Rochester Center for Economic Research

On the Axiomatics of Resource Allocation: Classifying Axioms and Mapping Out Promising Directions

William Thomson

Working Paper No. 596

October 2017; revised May 2019

UNIVERSITY OF ROCHESTER

On the axiomatics of resource allocation: classifying axioms and mapping out promising directions

William Thomson*

October 2017; revised version May 2019.

*wth2@mail.rochester.edu. I thank Shigehiro Serizawa for useful comments.

Abstract

The purpose of this note is to propose a two-way classification of the axioms of the theory of economic design, and to map out directions for future research that we perceive as particularly promising.

Axiomatic analysis holds a central place in economic design. It allows us to go beyond current practice and hypothetical examples. It provides explicit arguments for the use of particular allocation rules in terms of meaningful criteria of good behavior and it gives complete answers to design questions: Are those properties compatible? If yes, what rules satisfy them all?

The scope of the contexts and the richness of the requirements that have been explored keep expanding. What would the literatures on matching, queuing, school choice, organ allocation, auctions, to name a few examples, look like without the conceptual and technical apparatus that the axiomatic approach has provided?

This purpose of this note is to suggest how its impact could be increased further. We first argue for the need to organize the field of axioms and for that purpose offer a two-way classification, together with taxonomic suggestions. We then indicate several directions that we see as deserving particular attention in future axiomatic work.

1 Organizing the field of axioms.

What's an axiom? An **axiom** is the expression in mathematical form of the intuition we have about how a solution mapping could be required to behave in certain situations.[1]

The potential benefits of organizing the field of axioms, any field for that matter, should be obvious. Broad categories of axioms are sometimes noted in the literature and piecemeal and model-specific observations made; also, identifying logical relations between axioms is a routine part of axiomatic work. However, a comprehensive attempt at classification is lacking. Such a treatment should be useful, an inspiring precedent being the manner in which the invention of the periodic table of the elements set the research agenda of physicists for many years. We therefore begin with a nutshell summary of a two-way classification of axioms, in terms of scope and format on the one hand, and in terms of content on the other, developed in Thomson (2018b).

To illustrate our definitions, we refer to three classes of allocation problems. In a **classical economy** a group of agents equipped with preferences defined over some commodity space have endowments of these goods, and

[1]We are not addressing here the role of axioms in other areas of economics, such as decision theory, nor in other fields, such as mathematics. For a discussion of the axiomatic method in game theory and resource allocation, see Thomson (2001).

the issue is to redistribute these endowments among them. In a **problem of fair division**, a social endowment of resources has to be distributed among agents having equal rights on it. In a **claims problems**, a group of agents have one-dimensional claims over a single resource and there is not enough of the resource to fully honor all claims.[2]

A "solution mapping", a "solution" for short, associates with each allocation problem in some class one or several allocations, interpreted as recommendations for it. We focus on single-valued solutions, which we call "rules". An agent's "assignment" is what the rule chooses for him. Depending upon the model, it could be a commodity bundle, a welfare level, a contribution expressed in resources, or effort, or money ... to some collective enterprise.

1.1 Axiom scope and format.

The distinctions made below should be self-evident, and our main contribution may simply be to provide language that allows us to talk about them.[3]

• **Universal versus model-specific axioms.** Few axioms express ideas that are meaningful in all models. Conversely, when an axiom is tailored to a specific model, one can usually see in it some general principle at work. Thus, we are certainly not suggesting a partition here. Rather, one should think of axioms arranged in a spectrum from **universal** to **model-specific**.

• **Punctual versus relational axioms.** A **punctual axiom** applies to each problem separately, "point by point"; a **relational axiom** relates the choices made by a rule for problems that are related in a certain way.[4]

• **Pre-application versus post-application axioms.** In describing the relation between the problems that a relational axiom is concerned with, we may have to apply the rule to at least one of them (usually one); the rule appears on the hypothesis side. Such an axiom is a **post-application** axiom; the others are **pre-application** axioms.[5]

[2]For surveys, see Thomson (2003, 2019b).

[3]Yet, not keeping them in mind may be damaging. For example, the failure to distinguish between punctual and relational axioms has been responsible for an erroneous interpretation of the consistency principle and of some results involving consistency axioms. For a discussion of this point, see Thomson (2012).

[4]*Efficiency* is a punctual axiom. *Resource monotonicity*, which says that if the resources available to society become more abundant, every agent should end up at least as well off as he was made initially, is a relational axiom.

