

Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games[†]

VINCENT P. CRAWFORD*

Harstad and Selten's article in this forum performs a valuable service by highlighting the dominance of optimization-based models over boundedly rational models in modern microeconomics, and questioning whether optimization-based models are a better way forward than boundedly rational models. This article complements Rabin's response to Harstad and Selten, focusing on modeling strategic behavior. I consider Harstad and Selten's examples and proposed boundedly rational models in the light of modern behavioral economics and behavioral game theory, commenting on the challenges that remain and the most promising ways forward. (JEL B40, C72, D01, D03, D80)

1. Introduction

Although neoclassical microeconomics is one of social science's success stories, a growing body of experimental and empirical research has documented substantial deviations from its core behavioral assumptions and predictions (Selten 1990, 1998; Thaler 1992; Rabin 1998; Frederick, Loewenstein, and O'Donoghue 2002; Camerer 2003; Sobel 2005; DellaVigna 2009; Armstrong

and Huck 2010; Crawford, Costa-Gomes, and Iriberry 2013).¹

Because the deviations have identifiable systematic components, modeling them and integrating them into microeconomics has the potential to strengthen our models of individual decisions and strategic behavior, and the economic analyses that depend on them.

Attempts to model the deviations can be sorted into two categories, following Harstad and Selten's article in this forum. Most models in either category can be augmented to allow decision errors, which is usually essential in empirical applications. However, I will

*University of Oxford and All Souls College, Oxford, and University of California, San Diego. I am grateful to Ronald Harstad and Matthew Rabin for useful discussions, and to the University of Oxford and All Souls College for research support. I also owe thanks to my long-term coauthors Miguel Costa-Gomes and Nagore Iriberry. Some of the exposition here adapts brief passages from Crawford, Costa-Gomes, and Iriberry (2013).

[†]Go to <http://dx.doi.org/10.1257/jel.51.2.512> to visit the article page and view author disclosure statement(s).

¹ Here and below I make no attempt to give comprehensive citations, and I favor surveys when reasonably recent ones exist. My goal is to give readers a start in reading their way into an enormous body of relevant literature.

focus on the models' deterministic structures, mentioning decision errors only when they are important.

The first, "boundedly rational" category includes models that seek to improve upon neoclassical models by relaxing their core assumption that individuals optimize, in favor of various sensible, direct characterizations of individual decision behavior. Although boundedly rational models relax optimization, many of them maintain (if sometimes implicitly) the customary neoclassical assumptions about the domain of individuals' preferences, the coherence of their goals, and the accuracy of their models and the inferences they draw from them.

Examples of boundedly rational models of individual decisions can be found in Simon (1955), Cyert and March (1963), Newell and Simon (1972), Nelson and Winter (1982), and Rubinstein (1998), among many others. Thoughtful discussions can be found in Hogarth and Reder (1986), Selten (1990, 1998), Conlisk (1996), Munier et al. (1999), and Spiegler (2011a, 2011b).

Examples of boundedly rational models of strategic behavior can be found in Rosenthal (1989), Rubinstein (1998), Young (2004), and, with regard to "reinforcement" learning, Roth and Erev (1995) and Erev and Roth (1998).

Harstad and Selten also view Selten's notions of "learning direction theory" and "impulse-balance equilibrium" (Selten and Stoecker 1986; Selten and Buchta 1999; Selten, Abbink, and Cox 2005; Ockenfels and Selten 2005; Selten and Chmura 2008; Brunner, Camerer, and Goeree 2011; and Selten, Chmura, and Goerg 2011) as boundedly rational models and, in an important sense, that is correct. However, I will argue that there is also an important sense in which learning direction theory and impulse-balance equilibrium fall into the second category, which I will call "optimization-based."

Optimization-based models maintain the customary neoclassical assumption that individuals act as if to optimize *something*. Instead, such models seek to improve on neoclassical models by relaxing or replacing one or more of the other customary neoclassical assumptions.

As Harstad and Selten note (footnote 1 and pages 6–7), optimization-based models, broadly construed, have come to dominate modern behavioral economics, much as they have long dominated neoclassical microeconomics.

In one branch of behavioral decision theory, optimization-based models relax customary neoclassical assumptions about the domain of individual preferences to allow reference-dependent preferences, which respond to changes in consumption relative to a reference point as well as levels of consumption (Kahneman and Tversky 1979; Tversky and Kahneman 1991; Kőszegi and Rabin 2006). In another branch, optimization-based models relax the customary domain restrictions to allow social preferences and reciprocity, with individuals responding to others' outcomes and/or decisions as well as their own (Rabin 1993; Fehr and Schmidt 1999; Bolton and Ockenfels 2000; Andreoni and Miller 2002; Charness and Rabin 2002; Sobel 2005).

In still another branch of behavioral decision theory, predominantly optimization-based models relax customary assumptions about the coherence of individual goals to allow present-biased preferences with a time-inconsistent tension between the desire for immediate and future gratification (Laibson 1997; O'Donoghue and Rabin 1999; Frederick, Loewenstein, and O'Donoghue 2002). And in yet another branch, predominantly optimization-based models relax customary assumptions about the accuracy of individuals' models or inferences to allow various "heuristics and biases" (Tversky and Kahneman 1974; Rabin 2002).

Turning to behavioral game theory, optimization-based learning models relax the customary neoclassical assumption that players play a Nash equilibrium from the start of play, assuming instead that they follow simple adaptive rules that may converge to equilibrium decisions, if at all, only in the limit. In “beliefs-based” adaptive learning models—unlike the reinforcement adaptive learning models I list as boundedly rational—players’ adjustments are directly motivated by optimization, even though their beliefs are based on oversimplified models of others’ decisions (Woodford 1990; Milgrom and Roberts 1990, 1991; Selten 1991; Crawford 1995; Fudenberg and Levine 1998; Camerer and Ho 1999; Camerer, Ho, and Chong 2002).²

Learning direction theory and impulse-balance equilibrium postulate players’ adjustments that are motivated by optimization, even though they do not explicitly model players’ beliefs. In that sense, learning direction theory and impulse-balance equilibrium are more closely related to conventional beliefs-based learning models than they are to most bounded-rationality models (Ockenfels and Selten 2005, footnote 11).

