INTRODUCTION TO "GAME THEORY IN THE TRADITION OF BOB WILSON"

Bengt Holmstrom Paul Milgrom
Alvin E. Roth

One of the nicer events in academic life is when we pause to recognize a scholar whose work is unusually important and influential, whose work marks the start of a new tradition.

When that scholar is also a great teacher and advisor, his students have the added pleasure of recalling his influence on them, and how it is reflected both in their own scholarship, and in how they teach and advise their own students. Students are the generations through which traditions are transmitted. This volume of selected published papers by Bob's students, accompanied by new introductory essays, is a celebration of Bob's tradition, by those of us who had the exceptional good fortune to receive it at first hand.

And what is this tradition? Scholarship as varied and wide ranging as Bob's defies easy characterization. He was among the first to recognize that it was going to be of the utmost importance for game theorists to understand how information is distributed and manipulated, concealed, and revealed. He was among the first to emphasize the importance in strategic calculations of players' beliefs about what other players would do, even in situations that were not anticipated to arise. But what especially marks him as a leader among the great economists of his generation is his view of the role of theory. In his understated way, he wrote in the preface to his book *Nonlinear Pricing*: "**The value of theory is its usefulness in addressing practical problems**..." And he went on to reflect on the role of practical problems in his own scholarly development: "...**for the theorist, the problems encountered by practitioners provide a wealth of topics**."

So, game theory in the tradition of Bob Wilson is game theory in the service of economics as a confident, practical, useful discipline. And research in the style of Bob Wilson is work that takes its inspiration not only from a wide reading and deep understanding of the work of other academics, but also from the ordinary stuff of economic life. In this spirit, Bob's work has produced not only acute conceptual insights of great generality, but also advice about and solutions to knotty problems of strategy and design.

1 Bob's Work

Even a brief examination of Bob's long, productive, ongoing research career reveals some distinctive approaches, woven together in a signature style. This makes a chronological telling of his story difficult, as ideas and themes that he explored early in his career are developed throughout. But an indication of the evolution of his interests can be gleaned from the brief author biographies that accompanied his early papers in *Management Science*. In the September 1966 issue, his chief research interest is stated as "statistical decision theory, especially group decisions and applications in finance." In May 1967, his ongoing research activities "involve the economics of uncertainty, including group decision theory". In March 1972, he is "presently doing research in multi-person decision theory, including game theory".

In his 1963 doctoral dissertation, Bob pioneered the iterative use of a Quadratic Programming subproblem to solve general nonlinearly constrained convex programming problems. In Wilson's method, an inequality constrained QP is solved at each iteration, where the quadratic objective function is a quadratic approximation to the Lagrangian function and the linear inequality constraints are local linear approximations of the original nonlinear constraints. The Lagrange multiplier estimates in the objective function are based on the QP multipliers from the previous iteration. Sequential Quadratic Programming Methods (SQP), which are based on Wilson's original idea, are considered today to be the best algorithms for general nonlinear constrained optimization. Although Bob didn't further pursue this particular line of research, his interest in computational algorithms has been an enduring theme throughout his career. He has developed computational methods for bilateral complementarity problems, for general equilibrium problems, for auctions, for electricity markets and for computing Nash equilibria of games. Wherever possible, Bob wants to know what solutions look like and find a way to compute them; the blend of theory and algorithms is a big part of what makes his research unique.

Bob's 1968 paper, "The Theory of the Syndicates," influenced a whole generation of students from economics, finance, and accounting. The paper poses a fundamental question that cuts to the heart of many economic issues: Under what conditions will a group of Savage rational individuals behave like a single Savage rational decision maker (a syndicate)? In other words, when does the expected utility representation describe the behavior of a group of individuals who choose lotteries and share risk in a Pareto optimal way? As a simple example, one can think of a firm, owned and controlled by a group of active managers, who have to decide on an investment. What is the best course of action (in the Pareto sense) for this group? When can that choice be construed as made by a single individual? If the group behaves as a single expected utility maximizing individual, what is the firm's objective function and how does it relate to the individual utility functions?

