

Brave New World of Market Design

April 17, 2011

Yeon-Koo Che* and Erin Cho†

Running Head: Market Design

Abstract: Market Design is an emerging field that seeks to provide a practical solution for various resource allocation problems. The current article discusses several issues arising in real-world market design, innovative ideas proposed to overcome those issues, and their applications in areas such as spectrum licenses allocation, Internet sponsored-search advertising, housing allocation, public school placement, kidney exchange, course allocation, and online matchmaking.

JEL Codes: D02, D04

* Economics Department, Columbia University, 420 West 118th Street, IAB 1029, New York, NY 10027; e-mail: yk.che@columbia.edu.

† Design and Management, Parsons, The New School for Design, 66 5th Avenue, New York, NY 10011; e-mail: choje@newschool.edu.

1. Introduction

Market Design is an emerging field that seeks to provide a practical solution for various resource allocation problems. The need for designing a method of resource allocation arises in a variety of contexts. For instance, governments often assign their citizens public land, subsidized housing, the rights to use minerals and radio spectra, health care, public education, or civic duties such as military or jury services. Private firms also allocate budgets and equipments to different divisions and assign their employees to different tasks that need to be performed; religious organizations assign their ministers and priests to different parishes; and manufacturers organize ways to procure necessary materials and parts from subcontractors. Indeed, its ubiquity and significance make market design as old as economics itself.

For the most part, institutions allocating a particular good or service develop rather spontaneously, without any planning or designing. For many ordinary goods, markets emerge naturally as an allocation method. Typically, when the demand for a particular good arises, those with the means to produce the demanded good bring it to the market (they will be motivated by the profit incentive to do so), and the price adjustment will then take care of its allocation to the buyers, namely by assigning the good to those willing to pay the market-clearing price and denying access to that good to those unwilling to pay that price. There are other goods, however, that are either owned by the government, or the nature of the good is such that its provision cannot or should not be simply left to unregulated private markets. Examples of public goods include health care, education, military service procurement, clean air, and to some extent public housing and radio spectrum. For such goods, markets do not develop voluntarily (due to the government control of critical input), or sometimes markets are not viewed as appropriate ways to allocate resources; for instance, moral objections against “commoditizing” human organs, basic education, civic or civil rights eliminates their markets. For these objects, the method of (non-market) allocation must be designed.

While governments have long been involved in changing, improving and experimenting with different methods for allocating public resources, and economists have contributed to such an effort in various ways, it is relatively recent that practical design of a desirable allocation scheme has received rigorous analytical scrutiny and that the lessons from those analyses have been put to practical use. The use of auctions in 1994 to assign Federal Communication Commission (FCC)’s licenses for Radio Spectrum was the first significant large-scale application of economic theory in market design. Of course, there are other significant cases of market design, some of which will be discussed later.

At the same time, the advance in computer technologies and the widespread use of Internet have brought the application of market design to a new level. A design feature once considered too complex to implement is now routinely implemented via computer programming, and many relevant parties to transactions can be found and linked through the internet, without interference from temporal and spatial boundaries, thus opening up an immense scope of market transactions both in the B2C and B2B contexts.

The expansion of market size and scope has increased the need to design each particular market to meet the particular needs and challenges arising in that market, which has increased the demand for, and the growth of, market design as a new field within Economics.

In this paper, we discuss several market design mechanisms applied to different areas and the general guiding principles behind these designs as well as the challenges and limitation of these designs.¹

2. Why “Markets” may be useful as an allocation device

Suppose you are a government official in charge of designing a scheme for assigning a particular resource, say undeveloped public land, to private citizens to promote its development. It will seem quite natural to divide the land to pieces of an appropriate size and sell them at prices determined by the demand and supply. Creating a market in this way appears to be a natural way to assign the land (any resources, for that matter) given that many ordinary goods are indeed allocated through this channel. Markets have a number of beneficial properties that we will illustrate subsequently, but economic theory does not necessarily point to “markets” as an ideal method, at least for the initial assignment of ownerships.

The celebrated Coase theorem suggests that it does not matter how the ownership is assigned initially as long as those assigned are free to resell them to others. The idea is that the process of sale and resale mediated by voluntary negotiation will lead to the land being ultimately allocated those who value the most, no matter who is initially assigned it. In other words, even though the land was given away at random and for free to anybody who desires, the end outcome in terms of allocation will be efficient in this sense after the subsequent resale process is complete. For the efficient result to hold, however, the resale process must be efficient in the sense that those with highest value for the property must recognize the property at low search cost and be able to acquire the property without incurring much transaction costs.

