelihood of rain tomorrow) if his report will be used in making a decision that he cares about. The other problem is the negative connotation of the term “lie”, which would unfairly stigmatize people engaging in the behavior when it is recognized as a legitimate way in which someone is allowed to influence social choice. For example, it does not apply well to the action taken in the voting booth by a voter who does not check the name of his most preferred candidate because he expects this candidate not to win. Such a voter is indeed simply exercising his right to influence who is elected.

- *Ordinal efficiency versus stochastic dominance efficiency* The adjective “ordinal” is used in the context of probabilistic assignment to designate a concept of efficiency based on stochastic dominance comparisons. It is meant to indicate that “intensity” of preferences is not invoked in the definition: the concept only depends on the ranking of sure alternatives, that is, on ordinal information. The more standard efficiency notions when there is uncertainty are based on the assumption that preferences satisfy the von-Neumann–Morgenstern axioms and that they can be

given numerical representations satisfying a certain expectation property. These representations are not invariant under arbitrary monotonic transformations; thus, they contain “cardinal information”.

When discussing and comparing efficiency notions in this context, it is natural to use terminology that brings out the fact that one definition uses cardinal information and the other only uses ordinal information. A number of studies have addressed the issue of characterizing efficiency based on stochastic dominance (McLennan 2002; Abdulkadiroğlu and Sönmez 2003; Manea 2008). These studies had no other objective and did not bring in other considerations or axioms. Thus, calling “ordinal efficiency” the efficiency notion based on stochastic dominance comparisons was not a bad choice, although the expression does not say much about the manner in which these “ordinal” welfare comparisons are made.

However, in most of the recent literature dealing with probabilistic assignment, the search has been for rules satisfying criteria of desirability other than, or in addition to, efficiency. Fairness requirements and requirements of robustness under strategic behavior have been invoked too. Standard notions, such as no-envy and the individual-endowments lower bound for example, and variants of these notions, as well as implementability notions, such as *strategy-proofness* and *Nash implementability*, which had been mostly studied in deterministic settings, had to be reformulated so as to accommodate comparisons of assignments by means of stochastic dominance. Thus, one encounters theorems in which several properties of rules are listed, some of which being based on stochastic dominance comparisons and others not, but among the former, only efficiency is labeled “ordinal”, the others being given the names under which they are known in deterministic settings.

What terminology would be consistent and transparent? Let us discuss some options.

First, for consistency, we should probably tag all of the terms that are based on stochastic dominance comparisons, not only efficiency. Adding the adjective “ordinal” to the expressions “no-envy” or “individual-endowments lower bound”, would certainly achieve this goal, but it would be problematic because the standard, deterministic, notions that we call “no-envy” and the “individual-endowments lower bound” are ordinal too.

Besides, as already noted, the term “ordinal” is not very informative about the manner in which ordinal information is used; it does not say that the definitions are based on stochastic dominance comparisons.

A solution here is to use the prefix “sd” (abbreviation for “stochastic dominance”) to the terms “efficiency”, “no-envy”, “strategy-proofness”, and so on. Here are the benefits: the prefix is short, easy to pronounce, informative, and mnemonic (it takes no time for students in a class or for a seminar audience to know exactly what it means). It does seem to achieve consistency and transparency.

When we read the following theorem,

Theorem (Kasajima 2013) *No rule is sd-efficient, anonymous, neutral, and weakly sd-strategy-proof.*

we immediately see which properties are based on stochastic dominance comparisons, which would not be obvious in the following restatement:

Theorem *No rule is ordinally efficient, anonymous, neutral, and weakly strategy-proof.*

Here are additional examples, which involve variants of the sd definitions. The sd relation is incomplete and given two assignments p and p' , some definitions require that an agent be able to make the comparison in a particular direction, for example that p should sd-dominate p' for him, or simply that the domination not go in the opposite direction, (that it should not be the case that p' sd-dominates p). The adjective “weak” is usually added to the name of the less demanding version of the requirement. That is why in the above theorem, we distinguish between *sd-strategy-proofness* and *weak strategy-proofness*.

As an application of these definitions, we would write that the random priority rule satisfies *ex post efficiency*, *weak sd-no-envy*, and *sd strategy-proofness* and that the object-wise sequential equal-division rule (the BM rule; see our earlier discussion of the expression “probabilistic serial rule”) satisfies *sd-efficiency*, *sd no-envy*, and *weak sd-strategy-proofness*.