In strategic applications where lack of clear precedents makes simple learning models implausible, models of strategic thinking in initial responses to games relax Nash equilibrium in favor of “level- k ” (Crawford, Costa-Gomes, and Iriberry 2013) or “cognitive hierarchy” (Camerer, Ho, and Chong 2004) models. In those models, players’ decisions are optimal, given beliefs that are anchored in a simple model of other players’ instinctive reactions to the game and then adjusted via a small number (k) of iterated best responses. Other models of strategic thinking relax Nash equilibrium in favor of quantal response

equilibrium (QRE) (McKelvey and Palfrey 1995), in which individual players’ decisions are noisy with a specified distribution, and each player’s decision is a noisy best response to the other players’ decision distributions.³

Harstad and Selten’s article in this forum performs a valuable service by highlighting the dominance of optimization-based models over boundedly rational models in modern microeconomics, and questioning whether optimization-based models are a better way forward than boundedly rational models. In their words, “Yet [the optimization approach] can no longer be firmly defended as an appropriate first approximation to reality. Increasingly persuasive evidence has accumulated that the behavioral assumptions underlying the optimization approach are incorrect, and can point analysis in seriously wrong directions. As neoclassical theory has been more clearly elaborated, its serious deficiencies have become increasingly apparent” (2). As Harstad and Selten explain (footnote 1), their use of “neoclassical” here includes “behavioral” optimization-based models.

Although Harstad and Selten frame their proposed reorientation of microeconomics modestly, describing their article as “. . . intended much more as a welcome mat than as a critical commentary on the current state of affairs” (2), the questions they raise pose a challenge to the optimization-based approach that is of central importance for the future of microeconomics.

Rabin’s article in this forum responds to Harstad and Selten, focusing mainly on the questions they raise regarding modeling individual decisions. Rabin argues that, although microeconomics can indeed sometimes be improved by deviating from optimization-based models, maintaining optimization but

² I don’t discuss “rational learning” models, which assume equilibrium in the game that describes the entire learning process, here because adaptive learning models tend to be much more useful empirically.

³ McKelvey and Palfrey (1995) suggest using QRE for initial responses as well as steady states; but some researchers suggest reserving QRE to describe steady states.

relaxing customary neoclassical assumptions about preferences or judgment in specific, evidence-based ways—the approach of most modern behavioral economics—may be more productive in many applications than would a comprehensive switch to boundedly rational models whose connection to evidence is less direct.

The present article responds to Harstad and Selten, focusing on modeling strategic behavior. Like Harstad and Selten and most others who apply game theory in microeconomics, I adopt the viewpoint of noncooperative game theory, which starts with a complete model of the strategic situation and seeks to characterize how players respond to it.⁴ Because noncooperative games are simply interdependent individual decision problems, Harstad and Selten's and Rabin's points about modeling individual decisions apply with equal force to games. Even so, I will follow the usual division of labor in microeconomics (behavioral or neoclassical) by focusing on the issues Harstad and Selten raise that are unique to games.

Recall that the canonical neoclassical descriptive model of strategic behavior is Nash equilibrium, defined as a combination of strategies, one for each player, such that each player's strategy maximizes his expected utility or "payoff," given the others' strategies (Myerson 1999).

The assumption that players will play their Nash equilibrium (henceforth sometimes shortened to "equilibrium" when the meaning is clear) strategies can be justified in two ways (Crawford, Costa-Gomes, and Iriberry 2013, section 1). If players have enough

experience with closely analogous games, both theory and experimental results suggest that players will usually learn to predict each other's strategy choices well enough that their beliefs about others' strategy choices converge to some Nash equilibrium in the game currently being played.⁵ With such experience, equilibrium can be viewed as a model of players' steady-state behavior. If, however, the current game has only imperfect precedents, or none at all, then if assuming equilibrium is justified, it must be justified as a model of strategic thinking in players' initial responses to a game. Such thinking can lead players to a Nash equilibrium if they are rational in the decision-theoretic sense of optimizing expected utilities *and* if they have the same, "rational" (in the sense of self-confirming) expectations or beliefs about how players will play (Brandenburger 1992).⁶

Some of the notions that Harstad and Selten and I discuss refer to steady-state behavior in settings where learning is plausible, while others refer to strategic thinking in initial responses.

Harstad and Selten illustrate the need for boundedly rational models of strategic behavior and their potential benefits by giving several examples of observed phenomena that appear to resist Nash equilibrium explanations. As possible partial remedies for these apparent failings of an optimization-based approach, they also propose two illustrative models of strategic behavior.

Probably the most important of Harstad and Selten's examples is the occurrence of bubbles and crashes in financial markets

⁵ This statement omits some qualifications that are important only for extensive-form games.

⁶ Equilibrium yields much more precise predictions than assuming only that players optimize, given some beliefs. Without requiring beliefs to be self-confirming, even common knowledge that players optimize implies only that their strategies are rationalizable (Bernheim 1984; Pearce 1984), often leaving behavior completely unrestricted.

⁴ The alternative would be cooperative game theory, which sidesteps the need to specify the details of the strategic situation by assuming that whatever the details, players reach a Pareto-efficient agreement. An important exception to the dominance of noncooperative game theory in microeconomics is the (mostly) cooperative theory of matching markets, but the modeling issues matching raises do not figure directly in this forum's discussion.

(their section 4). Bubbles and crashes cannot occur in a conventional rational-expectations equilibrium (Milgrom and Stokey 1982; Tirole 1982). Yet they do occur in financial market experiments even when subjects are fully informed, so that informational inferences from others' decisions are not relevant (Smith, Suchanek, and Williams 1988; Lei, Noussair, and Plott 2001). On that basis, Harstad and Selten argue that bubbles and crashes cannot be explained by a model that is entirely optimization-based.

Another of Harstad and Selten's examples (section 4) is the "winner's curse," in which bidders in common-value auctions fail to adjust their value estimates for the information that would be revealed, via the logic of Nash equilibrium, if they won.⁷ Winning reveals that the bidder's own estimate was higher than other bidders' estimates and therefore likely to overestimate the common value. Bidders who fail to make this informational inference bid higher than in equilibrium, often losing money if they win (Capen, Clapp, and Campbell 1971; Milgrom and Weber 1982; Samuelson and Bazerman 1985; Selten, Abbink, and Cox 2005). As Harstad and Selten note, the winner's curse is linked to the more general phenomenon Eyster and Rabin (2005) call "cursedness" and Crawford, Costa-Gomes, and Iriberry (2013, section 5) call "informational naiveté," whereby people fail to draw correct inferences from the link between others' private information and their decisions. Through its effects on herding, informational naiveté may also exacerbate bubbles and crashes in financial markets (Eyster and Rabin 2010).

Harstad and Selten's final example (section 4) refers to games with strategy spaces large or complex enough to make a directly optimization-based model implausible.

⁷ In independent-private-value auctions, such informational inferences are not relevant.