Nevertheless, we shall argue that the resale process is unlikely to be efficient, in which case it is important to get the initial assignment “right” in the sense that those with highest valuation for the property should be obtain the property. From this perspective, the market mechanism tends to be desirable than non-market methods (such as lottery or other priority rule still found in many areas). Let us argue this point first by way of two brief case studies.

The first is the massive land assignment in late 19th and early 20th century by the Southwestern states of the US, most famously by the state of Oklahoma. These states had acquired a large swath of lands from the Native American tribes and subdivided the lands and assign the tracts of lands to individuals for “homesteading” or settlement. Initially in 1889, the method used was what is now famously known as the *Oklahoma Land Rush*. That is, the applicants ran a race to stake a claim on a 160acre tract, some in better quality than others.

¹ For a fuller and more complete survey on market design, one should consult Roth (2002).

There were several problems with this method. First, there were conflicting claims; for instance, some participants staked their claims only to discover another claimant on the same 160-acre claim but on the other side of a hill. The conflicts were resolved sometimes through violence, adding to the lawlessness in the frontier life. The more serious issue was the so-called “sooners.” These are the people who entered the lands illegally prior to the official openings. A large fraction (as much as 90%) of the land was occupied by the sooners. The problem of sooners led to a change in a rule to simple lotteries, starting with the *1901 Land Opening*. The use of lotteries effectively eliminated the problem of sooners; but a different kind of problem emerged. The integrity of the system was often questioned, for there was a suspicion of manipulation of the outcome to favor some participants. More serious was the ease of participation, leading to as many as 165,000 applicants registering for the program. One suspects that many of these participants did so for the purpose of reselling the assigned lot at a higher price rather than for the homesteading purpose. In 1906, the government finally decided to sell the land at market values, as it started auctioning off the land.

There is a clear advantage that auctions have over the two earlier methods. Auctions tend to attract serious and credible buyers and select as winners those who tend to value the object the most. The simple reason is that a buyer who values an object the most is willing to pay the most to acquire it, so in most standard auctions, they tend to come out as winners. Clearly, selling the resource at a market clearing price is an effective way to raise revenue for the government, but the market assignment has arguably a more important virtue: *Selecting the most deserving owner*. The basic insight is that requiring someone to outbid and outpay others is a simple, but credible, way to ascertain the social value of one’s ownership.

The process of trial and error exhibited in the assignment of public land was eerily repeated in the assignment of Radio spectrum. Initially, the FCC awarded the licenses to use Radio spectrum through comparative hearings (which evaluated the merits of various applicants). But by the early 1980s so many firms were applying for licenses that the system ground to a halt. The FCC then switched to lotteries; the thinking was essentially the Coase theorem: the telecom firms, the argument goes, will reallocate the licenses by selling and reselling them until the latter ends up in the hands of the most deserving parties. The problem is that FCC didn’t restrict who could participate in the lotteries. There were over 400,000 applicants, and there were some embarrassing incidents. One year, for instance, a group of dentists won a license to run cellular phones on Cape Cod, and quickly turn around and sold it to Southwestern Bell for \$41 million. In general, the voluntary reallocation process ended up taking a painfully long time, and ultimately delayed the mobilization of the cell-phone services in the US.

The FCC finally switched in 1994 to auctions.² The auctions designed with the advices of economists such as Paul Milgrom, Bob Wilson and Preston McAfee became a huge success. Not only did they generate the revenue close to \$1 billion, but the bids were also highly competitive, and there was

² See Kwerel and Williams (1993) for the discussion on the background on this policy.

hardly any resale afterwards. Ever since, the auctions have now become the most predominant way to assign spectrum licenses all around the world.

The experiences of the public land and Radio spectrum assignment teach us two important lessons. First, as Coase himself recognized, the method of initial assignment *does* often matter in the real world since the subsequent reallocation tends to be time consuming and entail nontrivial transaction costs; it is thus important to get the resources initially assigned to the individuals who can realize the most social value out of them. Second, using competitive markets or carefully designed auctions does a good job of identifying and selecting such individuals for initial assignment. We shall illustrate some of recent applications of “markets” and the some of the design issues arising in those applications.

