Tractable Choice

MODIBO K. CAMARA
University of Chicago

Models in economics and game theory often assume that people behave as if they can solve very complex problems, which can lead to misleading conclusions. To address this, I propose that we supplement the theory of rational choice with a theory of *tractable choice*. Tractable choice asks what an individual can accomplish using resources like time, memory, or data, which are often in short supply. The field of economics has been disciplined when it comes to insisting that choices in models be rational, but is less diligent in requiring that choices be tractable under reasonable assumptions about what resources are available. Fortunately, theoretical computer science has developed deep insights and powerful frameworks for understanding tractability. Using a recent paper as a case study, I argue that tractability is a first-order concern when studying behavior.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Economics, Human Factors, Theory

Additional Key Words and Phrases: Bounded Rationality, Computational Complexity

1. INTRODUCTION

Models in economics and game theory often assume that people behave as if they can solve very complex problems. This is concerning if it leads to incorrect predictions about people's behavior. It is also concerning when it comes to designing markets or policies, because the markets and policies that are optimal in our models may be too complicated for real-world actors to interact with. There is a need for a theory that can distinguish predictions and recommendations that are unrealistically complex from those that are at least plausible.

I propose that we supplement the theory of rational choice with a theory of *tractable choice*. To define tractable choice, it is helpful to think of theory as taking a stance on what kinds of behavioral predictions are credible. According to this view, rational choice says that a prediction that "individual i follows strategy s" is justified only if we can argue that individual i prefers s to any other strategy s'. Tractable choice says that such a prediction is justified only if we can argue that individual i is able to execute strategy s using the resources at her disposal.

Tractable choice asks what an individual can accomplish using resources like time, memory, communication channels, or data, which are often in short supply. Whereas rational choice relies on models of and assumptions about preferences, tractable choice relies on models of these resources and assumptions about their availability. Here, a choice is complex if making said choice requires a large amount of resources. But complexity is multi-faceted. For example, a choice may be complex insofar as the individual must deliberate for a long time, but simple insofar as the individual only needs to communicate a "yes" or "no" answer.

Authors' addresses: `mcamara@uchicago.edu`

Economics and closely-related fields have been disciplined when it comes to insisting that choices be rational, but are less diligent in requiring that choices be tractable under reasonable assumptions about what resources are available.[1] One reason for this is the lack of consensus on how to model boundedly-rational choice, which asks what people *will* do when faced with intractable problems (as opposed to tractable choice, which simply asks what people *can* do). However, it is not necessary to understand how individuals will respond to intractable problems in order to design markets and policies where optimization is tractable.[2] Immorlica et al. [2020] is a great example of this approach.

Fortunately, as many readers know well, theoretical computer science has developed deep insights and powerful frameworks for understanding tractability. These frameworks are often highly compatible with economic models, as they tend to be based on a similar foundation of optimization, probability, and logic. They involve general-purpose abstractions that seem capable of representing a wide range of phenomena, not only electronic computers or algorithms implemented in standard programming languages.

It is almost tautological to say that real choices must be tractable, but whether tractability (as understood by computer scientists) should be a first-order concern for the study of human behavior is not quite as obvious. There are three questions that we must ask ourselves:

(1) Are computational models compatible with and helpful for understanding human behavior?

(2) Do predictions that respect rationality and tractability look meaningfully different from predictions that only respect rationality?

(3) Is it really necessary for us to study rational and tractable choice at the same time, rather than having one community (e.g., economists) focus on rationality while another community (e.g., computer scientists) focuses on tractability?

Using my recent paper on "Computationally Tractable Choice" [Camara 2022a] as a case study, I argue that we should expect an affirmative answer to all three questions. My aim is to convince the reader that it is worth taking tractability as seriously as we take rationality, or risk reaching the wrong conclusions. For readers that are already convinced, I hope this case study will help them convince others. In future writing, I hope to address the natural follow-up question of how we can integrate tractability into economic models in a more systematic way.[3]

[1]To a lesser extent, this is also true in algorithmic game theory. For example, there are many models that study auctions or market design from an algorithmic perspective, insisting that allocations can be computed in polynomial time or that the designer's distributional knowledge come from sample data. But, when it comes to the market participants, many of these models still maintain assumptions like Bayes-Nash equilibrium that are hard to justify as tractable.