Harstad and Selten's first illustrative model (section 7) is a boundedly rational model of players' decisions in games with such large or complex strategy spaces, following Selten, Pittnauer, and Hohnisch (2012) and Arad and Rubinstein (2012). Harstad and Selten frame their model as one of strategic thinking in initial responses, although their approach might also be adapted to model the steady states of a learning process. In their proposed model, players first edit their strategy spaces into a manageable (although theoretically suboptimal) form. Players then follow a simple level- k (Crawford, Costa-Gomes, and Iriberry 2013) or cognitive hierarchy (Camerer, Ho, and Chong 2004) model of strategic thinking, but as applied to the edited game.

Harstad and Selten's second illustrative model (section 7) is Selten's notion of learning direction theory (Selten and Stoecker 1986; Selten and Buchta 1999; Selten, Abbink, and Cox 2005) and its steady-state counterpart, impulse-balance equilibrium (Selten, Abbink, and Cox 2005; Ockenfels and Selten 2005; Selten and Chmura 2008; Brunner, Camerer, and Goeree 2011; Selten, Chmura, and Goerg 2011). Impulse-balance equilibrium applies to games played with the ample precedents that make simple learning models plausible, and is thus an alternative to a Nash equilibrium or QRE model of steady states.

The rest of this article is organized as follows.

In section 2, I discuss impulse-balance equilibrium and learning direction theory, focusing on independent-private-value auctions, in which informational inferences from others' decisions are not relevant; and on simple matrix games. I compare the informational requirements and predictive performance of impulse-balance equilibrium with those of other steady-state notions, Nash equilibrium, and QRE; and with those of level- k and cognitive hierarchy models of

strategic thinking. Here the challenge is not to an optimization-based approach per se, but to find a way to extend impulse-balance equilibrium's blend of boundedly rational and optimization-based insights to a class of games more nearly as general as those to which Nash equilibrium applies.

In section 3, I discuss Harstad and Selten's example of bubbles and crashes in financial markets. I note that the results ruling out bubbles and crashes depend on extreme common-knowledge assumptions, over and above the assumption that people optimize. Relaxing common knowledge is empirically at least as plausible as relaxing optimization, nor is it hard to find optimization-based models that are consistent with the occurrence of bubbles and crashes (Barberis and Thaler 2003; Brunnermeier and Oehmke 2013; Xiong forthcoming). The challenge is to find a model (optimization-based or not) whose behavioral assumptions are well-grounded in evidence, and which is not merely consistent with bubbles and crashes, but which reliably explains their observed fact patterns. I close by suggesting one possible candidate for such a model, which is optimization-based but reflects evidence-based limitations on thinking.

In section 4, I turn to Harstad and Selten's example of the winner's curse in sealed-bid common-value auctions, where informational inferences from others' decisions are essential to optimal bidding (equilibrium or not). I also consider the more general phenomenon of informational naiveté in games where informational inferences are important. I compare Selten, Abbink, and Cox's (2005) explanation of the winner's curse in terms of learning direction theory and impulse-balance equilibrium, with optimization-based explanations via Eyster and Rabin's (2005) steady-state notion of "cursed equilibrium" in common-value auctions, and Crawford and Iriberry's (2007) level- k model of initial responses to common-value and

independent-private-value auctions. I next discuss Brocas et al.'s (2010) experimental and theoretical level- k analysis of zero-sum betting, for which a game-theoretic analogue of Milgrom and Stokey's (1982) "no-trade" theorem for market equilibria holds. I then discuss Eyster and Rabin's (2010) analysis of naive herding.

In section 5, I take up Harstad and Selten's first illustrative model, of behavior in games with strategy spaces large and complex enough to render optimization-based models implausible. I put Harstad and Selten's proposed boundedly rational model into context, and sketch the issues that must be resolved to enable further progress along these lines.

In section 6, I conclude with some general remarks.

2. *Learning Direction Theory, Impulse-Balance Equilibrium, and Steady States in Independent-Private-Value Auctions and Simple Matrix Games*

In this section, I discuss Selten's notion of learning direction theory (Selten and Stoecker 1986; Selten and Buchta 1999; Selten, Abbink, and Cox 2005) and its steady-state counterpart impulse-balance equilibrium (Selten, Abbink, and Cox 2005; Ockenfels and Selten 2005; Selten and Chmura 2008; Brunner, Camerer, and Goeree 2011; Selten, Chmura, and Goerg 2011). I focus on independent-private-value auctions, in which informational inferences from others' decisions are not relevant, and simple matrix games.

Learning direction theory is a simple model of learning in settings that permit direct learning from experience. A player is assumed to observe not only his own decision and its realized payoff, but also to have enough information about the game to assess the payoffs that would have resulted from alternative decisions. The feasible

decisions are assumed to be ordered along a single dimension, so that directions are well-defined. Learning direction theory then assumes that if a higher (respectively, lower) decision than the one actually chosen would have yielded a higher payoff, the player adjusts his decision upward (downward).

Learning direction theory's informational requirements are very close to those of the "what if?" reasoning that underlies beliefs-based learning, and far from the requirements of reinforcement learning, in which players simply react to the realized payoffs of their chosen actions, with no counterfactual reasoning (Selten, Abbink, and Cox 2005, 7). Learning direction theory can be viewed as a form of beliefs-based learning in which players estimate their optimal decisions non-parametrically (Ockenfels and Selten 2005, footnote 11). See also Crawford (1995), who takes a similar approach, borrowed from the engineering adaptive control literature via Woodford (1990), to model learning in coordination games.

Learning direction theory is qualitative in that it determines bounds on players' adjustments without specifying their magnitude. Impulse-balance equilibrium, by contrast, gives a quantitative characterization of learning direction theory's steady-state rest points, whose greater precision is helpful in applications. In Harstad and Selten's words (footnote 23), an impulse-balance equilibrium results when the chosen decision (stochastically) balances three impulses: "[a] when a higher action could have led to a better payoff, there is an impulse to adjust the chosen action upward, [b] when a lower action could have led to a better payoff, there is an impulse to adjust the chosen action downward, and [c] when an action led to a negative payoff that a lower action could have avoided, there is an impulse to adjust the chosen action downward." Harstad and Selten also argue that, if the three impulses [a], [b], and [c] are "treated as a priori

equally strong, impulse balance equilibrium can be considered a parameter-free model." However, the impulse weights play a role akin to loss aversion, and are sometimes estimated (Ockenfels and Selten 2005). With that interpretation, impulse balance equilibrium is equivalent to Nash equilibrium in a game with payoffs transformed to reflect loss aversion. There is then no a priori reason why the weights should be equal, but equal weights correspond to a loss aversion coefficient of 2, close to estimates from other settings (Brunner, Camerer, and Goeree 2011).

