

## COMPUTATIONALLY INTENSIVE ANALYSES IN ECONOMICS

KENNETH L. JUDD\*

*Hoover Institution, 434 Galvez Mall, Stanford, CA 94305, USA*  
*e-mail: [judd@hoover.stanford.edu](mailto:judd@hoover.stanford.edu) url: <http://bucky.stanford.edu>*

### Contents

Abstract	882
Keywords	882
1. Introduction	883
2. Computational tools	883
3. Weaknesses of standard models	884
4. Criticisms of computationally intensive research	886
5. Systematic approaches to computationally intensive research	887
5.1. Search for counterexamples	888
5.2. Sampling methods	889
5.3. Regression methods	890
5.4. Replication and generalization	891
5.5. Synergies with conventional theory	891
6. Conclusion	892
References	892

\* I thank Leigh Tesfatsion for helpful comments.

**Abstract**

Computer technology presents economists with new tools, but also raises novel methodological issues. This essay discusses the challenges faced by computational researchers, and proposes some solutions.

**Keywords**

computational economics, economic methodology

*JEL classification:* B49, C60, C63

## 1. Introduction

The growing power of computers gives economists a new tool to explore and evaluate both old and new economic theories. The essays in this handbook illustrate that potential in many parts of economics, and make clear that they have just scratched the surface of what can be done. The main goal of this handbook is to encourage new work. However, as with any new tool, there are many questions about how to use it wisely and effectively. Important methodological questions need to be addressed before computational tools can achieve their potential for contributing to economic science.

Conventional economics uses computation primarily for two purposes: empirical analysis of data and computing equilibria of conventional models. The primary computational tools for these activities are standard numerical analytic tools for solving optimization problems and nonlinear systems of equations. Agent-based computational economics (ACE) often takes us in new directions that focus on computer models of complex dynamical systems to analyze alternative theories of economic behavior. ACE research is often like theory since it studies the implications of alternative assumptions about economic systems, as described in the description of constructive theory in [Tesfatsion \(2006\)](#). Unfortunately, the complexity that is embraced by ACE research makes it difficult, if not impossible, to use conventional ways for describing theories such as stating and proving theorems, presenting cases with closed-form solutions, and proving comparative statics. Instead, much ACE research uses computer simulations to analyze complex dynamic models.

The computationally intensive approaches to economics research typified in ACE research (as well as some other economics research) presents us with basic questions about how they should be used and what we can learn from their results. Where does simulation of complex dynamic models fit into the general set of economic methodologies? When and how much can we rely on computational findings? What are the criticisms of computationally intensive work? How should we address the challenges raised by critics? This essay will examine these questions and offer some answers<sup>1</sup>.

## 2. Computational tools

Before discussing methodological issues, it is useful to recall why we are here. The key fact is that we now have increasingly powerful computational tools and rapid progress will continue. First, there has been and will almost surely continue to be tremendous progress on improving computer hardware. The progress of the past 40 years has been related to advances in semiconductors. We are all familiar with Moore's law declaring that "chip density doubles every 18 months." Of course, this cannot continue forever

<sup>1</sup> This essay updates [Judd \(1994\)](#) and [Judd \(1997\)](#) that also discusses similar questions. I also draw on the suggestions of [McFadden \(1992\)](#).

because the components of a chip cannot be smaller than a molecule. Even optimistic experts argue that this can continue for only another 10–15 years. However, that will likely not be the end of hardware progress. Current research on three-dimensional chips, asynchronous chips, spintronics, optical computing and other technological advances give us good reason to believe that computational speeds will continue to grow exponentially for at least a few more decades. Furthermore, the potential for computational work would explode if we are really lucky and quantum computing achieves just a fraction of its theoretical potential. While this is speculative, it indicates that progress will continue even after the end of Moore's law. While improvements in semiconductor technology have been immeasurably important, we get a better appreciation of historical trends when we remember that the rate of increase in computing speed due to semiconductors in the last half of the 20th century was no greater than improvement in computing speeds achieved in the first half of the 20th century using other technologies. Even if technological advances stopped today, the cost of hardware would continue to fall as we reap the benefits of learning curve effects and increasing returns to scale in the production of computer components. The cumulative impact of this progress will make computation increasingly efficient, cheap, and available to economists.