3. Applications of markets and auctions

3.1 FCC license auctions

A typical FCC spectrum auction assigns numerous licenses. For instance the 2006 3G auctions by the FCC sold more than 1000 licenses for the total gross bids of \$13 billion. The competition featured firms with different background, including T-Mobile backed by a Deutsche Telekom AG, wireless firms such as Verizon Wireless and Sprint Nextel Corp, and cable operators such as Comcast Corp. and Time Warner Inc. These bidders have different needs and market strategies in participating the auctions, some (such as T-Mobile) needing to acquire a sufficient number of licenses to build capacities to compete against its rivals, and others having preferences for serving certain regions.

To accommodate different needs for different participants and to realize the highest social aggregate values presents a difficult design challenges. What is desperately needed is to allow bidders to communicate to the auctioneer and one another on their preferences for each bundle of licenses, so that in the end the winning bidders end up creating the most valuable aggregation, or mosaic of license holding. The design adopted in 1994 and continued on by FCC attempt to solve this “efficient aggregation” problem by the *Simultaneous Sequential Sealed-Bid Auctions* (SSA), wherein each bidder at each round indicates whether or not to raise bid on each license separately, subject to a set of activity rules governing minimal bid increment and exit.³ The idea is that, by observing the dynamic progress of bid increments made by all bidders, each bidder makes inference about others’ demand/preference structure and makes efficient adjustment.

Despite the success of SSA, the efficient bid aggregation remains a challenge. One particular challenge is the so-called “exposure” problem illustrated as follows. Suppose two bidders say, 1 and 2, bid for two licenses, A and B. Bidder 1 values either license at \$2 million, but does not demand both A and B.

³ See the discussion by Milgrom (2004) for more details about SSA and other issues related to FCC license auctions.

(His value for the second license is zero.) Bidder 2 values each at nothing but A and B together at \$3 million. Suppose the bids for each license rise in price at the increment of \$0.1 million. The efficient aggregation in this case is for bidder 2 to acquire both A and B, but it is unlikely for this to occur from the SSA. For bidder 2 to win both, he should be able to beat bidder 1 in both auctions, and bidder 1 is willing to bid up to \$2 million each. But then bidder 2 will have to spend \$4 million, which is more than the combined value of the licenses. The dynamic feature of the SSA certainly facilitates efficient aggregation, but the separate auction of license cannot produce an efficient aggregation, since bidder 1 will certainly focus on the weaker flank of the two battle fronts, and there is no avoiding the payment by bidder 2 of bidder 1's full valuation of each license there.

A solution to this problem can be provided by a more powerful, but at the same time complicated, auction procedure known as a "Combinatorial Auction." In this auction procedure, bidders bid on each possible bundle and the auctioneer consider all possible aggregations and declare the highest value aggregation as winning bids. It is easy to illustrate why a combinatory auction can facilitate efficient aggregation. For instance, in the above example, each bidder will bid on three bundles: "A only", "B only," and "A and B together." A will bid at most \$2 million for each bundle, whereas bidder 2 will bid zero for the first two and high enough (possibly more than \$2 million) for A+B to win both licenses, which is what would be socially desirable in this example. In practice, there are likely to be many licenses (recall that the 3G auctions sold close to a thousand), so the number of bundles to bid on can easily be too numerous to be manageable. The complexity of the system can also challenge the bidders and may lead them to behave strategically.

Another challenging issue is collusion among bidders. Bidder collusion is a common problem arising in many contexts. Bidder coordination on their preferences is a double-edged sword in this regard. On the one hand, allowing and even facilitating bidders' communication of their preferred bundle is necessary for efficient aggregation to emerge. On the other hand, they may use this communication opportunity to collude, or bid non-competitively. There were some early incidences of possible bidder collusion in the initial FCC auctions, in which bidders communicated their preferred licenses by the last three digits of their bids in early rounds and subsequently refrained from bidding on the rivals' preferred licenses.

3.2 Design competition, Beauty contests and Scoring auctions

Another limitation of standard auctions is that they often focus solely on the price dimension and do not adequately provide incentives for the quality of services that the consumers may care about. Auctions inattentive to quality matters are likely to produce insufficient incentive for quality provision. In the European 3G license auctions, for instance, winning firms were criticized for lagging behind their timely

commitment and investment in network coverage and technological advance after winning their licenses. One important reason is that the auctions did not reward/incentivize quality investment properly by the winning firm. Furthermore, firms allegedly overexerted their resources in winning the license bidding, which left them with insufficient amount of resources for network coverage and technology investment. In fact, precisely this concern for investment and quality provision led several countries, such as Finland, Sweden, Norway, Spain, France to forego auctions in favor of the so-called “beauty contest” --- a bureaucratic procedure of selecting firms based on technical expertise, financial viability, network coverage, etc.