[2]Similarly, it is not necessary to understand precisely how individuals will respond to intractable problems in order to design markets and policies where *approximate* optimization is tractable. We can evaluate such markets and policies according to worst-case participant strategies, subject to the constraint that those strategies be approximately optimal.

[3]A recent line of work in data-driven mechanism design [Immorlica et al. 2020; Cummings et al. 2020; Camara 2022b; Camara et al. 2020] offers some guidance for modeling tractable choice when data is the limited resource. In addition, a recent line of work by Ryan Oprea develops an

2. COMPUTATIONALLY TRACTABLE CHOICE

Are computational models compatible with and helpful for understanding human behavior? I argue that the answer can be yes, using recent work that integrates computational constraints into decision theory [Camara 2022a]. Still, one must be thoughtful when applying computational models in economics. Developing frameworks that translate results in theoretical computer science to statements relevant for economists seems to require a nuanced understanding of both fields. Echenique et al. [2011] illustrate this point very well.[4]

The premise of Camara [2022a] is that (i) decision-makers have only a limited amount of time to make decisions, but (ii) making good decisions can be time-intensive. To explore the implications for choice, I propose an *axiom of computational tractability*. This axiom is weak: it only rules out behaviors that are thought to be implausible for any algorithm to exhibit in a reasonable amount of time.

I consider a model of choice under risk where the decision-maker has to make many different decisions. For example, consider a consumer choosing from the hundreds or thousands of products in a grocery store, or an investor purchasing shares among the thousands of firms listed on the New York Stock Exchange. The decisionmaker cares about high-dimensional random vectors, i.e.,

$$X = (X_1, \ldots, X_n)$$

For example, an investor cares about income X_i from assets $i = 1, \ldots, n$, while a consumer cares about consumption bundles, where X_i represents the quantity consumed of good i.

A choice correspondence c maps a menu of feasible options to the decisionmaker's choices X from that menu. The correspondence c is defined over a rich set of menus. This includes all *binary menus* where the decision-maker chooses between two lotteries X and X', as well as *product menus* where the decison-maker separately chooses each component X_i of the lottery.

I call the choices c *rational* if they maximize expected utility for some utility function u, a common assumption that is axiomatized by von Neumann and Morgenstern [1944]. Later on, I will return to this definition and evaluate its normative appeal in the presence of computational constraints.

I assume that the decisionmaker's choices can be generated by a Turing machine, a powerful model of computation used in computational complexity theory to study what algorithms can and cannot do. Given an appropriate description of a menu, the Turing machine outputs a choice from that menu within a certain amount of time. A choice correspondence is *tractable* if it can be generated by a Turing machine, within an amount of time that grows at most polynomially in the length

experimental paradigm that controls resource complexity while varying incentives, and vice-versa (see e.g., Oprea [2020]). This has already led to some remarkable results, many of which are not yet public, and suggests a path forwards for the empirical study of tractable choice.

[4]Echenique et al. [2011] also integrate computational constraints into decision theory. Their "revealed preference approach to computational complexity" shows that, in a model of consumer choice, any finite and rationalizable dataset can be rationalized by tractable preferences. This surprising result contrasts with the more naive conclusion that consumer choice is intractable because it resembles an NP-hard problem.

of the description.[5]

Having described the model, I can now address two common objections: that humans are not Turing machines, and that computational complexity theory has a misguided focus on worst-case runtime.

The first objection – that humans are not Turing machines – is not a problem in itself. Strictly speaking, it is not necessary for choices to be generated by a Turing machine for the results of Camara [2022a] to hold. All that is necessary is that people are unable to efficiently solve problems that are thought to be fundamentally hard. By contrast, suppose some person *can* make choices that maximize expected utility for a given utility function u, and that those choices are intractable. Then, using the algorithmic reductions developed in the paper, we could leverage that person's choices to efficiently solve problems that are thought to be fundamentally hard. That would be a surprising (and important) result.

The second objection – about worst-case analysis – is best understood as an issue with the definition of rationality, not with the definition of tractability. For context, it is common in computer science to evaluate algorithms by their runtime in the worst-case instance. Consider an algorithm A that takes one minute to solve 99% of inputs and one year for 1% of inputs, so that the worst-case runtime is one year. A decisionmaker that does not have a year to deliberate might use another algorithm A': see whether A returns an answer within a minute, otherwise choose something suboptimal. This is optimal 99% of the time, suboptimal 1% of the time, and always takes about a minute.