The informational requirements of impulse-balance equilibrium are close to those of Nash equilibrium, QRE, and level- k rules. Although Harstad and Selten view learning direction theory and impulse-balance equilibrium as boundedly rational theories, they are also optimization-based in the same sense as other beliefs-based learning models are.

The main empirical issue that models of steady-state behavior in games must address is the fact that people's decision distributions are normally sensitive to out-of-equilibrium payoffs. This sensitivity is usually in decision-theoretically intuitive directions in that increasing one or more of a decision's out-of-equilibrium payoffs makes it more likely to be chosen.

Nash equilibrium rules out such sensitivity to out-of-equilibrium payoffs a priori because they do not affect the expected payoffs of a player who expects others to play their Nash equilibrium strategies with certainty. QRE captures such sensitivity by assuming that a player's decision responds to the distribution of others' noisy decisions, which make out-of-equilibrium payoffs relevant to his expected payoffs. Level- k and cognitive hierarchy models capture sensitivity to out-of-equilibrium payoffs structurally, via a mixture of players who best respond to various nonequilibrium decisions

(Crawford, Costa-Gomes, and Iriberry 2013, sections 2.4–5).

Selten and Chmura (2008), Brunner, Camerer, and Goeree (2011), and Selten, Chmura, and Goerg (2011) ask which of these steady-state notions (and some others that are less relevant here) best describes subjects' behavior in a variety of experimental settings, including first- or second-price sealed-bid independent-private-value auctions, in which informational inferences from others' decisions are not relevant; and simple asymmetric matching pennies games.

In first- or second-price auctions with independent private values, experimental subjects have a strong, robust tendency to overbid, relative to the standard of the risk-neutral Bayesian Nash equilibrium (Kagel and Levin 1993; Ockenfels and Selten 2005). This tendency persists even with ample opportunity to learn. By contrast, there is no corresponding tendency to overbid in English auctions, the progressive analogues of second-price sealed-bid auctions.

Building on Kagel, Harstad, and Levin's (1987) suggestion, Harstad (2000) suggests an explanation of overbidding in second-price auctions with independent private values in the spirit of learning direction theory and impulse-balance equilibrium. Ockenfels and Selten (2005) suggest a similar impulse-balance equilibrium explanation of overbidding in first-price auctions.

If impulse-balance equilibrium is optimization-based, how can it explain systematic overbidding in second-price independent-private-value auctions, where bidding one's true value is a weakly dominant strategy? The key, in first-price as in second-price auctions, is the asymmetry of the out-of-equilibrium payoff feedback that players receive when their bids are higher or lower than their equilibrium bids. In second-price auctions, as Harstad (2000, 262) puts it, "When a subject loses money because both their bid and the second-highest bid exceed their value,

this is negative feedback suggesting they may be bidding too aggressively. However, in a second-price auction, a subject might overbid, win, and still make money: it may happen that no rival bids between his overbid and his value. Such an occurrence may be viewed (mistakenly) as positive feedback, serving to offset the negative feedback if overbids are not too large. No corresponding occurrence is possible in an English auction, as the question of whether to overbid does not really arise until the price has reached a subject's value and competitors remain."

In first-price auctions, the feedback asymmetry is different, but Ockenfels and Selten (2005) show that impulse-balance equilibrium also explains their subjects' tendency to overbid.

Thus, impulse-balance equilibrium provides a simple explanation of the systematic, persistent patterns of deviation from Bayesian Nash equilibrium in independent-private-value sealed-bid auctions, via the asymmetric effects of out-of-equilibrium payoffs. In second-price auctions, such effects can even explain systematic deviations from a dominant strategy.

Turning to asymmetric matching pennies matrix games, the comparative statics of the unique Nash equilibrium's mixed-strategy probabilities with respect to changing a payoff in both player roles (thus preserving the zero-sum property) are decision-theoretically intuitive in one player role, in that the "improved" pure strategy is played with higher probability; but counterintuitive in the other player role (Crawford and Smallwood 1984; Crawford, Costa-Gomes, and Iriberry 2013, section 4). This comparative statics implication has been tested in experiments, but usually only by varying only one player's payoff for one pure-strategy combination across treatments. In that case, the comparative statics are analogous because, in these games, each player's equilibrium mixed strategy is determined by the other's payoffs; but

the comparative statics may then be only weakly intuitive or counterintuitive.

In contrast to the theoretical results, experimental subjects' aggregate responses to such payoff changes across treatments tend to be at least weakly intuitive in both player roles; and Nash equilibrium therefore predicts poorly. QRE, impulse-balance equilibrium, and level- k or cognitive hierarchy models all avoid the counterintuitive implication. As a result, QRE and impulse-balance equilibrium fit approximately equally as well, and significantly better than Nash equilibrium (Brunner, Camerer, and Goeree 2011, section 2).⁸

Given Nash equilibrium's lack of sensitivity to out-of-equilibrium payoffs, it seems remarkable that a steady-state concept such as impulse-balance equilibrium can avoid the counterintuitive implications of Nash equilibrium and track subjects' observed sensitivity to out-of-equilibrium payoffs, without directly assuming that players best respond to decision noise.

However, more work is needed to learn whether impulse-balance equilibrium's fit reflects its accuracy in describing the structure of people's adjustments in response to feedback, or is simply due to the fact that in these games impulse-balance equilibrium mimics the effects of loss aversion. If the latter, loss aversion could easily be incorporated into a Nash equilibrium model, which would then fit approximately as well as impulse-balance equilibrium in this setting (Brunner, Camerer, and Goeree 2011, section 2).

More work is also needed to clarify whether the explanation for overbidding in second-price auctions is really as mechanical as impulse-balance equilibrium's success

suggests, or if the overbidding stems from deeper cognitive limits such as an inability or unwillingness to reason contingent on future events.

A final challenge is to find a way to define a notion like impulse-balance equilibrium for games whose feasible decisions are not ordered along a single dimension, so that its blend of boundedly rational and optimization-based insights can be extended to a class of games more nearly as general as that to which Nash equilibrium applies. If that problem proves to have a solution, I conjecture that it will be found by thinking of how players in such games might nonparametrically estimate their *optimal* decisions, adaptively or in steady state.