Second, there has also been significant progress in software engineering with many developments being particularly valuable for ACE modeling. Supercomputing used to mean vector processing, a technique of limited value for ACE modeling. The current strategies in high performance computing exploit massive parallelism and distributed computing. In these environments, many processors of possibly varying power are combined in a network and through communication, sometimes over the Internet, to work together to solve a problem. The value of parallelism depends on the problem. Fortunately, many of the problems discussed in the handbook, particularly those using Monte Carlo simulation, can easily make full use of the computational power of parallel and distributed computing. Of more specific value to ACE modeling has been the work on developing software tools for ACE models, such as the Sugarscape environment discussed in [Epstein and Axtell \(1996\)](#). Here, also, progress will continue and significantly reduce the human cost of doing computationally intensive economics research and make it easier for economists to profitably use ACE modeling.

This is all old news, but it bears repeating when we consider how computation could be used in economics. Some of the ideas I outline below will sound unreasonable and probably are infeasible today given current technology. However, we need to focus on how to proceed in the future, and that discussion should be mindful of the tools we will have then.

### **3. Weaknesses of standard models**

The other reason why economists are turning to ACE modelling is the dissatisfaction with conventional economic models and their frustration with the limitations of standard

research paradigms. Of course, all economics research is motivated by some dissatisfaction with the existing theories, and ACE modelling has been applied to many of the same questions, such as how an economic system gets to an equilibrium, that is studied by conventional means. What makes much of the ACE literature different from other research is its methodological novelty. Conventional economic theory, following the style of mathematics in general and real analysis in particular, begins with a set of definitions and assumptions, and proves theorems. The universe of models covered by the definitions is generally infinite. For example, general equilibrium theory begins with the concepts of preference orderings and feasible allocations. Theoretical models often make simplifying assumptions so that they can get clear, substantive results. In basic general equilibrium theory, we assume well-behaved excess demand functions and concave production functions in order to invoke the Brouwer fixed-point theorem. In other cases, such as the CAPM model of asset pricing or oligopoly models with linear demand and marginal cost curves, tractability considerations lead economists to make far more restrictive assumptions in order to get clean solutions. Furthermore, economists often examine simple models in the search for “the” cause of some economic phenomenon, and argue for a parsimonious explanation of their observations. This approach often ignores the possibility that the truth could be multidimensional, and that the multiple dimensions of reality could interact to produce phenomena that no one factor can explain. While we all like parsimony, true parsimony chooses a model as simple as possible without being too simple, and would not force our thinking into a conceptual straightjacket.

We often question the validity of the implications of these models because the elements which are sacrificed in the interest of simplicity are possibly of first-order importance. For many economists, this dissatisfaction with simple models is the main appeal of computational approaches. This dissatisfaction has moved economists in a variety of directions. For example, in public finance, economists often use computation to avoid the single-sector, representative agent models that are commonly used only because of their tractability.

The ACE literature generally aims at other weaknesses of standard models, often focusing on foundational problems instead of, for example, studying models with more goods. The chapters in this handbook study many models for which a computational approach is the only way to attain clear results. The models of social interactions presented in [Vriend \(2006\)](#), [Wilhite \(2006\)](#) and [Young \(2006\)](#) have combinatorial complexities that make it difficult (if not impossible) to attain closed-form solutions. The impact of learning on financial markets, discussed in [Hommes \(2006\)](#) and [LeBaron \(2006\)](#) also requires computational tools, particularly when individuals do not all follow exactly the same learning rules. Multiperson decision making, whether it is on the scale of a firm, as studied in [Chang and Harrington \(2006\)](#), or at the level of politics, as reviewed in [Kollman and Page \(2006\)](#), also involves complex patterns of learning and choices that are difficult to describe precisely without computation.

Many economists dismiss these complexities (along with many other features of real economic life glossed over in conventional models) arguing that they can't matter.

Some will point to convergence theorems and conclude that the convergence problem is “solved” and that factors affecting the convergence process cannot be important “in the long run.” Of course, Keynes’ observation that “in the long run we are all dead” forcefully reminds us that abstract convergence theorems tell us nothing about what happens at economic time scales. More generally, no matter how good our intuition is, we do not know which features of an economy are important and which are not until we examine them, and do so in a manner that reveals their quantitative importance. That focus on the quantitative properties of complex systems leads us to computationally intensive methods.