[5]*Resource monotonicity* (footnote 3) is a pre-application axiom. *Consistency*, which says that if some agents leave with their assignments, and the situation is reevaluated at

• **Individual-centric versus group-centric axioms.** An **individual-centric** axiom specifies how each individual's assignment should depend on a particular parameter of a problem, usually a parameter attached to that individual. A **group-centric** axiom specifies how the assignment of each member of a group, or of a group as a whole, should depend on a particular parameter, in most cases, a collective parameter.

• **Self-regarding versus other-regarding axioms.** When a model includes parameters attached to individual agents, our first inclination is usually to specify how the rule should take into account each particular agent's parameters when calculating *this agent's assignment*; we then have a **self-regarding** axiom. But we may also want to say something about how *the other agents' assignments* should depend on this parameter. This gives us an **other-regarding** axiom.[6]

• **Monotonicity versus invariance axioms.** When the space to which some parameter of a problem belongs and the space to which the outcome chosen by a rule belong are equipped with order structures, a **monotonicity** axiom requires the order structure of the former to be reflected in the order structure of the latter. An **invariance** axiom requires that certain changes in a parameter—these changes may or may not be evaluated in some order—not be accompanied by any changes in the chosen outcome.

• **Fixed-population versus variable-population axioms.** Even when a rule is defined over a domain that contains problems involving different populations, a **fixed-population** axiom specifies how the rule should behave on a subdomain of problems in which the population of agents is unchanged. A **variable-population** axiom specifies how a rule should respond to changes in populations.

1.2 Axiom content

The categories listed below are not meant to achieve a partition of the field of axioms.

---

this point, each of the remaining agents should be assigned the same thing as initially, is a post-application axiom.

[6]The *individual endowment lower bound*, which says that each agent should find his assignment at least as desirable as his private endowment, is a self-regarding axiom. *Other-regarding endowment monotonicity*, which says that when an agent's private endowment increases, each of the others should end up at least as well off as he was made initially, is an other-regarding axiom.

• **Efficiency and related properties.** These punctual properties are at the core of modern economic thinking.[7]

• **Symmetry, anonymity, order preservation.** The first two axioms are commonly invoked as embodying minimal objectives of punctual fairness. When the parameters attached to agents belong to a space that has an economically meaningful order structure, we may require that a rule "reflect" or "respect" that order. The idea can sometimes be applied to groups of agents when the parameters attached to their members can be aggregated, or when parameters are attached to groups themselves.

• **Lower bounds and upper bounds.** Bounds can be imposed on an individual's assignment (footnote 5), or on the assignments of groups, aggregated in some way, or on the group's assignment directly. Most commonly, it is the resources that each agent owns that are used to define a lower bound on his welfare. In fact, imposing such a bound can be interpreted as a way of making operational the abstract concept of ownership. Meaningful bounds can also be based on comparisons of the economy under consideration to counterfactual economies.

Each of the general principles listed next is the "root" of many specific relational axioms.

• **Solidarity.** The principle of solidarity says that if a parameter of an economy changes and no one (or no one in a particular group), is responsible for the change, the welfare of all agents (or of all agents in that group), should be affected in the same direction. A number of monotonicity requirements are obtained as special cases, when the space to which the parameter belongs and the space to which the outcomes belong have order structures. Considerations of solidarity also often underly other-regarding monotonicity requirements.[8]

• **Robustness with respect to changes in perspectives when evaluating a situation or a change in a situation.** There are frequently more than one way of looking at a situation, all equally valid. If they lead to outcomes that are not welfare-equivalent, adopting one of them and not the other(s) is bound to be seen as unfair by some agents. Thus, the interest of robustness requirements stating that the possible perspectives that one can take all result in the same, or in welfare-equivalent, outcomes.

• **Robustness to strategic behavior.** Strategic behavior on the part of agents may take multiple forms and to each type corresponds a robust-

[7] *Respect of unanimity* is another example.

[8] *Other-regarding endowment monotonicity* (footnote 5) is an example.

ness requirement. Misrepresentations of preferences or of some other private knowledge that agents may have, and manipulations of resources under their control (withdrawing, destroying, transferring, exaggerating . . .), are the underlying motivation for a number of axioms of robustness of rules, namely that individual agents or groups of agents never benefit from the behavior. A next step is the identification of properties of a rule guaranteed the existence of a "game form" so that for each economy in its domain, the set of equilibrium allocations of the resulting game (the game form augmented with the preference profile) coincides with the set of allocations the rule would select for that economy. This is the object of the vast "implementation" literature.

