Foreword

Robert J. Aumann

John von Neumann reportedly said that pure and applied mathematics have a symbiotic relationship: Not only does applied math draw heavily on the tools developed on the pure side, but correspondingly, pure math cannot exist in the rarefied atmosphere of abstract thought alone; if it is not somehow rooted in the real world, it will wither and die.

The work before us—which certainly qualifies as beautiful, subtle pure mathematics—is a case in point. It originated half a century ago, at the height of the Cold War between the United States and the Soviet Union, indeed as a direct result of that conflict. The US and SU were trying to keep the Cold War from going hot; to minimize the damage if it did; and to cut down the enormous expenses that the nuclear arms race entailed. To that end, they met repeatedly in Geneva to negotiate mutual reductions in their nuclear arsenals. Regarding these arsenals, both sides were in the dark. Neither knew how many weapons the other had; and clearly, it was the number retained, rather than destroyed, that mattered. In Princeton, Oskar Morgenstern and Harold Kuhn had just founded the mathematics consulting firm “Mathematica.” The United States Arms Control and Disarmament Agency (ACDA) was responsible for conducting the Geneva negotiations for the US; it turned to Mathematica to see whether the Theory of Games—created two decades earlier by John von Neumann and Oskar Morgenstern (the same as the Mathematica principal)—could help in addressing the strategic issues raised by these negotiations. Mathematica responded by assembling a team of theorists that included Gerard Debreu, John Harsanyi, Harold Kuhn, Mike Maschler, Jim Mayberry, Herb Scarf, Reinhard Selten, Martin Shubik, Dick Stearns, and the writer of these lines. Mike and I took charge of the informational aspect (Dick joined us later): whether one side could glean any information about the size of the other’s nuclear arsenal from its tactics in previous negotiation rounds. To get a handle on this problem, we started by looking at the simplest possible analogues: very simple-looking two-person zero-sum repeated games, in which one player knows the payoff matrix while the other does not, and each observes the action of the other at each stage of the repetition. In such games, can the uninformed player glean any information about the payoff matrix from the informed player’s actions at previous stages? Answering this question, even for the simplest 2x2 games, turned out to be surprisingly difficult—and challenging, fun! I vividly remember feeling that we were not working on a contrived, artificial problem, but were exploring the mysteries of the real world, like an astronomer or biologist. Thus was born the theory of repeated games of incomplete information.

What developed from that early work certainly *cannot* be considered applied math. To be sure, some insights may have been useful; for example, that in the context of a long series of repetitions, one cannot make use of information without implicitly revealing it. As a very practical corollary, we told the ACDA that it might be advisable to withhold some information from the ACDA’s own negotiators. But the lion’s share of the theory did not become directly useful, neither at that time nor subsequently. It really is pure mathematics:

though *inspired* by experience—by the real world—it is of no direct use, at least to date.
The theory born in the mid-to-late sixties under the Mathematica-ACDA project started to grow and develop soon thereafter. For many years, I was a frequent visitor at CORE—the Center for Operations Research and Econometrics—founded in the late sixties by Jacques Drèze as a unit of the ancient university of Leuven-Louvain in Belgium. Probably my first visit was in ’68 or ’69, at which time I met the brilliant, flamboyant young mathematician Jean-Francois Mertens (a little reminiscent of John Nash at MIT in the early fifties). One Friday afternoon, Jean-Francois took me in his Alfa-Romeo from Leuven to Brussels, driving at 215 km/hour, never slowing down, never sounding the horn, just blinking his lights—and indeed, the cars in front of him moved out of his way with alacrity. I told him about the formula, in terms of the concavification operator, for the value of an infinitely repeated two-person zero-sum game with one-sided incomplete information—which is the same as the limit of values of the *n*-times repeated games. He caught on immediately; the whole conversation, including the proof, took something like five or ten minutes. Those conversations—especially the vast array of fascinating, challenging open problems—hooked him; it was like taking a mountain climber to a peak in the foothills of a great mountain range, from where he could see all the beautiful unclimbed peaks. The area became a lifelong obsession with him; he reached the most challenging peaks.

At about the same time, Shmuel Zamir, a physics student at the Hebrew University, asked to do a math doctorate with me. Though a little skeptical, I was impressed by the young man, and decided to give it a try. I have never regretted that decision; Shmuel became a pillar of modern game theory, responsible for some of the most important results, not to speak of the tasks he has undertaken for the community. One problem treated in his thesis is estimating the error term in the above-mentioned limit of values; his seminal work in that area remains remarkable to this day. When Maschler and I published our Mathematica-ACDA reports in the early nineties, we included postscripts with notes on subsequent developments. The day that our typist came to the description of Zamir’s work, a Jerusalem bus was bombed by a terrorist, resulting in many dead and wounded civilians. By a slip of the pen—no doubt Freudian—she typed “terror term” instead of “error term.” Mike did not catch the slip, but I did, and to put the work in its historical context, purposely refrained from correcting it; it remains in the book to this day.

After finishing his doctorate, Shmuel—like many of my students—did a post-doctoral stint at CORE. While there, he naturally met up with Jean-Francois, and an immensely fruitful life-long collaboration ensued. Together they attacked and solved many of the central unsolved problems of Repeated Game theory.

