8 From Competitive Equilibrium to Mechanism Design in Eighteen Months

Sea Change

If I had to name one major shift in the sensibilities of economic theorists in the past half century, a prime candidate would be the way we conceptualize markets—from quasi-natural phenomena admired from afar to manmade institutions whose design can be tweaked by economist-engineers.

The traditional image of markets that reigned in the 1950s and 1960s was competitive equilibrium. Standard accounts of the evolution of modern economic thought emphasize the extreme assumptions of the competitive-equilibrium model, which were subsequently relaxed by the information-economics and game theory developments in the 1970s and 1980s. But what is perhaps most striking in the competitive-equilibrium model is its sense of the market as a *natural* phenomenon. No one designs it. No one runs it. It can be defined with no reference to an explicit mechanism. The Walrasian auctioneer is a fiction aimed at those of us who insist on a concrete mechanism, but it is not an intrinsic part of the model. Endowments are like initial conditions of a physical system, and prices react spontaneously to changes in these endowments until they somehow manage to equilibrate this system. Indeed, as historians of economic thought have pointed out, the conscious inspiration for this image was the physical theory of thermodynamics.[1] According to this point of view, economists study markets in a disengaged manner, as if they were observing a natural system. Sure, they can tweak the system's initial conditions, but then they let market processes unwind naturally.

Compare this view with the contemporary "market design" approach, which regards the economist as an "engineer" that designs the details of the market institution. A good guide into this new culture is Al Roth's popular book *Who Gets What and Why*.[2] A telling sign of the shift is that in Roth's book, there is a huge emphasis on market institutions in which prices play no role (for example, assignment of children to schools)—in marked contrast to the all-important role of prices in the competitive-equilibrium model.

Another manifestation of the sea change is the declining status of competitive equilibrium in the *education* of academic economists. Students are seeing less and less competitive equilibrium and more and more game theory during their basic graduate-level theory training. In some quarters, teaching competitive equilibrium is taking place in *macroeconomics* courses. The "market design" approach relies on game theory, which enables us to describe specific market rules and analyze participants' behavioral response to these rules. Competitive equilibrium cannot get to this level of resolution.

All this is familiar. The question is whether this great methodological shift affects our attitudes as economic theorists in ways that might be less salient. My goal in this chapter is to illuminate this question, making use of a personal experience of mine: a project I pursued in the last decade with my longtime coauthor, Kfir Eliaz. We started our project in a spirit of rebellion against the prevailing mechanism-design approach to the topic we were interested in, but eventually we capitulated to the prevailing paradigm and rewrote our paper as a mechanism-design exercise.

Why did we start out opposing the mechanism-design approach? And how did we end up accepting it as an effective mode for our exercise? This personal story can teach us a small lesson about the power that cultural currents in our profession exert over individual researchers—even those who make a deliberate effort to swim against the tide. Of course, I will make my best effort to make the reader curious about the substantive economic question that Kfir and I studied. But the real interest here is in the conscious trade-offs that we made between the two styles of research, our initial decision to adopt the outmoded style and the reasons behind it, and the forces that led to our eventual surrender.

What Do I Mean by "Competitive Equilibrium"?

Strictly speaking, the term refers to the classical model of a market (or a collection of markets), to which agents arrive with endowments or production technologies and make their individually optimal consumption or production decisions given market-clearing prices.

I have in mind a looser conception of the term, which describes an attitude to modeling rather than a specific set of models. This category includes the 1976 Rothschild-Stiglitz model of competitive screening in insurance markets that we encountered in chapter 7. In this model, market equilibrium is defined by a collection of insurance contracts, such that all contracts generate zero profits given consumers' behavior, and there is no scope for profitable entry of a new contract. This is quite different from the traditional formulation of competitive equilibrium. Nevertheless, the formalism is close in spirit. In particular, it eschews game theory: there is no clearly articulated protocol of trade. The competitive pressures that insurance companies face have the status of an "invisible hand"—they are in the background, captured by the zero-profit condition. By contrast, in game-theoretic models of market competition, intensity of competitive pressures is determined by the "rules of the game."

So, this is what I mean by a competitive-equilibrium approach: the "reduced form" manner in which the definition of market equilibrium captures an underlying competitive force. It offers a bird's-eye perspective into the market situation at hand. The culture of economic theory has grown rather hostile to this bird's-eye, "invisible hand" approach—favoring the game-theoretic, mechanism-design approach, with its explicit description of the "market protocol."

After this prologue, now to the story itself.

Can "the Market" Simulate an Ideal Search Engine?

Around 2008, keyword auctions were the rage in parts of economic theory. After years of operating without a clear method for monetizing its incredible search engine, Google had converged on "sponsored search" as a viable business model. Google maintained its so-called organic search engine, where answers to a user's query are determined by the engine's algorithm. However, side by side the organic search, Google started auctioning a chunk of the user's screen space to paying advertisers.

