

Market Design Using ACE Simulations

- I. Designer Markets**
- II. From Analysis to Design**
- III. Market Design**
- IV. Complexity of Design**
- V. Design Trade-Offs**
- VI. From Closed-Form ...**
- VII. ... To Agents**
- VIII. The Design Economist.**

Simon on Market Design

“The basic idea is that the several components in any complex system will perform particular subfunctions that contribute to the overall function. ... To design such a complex structure, one powerful technique is to discover viable ways of decomposing it into semi-independent components corresponding to its many functional parts. The design of each component can then be carried out with some degree of independence of the design of others... There is no reason to expect that the decomposition of the complete design into functional components will be unique.... Much of classical organization theory in fact was concerned precisely with this issue of alternative decompositions of a collection of interrelated tasks.” (Simon, 1996)

I. Designer Markets

Designing markets is a new discipline.

At least five examples of designed market can be identified:

- 1. simulated stock markets;**
- 2. emission markets;**
- 3. auctions for electro-magnetic spectrum;**
- 4. electricity markets; and**
- 5. on-line, e-commerce markets.**

Contract design is another area where agent-based modeling might be used, but negotiation and design of contracts by use of computer simulation and agent-based modeling is only now emerging from its infancy Jennings et al. (2001).

1. New Financial Markets

Markets for new financial derivatives were created and traded after Black, Scholes, and Merton solved the 70-year-old problem of pricing options.

Previously, financial traders knew that options were valuable, but not how to value them exactly.

More recently, research into the rules and micro-structure of stock markets, continuous double-auction trading, through the use of simulated markets (LeBaron (2005))

2. Markets for Environmental Amenity

Markets for pollution emissions, e.g. SO₂ and CO₂.

Realization that the emissions from anthropogenic processes were, at least potentially, altering the biosphere for the worse was followed only after a lag with the awareness by policy makers that market mechanisms could be harnessed to control such emissions, generally more efficiently than could other mechanisms.

Hailu and Schilizzi (2005), Janssen and Ostrom (2006)

3. Markets for Spectrum

Simultaneous ascending-bid auctions have recently been designed for selling bands of local electro-magnetic spectrum.

Use of auctions to choose the new owners and to value these assets slowly replaced so-called “beauty contests,” in which subject to certain technical requirements licenses were virtually given away.

But these new auction mechanisms at first did not allow for the complementary nature of bands in different localities.

Only after intensive efforts by economists advising governments and bidding companies did the successful “3G” auctions occur [Milgrom (2004)].

4. Markets for Electricity

A move away from centralized engineering-dominated means of allocating electricity load across generators and distribution networks to using market mechanisms of various kinds.

Since electricity cannot (easily or cheaply) be stored, previously existing market mechanisms were not appropriate.

Instead, several types of new market mechanisms have been introduced.

5. On-line markets.

With the growth of the use and extent of the Internet over the past ten years, and the dot-com boom, with buying and selling on-line, opportunities for designing on-line markets de novo, as opposed to trying to emulate existing face-to-face markets, have arisen.

In the last few years these opportunities have given rise to much work by computer scientists, as well as economists.

Indeed, there is a productive research intersection of the two disciplines, as revealed in some of the papers discussed below. MacKie-Mason and Wellman (2005)

Analysis First

Before design must come analysis.

Simulation allows analysis of systems that are too complex to analyze using traditional, closed-form techniques.

Once we understand through analysis how the elements of the phenomenon of concern work together, we can ask the question of how to improve its operation: how better to design it.

II. From Analysis to Design

Roth one of the first (and most prominent) economists to argue for the economist as designer: *the market engineer*.

Roth (1991) outlined the iterative process of market design using three possible approaches:

- 1. traditional closed-form game-theoretic analysis;**
- 2. experimental results from economics laboratories; and**
- 3. computational exploration of different designs.**

If the design criteria are clearly defined, some recent techniques of simulation and optimization from computer scientists and computational economists can be used to search for optimal market designs, directly and indirectly.

Markets Have *Emerged*

Historical market institutions have in general not been imposed from above (*top-down design*) but have emerged from the bottom up as a consequence of a multitude of actions and interactions of the myriad traders (McMillan 2002).