It would be meaningful to place in one group the first four categories of axioms just listed as they express normative or ethical concerns, and in a second group the fifth category, which has to do with strategic ones. Some axioms would belong to both groups. For example, *endowment monotonicity*[9] is clearly desirable from a normative viewpoint but it also has a strategic interpretation: if violated, an agent could gain by strategically destroying some of the resources he controls. The same comment applies to the *individual endowment lower bound*. It means "respecting" ownership but it can also be understood as providing agents the incentive to willingly participate in exchange.

• **Technical and operational axioms.** This last category stands somewhat apart from the others. Some axioms do not have a straightforward economic interpretation but are imposed for technical reasons, to allow the use of certain mathematical tools for example, or to impose some discipline on some complex class of rules. The term "operational" is also often used to designate an axiom that allows the real-world implementation of some socially meaningful objective that would be impractical otherwise. Such an axiom may concern the "engineering" or computational aspect of design, but is typically of limited normative or strategic interest. It should be noted however that ease of computation is a double-edged property: it facilitates the work of the designer but it also enables strategic agents, thereby hampering achieving the chosen social objectives.

[9]This says that when an agent's private endowment increases, he should end up at least as well as was made initially.

2 Directions in which deeper axiomatic analysis is needed or seems particularly promising.

Using the taxonomy of the previous sections, this section suggests directions that the axiomatic program would benefit from emphasizing in the future.

2.1 Being more systematic in formulating axioms.

When reading an axiomatic study, we often have the feeling that it would have benefited from their authors being more systematic in their formulation of axioms, that gaps in our knowledge would be filled faster and more thoroughly. For example, other-regarding versions of axioms are rarely explored, and group-centric versions of axioms are usually not paid much attention even though doing so could often yield some interesting insights.[10]

Let us take claims problems as an example. Because it could draw on a large body of existing axiomatic work, this literature quickly reached maturity, yet it is surprising that so many rather straightforward questions having to do with the dependence of rules on the claims vector have remained unexplored. Supposing that someone's claim increases, the first requirement that comes to mind is certainly that he not be awarded less than he was awarded initially and this self-regarding requirement has been discussed in a number of studies. However, there are multiple other ways, most of them either other-regarding or group-centric, in which a rule could be required to depend on claims: the increase in the award to an agent whose claim increases should be bounded by the amount by which his claim increases; the impact on each of the others' awards should be bounded by this amount; no one else's award should increase; the impact on that agent's award should be a monotonic function of the amount to divide; so should the impact on each of the other agents' awards; and similar bounds and monotonicity properties could be imposed when the focus is on groups of claimants (Thomson, 2019a).[11]

[10]Why did it take so long for a group version of *non-bossiness* (the requirement that, if an agent's assignment does not change as a response to a change in his preferences, no one else's assignment change either), to be formulated?

[11]These axioms are related in multiple ways. This should be expected. The richer one's formulation of axioms for a model, the more likely they will be logically related.

2.2 Interpreting and reinterpreting axioms.

Some axioms are sometimes unfairly perceived as being technical; conversely, economic interpretations are sometimes forced on axioms whose technical role should instead be openly acknowledged.

Technical axioms are clearly not what the focus of our investigations should be. However, upon examination, a technical axiom can often be seen to have compelling normative or strategic interpretations, sometimes both. Continuity requirements illustrate the point; they are usually thought of as technical, yet there is always some arbitrariness in the exact data of a problem and preventing small variations in these data from having a radical impact on outcomes seems normatively justified. Also, errors in measurement are unavoidable; on the other hand, errors can be corrected. It would not be right for welfare distributions to be overly sensitive to arguably irrelevant details of the situation at hand. Also, a discontinuous rule is more likely to be manipulable, so *continuity* is a step towards robustness to manipulation. (In fact, it is often a logical implication of requirements of immunity to manipulation.)

Another example is the ubiquitous requirement of *additivity* (coalitional form games, cost allocation, airport problems, claims problems). *Additivity* too is often described as a technical axiom, yet it belongs in the category of robustness axioms described in Subsection 1.2: when two equally valid perspectives can be taken in evaluating a situation—taking cost allocation as an example, whether or not two projects should be looked at separately, or consolidated as one big project—it seems natural to require that these two perspectives result in the same outcome.[12]

Consistency is sometimes called an operational axiom, but what is it supposed to make operational? Like all invariance properties, it too belongs to the category of axioms expressing the robustness of a rule to choices of perspectives in evaluating a situation, in this case a variation in population, specifically when some agents leave the scene with their assignments: should we stick with the initial choice for the remaining agents or should we look at

[12]Consider a society simultaneously facing two allocation problems; these problems are of the same type and they can be added, for example two cost allocation problems implicating the same agent set, such as a university contracting with the same list of suppliers for two construction projects. *Additivity* says that they can be handled separately, each agent receiving two bills, one for each of them, or consolidated into one and handled as one problem.

the situation anew, from their perspective, and reassess their opportunities? Again, these are two equally deserving viewpoints. The axiom says that it should not make any difference which is adopted. (It also has a fairness interpretation discussed below.)

