But First, Why is Mechanism Design a Big Deal?

Well, before mechanism design came along, there were *normative* descriptions of ideal economic policies or systems, but they were seldom mathematical. For example, people promoted socialism because it encouraged equality and community, or they promoted capitalism because it encouraged freedom and personal self-realization. On the other hand, mathematical economists typically did research that was either *exploratory* (suggesting basic economic patterns or hypotheses that seemed worth investigating), *descriptive* (describing the workings of economic systems as they exist), *analytical* (analyzing and explaining causal relationships within economic systems) and/or *predictive* (forecasting when, where and why certain economic phenomena might be likely to manifest in the future). All of these mathematical approaches took some economic system as a given and attempted to identify, describe or predict outcomes.

Mechanism design turns that on its head: It takes the *outcomes* as a given and uses mathematical techniques to generate models of economic institutions that produce those outcomes. It is sometimes referred to as the engineering side of economics, since it attempts to produce detailed, implementable designs for economic institutions, given detailed design specs.

Leo first articulated this idea (without using the term "mechanism design") in 1960: "In a broader perspective, these findings suggest the possibility of a more systematic study of resource allocation mechanisms. In such a study, unlike in the more traditional approach, the mechanism becomes the unknown of the problem, rather than a datum."[1] He expanded on these ideas in 1972,[2] while focusing less on resource allocation (flows of money, goods, capital and labor) and more on communication (information transfer and processing).

As early as 1977,[3] Eric Maskin came up with the important concept of monotonicity. The basic idea is that, if a goal is *monotonic*, there is a close relationship between (1) whether a particular outcome is consistent with that goal and (2) people's preferences for that outcome; the two must move together. So if the outcome is no longer consistent with the goal, someone must have changed his mind and now doesn't like the outcome as much as he previously did. Another way of putting it: If an outcome becomes more preferred (relative to other possible outcomes) and was previously consistent with the goal, then it must still be consistent. Monotonicity is a simple, intuitively appealing concept that turns

---

1 Leonid Hurwicz, "Optimality and Informational Efficiency in Resource Allocation Processes," *Mathematical Methods in the Social Sciences*, edited by Arrow, Karlin and Suppes, Stanford University Press, p. 28

2 "On Informationally Decentralized Systems," in *Decision and Organization*, edited by C.B. McGuire and R. Radner, North Holland, Amsterdam, pp. 297-336. Roger Myerson said of this article, "Around 1972, when he began to ask deep questions about people's incentives to communicate, then he made one of the great breakthroughs in the history of social science. When Leonid Hurwicz introduced the concept of incentive-compatibility, it was as if a pair of blinders had been removed. Suddenly, we could see how to analyze economic incentives without assuming any specific institutional structure, and even how optimal institutions might be characterized." (Appendix A)

3 Eric Maskin, "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, pp. 23-38, finally published in 1999

out to be very helpful in determining whether a particular goal can be implemented by some mechanism: if it is monotonic, it can be; otherwise, not.[4]

A few years later, Roger Myerson (and apparently about half a dozen other people)[5] demonstrated that, if a goal can be implemented by some mechanism, then it can be implemented by a mechanism in which participants have the incentive to reveal their true preferences. This fact, termed the *revelation principle*, makes life easier for mechanism designers, because it means that, for instance, before attempting to design a mechanism that copes with possible lying and cheating, the designer can do a sanity test by attempting the usually much easier task of creating a mechanism with the same outputs that assumes everyone is honest. If that's not possible, then there's no need to attempt the more difficult task of creating a mechanism that deals with lying and cheating.

So what exactly is a mechanism? My dad, apparently foreseeing that I might some day try to read the first few pages of his book, *Designing Economic Mechanisms*, included a plain English definition on the very first page:

> A mechanism is a mathematical structure that models institutions through which economic activity is guided and coordinated. There are many such institutions; markets are the most familiar ones. Lawmakers, administrators and officers of private companies create institutions in order to achieve desired goals. They seek to do so in ways that economize on the resources needed to operate the institutions, and that provide incentives that induce the required behavior.[6]

Mechanism design, then, consists of "systematic procedures for designing mechanisms that achieve specified performance, and economize on the resources required to operate the mechanism ..."