The omnipotent programmer can experiment with different market forms and different kinds of boundedly rational agents to discover sufficient combinations of each for specific behavior of the market,

But evolutionary computation raises the possibility of *bottom-up design*,

or emergence of market design through simulation.

Purpose of Agent-Based Simulated Market Design

Which?

- **as a model of human behavior (where analysis is followed by design, given the behavior of the agents and the emergent aggregate outcomes) — in which case it is an empirical question as to how boundedly rational the agents should be to best model human agents (Duffy, 2006) or**
- **as an end in themselves, because on-line it is possible to use agents (“buy-bots, sell-bots”) to buy and sell, without the errors that human agents are heir to.**

Two Deeper Issues

1. **to what extent are the learning processes of human participants in real-world markets mal-adapted to market institutions?**

Could the use of agent-based optimization tools improve human market behavior?

e.g. in eBay auctions, when bidders use software to enhance their chances of being the high bidder at the deadline.

2. to what extent have existing market protocols (or market designs) evolved or been designed to avoid the need for any great rationality on the part of market participants?

Gode & Sunder (1993) and others seek to answer this question for financial markets; their results may be valid for other markets.

Markets and the Participants

Market performance may depend on the degree of “intelligence” or “rationality” of the agents buying and selling.

What if a market design with agents of low degree of “intelligence” is found to be sufficient for a specific level of market performance?

Then less bounded participants would, through buying and selling, create an efficient market (Walia et al. 2003).

Unless a market loophole, or rent-seeking by more intelligent agents, lowers market efficiency (Arifovic 2001).

See Tesfatsion (2002), LeBaron (2006), Duffy (2006).

III. Market Design Difficulties

Design: a process of building directed by: the pre-specified design objectives, or an explicit how-to plan.

But specifying objectives does not → exactly how the model building should occur. Why?

Objectives specified in a performance space (or behavior space), but the building occurs in a design space.

Mapping: designed structure → the desired performance not clear.

In evolution: design in the genome space, while behavior or performance in the phenome space.

Designing Markets

In designer markets: policy-makers use theory, human experiments and computer simulations to help the mapping: design (structure and rules) → behavior of the economic actors (the performance of the system).

Where the mapping is sufficiently well understood, and where closed-form analytic solution is tractable, possible to describe not only *sufficiency*:

If the market has this structure, and the rules of trading are such and such and the traders are given this information, then this performance and behavior will follow

but also *necessity*:

If you want this performance and behavior, then this is the only set (or sets) of designs (combinations of structure and rules) that will produce it.

Sufficiency yes, but Necessity?

With no closed-form analytical solution, but with human experiments or with computer simulations: necessity is in general impossible, only sufficiency.

But with few degrees of freedom, necessity is close:

Watson & Crick (1953): simulated the structure of DNA, given its chemical properties (acid), known atomic composition (and electrical properties), and some X-ray diffractions

A Marketplace Design Framework

— MacKie-Mason & Wellman (2006)

Three fundamental steps that constitute a transaction:

- 1. the connection (searching for and discovering the opportunity to engage in a market interaction),**
- 2. the deal (negotiating and agreeing to terms),**
- 3. the exchange (executing a transaction).**

Define a “marketplace system” as consisting of: agents and the market mechanism through which they interact, all embedded in an environment of social institutions (language, laws, etc.).

M & W's Market Mechanism

The set of “rules, practices, and social structures of a social choice process, specifying —

- 1. permissible actions” (including messages), and**
- 2. market-based exchange transactions as outcomes of a function of agent messages.**

If there is some entity, apart from the participating agents, that manages any inter-agent communication and implements the mechanism rules, then the market mechanism is *mediated*.

Thus Several Design Decisions

M&W: this characterization of a marketplace → several design decisions:

- 1. the design of the market mechanism, which might be decomposed into the design of mechanisms for, successively, the connection, the deal, and the exchange phases of a transaction.**
- 2. design of agents to interact with the market mechanism, whether existing or newly designed.**

M&W define an agent

An agent: an “autonomous decision-making locus in a system of multiple decision-making entities”.

An agent has “type” attributes, such as:

- **preferences,**
- **beliefs,**
- **intentions, and**
- **capabilities.**

Want consistency between the agents’ behavior, beliefs, and preferences, consistent with some principle of rationality.