Alternative extensions have recently been proposed, under the names of “downward lexicographic”, “upward lexicographic” (Cho 2018), and “social welfare” (Doğan et al. 2018), which these authors have designated by the well-chosen prefixes of “dl”, “ul”, and “sw”, which go quite well with the sd prefix, and share its desirable attributes of being short, easy to pronounce, informative, and mnemonic (although we may have to rethink this entire naming scheme if they keep multiplying).

- *Maskin monotonicity versus Maskin invariance* Consider a mapping defined on some domain of preference profiles, which associates with each profile in its domain a nonempty subset of some set of alternatives. Say that a preference relation is obtained from some other relation by a monotonic transformation at a point if its lower contour set at that point contains the lower contour set of the other relation at that point. The mapping is said to be *Maskin monotonic* (Maskin 1999) if whenever it selects some alternative a for some profile of preferences, and each preference relation in the profile is subjected to a monotonic transformation at a , then it still selects a for the new profile.

As commonly understood in mathematics, a function is monotonic if both its domain of definition and its range are equipped with order structures, and whenever two elements of the domain can be related in the order relation defined on the domain, then their images are comparable in the order relation defined on the range; moreover, the comparison always goes in the same direction. (Thus, we speak of a function defined on the real line and taking its values in the real line, when both are equipped with the usual order structure of the real line, which happens to be a complete order, to be “monotone increasing” or “monotone decreasing”.) The order structure of the domain is reflected in the order structure of the range. This is definitely not the form of the definition under discussion. Thus, the term “monotonicity” does not seem appropriate.

Although the term “monotonicity” can be applied to the transformation itself, as the transformation does bring about an enlargement of something, namely the lower contour sets at a , the property is an invariance property. Expressions such as **invariance under monotonic transformations**, or **Maskin invariance**, seem preferable to

Maskin monotonicity, especially when we study how this property relates to other properties that are themselves (correctly) referred to as invariance properties.

But can't the argument be made that the term "monotonicity" does apply because the choice set for the second profile is required to contain the alternative that is the point of departure? There is indeed an inclusion, an enlargement.

That is true but this argument is far from providing a sufficient justification to use the term, for two reasons.

First, *Maskin invariance* is a meaningful property for functions (*single-valued* mappings) as well as for correspondences. When applied to a function, it certainly would not be very natural to interpret it as prescribing an expansion: describing an invariance requirement imposed on a *single-valued* mapping as an (albeit trivial) enlargement requirement (the set chosen for the second profile is a superset of the set chosen for the second profile, but since they are both singletons, they are actually the same) does not seem very helpful.

To explain the second point, let us ask: if *Maskin monotonicity* can be described as stating an inclusion requirement, exactly which sets does it say should be related by inclusion? The larger one is certainly the set of allocations the correspondence chooses for the second profile, but what is the smaller one? Presumably, it should be the set of allocations the correspondence chooses for some other profile, the most likely candidate being the initial profile. However, that is not what *Maskin monotonicity* says. Indeed, although the transformation to which preferences are subjected causes an enlargement of the lower contour sets at the alternative a that is taken as point of departure, the alternative a may be only one among several that the correspondence chooses for the initial profile. Except on some narrow and special domains, and for specific correspondences, the transformation need not cause an enlargement of the lower contour sets at any of the other alternatives that the correspondence chooses for the initial profile: thus *Maskin monotonicity does not say* that the set chosen for the second profile should contain the set chosen for the initial profile.¹²

Finally, if "monotonicity" is dismissed on these grounds, shouldn't "invariance" be dismissed for the same reasons? Isn't "invariance" also understood as requiring that for two profiles that are related in a certain way, the *sets* of allocations chosen for the two profiles should be the same? (Typically, the second profile would be obtained by subjecting the first profile to some operation.) When we say that a correspondence satisfies some invariance property, we think of the value taken by the correspondence for some initial profile (which is a set) to be the same as the value taken by the correspondence for some other profile (which is a second set) that is related to the initial profile in a particular way.

This is certainly one usage of "invariance" but the term should not necessarily be understood in this way. In fact, there is nothing wrong with calling "an invariance property" a property stating that the social desirability of a *particular* alternative for

¹² It does happen and the possibility underlies the proofs of some characterizations. For example, it is an important step in Schummer's (1997, 1999) characterizations of the *strategy-proof* selections from the *efficiency* correspondence on domains of private good economies (or domains of public good economies) with strictly monotonic and linear preferences. (Indeed, a monotonic transformation at an efficient point is a monotonic transformation at each other efficient point.)

a profile of preferences should be preserved, or unaffected, by certain operations performed on the profile (these operations may well depend on the alternative). Obviously, invariance properties vary in the scope of the change that they cover.¹³