The introduction to *Designing Economic Mechanisms* also has a couple of really cool drawings illustrating mechanism design. Not charts or diagrams or graphs. Actual sketches of physical things (boxes in this case) such as might be drawn by a typical fifth-grader to explain something to a typical hippie-techie-bum-folksinger. True, these boxes are decorated with mysterious Greek symbols, and the caption for the first drawing contains the words "nondeterministic algorithm." But, as it turns out, no harm was intended. (Check the drawings out in Appendix G.)[7]

The first box represents a mechanism. It has an input slot and an output slot. Two things go into the

---

4 Eric Maskin email to the author Oct. 28, 2018. For more on monotonicity, see Eric Maskin, "Mechanism Design: How to Implement Social Goals," Prize Lecture, December 8, 2007, https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2007/maskin-lecture.html

5 "But remarkably, the set of all possible equilibria of all possible games can be simply characterized by using the revelation principle, which many of us (Dasgupta, Hammond and Maskin 1979, Harris and 322 Townsend 1981, Holmström 1977, Myerson 1979, Rosenthal 1978) found independently, building on ideas of Gibbard (1973) and Aumann (1974)." Roger Myerson, "Perspectives on Mechanism Design in Economic Theory," Nobel Prize Lecture, December 8, 2007, https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2007/myerson_lecture.pdf, pp. 321-322

6 From *Designing Economic Mechanisms*, right inside the front cover, before the introduction and even before the title page. Titled "Designing Economic Mechanisms," the page is unnumbered.

7 Leonid Hurwicz and Stanley Reiter, 2006, *Designing Economic Mechanisms*, Cambridge University Press. Used with the generous permission of Cambridge University Press

input slot: 1) where the mechanism has to work (the "prevailing environment" chosen out of a set of possible environments) and 2) who the mechanism has to work for (a particular group of "agents" chosen out of one or more possible groups). The box is programmed with a "goal function" defining exactly how success is defined for this particular mechanism.

The output is simply a yes-no answer to the question, "Did we get a desired result as defined by the goal function?" Note that it's "a" desired result, not "the" desired result. That's why it's a *nondeterministic* algorithm: With a given input, you don't always get the same output. But you may be able to design a nondeterministic mechanism so that you always get a "good" output (however the goal function defines "good").

For instance, say you have three people and a mechanism that, when asked for a piece of fruit, hands out either a banana, an orange, or an apple – but you can't predict who's going to get what. As long as you (the mechanism designer) are happy no matter what each of the three gets (hopefully because they themselves are happy), you've got yourself a successful mechanism.

The prevailing environment, by the way, includes the preferences of the agents (e.g. the fact that they're all happy with any of the three types of fruit), but also any other economic factors that the mechanism designer can't control, such as perhaps the number of available oranges, apples and bananas, possibilities for consuming said fruit or putting it back in the hopper, possibilities for growing or importing more fruit, etc.

That's the mechanism box.

The second box might be called a "mechanism factory." It also has inputs and outputs. The output – no surprise – is a mechanism, or possibly a bunch of mechanisms.

Once again, the inputs define who the mechanism has to work for (the "agents") and where it has to work (the "environments"), but this time we may have multiple groups, each in a different environment, which could include different preferences. For example, perhaps we have: three people in Arizona who happen to like tacos, burritos and nachos; three people in California who would prefer oranges, bananas and apples; and three people in Idaho, who like latkes, french fries and tater tots. The correspondence between the state and the preferred foods is embodied in the goal function. The mechanism factory needs to turn out one or more mechanisms that pass the goal function's test by providing the right foods for each group.

This second box, in real life, would typically be an algorithm, that is, a set of mathematical instructions for producing the mechanism(s).