For an ordinary good sold in a market, there would be no concern about a firm’s incentive for quality provision, since consumers care about both quality and price in their purchasing decisions, and firms will then be forced to compete in both quality and price, or in whatever way to satisfy and win over their consumers. The same kind of incentive is lacking, however, in the standard price-only auctions: Firms have no reason to raise quality for price is only the criterion for winning. When the concern about quality is significant, thus a simple price-only auction will not work, and the multidimensional aspect of market competition needs to be explicitly incorporated into the design of an auction. An extremely simple way to do this would be to abandon the auction altogether, as has occurred with beauty contests. But beauty contests are unlikely to be efficient; for efficiency reason, some form of bidding competition would still be preferred. Indeed, a *scoring rule auctions*, suggested by Che (1993), can provide a practical solution to the problem. The idea behind the scoring rule auction is very simple. Instead of selecting a winning bidder solely based on his price offer (lowest price in the case of procurement), the auctioneer formulates a system of scoring a bundle of quality/design specifications and price, and selects the bidder that achieves the highest score. Since the firms will then compete to maximize the score, it will simply be a matter of designing the proper scoring function to create the incentive for quality provision. Some forms of scoring rule auctions have existed, used by such government agencies as Department of Defense or Department of Transportation, to procure weapons or transportation systems. The scoring functions used by these agencies are of a simple linear form which assigns weights to different aspects of performance measures and then aggregates them to produce the final score. With the improvement of computing and web technologies, the application of the scoring rule auctions to the B2B outsourcing context appears very promising.

3.3 Internet Ad (“Sponsored Links”) Auctions⁴

Internet search engines such as Google and Yahoo! provide the unprecedented opportunity for advertisers to identify and target their relevant customers. The advertisements placed in search engines, known as “sponsored links”, are the web links that appear on the top and on the right side in a highlighted box when

⁴ More detailed discussion of the Sponsored Links auctions can be found in Edelman et al (2007) and Varian (2007).

one enters a related term or keyword in a search engine. For instance, when a search term is entered in the Google, a number of sponsored links as well as relevant searched links appear.

The ads provided through sponsored links are closely related to the term searched, and thus likely to be of interest to the searchers. This keyword targeting provides a very powerful tool for advertisers to reach their most valued set of customers, much more easily and effectively than can be achieved by the traditional ads or even Internet banners. Indeed, much of the value of search engines as a business model comes from the sale of these sponsored links spaces. In 2007, 98% of the Google's revenue and 50% of Yahoo!'s revenue are attributed to these sponsored links.

How are these sponsored links spaces sold? Both Yahoo! and Google auction off these spaces, using similar formats. Several issues must be taken into consideration in the designing of such auctions. First, the position of where a given advertising link is placed matters; clearly, the advertising at the top position is most easily recognized by the searchers, and thus would be most valuable to the advertisers, followed by the second top, and the third, etc. A standard economic wisdom suggests that a better position should be sold to an advertiser likely to generate a larger revenue, for doing so will enable the search engine to more effectively extract large payments from the advertisers. This means that a bidding rule must create incentives for advertisers to compete for better positions. Second, it is important from the risk sharing perspective for the advertisers and the search engine to share the profit from the advertising. This means that the advertisers' payments must be correlated with not just the position but also with some measure of the value of advertising at that position, a good proxy of which will be the click-through rate, i.e., the number of times the advertised link is visited by the searchers. Third, the relevance and interest of a given keyword to searchers change over time; so is the value of advertising on the part of the advertisers. It would be useful to allow real time adjustment and repositioning of ads.