3. *Bubbles and Crashes in Financial Markets*

In this section, I discuss Harstad and Selten's example of bubbles and crashes in financial markets. Bubbles and crashes cannot occur in a conventional rational-expectations equilibrium (Milgrom and Stokey 1982; Tirole 1982). Yet they do occur in experiments whose subjects are fully informed (Smith, Suchanek, and Williams 1988; Lei, Noussair, and Plott 2001). On that basis, Harstad and Selten argue that bubbles and crashes cannot be explained by a model that is entirely optimization-based. In their words, bubbles and crashes in such markets are ". . . at odds with behavioral economics and level- k models just as [. . .] with mainstream microtheory" (8).

However, the results ruling out bubbles and crashes depend on the extreme common-knowledge assumptions that underlie the notion of a rational-expectations equilibrium, over and above the assumption that people optimize. Barberis and Thaler (2003), Brunnermeier and Oehmke (2013), and Xiong (forthcoming) survey models that have the potential to explain bubbles and crashes,

⁸ Brunner, Camerer, and Goeree did not consider level- k or cognitive hierarchy models, which as models of initial responses rather than steady states are less relevant here. But other work, on initial responses, suggests that level- k or cognitive hierarchy models would also fit approximately as well as QRE and impulse-balance equilibrium.

some of which are optimization-based but allow people to have diverse beliefs.

The challenge is to find a model (optimization-based or not) whose behavioral assumptions are well-grounded in evidence; and which is not merely consistent with bubbles and crashes, but makes specific predictions that help to explain their observed empirical regularities. In settings as rich as financial markets, specific predictions require a specific structure. I now sketch an illustrative optimization-based candidate model that has the potential to explain the empirical regularities of bubbles and crashes, by relaxing common-knowledge assumptions in favor of a structural model with diverse beliefs, which seems to accord with evidence.⁹

There is considerable evidence, going back to Selten and Stoecker's (1986) finitely repeated Prisoner's Dilemma experiments, that few people consider the future consequences of their current decisions far enough ahead to support the epistemic rational-expectations argument on which the no-bubbles results depend. The finitely repeated Prisoner's Dilemma has a unique Nash equilibrium (subgame-perfect or not), in which players defect in every period, without regard to the history of play. Yet Selten and Stoecker (1986) find that most subjects cooperate until close to the end of the ten-period horizons of their Prisoner's Dilemma games; that the timing of first defection varies across subjects, but defection tends to persist after a subject's first one; and that more and more subjects defect as the end comes closer. These patterns suggest a model in which people consider the future consequences of their current decisions only a small, heterogeneous

number of periods ahead, perhaps due to cognitive limitations, or to the perception that others have limitations, etc.

Finally, Selten and Stoecker (1986) find that, when the entire ten-period game is repeated, the timing of defections "unravels," with subjects starting to defect earlier and earlier. This last pattern is what originally suggested learning direction theory to Selten and Stoecker.

Johnson et al. (2002) find similar patterns of heterogeneity in the extent to which subjects considered the future consequences of their current decisions in experiments designed to elicit subjects' initial responses to three-period alternating-offers bargaining games. They explain their results via an explicit structural model, in which subjects "think ahead" a small, heterogeneous number of periods, following optimization-based rules that vary according to how many periods ahead a player thinks.

A similar optimization-based model of financial markets would have firmer behavioral foundations, and might help to explain the empirical regularities of bubbles and crashes. Although boundedly rational explanations are also possible, I am unaware of any evidence that would guide the choice among the many possible boundedly rational models one could specify.

4. *The Winner's Curse and Informational Naiveté*

In this section, I turn to Harstad and Selten's example of the winner's curse in common-value auctions, where informational inferences from others' decisions are essential to optimal bidding, equilibrium or not (Capen, Clapp, and Campbell 1971; Milgrom and Weber 1982; Samuelson and Bazerman 1985; Selten, Abbink, and Cox 2005). I also consider the more general phenomenon of informational naiveté in games like the "acquiring a company" game where

⁹ The model I have in mind is roughly similar to Morris, Postlewaite, and Shin's (1995) model, but replaces their epistemic arguments with a simpler model of behavior that I believe has stronger behavioral foundations. See also Abreu and Brunnermeier (2003). Harstad and Selten do not propose a specific boundedly rational model of bubbles and crashes.

informational inferences are important, which Harstad and Selten argue underlies much of the high volume of speculative trade in financial markets (Samuelson and Bazerman 1985; Charness and Levin 2009; Crawford, Costa-Gomes, and Iriberry 2013, section 5.3).

Selten, Abbink, and Cox (2005) propose an explanation of the winner's curse in terms of learning direction theory and impulse-balance equilibrium, much as Harstad (2000) and Ockenfels and Selten (2005) use them to explain overbidding in second- and first-price auctions with independent private values. Yet it remains unclear whether the curse is really a mechanical consequence of impulse balance, or if it stems instead from deeper facts about human cognition, such as an inability or unwillingness to reason contingent on future events.

As Harstad and Selten acknowledge, Eyster and Rabin (2005) give a credible optimization-based steady-state account of the winner's curse via their notion of "cursed equilibrium," whereby individuals do not fully attend to the possible correlations between others' decisions and others' private information, but otherwise follow the logic of Bayesian Nash equilibrium. Cursed equilibrium is a plausible candidate for a model of steady states if people really do not attend to the relation between others' decisions and private information. And Crawford and Iriberry (2007) use a level- k model in which players anchor their beliefs in a "level-0" that randomizes independently of its own private information to give a similar account of informationally naive initial responses in independent-private- and common-value auctions.

In other contributions, Brocas et al. (2010) report experiments on a zero-sum betting design, for which a game-theoretic analogue of Milgrom and Stokey's (1982) "no-trade" theorem for market equilibria holds. They show that their results are well explained

by a level- k model like that Crawford and Iriberry (2007) proposed (see also Crawford, Costa-Gomes, and Iriberry 2013, sections 2.4 and 5.1). Eyster and Rabin (2010) conduct an optimization-based theoretical analysis of "naive herding," in which people make informational inferences from others' decisions using rules that are informationally naive in ways that resemble Crawford and Iriberry's (2007) "random level-2" rules. Their results accord with intuitions about how herding works in practice better than those of the standard neoclassical analyses of herding.

In most settings, the winner's curse and informational naiveté seem to be reasonably well accounted for by optimization-based models. But cursed equilibrium and level- k models fit experimental data poorly for games like acquiring a company (Charness and Levin 2009), which suggests that closer attention to cognitive limitations is needed. Neither optimization-based nor boundedly rational models have yet to give a fully satisfactory account of cognitive limitations in such settings, which is a challenge and an opportunity for both approaches.

5. *Games with Large and Complex Strategy Spaces*

In this section, I take up Harstad and Selten's first illustrative model of behavior in games with strategy spaces large and complex enough to render optimization-based models implausible.