#### 4. Criticisms of computationally intensive research

Many economists are dissatisfied with conventional economic models, but have serious doubts about taking a computationally intensive approach to addressing fundamental issues. This is natural since any novel methodology and paradigm will be challenged and scrutinized before it is accepted. Economists using computationally intensive methods need to acknowledge this process and develop responses to the questions and criticisms raised by the status quo. Thinking about these issues will also help us construct more compelling formulations of our ideas.

First, critics point out that computational methods produce only examples, whereas conventional economic theory aims to produce theorems. This is true given the conventional use of the words “examples” and “theorem.” The usual theorem in economics, such as existence theorems in general equilibrium theory, will cover an infinite number of possible cases. However, the substantive gap between “examples” and “theorem” is less clear. In fact, isn’t “theorem” just a plural of “example”? Theories usually examine a continuum of examples but, in order to attain analytical tractability, that continuum often constitutes a measure zero set of economically plausible and interesting specifications. Assumptions made for reasons of tractability may miss many interesting phenomena. These assumptions may take the form of functional form specifications, such as the linear demand curves we often see in oligopoly theory, or may be qualitative assumptions such as the strong informational assumptions used in rational expectations analyses. While computations examine only a finite set of examples, that set can be taken from a much more robust set of possible specifications, allowing more flexible functional form specifications as well as more complex and realistic assumptions about the distribution of information and evolution of beliefs. The *relevance* and *robustness* of examples is more important than the *number* of examples, and computational methods allow one to examine cases that theory cannot touch. Furthermore, computation can often give us insights when there are no general theorems to be had. Simple general statements are not likely to be globally true, but there may still be patterns that are economically useful, such as statements about what is usually true over empirically plausible parts of the parameter space.

Second, critics point out that numerical results have errors. Again, this is a correct observation for most computational work. For example Monte Carlo simulations have nontrivial sampling error since  $N^{-1/2}$  convergence is slow. Many algorithms produce estimates of a bound on the numerical error, but this only reduces the uncertainty. Very few computational techniques produce error bounds along with the results. The presence of numerical errors is another distinction between theorems and computational results. However, these errors can be controlled by the application of sophisticated algorithms and powerful hardware. Careful simulation methods can reduce simulation error by increasing the sample size and by exploiting variance reduction methods. More generally, careful numerical work can reduce numerical errors. The problems of numerical errors in ACE models are no more difficult to handle (and often much easier) than the analogous numerical problems that arise in maximum likelihood estimation and other econometric methods.

Theoretical models may not have errors when they solve particular cases, but they often commit specification errors by focusing on tractable cases. In fact, computational work has an advantage here because numerical errors can be reduced through computation but correcting the specification errors of analytically tractable models is much more difficult. The issue is not whether we have errors, but where we put those errors. The key fact is that economists face a trade-off between the numerical errors in computational work and the specification errors of analytically tractable models. Computationally intensive approaches offer opportunities to examine realistic models, a valuable option even with the numerical errors. As Tukey (1962) put it, “Far better an approximate answer to the right question . . . than an exact answer to the wrong question. . . .”

Third, they argue that computational models are black boxes that offer few if any insights. This is an understandable reaction to a single computed example of a model, particularly one with many factors contributing to the result. A single example may show what is possible and an author may come up with an appealing story to explain the result, but one example cannot sort out the relative importance of a model’s various components. This is sometimes addressed by sensitivity analysis where a small number of alternative parameterizations are computed and the results are compared; this is essentially a computational version of comparative statics. However, it is unclear how much can be inferred from a few examples. Here, again, is a problem that can be addressed using computation. A few examples may not demonstrate much but a few thousand well chosen examples can be more convincing, and a few million examples may be as compelling as any theorem, as well as being less costly to produce. Of course, this presents us with a different problem: How do we communicate to a reader or listener the lessons learned from thousands of examples? We now turn to that issue.

## 5. Systematic approaches to computationally intensive research

An important advantage of conventional economic theory is that a theorem is an efficient means of communicating a result: it is a simple but informative statement of

a truth about a large set of examples. In contrast, computationally intensive papers in economics often focus on a few examples to show what can happen. Some papers will say “We have examined other cases and found similar results”; this statement may be true but falls far short of what is expected in a “scientific” paper. Readers of any kind of paper, theoretical or computational, want more than a couple of examples and unsupported assertions of generality. Sometimes demands by readers can be unreasonable (examples of which are related in [Axelrod \(2006\)](#)), particularly when the demand for robustness in computational models exceeds the demand for robustness in theoretical models. However, we need to develop tools for addressing reasonable demands.