Readers who object to worst-case analysis may point out that the algorithm A' is a perfectly reasonable solution. That may be true. But A' is not rational, insofar as standard definitions of rationality require choice to be optimal 100% of the time (e.g., von Neumann and Morgenstern [1944]). In contrast, A' is tractable because it makes a choice within the time constraint 100% of the time. Moving away from worst-case analysis requires a more flexible definition of rationality, rather than a different definition of tractability.

I use this framework of computationally tractable choice to obtain two kinds of results. First, I show that, under standard rationality assumptions, computational constraints necessarily lead to certain behavioral heuristics. Second, I use these results to give a formal justification for behavior that is not rationalizable by expected utility preferences. I describe these results in the next two sections.

3. FOUNDATIONS FOR BEHAVIORAL HEURISTICS

Do predictions that respect rationality and tractability look meaningfully different from predictions that only respect rationality? In Camara [2022a], I demonstrate that they do look meaningfully different. I show that, under standard rationality assumptions, computational constraints necessarily lead to forms of *choice bracketing*. These are heuristics that lead a decision-maker faced with many decisions

[5] In the paper, I distinguish between weak and strong tractability based on whether the Turing machine has access to polynomial-size advice. I ignore this distinction here and state results informally.

$i = 1, \ldots, n$ to focus on each decision i in isolation, without considering the rest.[6] Equivalently, I show that expected utility maximization is intractable unless the utility function satisfies a strong separability property.

I start by introducing a symmetry assumption that I will later relax. The decisionmaker's choices are *symmetric* if she is indifferent between vectors (X_i, X_j) and (X_j, X_i). Symmetry may be plausible for investors, where income from one asset i is interchangeable with income from another asset j.

Theorem 1 of Camara [2022a] shows – assuming the P $\neq$ NP conjecture holds – that rational, tractable, and symmetric choices c are observationally equivalent to *narrow choice bracketing*. This means that a decision-maker's choice of X_i in dimension i does not depend on what she chooses in other dimensions j.

More precisely, this result shows that expected utility maximization is intractable unless the utility function is *additively separable*, i.e.

$$u(x) = f(x_1) + \ldots + f(x_n)$$

In other words, Theorem 1 is a dichotomy theorem: it partitions a class of computational problems (parameterized by symmetric utility functions u) into polynomial-time (if u is additively separable) and NP-hard (if u is not additively separable).

Theorems 2 and 3 generalize Theorem 1 by dropping the symmetry assumption and strengthening P $\neq$ NP to the *non-uniform exponential time hypothesis*. They show that rational and tractable choice correspondences are observationally equivalent to *dynamic choice bracketing*, a larger class of heuristics that augment choice bracketing with ideas from dynamic programming. These heuristics preserve the computational advantages of choice bracketing while allowing for richer patterns of behavior.

As in Theorem 1, it is useful to restate this characterization in terms of a separability property. Theorem 2 shows that if expected utility maximization is tractable then u is *Hadwiger separable*. This property is a novel relaxation of additive separability that allows for some complementarity and substitutibility across dimensions, but limits their frequency. It is quite restrictive and rules out many common utility functions, such as

$$u(x) = f\left(x_1 + x_2 + \ldots\right)$$

where f is non-linear. More precisely, Hadwiger separability is defined using the notion of an *inseparability graph*. This is an undirected graph where nodes i and j are connected if and only if the utility function u can be represented as

$$u(x_1, x_2, \ldots) = f(x_i, x_{-ij}) + g(x_j, x_{-ij})$$

The utility function u is Hadwiger separable if the inseparability graph is sufficiently sparse. That is, if the graph's Hadwiger number grows at most logarithmically in the number of dimensions n.

Together, Theorems 1-3 describe the implications of computational constraints for behavior under standard rationality assumptions. In doing so, they demonstrate that certain behavioral heuristics are not only consistent with but *predicted by* an

[6]There is substantial empirical evidence for this kind of behavior. For example, see Tversky and Kahneman [1981] or Rabin and Weizsäcker [2009] for experimental evidence.

essentially standard model of choice with mild computational constraints. The strength of these results illustrate that tractability can significantly sharpen our predictions about behavior.