For economists interested in auction theory, the auction-design problem was novel. In this setting, the auction designer aims to allocate multiple "positions" on the user's screen. Different positions have different degrees of prominence. They differ in their ability to attract the user's attention, hence in their value for advertisers. Moreover, unlike traditional advertising, the search engine can partly monitor users' attention, by observing whether they click on ads. This means that the advertiser's bid and payment can be defined "per click." In 2007, Ben Edelman, Michael Ostrovsky, and Michael Schwarz published an influential paper that studied a specific auction format in this environment, known as the "generalized second-price auction."[3] This paper created a buzz in the auction theory community as well as in the subfield known as Econ-CS, which mixes researchers from economic theory and computer science. An important chunk of this multidisciplinary interaction has been taking place under the auspices of new research labs founded by internet giants.

The avalanche of papers in the wake of Edelman et al. (2007) was firmly in the mold of auction theory, itself part of the mechanism-design tradition. When Kfir and I became interested in search engines around 2008, we felt a need to try a different approach, which would abandon mechanism design toward a competitive-equilibrium approach.

Our starting point was the observation that the primary function of search engines is to bridge the gap between what users want and their limited ability to put it in words. When a consumer is perfectly able to describe a product she wants, all the search engine needs to do is provide a list of sellers that provide this product. These alternatives need not be identical; they may differ in price, quality, or certain idiosyncratic details that can be detected only by inspection. But this is not where search engines shine best. It is when they manage to take a vague query like "a tall, handsome Israeli economist" and spit out "Rani Spiegler" in reply, or suggest that when someone spells "Kefir Elias" he is really looking for Kfir Eliaz, that they take the search experience to the next level. Kfir and I referred to this function as "vocabulary expansion." Yet, all the papers that Kfir and I had seen in the literature on keyword auctions ignored this key aspect of search engines. They assumed that the user's query unambiguously defines the set of relevant objects, such that the search engine's job is merely to serve them to the user in a particular order.

Kfir and I posed the following question: Can a "sponsored search" engine, based entirely on the incentives of competing advertisers, mimic the vocabulary expansion performed by an ideal "organic" search engine? In other words, if Google ditched its stupendous algorithm and relied entirely on sponsored search, would the competitive forces that shape advertisers' bidding for keywords lead to a search environment that is as effective in fulfilling the vocabulary-expansion role?

As it happens, Google founders Sergei Brin and Larry Page had considered a similar question ten years before us, in an appendix to the famous 1998 paper that introduced their search engine:

> Currently, the predominant business model for commercial search engines is advertising. The goals of the advertising business model do not always correspond to providing quality search to users. . . . We expect that advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers. . . . We believe the issue of advertising causes enough mixed incentives that it is crucial to have a competitive search engine that is transparent and in the academic realm.[4]

Kfir and I have been unaware of this quote until very recently. What Brin and Page asserted back in 1998 (and ironically ignored when they later adopted the advertising business model), Kfir and I were asking from a purely academic point of view. Can "the market" offer "quality search" to users?

Why Competitive Equilibrium?

History is written by the victors, as the cliché goes. So are lecture notes and textbooks. The narrative that the currently dominant mechanism-design culture offers nowadays is that the "reduced form," "bird's eye," "invisible hand" aspects of the competitive-equilibrium paradigm make it inferior to the mechanism-design approach. The graduate-level microeconomic-theory textbooks by Kreps (1990) and Mas-Colell, Whinston, and Green (1995) were key moments in this development.

An extreme manifestation of the new culture is the intermediate microeconomics course that the leading theorist Jeff Ely has developed (building on the work of his former Northwestern University colleague, Kim-Sau Chung). In traditional courses, competitive equilibrium takes center stage. In contrast, Ely's course strictly adheres to the mechanism-design perspective, which he regards as more basic, leaving the competitive-equilibrium model (which he regards as more specific) to the very end of the course. As he writes on the website that presents his approach:[5]

> The goal is to study the main themes of microeconomics from an institution—and in particular market-free approach. To illustrate what I mean, when I cover public goods, I do not start by showing the inefficiency of market provided public goods. Instead I ask what are the possibilities and limitations of any institution for providing public goods. By doing this I illustrate the basic difficulty without confounding it with the additional problems that come from market provision. I do similar things with externalities, informational asymmetries, and monopoly.

All of this is done using the tools of dominant-strategy mechanism design. This enables me to talk about basic economic problems in their purest form.

Even if this is an extreme pedagogical approach at the moment, it may well anticipate the near future of economics pedagogy. It is an impressively unapologetic demonstration of power by the new, self-confident culture.

Like Ely, Kfir and I came of age intellectually when this new culture began its ascent. New graduate textbooks in the 1990s emphasized game theory at the expense of competitive equilibrium. Consequently, I feel more at home in Ely's radical course than in the old-style courses that place competitive equilibrium at the forefront. And yet, when Kfir and I wanted to study sponsored search engines, we were concerned about certain aspects of the mechanism-design approach.

First, when one formulates and analyzes a mechanism-design problem, one tacitly adopts the designer's point of view. But who is this designer and why should we empathize with him? For example, in a principal-agent model of the relationship between an employer and its worker, the theorist effectively identifies with the employer. The principal has a clear objective (maximizing profits or the worker's effort, attaining an efficient allocation), and the theorist finds herself pursuing it single-mindedly. In contrast, if the theorist abandons the designer's perspective and looks at the situation from a more disinterested point of view, she is more likely to explore the economic interaction from multiple perspectives and evaluate its outcome according to a more diverse set of criteria.