One final argument: couldn't we understand the expression "Maskin's monotonicity condition" to mean "Maskin's condition *about* monotonicity" (namely the monotonic transformations to which the lower contour sets are subjected in the hypotheses of the condition)? The model would be, for instance, the manner in which we refer to a step in a proof in which a function is proved to be monotonic as a "monotonicity step". We have other constructions of that type: the "Decomposition Lemma" (Alkan et al. 1991) for instance is about decompositions. In the scores of conversations that I have had about "Maskin monotonicity", many have defended the expression but no one has ever mentioned this possible justification for it. Nevertheless, for the sake of argument, let us explore it. If the condition is about something being monotonic, it should be the transformation to which lower contour sets are subjected, which has to do with the hypotheses of the property, but not the choice correspondences (or functions), whose arguments are the preferences that are modified at a particular allocation in this particular way. The adjective "monotonic" could not be applied to the choice correspondences themselves, contrarily to the way it is applied in the literature. The best way to describe this behavior is to say that the correspondences are invariant to the transformations.

The property had been discussed before Maskin brought out its central relevance to implementation. Muller and Satterthwaite (1977) consider it in the context of an abstract Arrovian social choice model with strict preferences, and call it *strong positive association*, explaining (footnote 2): "We have named this condition "strong positive association" because it is a straightforward strengthening of Arrow's condition of "positive association"." These expressions may be acceptable—at least they are not misleading—although it is hard to imagine that anyone new to the subject would be able to guess how to correctly transcribe them into mathematics: the expression "positive association" does not say much about the content of the property. But that is a common problem; there is simply too much to include in a name for it to fully describe most mathematical concepts. "Maskin invariance" is certainly not enough. "Invariance under monotonic transformations" is better. But transformations of what? Should we say "Invariance under monotonic transformations of preferences"? At what point? Should we be more explicit and write "Invariance under monotonic transformations of preferences at the allocations chosen for some initial economy"? Let us leave it at "Maskin invariance".

Besides Hurwicz's foundational papers, Maskin's paper is arguably the most central contribution to the theory of implementation. Thus, Maskin's name is deservedly linked to the property, but for someone who prefers descriptive terms, introducing it under the name of *invariance under monotonic transformations*, and subsequently

¹³ Perhaps we could distinguish between "local" invariance, and "global" invariance. For the property under discussion here, we could say that a rule is "locally invariant" if the following holds: at each point that it selects for a profile, if preferences are subjected to monotonic transformations *at this point*, then the point is still selected for the new profile. By contrast, a rule would satisfy a "global invariance" property if the entire set of allocations selected for a profile would still be chosen after preferences have been subjected to some transformation: the transformation would be defined in relation to that entire set.

referring to it under the simplified name of *invariance* (in a study that would involve no other invariance property), is a reasonable alternative.

This property actually illustrates how bad terminology may force subsequent authors to choices that make matters worse. Indeed, a property related to *Maskin invariance* is that if a new preference relation is contemplated by an agent that is obtained from his initial one by a monotonic transformation at his initial assignment, the new outcome should be ranked at least as high as the initial one according to this new relation. This property is a weakening of *Maskin invariance* and these authors were naturally led to refer to it as “weak Maskin monotonicity”. This is unfortunate because the weakening of the conclusion turns it into a true monotonicity property, whereas the term suggests that it is less of a monotonicity property than Maskin’s original formulation.

- *No justified envy* The expression “no justified envy” has appeared in the literature in two different contexts, and it has been given two different meanings. For both definitions, the starting point is the no-envy test, which involves checking whether at an allocation, some agent i would prefer some agent j ’s assignment to his own.

First, consider the problem of allocating objects when to each object is associated a priority order over its potential recipients. Suppose that at an allocation, agent i would prefer agent j ’s assignment to his own. Then, we say that **envy is not justified**, if in the priority order attached to the object assigned to agent j ’s, agent i is ranked lower than agent j (Balinski and Sönmez 1999). The envy constraints have to take into account, or respect, the priorities, and we could speak then of an allocation satisfying the **priority-respecting no-envy conditions**.

For the problem of allocating objects when objects are not indexed by priority orders over potential recipients, the starting point is also the comparison for each agent, of (i) the welfare that he would experience if he were to consume each of the other agents’ assignments and (ii) his welfare when he consumes his own assignment. However, this thought experiment is complemented with an attempt at correcting a preference going the wrong way: if agent i would prefer agent j ’s assignment to his own, the most natural thing to do would be to transpose their assignments. Let us then distinguish between two possible subtypes of situations.