Next, we will see that tractability can do more than refine rational choice; it can also highlight problems with how we define rationality in the first place.

4. CHOICE TRILEMMA

Is it really necessary for us to study rational and tractable choice at the same time or is it with minimal loss to have economists focus on rationality and computer scientists focus on tractability? I argue that it is important to study both at the same time. My evidence is the *choice trilemma* of Camara [2022a], which formally shows that incorporating tractability into our models highlights a problem with how we define rationality.

Suppose a decision-maker intrinsically wants to maximize the expected value of a given objective function $\bar{u}$. If $\bar{u}$ is not Hadwiger separable, Theorem 2 implies that the computationally-constrained decision-maker cannot make choices that exactly optimize the expected value of her objective function. Instead, she might turn to *approximation algorithms* that guarantee her a positive fraction of her optimal payoff. Will this decision-maker make choices that appear rational to an outside observer, insofar as they can be rationalized by some utility function u?

For many natural objective functions – and assuming NP $\not\subset$ P/poly – Theorem 4 of Camara [2022a] shows that a computationally-constrained decisionmaker *cannot* simultaneously (i) guarantee any non-zero fraction of her optimal payoff and (ii) be rationalized as maximizing the expected value of some utility function u.

Theorem 4 also shows that the decision-maker can guarantee approximately optimality (i) if she is willing to drop rationality (ii). That is, there do exist tractable algorithms that guarantee at least half of the decision-maker's optimal payoff. These algorithms do not satisfy the axiomatic definition of rationality of von Neumann and Morgenstern [1944], because they do not exactly maximize the expected value of any particular utility function u.

Altogether, my results imply a *choice trilemma* that relates rationality, tractability, and approximate optimality as properties of choice. For many objective functions $\bar{u}$, there exist choice correspondences that satisfy any two of these properties, but not all three. That is, a computationally-constrained decision-maker may be better off (according to her true objective function $\bar{u}$) if she is willing to make choices that an analyst would not be able to rationalize. This suggests that alternative definitions of rationality are needed.

5. CONCLUSION

Using Camara [2022a] as a case study, I argued that tractability should be a first-order concern for economists, and that tools from theoretical computer science can be useful for integrating tractability into economic models.

Specifically, I asked three questions. First, are computational models compatible with and helpful for understanding human behavior? Second, do predictions that respect rationality and tractability look meaningfully different from predictions that only respect rationality? Third, is it really necessary for us to study rational and tractable choice at the same time? The results in Camara [2022a] suggest that the answers to all three questions are affirmative.

REFERENCES

CAMARA, M. K. 2022a. Computationally tractable choice. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 28.

CAMARA, M. K. 2022b. Mechanism design with a common dataset. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. EC '22. Association for Computing Machinery, New York, NY, USA, 558.

CAMARA, M. K., HARTLINE, J. D., AND JOHNSEN, A. 2020. Mechanisms for a no-regret agent: Beyond the common prior. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 259–270.

CUMMINGS, R., DEVANUR, N. R., HUANG, Z., AND WANG, X. 2020. Algorithmic price discrimination. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '20.

ECHENIQUE, F., GOLOVIN, D., AND WIERMAN, A. 2011. A revealed preference approach to computational complexity in economics. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. EC '11. Association for Computing Machinery, New York, NY, USA, 101–110.

IMMORLICA, N., MAO, J., SLIVKINS, A., AND WU, Z. S. 2020. Incentivizing exploration with selective data disclosure. In *Proceedings of the 21st ACM Conference on Economics and Computation*. EC '20. Association for Computing Machinery, New York, NY, USA, 647–648.

OPREA, R. 2020. What makes a rule complex? *American Economic Review 110,* 12 (December), 3913–51.

RABIN, M. AND WEIZSÄCKER, G. 2009. Narrow bracketing and dominated choices. *American Economic Review 99,* 4, 1508–43.

TVERSKY, A. AND KAHNEMAN, D. 1981. The framing of decisions and the psychology of choice. *Science 211,* 4481, 453–458.

VON NEUMANN, J. AND MORGENSTERN, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.