This does not mean that there is no place in the theory of economic design for technical axioms. We do want all of our axioms to have a clear economic interpretation, but when we have trouble understanding the implications of some axiom system, invoking additional technical axioms may be a useful, hopefully temporary, step towards the answers we seek. However it is counterproductive to artificially claim normative or strategic interpretations for these axioms. Consider *non-bossiness* (footnote 9). This axiom does have multiple conceptual and logical connections to various normative and strategic axioms but these connections are really quite tenuous. In spite of its name, helping prevent what we call "bossiness" in common language is not a very convincing interpretation of the axiom. Its main rationale really is that it structures a class of rules in a way that makes it manageable (Thomson, 2016).

2.3 Parameterizing axioms.

An axiom is a yes-or-no proposition but we will argue here that there is much to gain from introducing measures of the extent to which axioms are satisfied or violated, and working with these measures.

We disqualify a rule as soon as we find one violation of an axiom that we have imposed. Of course, there are almost always other situations where violations occur and it is because our judgment is that the violations are "frequent enough" that we disqualify the rule. We rarely seek to identify all of those situations however, and most often that would be too challenging a task anyway. Usually, the most that is practical is identifying domains that are free of them, or sufficient conditions on domains for that to be the case. We propose here that parameterizing axioms would provide sensitive tools to evaluate rules and that more effort should be expanded in this direction.

In the study of *strategy-proofness* for discrete models a range of interesting measures have been proposed many of them based on counting profiles for which things go wrong, and some work has actually been done using such measures to compare rules. For example, one can calculate the proportions of profiles of preferences at which some agent can gain by misrepresenting his preferences. One may take into account the gain that he can achieve by optimally manipulating, or the average gain that a manipulating agent can achieve by manipulation. Orders can be defined on the space of rules based on inclusion of profiles for which manipulation is possible for someone. Finally, one can compare rules in terms of the equilibrium sets of the manipulation games associated with them. This involves characterizing these sets and relating them to the set of outcomes that would be selected under truthful behavior.

When allocation space is a subset of a Euclidean space, fine quantitative measures of the extent to which an axiom is violated can usually be defined although here too, one should not expect that there would be a unique way to parameterize violations. Consider the *no-envy* axiom[13]. Calculating the minimal radial expansion to which the bundle assigned to an agent who is envious (say agent 1) of some other agent (say agent 2) should be subjected for this instance of envy to disappear is an option. Another is to use the minimal amount of some reference good that should be added to agent 1's assignment for his envy of agent 2 to disappear. Alternatively we could subject agent 2's bundle to minimal radial contractions or minimal subtractions from his consumption of a reference good for agent 1 to stop envying him (Chaudhuri, 1985; Diamantaras and Thomson, 1990). For a relational axiom such as *resource monotonicity*, we can measure a violation by the minimal expansion to which the bundle assigned to an agent who is hurt by the augmentation of the social endowment should be subjected for him to return to his initial welfare level (Moulin and Thomson, 1988).

Once axioms are parameterized, numerical tradeoffs between them can be studied. For example, for fair division problems, a relationship can be identified between a parameter measuring the extent to which *resource monotonicity* and a parameter measuring the extent to which the *no-domination* requirement[14] can be jointly satisfied (Thomson,1987b).

For classical economies, a seminal result is that for two agents, no rule satisfies *efficiency*, the *individual-endowment lower bound*, and *strategy-proofness* (Hurwicz, 1973). A parameterized version of each of these axioms has been proposed, the parameter ranging from 0 (the axiom is vacuous then) to 1 (when the standard formulation is obtained). We now know that no matter how much each of these axioms is weakened and unless it is given up alto-

[13]This says that no agent should prefer someone else's assignment to his own

[14]This says that the assignment of no agent should dominate commodity by commodity the assignment of anybody else.

gether, an incompatibility with the other two still holds (Schummer, 2004; Cho, 2014). Thus, parameterizations of axioms have allowed to reach a much deeper understanding of the nature of the difficulty.

Conversely, finding ways of quantifying how well an axiom is satisfied may be worthwhile. Using the *no-envy* axiom to illustrate the point, the operations on bundles suggested above can also be used to quantify how far some agent is from envying someone else. Selections from the no-envy solution have been proposed based on measures of this type (Tadenuma and Thomson, 1995). Also, having a measure of the extent to which a bound is met allows a search for rules that best satisfy the bound.