Harstad and Selten organize their discussion of this issue around Arad and Rubinstein's (2012) experimental and theoretical analysis of "Colonel Blotto" games, where the strategy spaces are so large that it is difficult to imagine anyone optimizing in them. An even more striking illustration is implicit in Ewerhart's (2000) analysis of chess. Ewerhart shows that chess can theoretically be solved by applying two rounds of iterated elimination of weakly dominated

strategies. Thus, by the usual measures, chess has an exceedingly simple strategic structure. This highlights that much of the game's difficulty is due to its enormous and complex strategy spaces.

Although this conundrum could be posed with regard to modeling the steady states of behavior in games where learning from experience with closely analogous games is possible, Harstad and Selten pose it with regard to modeling strategic thinking in initial responses.¹⁰

Harstad and Selten sketch a behavioral model of strategic thinking following Selten, Pittnauer, and Hohnisch (2012) and Arad and Rubinstein (2012). Players first edit their strategy spaces into a few manageable dimensions, ignoring other aspects of their strategies. Players then follow a thinking model such as level- k (Crawford, Costa-Gomes, and Iriberry 2013) or cognitive hierarchy (Camerer, Ho, and Chong 2004), but dimension by dimension in the edited game. The behavioral model is thus a hybrid of boundedly rational and optimization-based approaches.

Harstad and Selten's proposal is a promising approach to the difficult problem of modeling people's strategic thinking in realistically complex applications. As they are aware, much more work is needed before it can be applied with the precision, generality, and portability across games that have made Nash equilibrium models so competitive. The most important gap regards the principles that govern players' editing. It might be possible to close their model by adding a theory of analogies, perhaps as in Mullainathan (2000), Gabaix (2011), or Samuelson (2001).¹¹ To me it seems no

accident that those theories of analogies all blend small measures of bounded rationality with large doses of optimization.

6. Conclusion

Harstad and Selten's article in this forum performs a valuable service by highlighting the dominance of optimization-based models over boundedly rational models in modern microeconomics, and questioning whether optimization-based models are a better way forward than boundedly rational models.

This article responds to Harstad and Selten's proposed reorientation of microeconomics, focusing on modeling strategic behavior. I comment on their examples of observed phenomena that appear to resist Nash equilibrium explanations, and the models of strategic behavior they propose as possible partial remedies. I try to put the models that Harstad and Selten propose into a broader context, to highlight the mostly optimization-based work that has already been done to address their examples, and to identify the challenges that remain and some ways forward.

Nash equilibrium dominates the analysis of strategic behavior in the social sciences because it has important advantages of tractability, precise predictions, generalizability, and portability across games (Myerson 1999). A successful competitor to Nash equilibrium must give a more accurate account of strategic behavior, implicitly or explicitly by identifying and modeling systematic deviations from Nash equilibrium, while preserving most of its advantages.

Of the several challengers to Nash equilibrium, Harstad and Selten and I discuss, level- k and cognitive hierarchy models, my sketched model of limited thinking ahead in financial markets (section 3), and Harstad and Selten's sketched model of strategic thinking in games with large or complex strategy spaces (section 5) all maintain the

¹⁰ If it is desired to model steady-state behavior instead of strategic thinking, genetic algorithms (Holland 1992), which model evolution in complex strategy spaces, might yield insights into how people learn in complex games.

¹¹ Jehiel (2005) suggests a way to integrate a model of analogies into a general model of strategic behavior.

assumption that players optimize. Learning direction theory and impulse-balance equilibrium are optimization-based in the weaker sense that they use optimization as the motivation for players' strategy adjustments without directly imposing it on their strategy choices. QRE and cursed equilibrium are also optimization-based, while maintaining rational expectations in different senses.

It is noteworthy that all of these models, including those that Harstad and Selten propose and those I propose or reference in response to their examples, are to some degree optimization-based, though often with elements of bounded rationality. Although Harstad and Selten might well view the predominance of optimization-based models, even in a forum on bounded rationality, as simply proving their point, I think it follows from the logic of what we are all trying to accomplish, and is therefore not a bad thing.

In most settings, there is an enormous number of logically possible models, optimization-based or boundedly rational, that deviate from neoclassical models. In attempting to improve upon neoclassical models, it is essential to have some principled way of choosing among alternatives. Modern behavioral economics and behavioral game theory guide model selection by grounding assumptions firmly in experimental and empirical evidence on preferences, judgment, learning, strategic thinking, and cognition more generally.

Of course, the direct characterizations of individual behavior on which boundedly rational models are built also reflect evidence on cognition and behavior. But to move forward, those who advocate bounded rationality modeling must find comparably convincing, evidence-based ways to choose among the equally enormous number of possible nonoptimizing models. This, I think, will prove a more difficult task than finding empirical support for optimization-based

deviations from neoclassical models, because the benchmark that optimization provides seems to aid the evidence-gathering process.

To improve on a neoclassical model, one must identify *systematic* deviations; otherwise one would do better to stick with a noisier neoclassical model. Behavioral decision theory is built on such systematic deviations—the empirical regularity that a great majority of people are either “neoclassical” or present-biased, very rarely future-biased; the regularity that a great majority of people are either neoclassical or loss-averse, very rarely gain-averse; and so on. These behavioral “biases” would likely have been invisible to researchers without a neoclassical, optimization-based benchmark to measure them against.

Behavioral game theory is also built on systematic deviations. To take an illustration that is connected with many of the issues that Harstad and Selten and I discuss, level- k and cognitive hierarchy models of strategic thinking rest on the empirical regularity that people seem unable or unwilling to do the fixed-point or indefinitely iterated reasoning that a thinking justification for Nash or Bayesian equilibrium often requires. The experimental and some field evidence suggests that people are therefore often driven to use level- k rules that anchor beliefs in a simple model of others' instinctive reactions to the game and then adjust them via a small number (k) of iterated best responses (Costa-Gomes and Crawford 2006; Crawford, Costa-Gomes, and Iriberry 2013, section 3). It seems unlikely that researchers would have looked for evidence on people's ability or willingness to do fixed-point or indefinitely iterated reasoning if it did not figure in the theory of Nash equilibrium. Yet the avoidance of such reasoning has important implications for strategic behavior, and using a specific behavioral model yields insights in settings where Nash equilibrium does not describe observed

behavior (Crawford, Costa-Gomes, and Iriberri 2013).

For all of these reasons, it seems to me that the most promising way forward is to learn more about how best to specify optimization-based models, and how to use them in applications. That said, Harstad and Selten are right to point out that the things optimization-based models get wrong leave them open to challenge by boundedly rational models. The reciprocal challenge for strategic modeling is to find evidence-based boundedly rational models that have the precision, generalizability, and portability to be worthy competitors to Nash equilibrium.