This distinction is not clear-cut. For example, Myerson and Satterthwaite's (1983) classic impossibility result employs the mechanism-design approach to obtain a result about the inefficiency of bilateral trade in Nash equilibrium under asymmetric information, which applies to arbitrary trading mechanisms, and regardless of the designer's motivations. Thus, on one hand, it is firmly in the mechanism-design tradition, yet on the other hand, it does preserve something of the outside perspective that I associated with the competitive-equilibrium tradition. However, the Myerson-Satterthwaite paper is an exception in this regard. The typical mechanism-design exercise does not pretend to describe a situation from a disinterested position; it aims to solve a problem. And typically, it is the specific problem of a specific economic agent (an employer, a seller). The single-mindedness of this endeavor can blind the analyst to other aspects of the situation. The analytical identification with the designer risks becoming an emotional one. In the context of the sponsored search problem, it meant identifying with Google and its objectives. And Kfir and I didn't want to act like unpaid Google employees.

Second, the single-minded pursuit of the designer's objective usually comes with fierce determination to relax any constraint on the instruments at the designer's disposal, even if this leads to artificial mechanisms. From this perspective, every limitation on the space of feasible mechanisms feels ad hoc. When theorists carry out a mechanism-design exercise, they tend to have fewer inhibitions when introducing unrealistic instruments, compared with descriptively motivated theoretical exercises. A classic example is the role of integer games in Eric Maskin's (1999) canonical Nash-implementation mechanism.[6] In contrast, the competitive-equilibrium approach is more at ease with restricting contract spaces. The most obvious example is the classic textbook model, which restricts itself to linear prices without trying to derive this restriction from first principles.

To summarize, focus on a single objective and reluctance to limit contractual instruments are two characteristics of the mechanism-design approach. Kfir and I felt that when trying to understand whether "the market can attain quality search," these two characteristics would derail us from our mission and blind us to interesting aspects of the problem. We found competitive equilibrium's disinterested "bird's eye" mentality more fitting. Or was it no more than a childish reaction to the "Econ-CS" papers that gave us the impression that their authors were about to the enter the executive suite of some internet giant (some did)?[7]

And so now I'll tell you the story of our research, how we were determined to follow the competitive-equilibrium style, how we failed, how we ended up adopting the mechanism-design language, how our initial motivation did leave its imprint on the final product (published as Eliaz and Spiegler 2016), and why I think there are general lessons to be learned from this otherwise idiosyncratic experience.

Our theoretical argument can be conveyed with a simple example, so I'll stick to it throughout the chapter, resisting the temptation to explain how much more general our exercise really is. (I should add that my presentation here cheats a little bit; I have tweaked the formal exposition relative to our original working paper, for the sake of clarity.)

Mozart or Stravinsky

Suppose our search-engine users are interested in classical music. They heard a piece of music on the radio while driving, or possibly as part of a film soundtrack. Having liked the piece, they would like to retrieve it. Unfortunately, they don't know the name of the piece; they can barely hum it. They will recognize it if they get to hear it again, but they cannot describe it.

However, our users are not all helpless. Some of them know the *name* of the piece's composer. For the sake of our example, suppose the universe of classical-music composers consists of only two, Mozart and Stravinsky. Some of our users are looking for a Mozart piece, while others are looking for a Stravinsky piece. Within each group, some can name the composer of the piece they are looking for. Accordingly, there are four types of users:

1. Type (Mozart, MOZART): The composer of this type's favorite piece is Mozart and the user can name him, and therefore he submits the specific query "MOZART."
2. Type (Stravinsky, STRAVINSKY): The composer of this type's favorite piece is Stravinsky and he can name him, and therefore he submits the specific query "STRAVINSKY."
3. (Mozart, CLASSICAL MUSIC): The composer of this type's favorite piece is Mozart but he can't name him, and therefore he submits the generic query "CLASSICAL MUSIC."
4. (Stravinsky, CLASSICAL MUSIC): The composer of this type's favorite piece is Stravinsky but he can't name him, and therefore he submits the generic query "CLASSICAL MUSIC."

I am using capital letters to indicate queries and lowercase letters to indicate the composer of the piece our user is looking for. I will also use the terms "queries" and "keywords" interchangeably.

The economic allocation problem is thus to give each of these types access to a "search pool" from which they can repeatedly sample specimens until they find what they are looking for. Formalize an "ideal search engine" as a function that assigns such a search pool to each query. Note that the search engine cannot distinguish between types 3 and 4 because they submit the same query. Therefore, there will be only three search pools, one for each of the queries MOZART, STRAVINSKY, and CLASSICAL MUSIC.

How does the user navigate inside his assigned search pool? Take the simplest model in the textbook: *random sequential search*. The user samples specimens in random order, until he finds what he likes. (Since we will deal with search pools that contain an infinite number of any given type of seller, the distinction between sequential search with and without replacement is irrelevant.) A user whose favorite composer is Stravinsky will repeatedly sample pieces of music from the pool. He will never choose a Mozart piece; and conditional on drawing a Stravinsky piece, there is some constant probability q that he will select it and terminate the search.