The auctions adopted by Google and Yahoo largely accommodate these issues. They both use real-time dynamic auctions that assign the links spaces of differing values in the order of bids, i.e., assigning the most valuable position to the top bidder, the second most valuable position to the second highest bidder, and so on and so forth. The winning bidders are then charged fees based on their bids per click. In other words, during any given period, an advertiser of a winning bid would be charged a fee based on his bid multiplied by the number of times the linked site was visited. Up until late 1990s, the per-click fee charged by Yahoo was exactly the same as the advertiser's bid. For instance, the following would be an example of this so-called *Generalized First-Price auctions*:

Table 1: Generalized First-Price Auctions

Positions	Bidders	Bids	Payment per click
1	A	\$4	\$4
2	B	\$3	\$3
	C	\$1	

The problem with this fee system is that it creates an incentive for advertisers to constantly monitor others' bids and revise their own bids in response.⁵ To see this, notice that bidder A can keep the same top position for any bid greater than \$3, so he will try to lower the bid down to \$3.01, just enough to keep his top position. If A does so, then B will try to jump over to \$3.02 and get a better position without paying much more what he used to. B doing this in turn will induce A to raise his bid over B's, who then follows suit again, raising his own bid over A's; and the bids will rise until the price for the top position becomes so high that B decides rather to save money by jumping down his bid to \$3. But this is precisely what retriggers A's initial deviation! So, they will simply repeat the same cycle of price movement, which clearly requires their constant monitoring and adjustment of their bids. Clearly, the constant monitoring and revising of bids adds to transaction costs to the advertisers and thus adversely affects the appeal and value of the ad spaces sold by the search engines.

These problems were later fixed when the auction rule changed to the **Generalized Second-Price Auction** (henceforth, **GSP**).⁶ Under this latter rule, a winning bidder for a given position is charged per click a fee equal to the winning bid for a position one-step lower than that position. For example, the same bids as in the example will lead to the following per-click payments:⁷

Table 2: Generalized Second-Price Auctions

Positions	Bidders	Bids	Payment per click
1	A	\$4	\$3.01
2	B	\$3	\$1.01
	C	\$1	

⁵ The problem discussed here follows Edelman and Schwarz (2007)

⁶ Google's fee system and the ranking of bids are more complicated, as they were adjusted by the likelihood of visit at the given link, summarized by the estimated click-through rates. The bids are then indexed by the associated scores, namely the per-click bid multiplied by the estimated click-through rate. And the per-click price of a winning bidder is then determined by the price that will make his/her score the same as the next highest score.

⁷ Readers should be warned that the advertiser in general will change their bids (in particular most likely raise their bids) when the rule changes from the first-price to the second-price auction. The same bids being assumed here were thus just a matter of expositional simplicity.

This auction rule would be clearly more stable since bidder A now no longer faces the temptation to lower his bid, since his payment will not go down when he lowers his bid above \$3. Likewise, B will not be able to win a better position by raising his bid slightly above \$3.

Aside from ensuring the stable bidding behavior, the GSP auctions tend to allocate the positions efficiently in the sense of maximizing the total values from assignment. Edelman et al (2007) and Varian (2007) argue that the GSP auctions entail multiple equilibria in bidding behavior, but all of them support the efficient allocation of advertising positions. The efficiency claim is largely supported in a recent experimental study by Che, Choi and Kim (2011).

4. Why “Markets” may be feasible or desirable

Despite the increasing popularity of markets/auctions as a method of assigning resources, non-market assignment methods such as first-come-first-served, lotteries, and other priority rules, have a long history, and their use remains widespread. Human organs are assigned to transplant patients by a priority rule that depends on factors including the recipient's age, the severity of the condition, and the distance between the donor and recipient. Typically, assignment to public school is based on one's residence or test, rather than one's willingness to pay for a market-clearing price. In Korea and Singapore, substantial numbers of new housing units are subject to price caps that are well below market prices, and the units are assigned by lottery. In addition, lotteries have been used to assign immigration visas and jury duty, and to select conscripts for military service.

At least in some of these examples, the use of non-market method appears to originate from the reluctance by the society to put a “price” on something as sacred or inalienable as human organs, elementary public education, immigration rights, basic health care, and housing. This attitude may change, though, as the efficiency benefit from markets may ultimately prove too powerful to be held back by whatever moral objections currently swaying the mind. Indeed, the debate continues on the merit of introducing markets for things that were once considered outside the realm of the so-called ordinary economic goods. A recent one is the proposal by Gary Becker, a distinguished member of the Chicago School and a Nobel Laureate, to auction off 50,000 immigration visas that are currently assigned by lotteries to randomly selected applicants all over the world.⁸ The idea of using markets here is not to raise revenue for the government, but rather to select the immigrants more efficiently. The reasoning is that those who will add the most by immigrating will stand to reap the most from immigrating, and will thus be most eager to bid large sums of money. We will come back to Becker's proposal later, but it is useful to make the following two points.