A computational economist can easily offer up many examples, but it is not obvious how to communicate his findings in a compact and informative manner. For example, space limitations mean that a paper can present only a few graphs of time series generated by simulations of a dynamic process. Tables can summarize results for several cases, but they are often harder to quickly digest than a good graph, and space limitations again will limit the amount of information that can be conveyed.

Research that relies on computationally intensive methods needs to find effective ways to communicate its findings, and it needs to develop its own style. It cannot necessarily follow what, for example, physicists do. For example, if a physicist wants to simulate the collision of two black holes, he writes down the relevant equations from general relativity, uses the constants of nature that have been precisely estimated by experimentation, and uses astronomical observations to judge what size of black holes he needs to consider. A few examples will suffice for his purposes. Economics is a much less clear mixture of the quantitative and qualitative. We often make qualitative restrictions, such as concave utility, but we do not want to make inflexible functional form assumptions. When we do compute something, we have to make functional form assumptions that we acknowledge are only approximations, and calibrate them with imprecise estimates of parameters of functions that are themselves just approximations to true functions.

In this section, I will discuss some approaches that computationally intensive work may take to address the critical issues.

### *5.1. Search for counterexamples*

While a computer cannot prove a theorem<sup>2</sup>, it can help us look for falsifying examples. Suppose we have a model with parameters  $\theta$  and we have a conjecture that can be expressed as a proposition  $P(\theta)$ . For the sake of specificity, suppose that the proposition is true if  $P(\theta) \geq 0$ . For example, suppose we are examining one of the asset market models described in [Hommes \(2006\)](#) and want to test a hypothesis about the relation between price volatility and the parameters describing learning rules or agent heterogeneity. In

<sup>2</sup> Of course, computers have occasionally proven nontrivial theorems in mathematics, but we are a long way from computer theorem-proving being a common tool in economics.



this case,  $P(\theta)$  would be a statement about measures of volatility (some moments or a measure of chaoticity) and  $\theta$  would include the exogenous parameters. Even if we could not prove the truth of  $P$  (that is, the global nonnegativity of  $P(\theta)$ ) we could assess the likelihood of its truth by searching for counterexamples, that is, values of  $\theta$  such that  $P(\theta) < 0$ .

Global optimization software could be used for testing  $P(\theta)$  by finding the global minimum of  $P(\theta)$ , and determining if it were ever negative. The choice of global optimization software would depend on the nature of the function  $P(\theta)$ . If  $P(\theta)$  were a rough function, we would have to use methods like genetic algorithms or simulated annealing. If  $P(\theta)$  were piecewise continuous, then we would want to combine a global strategy (such as in GA or simulated annealing) with a more conventional optimization method, such as Nelder–Mead, to take the guesses generated by the global strategy and find nearby local optima. If  $P(\theta)$  were a smooth, but possibly multimodal, function, we could even combine a Newton-style method with a global strategy. Once we exploit the properties of  $P(\theta)$ , we could formulate an efficient as well as systematic approach for finding counterexamples.

If we find a counterexample then we will have learned something about the model. Also, the counterexample, or counterexamples, will give insight about when and why a proposition fails to hold. Failure to find a counterexample would not prove the conjecture, but would be strong evidence for its truth. If high-quality global optimization software is used, then this would be even more compelling evidence.

If we do find counterexamples, we would like to find ways to describe when  $P$  is true. If we fail to find a counterexample, we may want to consider alternative ways to express the apparent global validity of proposition  $P$ . We next turn to methods that help us in those tasks.

## 5.2. Sampling methods

If we are convinced of a proposition's truth, then we would want to express that in some compact way. Various sampling schemes can be developed for this purpose and produce statements using standard language from statistics or analysis.