2.4 Beyond one-dimensional parameterizations.

A "price of anarchy" measure of the type formulated by computer scientists is based on worst-case scenarios. Such a measure summarizes the behavior of a rule by means of a single number. We argue here that this may be too blunt of a tool and that more sensitive evaluations can sometimes be obtained by partitioning the space of problems, measuring on each component separately how well-behaved the rule is, and working with the resulting lists of answers.

To illustrate, consider the notion of a *guarantee structure* that was proposed to evaluate how bargaining solutions respond to population increases as a function of two parameters, the size of the initial population and the size of the additional population (Chun and Thomson, 1989; Lensberg and Thomson, 1983; Thomson, 1983, 1987a). When new agents arrive without the opportunities open to society expanding, the agents initially present will typically be affected negatively as a group, but how badly they will be affected individually should be a concern. If a rule is *population monotonic*, each of the incumbents will have to make some sacrifice. How much of a sacrifice? How disruptive will this arrival be? When incumbents are allowed to express an opinion as to whether potential newcomers should be allowed in, they will want to know what it will cost them.

So for a generic agent, let us identify the minimal ratio of the final to initial numerical welfares that a rule assigns to him when such events occur. Seen positively, the resulting number can be understood as a *guarantee* to the agent that when population expands, his final welfare will reach at least a certain proportion of his initial welfare. A *guarantee structure* gives the list of guarantees to a generic agent involved in a generic problem as a function of the numbers of incumbents and newcomers. Higher guarantees are of course more desirable. These notions have recently been successfully adapted to claims problems (Harless, 2017a,b). A disadvantage they have is that they only provide a partial ranking of rules. Yet, in both applications, bargaining problems and claims problems, broad rankings of rules have been obtained based on guarantee structures, and maximal elements within large classes have even been identified, and characterized.

References

Chaudhuri, A., “Some implications of an intensity measure of envy,” *Social Choice and Welfare* 3 (1986), 255-270.

Chun, Y. and W. Thomson, “Bargaining solutions and relative guarantees”, *Mathematical Social Sciences* 17 (1989), 295-295.

Cho, W., “Impossibility results for parametrized notions of efficiency and strategy-proofness in exchange economies”, *Games and Economic Behavior* 86 (2014), 26-39.

Diamantaras, D. and W. Thomson, “An extension and refinement of the no-envy concept”, *Economics Letters* 33 (1990), 217-222.

Harless, P., “Wary of the worst: maximizing award guarantees when new claimants arrive”, *Games and Economic Behavior* 105 (2017a), 316-328.

——, “Guarantee structures for claims problems; properties and comparisons among rules”, mimeo, 2017b.

Hurwicz, L., “On informationally decentralized systems”, Chapter 14 in *Decision and Organisation* (C.B. McGuire and R. Radner, eds) University of Minnesota Press, Minneapolis, 1972, 297-336.

Lensberg, T. and W. Thomson, “Guarantee structures for problems of fair division,” *Mathematical Social Sciences* 3 (1983), 205-218.

Moulin, H. and W. Thomson, “Can everyone benefit from growth? Two difficulties”, *Journal of Mathematical Economics* 17 (1988), 339-345.

Schummer, J., “Almost-dominant strategy implementation: exchange economies”, *Games and Economic Behavior* 30 (2004), 154-170.

Tadenuma. K. and W. Thomson, “Refinements of the no-envy solution in economies with indivisible goods,” *Theory and Decision* 39 (1995), 189-206.

Thomson, W., “Collective guarantee structures,” *Economics Letters* 11 (1983), 63-68.

——, “Individual and collective opportunities,” *International Journal of Game Theory* 16 (1987a), 245-252.

——, “Can everyone benefit from growth? Another impossibility,” University of Rochester mimeo, 1987b.

——, “On the axiomatic method and its recent applications to game theory and resource allocation”, *Social Choice and Welfare* 18 (2001), 327-387.

——, “Axiomatic and game-theoretic analysis of bankruptcy and taxation problems: a survey”, *Mathematical Social Sciences* 45 (2003), 249-297.

——, “On the axiomatics of resource allocation: Interpreting the consistency principle”, *Economics and Philosophy* 28 (2012), 385-421.

——, “Non-bossiness”, *Social Choice and Welfare* 47 (2016), 665-696.

——, *How to Divide When There Isn’t Enough: from Aristotle, the Talmud, and Maimonides to the Axiomatics of Resource Allocation*, Monograph of the

Econometric Society, Cambridge University Press, forthcoming, 2019a.
——, *Let's Talk about Economic Design*, 2019b.