Here: focus on design of MacKie-Wellman’s market mechanism, specifically, the deal negotiation task, that govern the settlement from allowable actions.

Market Mechanisms

Mechanisms specify:

- 1. the agents' concerns that are recognized, and,**
- 2. rules mapping actions into allocation outcomes.**

A rule might:

- specify which actions are permissible, or**
- the procedure for choosing a settlement of agents' concerns based on observable actions.**

E.g., auctions have rules governing allowable actions, and rules governing settlement.

Designs are constrained

Design of the market mechanism must be measured,

usually consists of a constrained optimization, even if not explicitly or directly.

E.g. “No external subsidies” or “maintain horizontal equity”

— are two possible constraints.

The General Market Design Issue:

has become designing a market mechanism that:

- includes defining a set of concerns over which agents can interact,**
- specifying rules of permissible actions, and**
- specifying rules for mapping from actions to settlement and outcomes.**

IV. Complexity of Design

“... it is typical of many kinds of design problems that the inner system consists of components whose fundamental laws of behavior ... are well known. The difficulty of the design problem often resides in predicting how an assemblage of such components will behave.” — Simon (1981).

Simon is speaking of complex systems, in which emergence may occur.

Syntactic Complexity

The *syntactic complexity* of design (Edmonds & Bryson 2003):

the lack of a clear mapping from design to behavior:

the only way to know the system's outcomes is to run the system, and observe the emerging performance: analysis is not able to predict the outcome.

Reasons for Design Complexity

**One reason why analytical methods of analysis might fail: the mapping: initial conditions of structure and rules → behavior and performance is not smooth or continuous,
∴ is not amenable to calculus-based tools.**

The rugged landscape (Kauffman 1995) is even more difficult if it too is changing, perhaps because of the strategic complexity of other players' actions and learning: *co-evolution*.

Partly because of these complexities, direct design of markets is hardly ever attempted. But see Byde (2006) next lecture.

V. Design Trade-offs

Where there are several design criteria, the possibility arises of trade-offs between the criteria.

E.g., if a firm has market power, it can maximize its seller revenue, but at the cost of market efficiency, as measured by the sum of seller (or producer) surplus and buyer (or consumer) surplus.

Or it might be possible to improve the fairness of a market outcome, but at the cost of market efficiency.

Such trade-offs must be explicit.

Satisficing in Market Design?

Possible to use a version of Simon's (1981) satisficing: so long as the other criteria are met (above some target level), the remaining criterion is used to rank designs.

Or different criteria could be weighted to derive a single, scalar maximand.

How good is a designed auction market?

(Phelps et al., 2002, 2004) suggest eight possible criteria:

1. **maximizing seller revenue:**
one of the main criteria in the design of the spectrum auctions, such as the 3G auctions (Milgrom, 2004);
2. **maximizing market allocative efficiency:**
a desirable policy attribute of a marketplace system;
3. **discouraging collusion**
to attaining the first and second criteria;
4. **discouraging predatory behavior**
to help to maximize efficiency;

-
- 5. discouraging entry-detering behavior to maximize seller revenue;**
 - 6. budget balance:
no third-party payments for a deal to be reached;**
 - 7. individual rationality:
the expected net benefit to each participant from the market mechanism should be no less than the best alternative; and**
 - 8. strategy-proofness:
participants should not be able to gain from non-truth-telling behavior.**

A Design Impossibility Theorem

Myerson & Satterthwaite (1983) derived an impossibility result:

No double-sided auction mechanism with discriminatory pricing can be simultaneously efficient, budget-balanced, and individually rational.

In discriminatory-price auctions (or “pay-as-bid” auctions), distinct trades in the same auction round occur at distinct prices;

In uniform-price auctions, all trades in any given auction round occur at the same price.

Talukdar's requirements:

Talukdar (2002): before the market can be designed (solved), the design problem must be well posed, that is, complete, feasible (all constraints can be satisfied), and rich (allows for innovative and desirable solutions).

To be complete, the design problem must contain:

- 1. the attributes to be used in characterizing behavior of the market;**
- 2. the decision variables to be used to characterize the structure;**
- 3. the goals to be attained (desired behaviors, laws, regulations); and**

-
- 4. a computable mapping of decision variables into goals (does each point in decision space meet the goals?).**

This is achieved for complex design problems by iterative analysis, achieved using agent-based simulation tools and agent-based verification tools, since such tools are open and modular

Blake's Seven

Blake LeBaron (2004), in examining the use of agent-based models of financial markets, discusses seven basic design questions for his models, which translate across to more general models.