Monte Carlo sampling offers one simple procedure. Suppose we want to investigate a set of models where we have imposed a probability measure,  $\mu$ , over the parameter space  $\theta$ . Suppose we want to evaluate our proposition  $P(\theta)$  over a set  $\theta \in \Theta$ . We could draw  $N$  models at random from  $\Theta$  according to the measure  $\mu$ , and use computation to determine the truth of the proposition in those cases. If computation showed that proposition  $P$  held in each case, then we could say "We reject the hypothesis that the  $\mu$ -measure of counterexamples to proposition  $P$  exceeds  $\epsilon$  at the confidence level of  $1 - (1 - \epsilon)^N$ ." Note the crucial role of the randomization; the fact that we randomly drew the cases allows us to use the language of classical statistics.

We could also use Bayesian methods to express "posterior beliefs" after several computations. Let  $p$  be the probability that a  $\mu$ -measure randomly drawn point satisfies proposition  $P$ , and suppose that we have a uniform prior belief about the value of  $p$ .

Then our posterior belief about  $p$  after  $N$  draws which satisfy proposition  $P$  can be directly computed.

The advantage of Monte Carlo sampling methods is the ease of expression using language from either classical or Bayesian statistics. There is little question about the meaning of these statements since independent draws are easy to implement and well-understood.

Some have told me that they would prefer to use a prespecified, uniform grid of cases for this task instead of random draws. The idea is to examine a set of examples such that each possible case is within some distance  $\delta$  of one of the cases computed. The uniform grid approach has an advantage over Monte Carlo in that it avoids the clumping and gaps that, due to the Central Limit Theorem, must occur with Monte Carlo sampling. However, uniform grids are inefficient ways of sampling in a multidimensional space. Fortunately, there are quasi-Monte Carlo sampling methods, such as low discrepancy sets, that use far fewer points than the Cartesian grid and accomplish the same goal. With deterministic grids, one cannot use the statistical concept of “confidence levels” to summarize a result. The alternative statement would be based on the maximal size of a ball or cube of counterexamples; that is, if proposition  $P$  is true at each point on a grid and the largest ball which can miss each point on the grid is of diameter  $\delta$ , then  $\delta$  could be used as a measure of the strength of proposition  $P$ .

One advantage of all sampling methods is the ease of implementation. If you can compute  $P(\theta)$ , then you can execute a sampling method. Sampling methods can efficiently use any computer environment. In particular, because there is little interdependence across different points in a grid, sampling methods can be directly implemented in all distributed computing environments, such as massively parallel supercomputers or grid computing systems. The global optimization approach could also exploit a parallel environment but would require some coordination.

### 5.3. Regression methods

Instead of trying to prove that  $P(\theta)$  is always nonnegative, we may instead try to find the shape of  $P$ . Judicious use of computational power can help us here as well. A computational study can compute  $P(\theta)$  for a large number  $\theta$  values, use approximation methods, such as regression, neural nets, or radial basis functions, to express how  $P$  depends on the exogenous parameters,  $\theta$ . The approximation results would then tell us how a model depends on its parameters. If a simple functional form, such a low order polynomial in the components of  $\theta$  (or  $\log \theta$ ), could fit the data, then the fitting function would cleanly express our findings.

While this looks a lot like statistics, our task would be easier than standard econometrics. First, we can define the set of sample points  $\theta$ . Econometricians are stuck with the  $\theta$ 's nature gives them. We instead can control the number and distribution of  $\theta$ 's so as to maximize the information we get from our computations. Second, the error in computing  $P(\theta)$  can often be controlled much better than an econometrician can control measurement errors. Third, because we have control over the measurement errors,

sampling errors and sample size, we can be more flexible in terms of functional form specifications and get more information out of our data. In particular, we could focus on finding functional forms that can compactly express the patterns we find.

There are many ways to accomplish this. The main point is that approximation methods, and data mining in general, could be used to summarize results of a computational study and test hypotheses one has for a model.

#### *5.4. Replication and generalization*

We have discussed ways that computational work could be conducted and expressed to produce conclusions that are clear to a reader, and, in some cases, nearly as compelling as a proof. Computational methods have one potentially important advantage over theorem-proving. Suppose that your paper did not examine a case or class of models that some reader cares about. If you have proved a theorem that does not include that case, the reader has to work hard to see what happens for his case, and will usually fail to find the answer unless he has expertise comparable to the author. In contrast, in computational models the reader could just take the computer program you wrote and apply it to the parameter values he wants to examine. If his case is qualitatively different, he could perhaps make modest changes in the (hopefully well-written) code and then run the program. In either case, he can quickly find the answers to his questions in a way that is not possible for theoretical work.