REFERENCES

- Abreu, Dilip, and Markus K. Brunnermeier. 2003. "Bubbles and Crashes." *Econometrica* 71 (1): 173–204.
- Allen, Franklin, Stephen Morris, and Andrew Postlewaite. 1993. "Finite Bubbles with Short Sale Constraints and Asymmetric Information." *Journal of Economic Theory* 61 (2): 206–29.
- Andreoni, James, and John Miller. 2002. "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica* 70 (2): 737–53.
- Arad, Ayala, and Ariel Rubinstein. 2012. "Multi-dimensional Iterative Reasoning in Action: The Case of the Colonel Blotto Game." *Journal of Economic Behavior and Organization* 84 (2): 571–85.
- Armstrong, Mark, and Steffen Huck. 2010. "Behavioral Economics as Applied to Firms: A Primer." *Competition Policy International* 6 (1): 3–45.
- Barberis, Nicholas, and Richard Thaler. 2003. "A Survey of Behavioral Finance." In *Handbook of the Economics of Finance. Volume 1B. Financial Markets and Asset Pricing*, edited by George M. Constantinides, Milton Harris, and René Stulz, 1053–1123. Amsterdam; London and New York: Elsevier, North Holland.
- Bernheim, B. Douglas. 1984. "Rationalizable Strategic Behavior." *Econometrica* 52 (4): 1007–28.
- Bolton, Gary E., and Axel Ockenfels. 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review* 90 (1): 166–93.
- Brandenburger, Adam. 1992. "Knowledge and Equilibrium in Games." *Journal of Economic Perspectives* 6 (4): 83–101.
- Brocas, Isabelle, Juan D. Carrillo, Stephanie W. Wang, and Colin F. Camerer. 2010. "Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games." <http://www.hss.caltech.edu/~sweiwang/papers/Betting.pdf>.
- Brunner, Christoph, Colin F. Camerer, and Jacob K. Goeree. 2011. "Stationary Concepts for Experimental 2 x 2 Games: Comment." *American Economic Review* 101 (2): 1029–40.
- Brunnermeier, Markus K., and Martin Oehmke. 2013. "Bubbles, Financial Crises, and Systemic Risk." In *Handbook of the Economics of Finance, Volume 2, Part B*, edited by George M. Constantinides, Milton Harris, and René M. Stulz, 1221–88. Amsterdam and Boston: Elsevier, North-Holland.
- Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press; New York: Russell Sage Foundation.
- Camerer, Colin F., and Teck-Hua Ho. 1999. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica* 67 (4): 827–74.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. 2002. "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games." *Journal of Economic Theory* 104 (1): 137–88.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics* 119 (3): 861–98.
- Capen, Edward C., Robert W. Clapp, and William M. Campbell. 1971. "Competitive Bidding in High-Risk Situations." *Journal of Petroleum Technology* 23 (6): 641–53.
- Charness, Gary, and Dan Levin. 2009. "The Origin of the Winner's Curse: A Laboratory Study." *American Economic Journal: Microeconomics* 1 (1): 207–36.
- Charness, Gary, and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117 (3): 817–69.
- Conlisk, John. 1996. "Why Bounded Rationality?" *Journal of Economic Literature* 34 (2): 669–700.
- Costa-Gomes, Miguel A., and Vincent P. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review* 96 (5): 1737–68.
- Crawford, Vincent P. 1995. "Adaptive Dynamics in Coordination Games." *Econometrica* 63 (1): 103–43.
- Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. 2013. "Structural Models of Non-equilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51 (1): 5–62.
- Crawford, Vincent P., and Nagore Iriberri. 2007. "Level- k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica* 75 (6): 1721–70.
- Crawford, Vincent P., and Dennis Smallwood. 1984. "Comparative Statics of Mixed-Strategy Equilibria in Noncooperative Two-Person Games." *Theory and Decision* 16 (3): 225–32.
- Cyert, Richard M., and James G. March. 1963. *A Behavioral Theory of the Firm*. New York: Prentice Hall.
- DellaVigna, Stefano. 2009. "Psychology and Economics:

- Evidence from the Field." *Journal of Economic Literature* 47 (2): 315–72.
- Erev, Ido, and Alvin E. Roth. 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88 (4): 848–81.
- Ewerhart, Christian. 2000. "Chess-like Games Are Dominance Solvable in at Most Two Steps." *Games and Economic Behavior* 33 (1): 41–47.
- Eyster, Erik, and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica* 73 (5): 1623–72.
- Eyster, Erik, and Matthew Rabin. 2010. "Naive Herding in Rich-Information Settings." *American Economic Journal: Microeconomics* 2 (4): 221–43.
- Fehr, Ernst, and Simon Gächter. 2000. "Fairness and Retaliation: The Economics of Reciprocity." *Journal of Economic Perspectives* 14 (3): 159–81.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–68.
- Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40 (2): 351–401.
- Fudenberg, Drew, and David K. Levine. 1998. *The Theory of Learning in Games*. Cambridge and London: MIT Press.
- Gabaix, Xavier. 2011. "A Sparsity-Based Model of Bounded Rationality." National Bureau of Economic Research Working Paper 16911.
- Harstad, Ronald M. 2000. "Dominant Strategy Adoption and Bidders' Experience with Pricing Rules." *Experimental Economics* 3 (3): 261–80.
- Harstad, Ronald M., and Reinhard Selten. 2013. "Bounded-Rationality Models: Tasks to Become Intellectually Competitive." *Journal of Economic Literature* 51 (2): 496–511.
- Hogarth, Robin M., and Melvin W. Reder, eds. 1986. *Rational Choice: The Contrast between Economics and Psychology*. Chicago and London: University of Chicago Press.
- Holland, John H. 1992. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, Second edition. Cambridge and London: MIT Press.
- Jehiel, Philippe. 2005. "Analogy-Based Expectation Equilibrium." *Journal of Economic Theory* 123 (2): 81–104.
- Johnson, Eric J., Colin F. Camerer, Sankar Sen, and Talia Rymon. 2002. "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining." *Journal of Economic Theory* 104 (1): 16–47.
- Kagel, John H., Ronald M. Harstad, and Dan Levin. 1987. "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica* 55 (6): 1275–1304.
- Kagel, John H., and Dan Levin. 1993. "Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders." *Economic Journal* 103 (419): 868–79.
- Kagel, John H., Dan Levin, and Ronald M. Harstad. 1995. "Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common Value Auctions." *International Journal of Game Theory* 24 (3): 293–319.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–92.
- Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65.
- Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–77.
- Lei, Vivian, Charles N. Noussair, and Charles R. Plott. 2001. "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality." *Econometrica* 69 (4): 831–59.
- McKelvey, Richard D., and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10 (1): 6–38.
- Milgrom, Paul R., and John Roberts. 1990. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica* 58 (6): 1255–77.
- Milgrom, Paul R., and John Roberts. 1991. "Adaptive and Sophisticated Learning in Normal Form Games." *Games and Economic Behavior* 3 (1): 82–100.
- Milgrom, Paul R., and Nancy Stokey. 1982. "Information, Trade and Common Knowledge." *Journal of Economic Theory* 26 (1): 17–27.
- Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–1122.
- Morris, Stephen, Andrew Postlewaite, and Hyun Song Shin. 1995. "Depth of Knowledge and the Effect of Higher Order Uncertainty." *Economic Theory* 6 (3): 453–67.
- Mullainathan, Sendhil. 2000. "Thinking through Categories." Unpublished.
- Munier, Bertrand, et al. 1999. "Bounded Rationality Modeling." *Marketing Letters* 10 (3): 233–48.
- Myerson, Roger B. 1999. "Nash Equilibrium and the History of Economic Theory." *Journal of Economic Literature* 37 (3): 1067–82.
- Nelson, Richard R., and Sidney G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge and London: Harvard University Press.
- Newell, Allen, and Herbert A. Simon. 1972. *Human Problem Solving*. New York: Prentice Hall.
- Ockenfels, Axel, and Reinhard Selten. 2005. "Impulse Balance Equilibrium and Feedback in First Price Auctions." *Games and Economic Behavior* 51 (1): 155–70.
- O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review* 89 (1): 103–24.