When individuals hold identical beliefs, a condition the paper considers rare, it is easy to see that group behavior satisfies the expected utility theorem. The Pareto set of lotteries and sharing rules is convex so that each point on the

Pareto frontier is picked out by maximizing some weighted average of individual expected utilities. Because beliefs are shared, the syndicate utility is simply the weighted average of individual utilities. Risk is shared optimally according to Borch's well-known formula, which, when utilities satisfy constant absolute risk aversion, implies sharing rules that are linear.

The main focus of the paper is on the case with different beliefs. With different beliefs, Pareto optimal group behavior rarely satisfies Savage rationality. The only exception is when optimal risk sharing results in linear sharing rules, that is, when utilities exhibit constant absolute risk aversion. The paper views this result as rather discouraging, though now that behavioral economics is staging a comeback, this finding may get more attention. On the other hand, when preferences can be aggregated into a syndicate utility and a syndicate belief, the results are both intuitive and analytically useful. Risk is shared in proportion to individual risk tolerances and differences in beliefs get resolved through natural side-bets. Most interesting is the fact that the syndicate's utility will coincide with each individual's utility induced by the sharing rule. Once risk is shared optimally (including the side-bets), the group is in full agreement on the most desirable course of action; unanimity prevails. Therefore, any given individual can represent the syndicate and decide on its behalf; there is no further need to negotiate over decisions once the shares have been negotiated (this assumes, of course, that all conceivable decisions are envisioned in advance).

The unanimity result bears directly on two important literatures — the agency literature and the literature on the optimal objective function of the firm. In "The structure of incentives for decentralization under uncertainty", published in 1967, Bob presented the first formal analysis of a principal-agent model, using the syndicate analysis to show when an agent can be trusted to act in the interest of a principal. In his paper with Steinar Ekern (in this volume) the intuition of the syndicate theory was used to provide sufficient conditions for unanimity in an incomplete market and to show that the implied market risk coincides with the market syndicate's coefficient of risk aversion.

The influence of the syndicate paper in the 1970s can largely be traced to these two applications. The paper was written at a time when basic asset-pricing models were still being developed. Bob came to these questions from a very different background; the influence of his Harvard Business School training, particularly with Raiffa, is easy to detect. Yet, he saw quickly the connections to the emerging finance literature and in no small way influenced the direction this literature took, particularly through its Stanford students.

With the advent of asymmetric information, adverse selection and common knowledge it became unfashionable to assume that individuals hold different beliefs. For the past two decades the syndicate paper has been living a quieter life. But with behavioral economics coming into fashion, the paper may see a revival of interest. The question of group behavior is fundamental to our understanding of organizations like firms and it is intriguing that belief differences alone can cause a departure from rational behavior. And why not think of individuals as being of many minds?

Bob's first papers on auctions date from around the same time. His earliest

work on Bayesian auction models dates to 1967. That is so early that it predates Harsanyi's famous papers setting the foundation for Bayesian games.

Every area of modern auction theory bears Bob's mark. He was the first to study models with what we now call common values and the first to treat models in which types are not independently distributed. He studied symmetric models and asymmetric ones that represent bidding on drainage tracts, where one bidder is especially well informed. The drainage tract model evolved into the foundation for important empirical work by Hendricks and Porter, confirming in the data some of the most surprising predictions of the game theoretic solution.

While most students of auction theory had limited their attention either to models in which each bidder demands a single good or to models with multiple goods but complete information, Wilson's "Auctions of Shares" broke the mold. It initiated the analysis of the auctions for divisible goods—a model that is useful for evaluating sales of Treasury bills or electric power. It immediately identified the possibility that such auctions would increase the bidders market power, leading to unfavorable terms for the auctioneer. The recent history of power markets shows the importance of that early insight for today's practical reality.

Starting from his early auction work, Bob's game theoretic research brought into clear focus the importance of differential information, and of agents' beliefs. Among Wilson's most famous and influential papers, none is more widely cited than "Sequential Equilibria," which he coauthored with David Kreps. An on-line search shows hundreds of articles that include the phrase "sequential equilibria" or "sequential equilibrium" *in the title of the article*, and there are many more papers that employ the idea.