This observation also points out how replication of computational work could be done. Of course, this assumes that the software you use is easily transportable and flexible to use. This is often not true today, partly because there is little incentive to write software that can be used by others. This is a not as bad a problem in empirical work since people often use common data sets and common econometric software. The lack of similar software is holding back the potential of ACE approaches, but that is hopefully only a temporary problem.

#### *5.5. Synergies with conventional theory*

The observations above have been of an “us versus them” nature. While this is a useful framework to use when discussing these issues, it is counterproductive for us to view this as a zero-sum contest between two methodologies. The ultimate aim is for computational and theoretical tools to interact in a fruitful manner. We have already seen some examples of this. For example, [Arthur \(1994\)](#) posed the El Farol Problem where customers want to predict the number of people at the El Farol bar because they want to avoid times when it is crowded. [Arthur \(1994\)](#) used computational methods to study the implications of inductive inference by patrons. Motivated by the insights in [Arthur \(1994\)](#), [Zambrano \(2004\)](#) reexamined the problem analytically and arrived at important insights that went beyond both the computational results and the related game theory.

I have laid out several distinctions between conventional theory and the kind of computationally intensive approaches to studying economic models advocated in ACE.

These distinctions will remain in the future even after we have refined our computational tools, and both approaches will be used, each exploiting its own unique strengths. The aim of this essay, as well as much of this handbook, is to highlight the distinct nature of computationally intensive research tools. However, there is no desire that the economics community be divided between computational economists and practitioners of conventional theory. Instead, the hope here is that a clear understanding of alternative methodologies will foster vigorous interactions where each approach benefits from the insights of the other.

## 6. Conclusion

Any time a new tool is introduced into economics, economists need to decide how best to use it to produce insights about economic problems. Economists who were trained in the literary tradition of classical economics were troubled by the infusion of mathematics in the middle of the 20th century. Likewise, the infusion of computationally intensive approaches will push economists to learn new tools, and raise questions about how best to use them in economic analysis. These questions need to be addressed in a systematic manner. This essay proposes some ideas that could be used more in economics, and I am sure that others will offer suggestions as we think about these issues. The potential usefulness of computational methods is enormous. I am confident that we will find suitable answers if we think carefully about the tradeoffs between conventional approaches and the ACE tools presented in this handbook, and if we make efficient and full use of the computational resources that will become available in the future.

## References

- Arthur, W.B. (1994). "Bounded rationality and inductive behavior (the El Farol Problem)". *American Economic Review Papers and Proceedings* 84, 406–411.
- Axelrod, R. (2006). "Agent-based modeling as a bridge between disciplines", this handbook.
- Chang, M.-H., Harrington, J. Jr. (2006) "Agent-based models of organizations", this handbook.
- Epstein, J., Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press/Brookings, MA.
- Hommes, C. (2006). "Heterogeneous agent models in economics and finance", this handbook.
- Judd, K.L. (1994). "Comment". In: Sims, C. (Ed.), *Advances in Econometrics Sixth World Congress*, vol. II. Cambridge University Press, Cambridge, UK, pp. 261–274.
- Judd, K.L. (1997). "Computational economics and economic theory: complements or substitutes?". *Journal of Economic Dynamics and Control* 21, 907–942.
- Kollman, K., Page, S.E. (2006). "Computational methods and models of politics", this handbook.
- LeBaron, B. (2006). "Agent-based computational finance", this handbook.
- McFadden, D. (1992). "Simulation methods for unemployed theorists". Santa Fe Institute Conference.
- Tesfatsion, L. (2006). "Agent-based computational economics: A constructive approach to economic theory", this handbook.
- Tukey, J.W. (1962). "The future of data analysis". *The Annals of Mathematical Statistics* 33 (1), 1–67.
- Vriend, N. (2006). "ACE models of endogenous interactions", this handbook.

Wilhite, A. (2006). "Economic activity on fixed networks", this handbook.

Young, H.P. (2006). "Social dynamics: Theory and applications", this handbook.

Zambrano, E. (2004). "The interplay between analytics and computation in the study of congestion externalities: The case of the El Farol Problem". *Journal of Public Economic Theory* 6, 375–395.