Suppose first that ownership is mixed, namely some objects are initially owned by particular agents, the others being collectively owned. In such a situation, a natural way to respect private ownership, in fact, the most natural way to make the expression “private ownership” operational, is to require of an allocation that it meets the *individual-endowment lower bound*: any agent who is endowed with an object should be assigned one that he finds at least as desirable as the one he owns. Now, if at an allocation, agent i would prefer agent j ’s assignment to his own, the transposition of their assignments may cause a violation of the *individual-endowments lower bound* for agent j , and this would get in the way of ownership rights. So, the violation of envy would be acceptable: envy would not be “justified” (Yılmaz 2010). It is a definition that give primacy to ownership rights over fairness. It enforces fairness only to the extent that it does not interfere with property rights. Here the envy constraints are weakened to accommodate ownership rights, and we could speak of an allocation satisfying the **ownership-respecting no-envy conditions**.

- *Law of demand versus size monotonicity versus inclusion monotonicity* In classical demand theory, the expression “law of demand” refers to the negative impact that the price of a good has on the demand for the good: the higher the price, the less of it a consumer buys. In the context of object allocation when each agent may consume several, it has been used to mean that confronted to two choice sets that are related by inclusion, a consumer would choose from the larger set a set that contains the set it would choose from the smaller one. Thus, the standard meaning¹⁴ has to do with an economic institution; it presumes that a specific allocation rule is used, that resource allocation is mediated through prices. Also, to the extent that we think of prices as determining choice sets, a rise in the price of a good implies a particular way in which the choice set expands. However, the intended meaning in object allocation theory has to do with preferences, with the psychology of the individual consumer, before the economist steps in so to speak.

The similarity between the property we need to name and the property of demand theory seems much too tenuous to justify that the same names or expressions be used for both. An alternative proposal that had been made, **demand monotonicity**, seems to be well adapted to that property. The property that the number of elements in the set chosen from the larger opportunity set contains more elements than the set chosen from the smaller opportunity set could be called **size monotonicity** (Alkan 2002, uses the expression “cardinal monotonicity”). A related property would say that in the same circumstances, the subset chosen from the larger opportunity set contains the subset chosen from the smaller opportunity set. The expression **inclusion monotonicity** could be used to designate it.

References

- Abdulkadiroğlu A, Sönmez T (1999) House allocation with existing tenants. *J Econ Theory* 88:233–260
- Abdulkadiroğlu A, Sönmez T (2003) Ordinal efficiency and dominated sets of assignments. *J Econ Theory* 112:157–172
- Adachi H (2000) On a characterization of stable matchings. *Econ Lett* 68:43–49
- Adachi T (2010) The uniform rule with several commodities: a generalization of Sprumont’s characterization. *J Math Econ* 46:952–964
- Alcalde J (1996) Implementation of stable solutions to marriage problems. *J Econ Theory* 69:240–254
- Alkan A (2002) A class of multipartner matching markets with a strong lattice structure. *Econ Theory* 19:737–746
- Alkan A, Demange G, Gale D (1991) Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59:1023–1039
- Amorós P (2002) Single-peaked preferences with several commodities. *Soc Choice Welf* 19:57–67
- Balinski M, Sönmez T (1999) A tale of two mechanisms: student placement. *J Econ Theory* 84:73–94
- Barberà S, Sonnenschein H, Zhou L (1991) Voting by committees. *Econometrica* 59:595–609
- Barberà S, Jackson M, Neme A (1997) Strategy-proof allotment rules. *Games Econ Behav* 18:1–21
- Bénassy JP (1982) *The economics of market disequilibrium*. Academic Press, Cambridge
- Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. *J Econ Theory* 100:295–328
- Border K, Jordan J (1983) Straightforward elections, unanimity and phantom voters. *Rev Econ Stud* 50:153–170

¹⁴ In the context of classical demand theory, the expression “the law of demand” does not seem well justified either since preference relations that violate it are certainly not thought to be pathological.