Before "Sequential Equilibria," economists had struggled to explain why certain behaviors occurred at Nash equilibrium, sometimes referring to what players might think or believe as the game was being played, but such discussions were usually limited both by being informal and by focusing solely on play along the equilibrium path. "Sequential Equilibria" changed that by making the beliefs that were the underpinning of players' choices part of the formal description of play, both on and off the equilibrium path. The new concept found immediate applications in bargaining theory, contract theory, and theoretical industrial organization and labor economics. It also helped to inspire developments in pure game theory, as further attempts were made to refine equilibrium concepts or to criticize the whole notion of equilibrium as involving strong—and now explicit—assumptions about what the players believe.

Among the first applications of the sequential equilibrium were Wilson's own, including his papers about the reputation phenomenon. Two coauthored papers, "Reputation and Imperfect Information" and "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," revealed the sensitivity of certain equilibrium calculations to extremely precise information about the payoffs or strategies of other players. In the Chain Store game, if the probability that the incumbent store is irrationally over-committed to defending its markets is $\epsilon > 0$, then if the game is repeated more than $-\alpha \cdot \log(\epsilon)$ times, where α depends on the payoffs, then initial plays of the game will differ from the full information

equilibrium prediction with probability *one*.

These celebrated and oft-reprinted papers pose important challenges both for game theory itself and for its applications in economics. In terms of applied economics, the notion that equilibrium predictions in some dynamic games could depend sensitively on beliefs was a challenge to the whole underpinning of industrial organization. The traditional structure-conduct-performance paradigm had tried to derive predictions about conduct from the supply and demand structure. This involved data about costs, scale economies, demand elasticities, the segmentation of demand, and so on. Now, Wilson and his fellow game theorists were saying that wasn't enough. Practitioners began to worry that barely discernible changes in beliefs could have dramatic effects on equilibrium behavior. The most pessimistic practitioners began to say that a game theorist could reach any conclusion at all by simply playing with minor assumptions. Of course, this critique misses the main point: when outcomes depend sensitively on expectations, as they sometimes do, assumptions about expectations are not minor and outcomes cannot be predicted from structural variables alone.

For theorists, the possibilities opened by the reputation theory led to developments in several directions. Some researchers found in these results an impetus to study the effect of general kinds of behavioral uncertainty in games. That leads to questions about what kinds of reputations a chain store or a company-owner or a bargainer want to develop. Any such endeavor assumes that players use past play to forecast future play, that is, that they expect some consistency over time. Within abstract game theory, interactive epistemology models look at the foundation for such beliefs and, in arguments reminiscent of David Hume's *Skeptical Doubts Concerning Human Understanding*, finds none. The implications of the line of thought initiated in these models continue to reverberate within economics.

Another important application of the sequential equilibrium approach was to dynamic monopoly, bargaining, and the "Coase conjecture." Coase had studied the pricing policy of a monopolist selling a new kind of durable good and conjectured that the monopolist would have little power to maintain the monopoly price. For if customers believed it did have such power and purchased at that high price, the monopolist would have an incentive to cut price immediately afterward to make sales to additional customers.

Bob, together with Faruk Gul and Hugo Sonnenschein, formalized and proved the Coase conjecture, extending it also to certain bargaining problems. In bargaining, if I convince you that I am committed to holding out for a price of $10 for my product (which costs me $5 to produce) but you don't buy at that price, then the item must not be worth $10 to you. At that point, it will be in my interest to cut my asking price below $10. If you anticipate that, you won't agree to pay $10 in the first place, since you'll expect me to cut my price in short order. It turns out that this effect is so strong that if offers can be made frequently, the seller can't hold out for any price above $5. Obtaining this result was one of the main developments in non-cooperative bargaining theory.