This image is quite different from Google search, where alternatives are presented in a certain order, which affects the order by which the user inspects the alternatives. In this sense, the modeling approach I am describing here does *not* aim at a faithful description of Google search. Instead, it starts from the abstract notion of an "ideal search engine" as a function that assigns a search pool to search queries. It then looks at the simplest textbook model in the search theory literature: random sequential sampling in a stationary environment. This is a difference between "applied theory" and "pure theory" sensibilities.

Given the user's behavioral model, a search pool is fully described by its shares of Mozart and Stravinsky pieces. An ideal search engine will choose a composition for each search pool in a way that *minimizes the user's expected search time*.

The optimal search pool in response to the query MOZART will consist of Mozart pieces only. Likewise, the optimal search pool in response to the query STRAVINSKY will consist of Stravinsky pieces only. Recall that when a user encounters a piece by the right composer, he stops the search with probability q. The user's expected search time in either of these two pools will be

$$q \cdot 1 + q(1-q) \cdot 2 + \cdots = \frac{1}{q}$$

What about the search pool that the search engine designs in response to the generic query CLASSICAL MUSIC (I'll use CL as a convenient abbreviation for this query)? This pool should include pieces from both composers. The optimal composition will minimize the expected search time of users who submit this query:

$$\Pr(Mozart \mid CL) \frac{1}{q \cdot Share(Mozart)} + \Pr(Stravinsky \mid CL) \frac{1}{q \cdot Share(Stravinsky)}$$

In this formula, $Share(x)$ means the fraction of x pieces in the search pool, and $\Pr(x \mid w)$ is the probability that the user wants x conditional on him submitting the query w. Solving this minimization problem is straightforward. The solution satisfies:

$$\frac{Share(Mozart)}{Share(Stravinsky)} = \sqrt{\frac{\Pr(Mozart \mid CL)}{\Pr(Stravinsky \mid CL)}}$$

This is a key observation. When the user submits the generic query, the ideal search engine faces uncertainty regarding the user's taste. To minimize the user's expected search time, the search engine should design a search pool in which the share of a composer is proportional to the *square root* of the fraction of this composer's fans in the population of users who submit the query. In particular, this means that the minority taste group in this population should be *overrepresented*. For example, when the fraction of Stravinsky fans among the users who submit CLASSICAL MUSIC is 20%, the share of Stravinsky pieces in the pool should be *one-third*:

$$\sqrt{\frac{20}{80}} = \frac{1/3}{2/3}$$

A Competitive Market for Keywords

Suppose there are many sellers of Mozart pieces and *just as many* sellers of Stravinsky pieces. (We can think of a seller as a group of musicians who made a recording of a particular piece of classical music.) Each seller can serve any number of customers, but no seller can provide both types of product. The value of a successful transaction is 1 for all sellers.

The allocation problem is to assign sellers to search pools. This is a many-to-many allocation: multiple sellers are admitted to a given pool, and a given seller can be assigned to multiple pools. The allocation determines the composition of each search pool. For example, if 75% of Mozart sellers and 50% of Stravinsky sellers are allocated to the search pool corresponding to the query CLASSICAL MUSIC, then this pool will consist of 60% Mozart (because $75/(50 + 75) = 0.6$) and 40% Stravinsky.

We want to conceptualize a "sponsored search" engine that performs the allocation task via some kind of "competitive market." For every query there is a market price that a seller needs to pay if she wants to enter the search pool associated with this query. The price can be defined as a fixed entry fee, or equivalently—in the spirit of real-life sponsored search—as a *price per impression* that the seller will pay each time she is examined by a user in the pool she was admitted into. (The search engine cannot monitor whether an impression results in a transaction. Therefore, a price-per-transaction is infeasible.)

A *market equilibrium* consists of a price-per-impression for each keyword and a decision for each seller as to which keywords to pay for (recall that paying the price of a keyword is a necessary and sufficient condition for entering its associated search pool). There are two conditions for this pair to constitute a market equilibrium. First, all sellers earn zero profits. Second, no seller can earn strictly positive profits by paying equilibrium prices for some other bundle of keywords.

The zero-profit condition captures the intense competition among sellers as they attempt to enter search pools. The definition doesn't ask who receives the sellers' payments. We can assume it is the search engine, but the definition is silent on this issue. This is unlike the mechanism-design approach, which puts the designer at the forefront.

Kfir and I thought of this kind of definition as belonging to the tradition of competitive equilibrium. Not because there are clearly articulated supply and demand. Neither did our model introduce an explicit scarcity that presumably calls for a market allocation. The scarce resource is the users' time, but it lies in the background. Like the 1976 Rothschild-Stiglitz concept that served as our inspiration, the definition captures in "reduced form" the competitive pressures that sellers experience when trying to enter users' consideration sets.

An immediate consequence of the two conditions is that we can apply the zero-profit condition to each query separately. The equilibrium price-per-impression of each query is equal to the *conversion rate* experienced by the sellers who are admitted into the query's pool—namely, the probability that an impression in that pool will result in a transaction.

Kfir and I focused on *symmetric market equilibria*. In a symmetric equilibrium, all sellers of a given type make the same choices. Asymmetric equilibria were less appealing in our model, as they could "weaponize" the zero-profit condition to obtain results that disappear with various perturbations of the model.

A First Welfare Theorem?