- Camus A (1944) Sur une philosophie de l'expression. Poésie
- Chen Y, Kesten O (2017) Chinese college admissions and school choice reforms: a theoretical analysis. *J Polit Econ* 125:99–139
- Cho J (2018) Probabilistic assignment: an extension approach. *Soc Choice Welf* (**forthcoming**)
- Cho J, Thomson W (2011) On the existence of the uniform rule to more than one commodity: existence and maximality results. Mimeo, New York
- Dasgupta P, Hammond P, Maskin E (1979) The implementation of social choice rules. *Rev Econ Stud* 46:153–170
- Doğan B, Doğan S, Yıldız K (2018) A new ex-ante efficiency criterion and implications for the probabilistic serial mechanism. *J Econ Theory* 175:178–200
- Dur U (2015) The modified Boston mechanism. Mimeo, New York
- Ehlers L, Peters H, Storcken T (2002) Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *J Econ Theory* 105:408–434
- Farquharson R (1956) Straightforwardness in voting procedures. *Oxf Econ Pap* 8:80–89
- Feldman A (1979) Manipulation and the Pareto rule. *J Econ Theory* 21:473–482
- Gärdenfors P (1977) A concise proof of theorem on manipulation of social choice functions. *Pub Choice* 32:137–142
- Harless P (2015) A school choice compromise: between immediate and deferred acceptance. Mimeo, New York
- Hurwicz L (1972) On informationally decentralized systems. In: McGuire CB, Radner R (eds) Chapter 14 in decision and organisation. University of Minnesota Press, Minneapolis, pp 297–336
- Ju B-G (2005) A characterization of plurality-like rules based on non-manipulability, restricted efficiency, and anonymity. *Int J Game Theory* 33:335–354
- Kasajima Y (2013) Probabilistic assignment of indivisible goods with single-peaked preferences. *Soc Choice Welf* 41:203–215
- Klaus B (2008) The coordinate-wise core for multiple-type housing markets is second-best incentive compatible. *J Math Econ* 44:919–924
- Manea M (2008) Random serial dictatorship and ordinally efficient contracts. *Int J Game Theory* 36:489–496
- Manjunath V (2014) Efficient and strategy-proof social choice when preferences are single-dipped. *Int J Game Theory* 43:579–597
- Maskin E (1999) Nash equilibrium and welfare optimality. *Rev Econ Stud* 66:83–114 (**first circulated in 1977**)
- McLennan A (2002) Ordinal efficiency and the polyhedral separating hyperplane theorem. *J Econ Theory* 105:29–54
- Mennle T, Seuken S (2014) The naive versus the adaptive Boston mechanism. Mimeo, New York
- Miyagawa E (2004) Strategy-proofness for the reallocation of multiple types of indivisible goods. Mimeo, New York
- Monte D, Tumennasan N (2013) Matching with quorums. *Econ Lett* 120:14–17
- Morimoto S, Serizawa S, Ching S (2013) A characterization of the uniform rule with several commodities and agents. *Soc Choice Welf* 40:871–911
- Moulin H (1980) On strategy-proofness and single peakedness. *Pub Choice* 35:437–455
- Moulin H (1984) Generalized Condorcet-winners for single-peaked and single-plateau preferences. *Soc Choice Welf* 1:127–147
- Muller E, Satterthwaite M (1977) The equivalence of strong positive association and strategy-proofness. *J Econ Theory* 14:412–418
- Murakami S (1968) Logic and social choice. Routledge & Kegan Paul, London
- Nash JF (1950) The bargaining problem. *Econometrica* 28:155–162
- Pazner E, Wesley E (1977) Stability of social choices in infinitely large societies. *J Econ Theory* 14:252–262
- Pinker S (2014) The sense of style: the thinking person's guide to writing in the 21st century. Penguin, New York
- Schummer J (1997) Strategy-proofness versus efficiency on restricted domains of exchange economies. *Soc Choice Welf* 14:47–56
- Schummer J (1999) Strategy-proofness versus efficiency for small domains of preferences over public goods. *Econ Theory* 13:709–722
- Shapley L, Scarf H (1974) On cores and indivisibility. *J Math Econ* 1:23–28
- Shapley LS, Shubik M (1972) The assignment game I: the core. *Int J Game Theory* 1:111–130

- Su F (1999) Rental harmony: Sperner's lemma in fair division. *Am Math Mon* 106:930–942
- Tadenuma K, Thomson W (1993) The fair allocation of an indivisible good when monetary compensations are possible. *Math Soc Sci* 25:117–132
- Thomson W (1992) Consistency in exchange economies. Mimeo, New York
- Thomson W (2006) Airport problems and cost allocation. Mimeo, New York (**revised May 2017**)
- Thomson W (2012) On the axiomatics of resource allocation: interpreting the consistency principle. *Econ Philos* 28:385–421
- Thomson W, Myerson RB (1980) Monotonicity and independence axioms. *Int J Game Theory* 9:37–49
- Velez R (2016) Fairness and externalities. *Theor Econ* 11:381–410
- Yılmaz Ö (2010) The probabilistic serial mechanism with private endowments. *Games Econ Behav* 69:475–491