A defining feature of mechanism design is its mathematical specificity. It requires you to explicitly define, in mathematical terms, what you are talking about. In discussing things in less mathematical terms, my father would occasionally use a phrase equivalent to "whatever that may mean" – highlighting the fact that a discussion was taking place without ever really defining what was being discussed. For instance:

> *Let me make it clear at this point that I do not intend to argue the advantages or disadvantages of whatever may be meant by "central planning" or "industrial policy." Rather, my purpose is to*

instill some skepticism with respect to oversimplified arguments sometimes used in this area.[8]

Of course he wasn't opposed to thinking about the fundamentals of economic systems: His work was guided by thinking about fundamental concepts such as *feasibility*, *efficiency*, *equilibrium* and *incentive-compatibility*. For example, given *informational decentralization* (the fact that pertinent information is dispersed among economic agents, who don't have to share it and can even deliberately mislead others about it), a mechanism may only be feasible (i.e. operable) if agents actually can and do share information honestly and are then able to process the information they receive. This means that agents have to be motivated to share information honestly (*incentive-compatibility*)[9] and able to transfer and process information efficiently (*informational efficiency*). Even then, if a solution is so unstable that it is unlikely to last for any significant period of time (i.e. the solution does not represent a point of equilibrium), then it may not be much of a solution after all.

But such foundational ideas were just a starting point. His ultimate goal was to embody whatever understanding or intuition he had in a mathematical form that permitted no ambiguity, to reason from basic premises to logical conclusions and produce logically consistent systems of postulates and theorems. This has sometimes been characterized as an "abstract" approach. In its own hyper-mathematical way, it is actually a highly tangible and concrete one, not allowing for any fuzziness or doubt as to exactly what it is you're talking about.

Of course, mechanism design probably wouldn't be as important as it is if it weren't down-to-earth and concrete in another way, namely, that it has actually helped solve some important practical problems. In fact, perhaps one of the reasons it took the good folks at the Royal Swedish Academy of Coolness so long to award this particular prize is that applications emerged over a long period of time.

For instance, one of area where mechanism design has been successfully applied is in the design of auctions. A commonly-cited exampled is the Vickrey auction, or "second-price" auction, first described academically by William Vickrey, a professor at Columbia University, in 1961.[10] In this type of auction, the highest bidder gets the item, but only pays the second-highest bid. The advantage is that bidders have no incentive to submit bids below the actual value they place on the item, because lowering your bid can only reduce your chances of winning, not the amount you pay if you do win. It is also not smart to bid more than the item is worth to you, since you could then end up paying more than it's worth. This is an example of an incentive-compatible mechanism, one that rewards those who honestly reveal their private information.

Vickrey's clever idea was "largely ignored for a decade."[11] Then in the 70's, the Vickrey-Clarke-

---

8 In a comment on *Economic Planning and the Knowledge Problem*, by Israel M. Kirzner, in *Cato Journal*, 1984, vol. 4, issue 2, p. 420

9 Leonid Hurwicz, "Optimality and Informational Efficiency in Resource Allocation Processes," *Mathematical Methods in the Social Sciences*, edited by Arrow, Karlin and Suppes, Stanford University Press, 1960, p. 28. A later paper went into depth on the idea: Leonid Hurwicz, "On informationally decentralized systems," in *Decision and Organization: A Volume in Honor of Jacob Marschak*, ed. R. Radner and C.B. McGuire, Amsterdam: North-Holland, 297–336.

10 William Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, 16(1):8–37, 1961

11 Lawrence M. Ausubel and Paul Milgrom, "The Lovely but Lonely Vickrey Auction," Discussion Papers 03-036, Stanford Institute for Economic Policy Research, p. 1,

Groves (VCG) auction was devised, which has similar truth-inducing characteristics for multiple bidders competing for multiple items, and is more generally applicable in other ways as well. Even that has been used more as a "lovely and elegant reference point" than as a blueprint for actual auctions.[12] Among other possible drawbacks, VCG requires bidders to bid on every possible combination of goods.[13] By the 1980's, one could reasonably have concluded that the auction design branch of the mechanism design tree had matured, bearing some interesting but perhaps not terribly edible fruits.