Does market equilibrium induce an optimal composition of the search pools associated with the various queries? If so, this would be a "first welfare theorem" specialized for our "market for keywords" setting, and sort of a theoretical "yes" to Brin and Page's 1998 question posed above.

The answer is trivially affirmative as far as the specific queries MOZART and STRAVINSKY are concerned. Users who submit these queries are homogeneous: all users who submit MOZART are looking for a Mozart piece, and all users who submit STRAVINSKY are looking for a Stravinsky piece. We can therefore allocate Mozart sellers to the MOZART pool and Stravinsky sellers to the STRAVINSKY pool. The market price-per-impression that induces zero profits for these sellers is q for each of the two queries. No Stravinsky seller wants to pay this price for MOZART because the conversion rate it will experience in the MOZART pool is zero. Likewise, no Mozart seller wants to pay the market price of STRAVINSKY. So, the equilibrium conditions hold for the specific queries.

The situation is different with the generic query CLASSICAL MUSIC. The optimal search pool associated with this query must include both types of sellers. If we focus on *symmetric* equilibria, then *all* sellers should be in the pool, such that the fraction of Mozart sellers in it will be 50%. This will almost never be the optimal composition.

For the sake of the argument, suppose we *did* allow for asymmetric equilibria, such that not all sellers of a given type act the same. In particular, suppose that m Mozart sellers and s Stravinsky sellers choose to pay the market price of CLASSICAL MUSIC and thus enter its search pool. The conversion rate that each seller experiences in this pool should be equal to the market price. Let us see what this entails.

The random-sequential-search assumption means that every seller in the pool gets the same number of impressions. In order for all sellers' conversion rate to be the same, the number of transactions that each of them completes must be identical as well. Let's calculate this number for each type of seller. Suppose that the total number of users who submit the query CLASSICAL MUSIC is n (m, s and n are all large numbers). Then, a Mozart seller gets the following number of transactions:

$$\frac{n \cdot \Pr(Mozart \mid CL)}{m}$$

Likewise, a Stravinsky seller gets the following number:

$$\frac{n \cdot \Pr(Stravinsky \mid CL)}{s}$$

The requirement that these numbers are identical translates to the following equation:

$$\frac{m}{s} = \frac{\Pr(Mozart \mid CL)}{\Pr(Stravinsky \mid CL)}$$

But the ratio m/s is exactly the ratio between the shares of Mozart and Stravinsky sellers in the pool. That is,

$$\frac{m}{s} = \frac{Share(Mozart)}{Share(Stravinsky)}$$

We see that $Share(x)$ must be *proportional* to the fraction of x fans in the population of users who submit CLASSICAL MUSIC. But recall that an ideal search engine requires $Share(x)$ to be proportional to the *square root* of this fraction. In other words, if m and s were selected to implement the optimal pool composition, Mozart sellers would be earning more than Stravinsky sellers, because the latter would be overrepresented in the pool relative to their fan base. Competitive forces would then lead Mozartians to crowd out the Stravinskians.

The conclusion is that a competitive market equilibrium cannot sustain the search pool that an ideal search engine would generate—even if we allow for asymmetric equilibria. We can't get our first welfare theorem. There is a fundamental tension between search-time minimization and the zero-profit condition that dictates competitive allocation of sellers into search pools.

Broad Match

Let us now tweak our competitive market for keywords, by redefining the entitlement that paying for a keyword gives. Suppose that when a seller pays the market price for the keyword MOZART or STRAVINSKY, she is also granted *probabilistic entry* into the search pool CLASSICAL MUSIC.

This indirect access is in the spirit of the "vocabulary expansion" function of search engines, bridging between supply and users' imperfectly described demand. When someone submits the query CLASSICAL MUSIC, there is some probability that what he is looking for is a Mozart piece. Just as when someone Googles "Kefir Elias," there is a good chance that he is actually looking for material on Kfir Eliaz. It is therefore helpful to form some linkage between the query "Kefir Elias" and the objects that are more obviously associated with the query "Kfir Eliaz."

Linkages of this sort are known in the industry as "broad match." Let us hijack this term and apply it to our model of a competitive market for keywords. This means adding the following component to the model: a function that assigns to every pair of queries w and v a number

$b(w, v)$ between 0 and 1. The interpretation is that if a seller pays the market price for w, she is admitted into the search pool of query v with probability $b(w, v)$.

This is the only change in the model: an extended definition of what paying the market price of a keyword entails in terms of admission into search pools. Under the original definition, a seller is admitted into the pool associated with a query if and only if she pays the market price for that same query. The broad match function means that a seller can enter a search pool associated with one query even if she pays the market price of a different query.

For example, suppose that all Mozart sellers pay the market price for MOZART and CLASSICAL MUSIC, whereas all Stravinsky sellers pay the market price for STRAVINSKY only. Then, the fraction of Stravinsky sellers in the search pool of CLASSICAL MUSIC will be (using abbreviations for the three queries):

$$\frac{b(STR, CL)}{b(STR, CL) + b(MOZ, CL) + b(CL, CL)}$$

The remaining fraction will consist of Mozart sellers.

The definition of market equilibrium remains the same: sellers must earn zero profits, and they should not be able to find a more profitable bundle of keywords. The only thing that changes is the entitlement that paying the market price of a keyword gives. This modified entitlement is defined by the broad match function b. In this sense, the broad match function is analogous to the endowments in the classical exchange-economy model.