- 1. the economic environment itself needs to be resolved: what will be traded? what is the scope of the market?**
- 2. how are agents' preferences to be modelled: with particular functional forms such as mean–variance, Constant Absolute Risk Aversion, myopic or inter-temporal, or perhaps just using evaluation of specific behavioral rules.**

- 3. modelling of market clearing and price formation.**
- 4. evaluating the fitness of the model: wealth or utility? And whether the evolving rules are forecast-based (what will the price be at time t ?) or demand- and action-based.**
- 5. how information is precessed and revealed.**
- 6. how learning occurs: is it social and direct or at arm's length; is it individual?**
- 7. how is benchmarking to be undertaken?**

While these questions relate to the models used to design markets, they may also reflect on the design criteria for the final designer markets.

VI. From Closed-Form Equilibria ...

Since Samuelson, economists have sought closed-form solutions to understand the performance of economic institutions.

Economic actors assumed to be perfectly rational, with the means to solve for equilibria outcomes in complex situations.

Economists have examined the existence, uniqueness, and stability of equilibria of economic interactions.

When the interactions among economic actors are strategic, the equilibria become Nash equilibria.

Disequilibrium etc.

But in an operating, real-time actual market, we are not interested just in equilibrium characterization:

continual shocks might never allow the system to approach, let alone reach, the equilibrium.

Moreover, in a repeated interaction almost any individually rational outcome for each player can be supported as an equilibrium. (The Folk Theorem of repeated games.)

Particularly so for interactions which have the general character of the iterated Prisoner's Dilemma (IPD).

Four Reasons for other means of market design

- 1. tractability:**
still very difficult to obtain solutions to the design of some markets, such as continuous double auctions (CDAs).
- 2. like to characterize out-of-equilibrium behavior, and especially the dynamic behavior of an operating market with fluctuating demand, and perhaps varying numbers of sellers, with unpredictable, varying costs.**

- 3. the assumption of perfect rationality and unlimited computational ability on the part of human traders is unrealistic, and not borne out by laboratory experiments with human subjects.**

Instead, using computer models of trading agents, should like to model economic actors in markets as “boundedly rational” — bounded computational ability, or bounded memory, or bounded perception (Marks 1998).

Conlisk’s (1996) four reasons for using bounded rationality in economic models: (1) evidence of limits to human cognition, (2) successful performance of economic models with bounded rationality, (3) sometimes unconvincing arguments in favor of unbounded rationality, and (4) the costs of deliberation.

4. to model learning:

Two reasons to include learning in any models used to design markets:

- a. Individuals and organizations learn: a model without learning is not as realistic as one incorporating learning.**

Bunn & Oliveira (2003) note that many researchers (including Erev & Roth 1998) have shown that learning models predict people's behavior better than do Nash equilibria.

-
- b. Moreover, learning can help to eliminate many otherwise legitimate Nash equilibria.**

Indeed, evolutionary (or learning) game theory has been seen as a solution to the multiplicity of Nash equilibria that occur in closed-form game-theoretic solutions: a priori, all are possible, but to see which are likely in reality, see how players learn and choose amongst them.

VII. ... Enter the Agents

**Can design without the use of agents:
given a market with demand and supply schedules,
economic efficiency is maximized at the output
level where marginal value equals the marginal unit
cost, no matter how the social surplus is divided
between buyers and sellers.**

**But such direct design (optimization) requires a
well defined problem.**

**With several design trade-offs and the possible
emergence of unforeseen performance in the
system:**

**enter agent-based analysis and design:
models the market system as “evolving systems of
autonomous, interacting agents with learning
capabilities“ (Koesrindartoto & Tesfatsion, 2004)**

LeBaron argues ...

LeBaron (2004) places some weight on how actual trading occurs: the institutions under which trading is executed.

He argues that agent-based models are well suited to examining market design and micro-structure questions because:

- 1. they can produce a large amount of data, and**
- 2. they allow testing of market design in a heterogeneous, adaptive environment.**

Examples.