Much of Bob's theoretical work arose from practical considerations. In 1979, Bob was retained as a consultant to the Analysis Research Group at the Xerox

Palo Alto Research Center (PARC) . Bob was brought to the group by Michael Rothkopf who was interested in auction theory. However, at that time Xerox entered the broad band communication business, and Bob started to interact with Shmuel Oren and Stephen Smith who were involved in the study of network externalities and in the development of pricing strategies for telecommunication services. Consequently Bob became interested in the pricing of differentiated products and nonlinear pricing which led to his 1982 papers with Oren and Smith. Bob continued his collaborative work on nonlinear pricing with his work on electricity pricing, motivated by a research initiative led by Hung Po Chao at the Electric Power Research Institute (EPRI). With Oren and Smith he extended the nonlinear pricing framework to a two dimensional setting that addressed joined pricing of capacity and quantity enabling firms to recover the cost of capacity installed to meet customers' peak loads. Reliability based differentiation of electricity service and pricing of interruptible service contracts was another focus area of the EPRI work. Bob's work over many years culminated in the publication of his prizewinning book *Nonlinear Pricing* (1993), which became an instant classic on tariff design and related topics for public utilities. His interest in practical solutions is clearly evident in this work, and many of his theoretical innovations are in directions that facilitate computation.

The design of nonlinear pricing policies seems to have led Bob naturally into full fledged market design. In 1992 he became involved in the evaluation and design of practical market mechanisms, including emissions allowances markets and electricity markets. He was involved in the design of the California Power Exchange energy auction and in the design of the California Independent System Operator ancillary service auction. Unlike most theorists, his proposals are not based on theory alone. He collaborated with Charlie Plott to test his proposed activity rules for electricity markets, in the Caltech experimental laboratory.

His review of the proposed rules for wholesale electricity markets in New England (with Cramton in 1998) had significant impacts on the subsequent developments of the energy, capacity and ancillary services markets in New England. He has continued to work with EPRI in advising institutions around the world, such as Taiwan, Poland and Japan, involved in the design of wholesale electricity markets. Bob was prevailed upon to distil some of his electricity market design experience into his presidential address to the Econometric Society, which became his paper called "Market Architecture."[1]

Perhaps Bob's most widely known design work has been for the radio spectrum auctions conducted by the Federal Communications Commission. He was one of the main thinkers behind the FCC auction design, which has evolved into a world standard for sales of radio spectrum. A reading of the original FCC Report and Order adopting the auction design shows his names mentioned more than

[1] A telling sign of Bob's modesty was his reluctance to write about his design work. Although he regarded it as presenting problems as subtle and interesting as any he had tackled, the absence of general theoretical results made him suspect that others would not be very interested. (This while he was the incoming president of the Econometric Society.) What became his Presidential lecture to the econometric society was originally submitted to, and rejected from, a relatively obscure journal whose name is omitted to protect the guilty.

thirty times, including an entire section of the report devoted to his proposal with Milgrom. This work has given new life to auction theory as well, inspiring increased attention to the practical handling of complexities such as auctioning multiple items for which bidders' values may exhibit strong complementarities.

2 Mentoring in the Tradition of Bob Wilson

There are many personal relationships that shape a life. For good or ill, the relationships between parents and children, between brothers and sisters, between spouses or other life partners, are as important in life as they are celebrated in literature and biography. Scholars, if they are lucky, have an additional set of important, life shaping relationships, with their advisors and students. These relationships haven't escaped notice in the rabbinic literature: the Talmud advises us, for example, to choose a teacher and acquire a friend.

An important part of Bob Wilson's tradition was his kindness and helpfulness to his students and young colleagues, and his friendship. This was brought home to us with force at the funeral of Bob Wilson's student Bob Rosenthal, on March 4, 2002. Bob R. made clear how this tradition is passed on. Prominent among Bob R's mourners were students whose lives he had touched, who had become close to both Bob and Liz, and who, it was apparent from their words and their anguish, had been deprived by his unseasonable death of a close friend. Listening to their memories of Bob R's kindness and helpfulness to them, we, his fellow students of Bob Wilson, could not help but be reminded of our own relationships with our advisor. We were also reminded of Bob Rosenthal's kindness to us, when, starting early in our careers, he treated us as his academic younger brothers. Bob R's life was the finest example of how the mentoring tradition is passed on from one generation to the next.

Those of us who have had the good fortune to have students and young colleagues of our own to advise and befriend have an opportunity to pass on this tradition. Just as children of happy families have an enhanced opportunity to be spouses and parents in happy families, so too, it seems to us, do students of good advisors have an opportunity to reap some advantage from this good example. Bob had this advantage himself, as the tribute in this volume by Howard Raiffa, one of his advisors at Harvard, makes plain.