A "Second Welfare Theorem"?

The introduction of broad match into the model raises a question in the spirit of the second welfare theorem: Is there a specification of the broad match function b for which the optimal composition of the ideal search engine can be sustained in *symmetric* market equilibrium?

The answer turns out to be—it depends. More precisely, the following inequality is a necessary and sufficient condition. Suppose that in the general population of users, there are more Mozart fans than Stravinsky fans. Then, this is what the inequality looks like:

$$\frac{\textit{Share of Mozart fans}}{\textit{Share of Stravinsky fans}} \cdot BC \leq 1 \qquad (*)$$

The first term on the left-hand side is clear: that is the ratio of Mozart and Stravinsky fans in the general user population. But what is BC? This is a measure of the similarity between the query distributions that characterize Mozart and Stravinsky fans. It is known as the *Bhattacharyya Coefficient*, following Bhattacharyya (1943)—hence the abbreviation BC—and defined as follows:[8]

$$\left(\sum_{w}\sqrt{\Pr\left(w\,|\,Mozart\right)\Pr\left(w\,|\,Stravinsky\right)}\right)^{2}$$

In this definition of BC, w is a query; $\Pr(w\,|\,x)$ is the probability that a user who wants x submits the query w. BC takes values between 0 and 1. It increases with the similarity between the query distributions of Mozart and Stravinsky fans. Put differently, BC measures how informative users' queries are of their preference type. For example, when the query is fully informative, $\Pr(w\,|\,Mozart) = 0$ or $\Pr(w\,|\,Stravnisky) = 0$ for every w, such that BC is zero. At the other extreme, if the query is entirely uninformative, $\Pr(w\,|\,Mozart) = \Pr(w\,|\,Stravinsky)$ for every w, such that BC becomes

$$\sum_{w}\Pr(w\,|\,Mozart) = 1$$

The BC formula applies to any conditional query distribution. In our simple example, BC is reduced to the simple product

$$\Pr(CL\,|\,Mozart)\cdot\Pr(CL\,|\,Stravinsky)$$

Inequality (*) conveys the following lesson. If users' queries are sufficiently informative about their preferences, and if the preference distribution is not too skewed, then the search pools that an ideal search engine would devise can be sustained in symmetric market equilibrium, provided that we design the broad match function appropriately. In this sense, we have a qualified second welfare theorem.

The reason broad match can help is that it dissociates the considerations that dictate the optimal composition of search pools from the competitive pressures that govern sellers' access to the pools. Recall that the former consideration requires overrepresentation of Stravinsky sellers in the CLASSICAL MUSIC pool. Stravinsky sellers who pay for STRAVINSKY and enter the CLASSICAL MUSIC pool thanks to broad match give the necessary boost to the representation of Stravinsky in that pool. At the same time, the conversion rate that these sellers get from STRAVINSKY is lower than in the previous, "narrow match" case because now they also encounter Mozart fans who submitted the query

CLASSICAL MUSIC. If there are too many of those users, then Mozart sellers, exploiting broad match, will become too eager to pay the market price for the keyword STRAVINSKY. As a result, they will disrupt the optimal allocation by diluting the STRAVINSKY pool. Inequality (*) prevents this scenario from materializing.

Kfir and I went further and constructed the broad match function and the equilibrium prices-per-impression of all queries in the market equilibrium that sustains the optimal search pools, when inequality (*) holds. For example, we showed how BC enters the equilibrium price formula: as users' queries become more informative, the price-per-impression of keywords rises across the board.

A happy ending?

Primitives and Solution Concepts

Only as far as Kfir and I in our "research lab" were concerned. But in the summer of 2013 came the time to tell other people about it in seminar or conference talks, as well as in ten-minute personal chats over coffee.

This proved to be not so happy. Time and again, we found it hard to communicate our model and our findings. The market equilibrium concept itself was difficult to get across. Not because it is formally complicated; it is in fact quite simple to describe, as we saw earlier. It is certain that our written exposition back then was muddled. But we had enough experience with seminar presentations and one-on-one chats to be able to explain what we were doing in an intelligible manner. Yet, unlike other occasions with other papers, we felt we were failing at that.

Of course, this is not interesting by itself, but I believe our communication failure had a broader significance. The difficulty in getting the message across efficiently—"twenty five words or less," as the villainous movie producer Griffin Mill demands from screenwriters who try to pitch him a story in Robert Altman's fabulous 1992 film *The Player*—seemed to lie in economic theorists' guarded attitude to any theoretical exercise that involves a nonstandard solution concept.

Most economic models have two built-in parts: a description of the model's *primitives* and the *solution concept* that one applies to these primitives. The culture of economic theory welcomes a proliferation of primitives but dislikes a proliferation of solution concepts; it prefers to take solution concepts off the shelf. If a paper wants to introduce a new solution concept, it must make a big fuss around it. The concept is expected to apply to a general, abstract class of models. As a result, the paper ends up being *about* that solution concept. Casually introducing a new solution concept into a paper that is nominally about something else is considered poor form.

Economic theorists have become extremely efficient in communicating models to each other. This efficiency depends on having an effective language for defining primitives and a small set of handy solution concepts. Confronted with a new solution concept, theorists switch to a slower, more critical, and less efficient reception mode.