Audet, Gravelle & Yang (2001) is an agent-based study of micro-structure (order books v. dealers)

Bottezzi, Dosi & Rebesco (2003) is another that examines tick sizes (and unexpectedly determines that smaller tick sizes do not necessarily improve the market's efficiency).

Chan & Shelton (2001) examines how a model behaves with different reinforcement-learning mechanisms, all of which enable the optimum policy function for a market-making broker to be found.

Marks et al. (1995), Arifovic (1994), Midgley et al. (1997) modelled the interactions using a GA.

Computer Scientists Designing ...

Designing systems of exchange, of markets, of distributed computing systems, and on-line trading in real time.

But the equilibrium characterizations of mathematical economics do not provide the answers they need:

their on-line markets are in disequilibrium almost always if trading in real time.

The adjustments of the operation of the markets to the current equilibrium (or attractor) will almost never happen fast enough to reach equilibrium, especially when the location of the attractor is continuously changing.

VIII. The Design Economist

Roth (2000, 2002) points the way forward for market design, with the economist as engineer.

Roth: as engineers have learnt to borrow from the insights of physics, the design economist can use insights from equilibrium mathematical economics, and from computer science.

When these insights are curtailed, perhaps by the tractability of closed-form analytical methods, both economists and software engineers have been using simulation in analysis, to obtain sufficient, but rarely necessary, conditions.

Simulation has occurred using GAs, numerical solutions, and explicit agent-based models, with iteration.

Some Criticisms of the ACE approach to modelling markets

- **Too many parameters,**
- **Questions about the stability of trading to the introduction of new trading strategies, sensitivity to the number of agents trading,**
- **Over-reliance on inductive models of agents, which respond to past rules and forecasts and**
- **Not enough on deductive models which might learn commonly held beliefs about how markets work.**

Addressed in the two areas of market design that we now consider: electricity markets and automated markets.

References

This is incomplete. For missing references, see Marks (2006), at <http://www.agsm.edu.au/~bobm/teaching/Taiwan.html>

- [1] Duffy, J., (2006), Agent-based models and human-subject experiments, *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science, forthcoming).
- [2] Edmonds, B., and J. J. Bryson (2003), Beyond the Design Stance: The Intention of Agent-Based Engineering, Centre for Policy Modelling, CPM Report No.: CPM-03-126.
<http://cfpm.org/papers/btds.pdf>
- [3] LeBaron, B., (2006), Agent-based computational finance, *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science, forthcoming).
- [4] MacKie-Mason, J.K., and M.P. Wellman (2006), Automated markets and trading agents *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science, forthcoming).
- [5] Marks R.E., 2006, Market Design Using Agent-Based Models, Ch 12 in the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science, forthcoming).
- [6] McMillan, J. (2002), *Reinventing the Bazaar: A Natural History of Markets*, NY: Norton.
- [7] Phelps, S., P. McBurney, S. Parsons, and E. Sklar (2002a), Co-evolutionary auction mechanism design: a preliminary report. *Lecture Notes In Computer*

Commerce IV, Designing Mechanisms and Systems, ed. by J. A. Padget, O. Shehory, D. C. Parkes, N. M. Sadeh, and W. E. Walsh, Berlin: Springer, pp. 123–142.

- [8] Phelps, S., S. Parsons, and P. McBurney (2004), An evolutionary game-theoretic comparison of two double-auction market designs, presented at the Workshop on Agent-Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems, New York, July 19.
- [9] Roth, A.E., (1991), Game Theory as a Part of Empirical Economics, *Economic Journal*, 101(401): 107–14.
- [10] Roth, A.E., (2000), Game theory as a tool for market design, in F. Patrone, I. García-Jurado, S. Tijs (eds.) *Game Practice: Contributions from Applied Game Theory*, Kluwer, pp. 7–18.
- [11] Roth, A.E., (2002), The economist as engineer: game theory, experimentation, and computation as tools for design economics, *Econometrica*, 70(4): 1341–1378.
- [12] Simon, H., (1981), *The Sciences of the Artificial*, 2nd ed., Cambridge: The MIT Press.
- [13] Tesfatsion, L., (2002), Agent-Based Computational Economics: Growing Economies from the Bottom Up, *Artificial Life*, 8(1): 55–82, October 2001, pp. 504-523.