

PREFACE TO “THE REDESIGN OF THE MATCHING MARKET FOR AMERICAN PHYSICIANS: SOME ENGINEERING ASPECTS OF ECONOMIC DESIGN”

Alvin E. Roth

I met Bob Wilson in the early 1970's, when I was a graduate student in Operations Research, starting to worry that I had no passion for queuing theory, inventory theory, reliability theory. So I felt like a man going down for the third time when Bob rescued me by agreeing to supervise a dissertation in game theory.

Bob was on sabbatical the year he agreed to supervise my dissertation. But I saw him an hour a week, and in those hours he helped me find a problem, and struggle to solve it. In memory, those meetings always began with my recitation of what I had tried to do that week, and the obstacle that had prevented me from succeeding. He would then spend some time reassuring me that research was hard, and that I was doing fine. Sometime towards the end of the meeting, he would suggest, sometimes in the most offhand of manners, a paper that I might profitably read. I would go straight to the library to get it, and often spent much of the following week trying to understand it. I frequently began by suspecting that Bob had made a mistake in suggesting the paper at hand, which seemed to have nothing to do with my problem. But in those cases, somewhere in the middle of the paper would be a minor lemma, or a method of proof, that would help me around my obstacle, and on to the next one.

It was already clear that game theory was on its way to opening up new vistas in economics, but, in those early days, game theory was still as much at home in operations research. Bob's own background in OR influenced his taste in game theory. He looked to game theory not merely as the source of deep conceptual insights, such as those through which he had already begun to make his mark on the study of auctions, but also for solutions to practical problems. His longstanding interest in computational algorithms was testimony to his desire to know not only some properties of equilibria, but to be able to compute them exactly. When his auction work led him in applied directions, first as an early advisor on bidding in high value auctions, and then, in the 1990's, as a pioneering auction designer, he helped open up a new way for game theorists to earn their livings, as consulting engineers for the market economy.

None of this was clear to me when I was a graduate student. In fact, for my dissertation, I wrestled with a classical problem in cooperative game theory,

trying to fix a failed conjecture by von Neumann and Morgenstern about how to usefully formulate a notion of stable sets of outcomes. But I must have absorbed, subliminally, much of Bob's taste in practical problems, because not too many years after I finished my dissertation (in 1974), I found myself thinking about markets as algorithms that compute outcomes from the decentralized information revealed by buyers and sellers. But rather than focus on how the market determines a price, I found myself thinking about how, when there are many buyers and sellers, the market determines which buyers are matched with which sellers. This also led me to think about the timing of transactions, but that is another story.

In particular, I found myself thinking about labor markets, and market failures, and searching both for conceptual insights into matching markets in general, and for an understanding of the detailed functioning of specific markets. The first of these that I studied was the entry level market for new doctors in the United States, and I spent the next decade and more happily observing many entry-level professional labor markets, and theorizing about how they worked.

The present paper grew out of a phone call that came in 1995, after I had grown fat and happy, and, it seemed to me, old and wise, in this work. The market for new doctors was suffering a mid-life crisis, said the voice on the phone; would I be prepared to lead an effort to redesign it?

I still recall the visceral feeling of that call—why me? I knew why, of course; I had coauthored a book on matching, and studied many market failures, and if they couldn't call me when a new market design was needed, what was this kind of work good for? But I also knew that the only parts of the book that would apply directly to the complexities of the medical market were the counterexamples; all the theorems were about models that assumed away some of the complications of the actual market.

So, this paper, which records how we overcame the gaps in what we knew, to design a market clearinghouse that would deal with the medical market in all its complexity, is also the story of how I took up the burden that my dissertation advisor had taken upon himself, to make theory that would be a reliable guide to market design, and would allow economists to take on the responsibilities of engineers.