- Pearce, David G. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52 (4): 1029–50.
- Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83 (5): 1281–1302.
- Rabin, Matthew. 1998. "Psychology and Economics." *Journal of Economic Literature* 36 (1): 11–46.
- Rabin, Matthew. 2002. "Inference by Believers in the Law of Small Numbers." *Quarterly Journal of Economics* 117 (3): 775–816.
- Rabin, Matthew. 2013. "Behavioral Optimization Models versus Bounded-Rationality Models in Decisions." *Journal of Economic Literature* 51 (2): 528–43.
- Rosenthal, Robert W. 1989. "A Bounded-Rationality Approach to the Study of Noncooperative Games." *International Journal of Game Theory* 18 (3): 273–91.
- Roth, Alvin E., and Ido Erev. 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior* 8 (1): 164–212.
- Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge and London: MIT Press.
- Samuelson, Larry. 2001. "Analogies, Adaptation, and Anomalies." *Journal of Economic Theory* 97 (2): 320–66.
- Samuelson, William F., and Max H. Bazerman. 1985. "The Winner's Curse in Bilateral Negotiations." In *Research in Experimental Economics, Volume 3*, edited by Vernon Smith, 105–37. Greenwich, Conn.: JAI Press.
- Selten, Reinhard. 1990. "Bounded Rationality." *Journal of Institutional and Theoretical Economics* 146 (4): 649–58.
- Selten, Reinhard. 1991. "Anticipatory Learning in Two-Person Games." In *Game Equilibrium Models I: Evolution and Game Dynamics*, edited by Reinhard Selten, 98–154. New York; Berlin; London and Tokyo: Springer.
- Selten, Reinhard. 1998. "Features of Experimentally Observed Bounded Rationality." *European Economic Review* 42 (3–5): 413–36.
- Selten, Reinhard, Klaus Abbink, and Ricarda Cox. 2005. "Learning Direction Theory and the Winner's Curse." *Experimental Economics* 8 (1): 5–20.
- Selten, Reinhard, and Joachim Buchta. 1999. "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, edited by David V. Budescu, Ido Erev, and Rami Zwick, 79–104. Mahwah, N.J.: Lawrence Erlbaum Associates.
- Selten, Reinhard, and Thorsten Chmura. 2008. "Stationary Concepts for Experimental 2x2-Games." *American Economic Review* 98 (3): 938–66.
- Selten, Reinhard, Thorsten Chmura, and Sebastian J. Goerg. 2011. "Stationary Concepts for Experimental 2x2 Games: Reply." *American Economic Review* 101 (2): 1041–44.
- Selten, Reinhard, Sabine Pittnauer, and Martin Hohmann. 2012. "Dealing with Dynamic Decision Problems When Knowledge of the Environment Is Limited: An Approach Based on Goal Systems." *Journal of Behavioral Decision Making* 25 (5): 443–57.
- Selten, Reinhard, and Rolf Stoecker. 1986. "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach." *Journal of Economic Behavior and Organization* 7 (1): 47–70.
- Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69 (1): 99–118.
- Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams. 1988. "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets." *Econometrica* 56 (5): 1119–51.
- Sobel, Joel. 2005. "Interdependent Preferences and Reciprocity." *Journal of Economic Literature* 43 (2): 392–436.
- Spiegler, Ran. 2011a. *Bounded Rationality and Industrial Organization*. Oxford and New York: Oxford University Press.
- Spiegler, Ran. 2011b. "But Can't We Get the Same Thing with a Standard Model? Rationalizing Bounded-Rationality Models." *Economics and Philosophy* 27 (1): 23–43.
- Thaler, Richard H. 1992. *The Winner's Curse: Paradoxes and Anomalies of Economic Life*. Princeton and Chichester, U.K.: Princeton University Press.
- Tirole, Jean. 1982. "On the Possibility of Speculation under Rational Expectations." *Econometrica* 50 (5): 1163–81.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–31.
- Tversky, Amos, and Daniel Kahneman. 1991. "Loss Aversion in Riskless Choice: A Reference-Dependent Model." *Quarterly Journal of Economics* 106 (4): 1039–61.
- Woodford, Michael. 1990. "Learning to Believe in Sunspots." *Econometrica* 58 (2): 277–307.
- Xiong, Wei. Forthcoming. "Bubbles, Crises, and Heterogeneous Beliefs." In *Handbook on Systemic Risk*, edited by Jean-Pierre Fouque and Joseph A. Langsam. Cambridge and New York: Cambridge University Press.
- Young, H. Peyton. 2004. *Strategic Learning and Its Limits*. Oxford and New York: Oxford University Press.

This article has been cited by:

1. Ronald M. Harstad,, Reinhard Selten. 2013. Bounded-Rationality Models: Tasks to Become Intellectually Competitive. *Journal of Economic Literature* 51:2, 496-511. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]