But the best was yet to come. Precipitated, in part at least, by the worst. In 1991, it was discovered that some major cheating had been going on in auctions for U.S. government securities: No single bidder was supposed to be able to bid on more than 35% of the total supply, but the Salomon Brothers investment bank had tried to corner the market by bidding on amounts at times *exceeding* the total supply. Salomon was fined $190 million and required to set aside another $100 million for restitution to injured parties. The firm's reputation was seriously hurt by the scandal, and it ended up getting acquired by Travelers Group. This raised awareness among those who followed such things that the good old sealed-bid, pay-what-you-offer auction had some shortcomings.

So in 1994, when the FCC started auctioning off communication spectrum (rather than assigning it by previously preferred means such as lottery or comparative hearings sometimes referred to as "beauty contests"), they consulted experts on how to devise an auction mechanism that would be transparent, fair, foolproof, and allocate the spectrum to the companies that valued it most. A very public test of the concept and potential of mechanism design, it was a huge success.

About this same time, the public Internet began its explosive expansion, and a need emerged to find ways to route traffic across networks owned by many different entities, serving all kinds of different customers, in a way that was both globally efficient and individually incentive-compatible. This is a classic mechanism design problem, and it was recognized that some of the same ideas that worked for auctions could be applied here: essentially, nodes and links would "auction off" their services and capacity, and traffic would "bid" on it, using least-cost routing algorithms to make decisions.

By the time all this was happening, though, it seemed that awareness of Leo's foundational contributions might have been fading. In fact, famously, when he got the early-morning call telling him he had won the Nobel Prize, he hung up, thinking it was a "stupid joke." (Well, actually, he had my mom hang up. She had answered the call, because he was hard of hearing.)

"There were times when other people said I was on the short list," he commented, "but as time passed and nothing happened, I didn't expect the recognition would come because people who were familiar with my work were slowly dying off."[14]

So on that chilly Minneapolis[15] morning, December 10, 2007, in the Ted Mann Theater at the

---

https://web.stanford.edu/~milgrom/publishedarticles/

12 *ibid.*, p. 35

13 Eric Maskin email to the author Oct. 28, 2018. For more on the VCG auction, see Eric Maskin and Tomas Sjöström, "Implementation Theory," *Handbook of Social Choice and Welfare*, Volume 1, Edited by K.J. Arrow, A.K. Sen and K. Suzumura, 2002, p. 241

14 William Grimes, "Leonid Hurwicz, Nobel Economist, Dead at 90," June 26, 2008, https://www.nytimes.com/2008/06/26/world/americas/26iht-obits.1.14006616.html

15 He was receiving dialysis several times a week and was unable to travel to Sweden. My sister Ruth

University of Minnesota, when my father slowly rose from his chair on the stage, slightly hunched over, leaning on his cane, and Jonas Hafström, Swedish Ambassador to the U.S., presented the red leather box containing the gold medal with the embossed profile of Alfred Nobel, those who *hadn't* forgotten knew he was harvesting the fruit of a tree he had planted half a century earlier.

Even those most in the know, however, may not have known much about the soil that nurtured that tree: our family's Polish-Jewish background. That will be the focus of the next chapter.

I may have left out one or two technical details in this discussion of mechanism design. If you're interested in going a little deeper, I recommend the short section from the introduction to *Designing Economic Mechanisms* by my father and his partner-in-design Stan Reiter, reprinted as Appendix G, with the generous permission of Cambridge University Press. Eric Maskin's[16] and Roger Myerson's[17] Nobel lectures are also great introductory material.

(with her family) traveled to Sweden to represent him.

16 Eric Maskin, "Mechanism Design: How to Implement Social Goals," Prize Lecture, December 8, 2007, https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2007/maskin-lecture.html

17 Roger Myerson, "Perspectives on Mechanism Design in Economic Theory," Prize Lecture, December 8, 2007, https://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2007/myerson-lecture.html