Our solution concept was new-ish. Conceptually, it was traditional. On the other hand, it wasn't bread-and-butter competitive equilibrium because there were no clearly articulated supply, demand, or endowments. Rather, our approach was inspired by Rothschild and Stiglitz's 1976 paper: they had invented a new solution concept in the *spirit* of competitive equilibrium, specialized for an insurance market model that departs from linear prices. Their concept was later applied to other settings, but it was initially introduced in that specific context of an insurance market. Likewise, Kfir and I wanted to define a solution concept in the spirit of competitive equilibrium, specialized for a "market for keywords" model. In the contemporary culture of economic theory, an exercise that introduces new primitives *and* a new solution concept is harder to get across.

From a broader perspective, which goes far beyond our little study of search engines, the competitive-equilibrium approach has a built-in limitation. Historically, the textbook definition was developed by Arrow, Debreu, and others in the context of rigidly structured market models. Agents in these models arrive with endowments or technologies, and prices must be linear. At the time, this seemed like the entire universe of economics, so economists didn't internalize its narrow scope. In this sense, game theory is much more flexible: its language can accommodate more diverse primitives. As a result, when theorists like Rothschild and Stiglitz wanted to carry the spirit of competitive equilibrium over to insurance markets with nonlinear pricing, they felt a need to invent a new solution concept, whereas a game theorist would need only write a new game form.

Recently, Michael Richter and Ariel Rubinstein developed an interesting research agenda that addresses this limitation. They argue that "competitive equilibrium" is a modeling attitude with much broader scope than what we have been conditioned to believe. Their approach (which they launched in Richter and Rubinstein 2015) was to present an abstraction of the competitive-equilibrium model that would go beyond the standard exchange/production economy and beyond prices as regulators of economic activity. This was more grandiose than our own exercise, but it had a similar underlying impulse.

Our audiences' distrust of our nonstandard solution concept was accompanied by the more general suspicion of models that involve hand-waving about invisible hands. If some competitive force is going to bring our sellers' profits to zero, then why won't we write down a game-theoretic model in which sellers have explicit moves that capture this force?

Ultimately, it wasn't clear to our audiences why we were so keen to abandon the mechanism-design approach. Under that approach, we would need to innovate only at the level of primitives: a new allocation problem that is concerned with assigning sellers to search pools. We wouldn't need to innovate at the level of the solution concept. Given the presentational advantages of this approach, our willful refusal to take it seemed annoying.

Conversion

Faced with such a clear communication failure, Kfir and I didn't even bother to submit our paper for publication. Instead, we decided to soften our intransigence and describe our exercise entirely in mechanism-design terms.

We started with so-called (anonymous) direct mechanisms. In this telling, every seller reports whether she is of a Mozart or Stravinsky type. The mechanism responds to every report with a probabilistic assignment to search pools and a per-impression fee. Given the reporting strategies of other sellers, each seller can compute the number of impressions and transactions she will experience in each pool, and use it to evaluate her possible reports. The mechanism is incentive-compatible if no seller wants to misreport her type given that all other sellers are truth-telling.

And that's it! Of course, this description seems so much shorter than my earlier pitch in this chapter because I've already presented the primitives. So the proper comparison is between this conventional definition of incentive-compatible direct mechanisms and our earlier definition of market equilibrium. The familiarity of standard mechanism-design definitions contributes to the efficiency in communicating our idea.

The most interesting comparison between the two approaches concerns the *questions* that each of them generates. The mechanism-design description immediately suggests the following question: Are optimal search pools implementable by an incentive-compatible mechanism, even without imposing the restriction that sellers earn zero profits? The answer turns out to be an unqualified *yes*, using a basic result from the theory of mechanism design due to Rochet (1987). This question had no counterpart in the competitive-equilibrium version.

When we ask whether the search engine can implement the optimal search pools and at the same time extract the sellers' entire profits, we get a problem that is formally equivalent to the question we posed to our competitive-equilibrium model. And indeed, the answer is the same: the twin objectives of implementing optimal search and extracting sellers' profits are attainable if and only if the inequality (*) holds. But these twin objectives sound rather strange in a mechanism-design context. If the designer is the search engine, why should it regard search-time minimization as a primary motive? And if the designer is a benevolent social planner, why should it care about extracting advertisers' surplus? This is a neat demonstration of how the mechanism-design approach leads the analyst to identify with a specific designer and refrain from asking questions that do not correspond to the designer's natural motivations.

To maintain the interpretation of the designer as a search engine and to rationalize its interest in implementing optimal search, we can assume that it also collects access fees from the search engine's *users*. Their willingness to pay these fees will depend on the quality of their search experience. The lower the search time, the higher the access fee the search engine can charge them. This is what Kfir and I did. We rewrote our model of user behavior as a rational search model (instead of the mechanical sequential search process that we originally assumed) and introduced user access fees. In this manner, the search engine doesn't have twin objectives but a single, conventional one: maximizing profits. But that's exactly the kind of modeling strategy we had originally wanted to avoid. Not only does it carry the "unpaid Google employee" mindset further, but it also leads to something that we *don't* see in reality: user access fees. It is a feature that is demanded not by realism but by the mechanism-design modeling strategy.