One of their beautiful results concerns the limit of values of *n*-times repeated two-person zero-sum games with incomplete information on *both* sides—like the original repeated Geneva negotiations, where neither the US nor the SU knew how many nuclear weapons the other side held. In the Mathematica-ACDA work, Maschler, Stearns, and I had shown that the infinite repetition of such games need not have a value: The minmax may be strictly greater than the maxmin. Very roughly, that is because, as mentioned above, using information involves revealing it. The minmax is attained when the maximizing player uses his information, thereby revealing it; but the minimizing player refrains from using her information until she has learned the maximizing player's information, and so can use *it,* in addition to her own. The maxmin is attained in the opposite situation, when he waits for her. In the infinitely repeated game, no initial segment affects the payoff, so each side waits for the other to use its information; the upshot is that there is *no* value—no way of playing a "long" repetition optimally, if you don't know *how* long it is.

But in the *n*-times repeated game, you can't afford waiting to use your information; the repetition will eventually end, rendering your information useless. Each side must use its information gradually, right from the start, thereby gradually revealing it; simultaneously, each side gradually learns the information revealed by the other, so can—and does—use it. So it is natural to ask whether the values converge—whether one can speak of the value of a "long" repetition, without saying *how* long. Mike, Dick and I did not succeed in answering this question. Mertens and Zamir did: They showed that the values indeed converge. Thus one *can* speak of the *value* of a "long" repetition without saying how long, even though one cannot speak of optimal *play* in such a setting. This result was published in the first issue—Vol. 1, No. 1—of the International Journal of Game Theory, of which Zamir is now, over forty years later, the editor.

The Mertens-Zamir team made many other seminal contributions. Perhaps best known is their construction of the complete type space. This is not directly related to repeated games, but rather to all incomplete information situations—it fully justifies John Harsanyi's ingenious concept of "type" to represent multi-agent incomplete information.

I vividly remember my first meeting with Sylvain Sorin. It was after giving a seminar on repeated games (of complete information, to the best of my recall) in Paris, sometime in the late seventies, perhaps around '78 or '79. There is a picture in my head of standing in front of a grand Paris building, built in the classical style with a row of Greek columns in front, and discussing repeated games with a lanky young French mathematician who actually understood everything I was saying—and more. I don't remember the contents of the conversation; but the picture is there, in my mind, vividly.

There followed years and decades of close cooperation between Sylvain, Jean-Francois, Shmuel, and other top Israeli mathematical game theorists. Sylvain and Jean-Francois came to Israel frequently, and the Israelis went to France and Belgium frequently. One winter, Sylvain and his family even joined me and my family for a few days of skiing in the Trois Vallées. During those years, Sylvain succeeded in attracting an amazing group of students, which became today's magnificent French school of mathematical game theory. One summer, he came to the annual game theory festival at Stony Brook University with twelve doctoral students; "Sylvain and his apostles" were the talk of the town.

Of the book's three authors, only Sylvain actually conducted joint research with the writer of these lines. We conjectured a result during the conference on repeated games organized by Abraham Neyman at the Israel Academy of Sciences in the spring of 1985; concentrated work on it started at the 1985-6 emphasis year in Math Econ and Computation organized by Gerard Debreu at the Mathematical Sciences Research Institute in Berkeley, in which we both participated; it continued by correspondence after we each returned to our home bases; finally, we succeeded in proving the conjecture, and in 1989 published it as the first paper in Vol.1, No.1 of the journal “Games and Economic Behavior.” The result concerns endogenous emergence of cooperation in a repeated game, and perhaps that is a good place to wrap up this preface. The book before us has been in the making, in one sense or another, for close to half a century; so its production may well be viewed as a repeated—or dynamic—game. And, both the production of the book itself, and the work described therein, have been highly cooperative ventures, spanning decades and continents.

The above has been a highly personal account of my involvement with the people and the work that made this extraordinary book happen. I have not done justice to the book itself. Perhaps the best way to do so is to quote from the reports of the anonymous readers who were asked by the publisher to report on the book. These reports are uniformly excellent and highly enthusiastic—I wish my work got reports like that. We here content ourselves with the opening paragraph of just one of those reports; the enthusiastic tone is typical:

“The results and proofs in this text are the foundations on which modern repeated-game theory is built. These are results that apply to zero-sum games, stochastic games, repeated games of incomplete information, spaces of beliefs, stochastic processes and many many other topics. It is impossible to find these results together in one place except in this volume. Existing texts and monographs cover some of them, but none covers anything like all of these topics. However, it is not the coverage of foundational material that makes this text one of a kind; it is the generality and the breadth of vision that is its most special feature. In virtually every section and result the authors strive to establish the most powerful and most general statement. The intellectual effort required to produce this work is huge. It was an enormous undertaking to have brought these results together in this one place. This makes the work as a whole sound leaden and dull; however, it is anything but that. It is filled with an intellectual *joie de vivre* that delights in the subject. This is epitomized by the astonishing links between disparate topics that are casually scattered throughout its pages–the Minmax Theorem used to prove the Peron-Frobenius Theorem; the Normal distribution arising in repeated games with incomplete information; the use of medial limits as a way of describing payoffs, …”

It should be added that the book provides encyclopedic coverage of the area of repeated games—with and without complete information—as well as of stochastic and other dynamic games. The main emphasis is on developments during the classical period—the second half of the twentieth century—during which the theory took shape. Later developments—right up to the present—are also thoroughly covered, albeit more briefly.

In short, the work before us is an extraordinary intellectual tour de force; I congratulate and salute the authors, and wish the reader much joy and inspiration from studying it.

Jerusalem, January 2014