Another problem that the mechanism-design pitch demands is finding "indirect," auction-like mechanisms, to complement the analytically convenient but unrealistic direct mechanisms. As we had expected, analyzing Nash equilibria in such auction games forces the analyst to consider all kinds of deviations from equilibrium play. Some of these deviations make sense, while others are excess baggage that one must accept if one is going to apply Nash equilibrium to the game. It's part of the deal, but it's a distraction that the competitive-equilibrium approach avoided.

Finally, once we learn that the search engine cannot always implement the "first-best" (maximizing and fully extracting social surplus), it is conventional in mechanism design to look for the "second-best": what is the maximal profit that the search engine can generate with an incentive-compatible mechanism? Once again, it is a question that doesn't arise in the competitive-equilibrium framework. This is a question we knowingly left open.

After we rewrote our paper as a strict mechanism-design exercise, without any reference to competitive equilibrium, it became a much easier expositional task. There were fewer misunderstandings, fewer lapses of communication. The stark assumptions we made exposed us to valid criticisms about the limited scope of our exercise—especially if one expected it to be a model of Google search, an expectation abetted by the mechanism-design style, unlike the more detached competitive-equilibrium approach. But it was easier to convey what we were doing.

Recap: Why Did We Eschew Mechanism Design in the First Place?

The research directions that the switch to a mechanism-design approach forced us to consider demonstrate the reservations about the approach that had motivated us from the beginning. Let us recall these concerns.

First, the mechanism-design approach leads the analyst to identify with the designer and therefore ask questions that are of importance for the designer, possibly at the expense of questions that are more interesting for an outside observer. And indeed, we found ourselves reformulating the problem in terms of a profit-maximizing designer.

Second, according to this mechanism-design protocol, when the designer is unable to implement her "first-best," figuring out the "second-best" seems like an obvious next step. The pressure to pose such a question under a competitive-equilibrium approach is weaker.

Third, another question that arises naturally under the mechanism-design approach is the quest for realistic, auction-like indirect mechanisms. But Nash equilibrium analysis of such mechanisms forces the analyst to consider stability with respect to deviations that aren't always natural or interesting.

Fourth, the mechanism-design perspective impels the analyst to expand the set of instruments at the designer's disposal, even if these instruments aren't realistic. And indeed, to reformulate our original question as a profit maximization problem for the search engine, we were led to introduce unrealistic user access fees into our model.

As to the crowding out of interesting questions, the very notion of broad match may be a case in point. The large Econ-CS literature on sponsored search, with its mechanism-design orientation, almost never addressed broad match. The overwhelming majority of works on auction-theoretic aspects of keyword pricing focus on the case of a *single* query. Very few address the case of multiple queries and the notion of broad match that arises in them. Those that do typically view broad match as a means for thickening auction markets, or examine the computational complexity of bidding in a multi-keyword environment.[9] Broad match as a fundamental feature of search engines that enables them to fulfill the "vocabulary expansion" role—we found none of that in this sizeable literature. I believe that the mechanism-design perspective is responsible for this neglect. It creates this tunnel vision that pushes the researcher to go deep in certain directions at the expense of others.

Epilogue

After the mechanism-design expositional overhaul, Kfir and I submitted our paper to the *American Economic Review* at the start of 2015, and our paper was rather quickly accepted for publication. In material terms, this was a successful outcome.

And yet, looking back, I think I still prefer the older, competitive-equilibrium version. The mechanism-design gambit led us to ask questions that appear "natural" through that particular prism. However, it may have muted other, more interesting problems that revolve around a question I still find fascinating: *Can market forces regulate effective organization of human knowledge?*

By accepting the mechanism-design mindset, Kfir and I were eventually, despite our initial intentions, acting like Google minions. The detached, bird's-eye view of competitive equilibrium seems to buy you a certain freedom and independence that the more practically minded "market design" culture doesn't. When the great transformation from competitive equilibrium to market design is complete and the dust settles, we should recognize this subtle cultural change as an important by-product.

The Curious Culture of Economic Theory

By: Ran Spiegler

Citation:

The Curious Culture of Economic Theory
By: Ran Spiegler
DOI: 10.7551/mitpress/14884.001.0001
ISBN (electronic): 9780262379038
Publisher: The MIT Press
Published: 2024

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Palatino by Westchester Publishing Services.

The author's preference is to use binary pronouns in this work. The MIT Press recognizes all gender identities and recommends the use of gender-neutral pronouns.

Library of Congress Cataloging-in-Publication Data

Names: Spiegler, Ran, author.
Title: The curious culture of economic theory / Ran Spiegler.
Description: Cambridge, Massachusetts : The MIT Press, [2024] | Includes bibliographical references and index.
Identifiers: LCCN 2023028463 (print) | LCCN 2023028464 (ebook) | ISBN 9780262548229 (paperback) | ISBN 9780262379021 (epub) | ISBN 9780262379038 (pdf)
Subjects: LCSH: Economics. | Economics—Sociological aspects.
Classification: LCC HB171 .S7155 2024 (print) | LCC HB171 (ebook) | DDC 330.01—dc23/eng/20230713
LC record available at https://lccn.loc.gov/2023028463
LC ebook record available at https://lccn.loc.gov/2023028464