⁸ See “The Price of Entry: a new proposal from Gary Becker to make a market in immigration,” *The Economist*, 2010.

First of all, we have to ask if, moral or fairness objections aside, markets can be fully justified at least from the efficiency perspective for the kinds of goods mentioned above, namely those for which markets have not yet developed or not fully embraced by the general public. Second, regardless of whether markets will ultimately be a better solution, it is worth investigating how to design the effective and desirable assignment schemes in the interim when the absence of explicit market taken as a given constraint. We are going to discuss the first issue here, while the second issue will be discussed later in the context of several specific applications through the remainder.

On the first issue, we shall argue that the answer to this question depends on the wealth inequality of individuals as well as the availability of well-functioning capital markets.⁹ The simple intuition is that a full-fledged market need not be efficient if individuals have different wealth levels and the access to capital markets is severely limited for some members of the society. As mentioned, the virtue of markets is that individuals are asked to credibly “demonstrate” their social value of ownership by their willingness to outpay all others for the last unit of the resource available. That the markets rely on one’s purchasing power to demonstrate his social value of ownership can be its drawback, however, if individuals differ in their wealth levels and the capital markets are not perfect. If capital markets were perfect, then the markets would still function effectively to screen for individuals with the high social value of ownership, for an individual will be able to convince the capital market to lend as much fund as would equal the social value generated by the resource.

Realistically, however, capital markets are not perfect, so individuals with low wealth will not be able to spend as much as the social value of resource under his ownership. This also means that market assignment --- i.e., the screening based on one’s payment --- will likely be distorted in favor of high wealth individuals, and will thus likely be inefficient. In other words, with imperfect capital markets, markets may rather select rich rather than the deserving individuals for ownership assignment. To illustrate this point, imagine a hypothetical economy with 300 inhabitants. There is a valuable resource, say an apartment, of 100 units available to be assigned to the people. The breakdown of the population in terms of their valuation of the apartments and their wealth levels is as follows:

Table 3: Hypothetical Economy with Liquidity Constrained Agents

Valuations\Wealth	\$5	\$20
\$20	50	50
\$10	50	50
\$0	50	50

⁹ This part of the current section is based on Che, Kim, and Gale (2011).

The figure in each cell represents the number of population with the particular wealth-valuation profile. That is, there are 50 people who value the apartment at \$20 but have wealth equal to \$5.¹⁰ The valuations represent individuals' user values and do not include investment values; in case of an apartment, those with zero valuation could be interpreted as individuals who already own houses and whose value from having second houses is zero.

It would be socially desirable for high value individuals get the apartment before the low value individuals do. In particular, since there are only 100 units, it should be those 100 individuals with valuation \$20, regardless of their wealth levels. But this assignment will not arise in the competitive market. These people with wealth \$5 only spend only up to \$5 and thus will not be able to compete effectively on the market. It is easy to see the market-clearing price must exceed \$5. If not, all those individuals with valuations \$10 or higher will each demand one unit of apartment, and there are 200 of those buyers, and we have only 100 units of the apartment in supply, so the market cannot clear at \$5. In other words, the market clearing price must be higher than \$5. This means that in equilibrium only the rich (i.e., people with wealth level \$20) with valuations \$10 or \$20 will get the apartment. Clearly this is not efficient since the poor people with valuation \$20 will not get the apartment.

The crucial assumption for this result is the capital market imperfections. If those who are liquidity-constrained can finance their purchase, they will be able to purchase payment as long as they have enough future earning power. Capital markets imperfections are likely to be more pronounced for those resources whose potential value are very large (relative to liquid assets average individuals hold), and the goods mentioned above arguably fall into this category. If markets do not work very well for such goods, then the real question is whether there is any that does work well, or perhaps more relevantly, better than the market. We answer this question with our first application of non-market mechanisms.

4.1 Lotteries with Resale: The Korean Housing Market Approach

An alternative approach to markets is to regulate the price to be capped below the equilibrium level and use lotteries or some other eligibility rules. It would be most desirable if the government can actually observe the user values of the individuals, in which case the socially optimal allocation can be simply enforced. It is unlikely that the government has such an ability to observe individuals' valuations of a good (which can be often subjective), and it cannot rely on the individuals to report them truthfully since they will have the incentives to lie about their valuations.

Alternatively, if the government can observe individuals' wealth levels, the socially efficient outcome can be achieved. To see this, suppose in the above example, the government sells the apartment at

¹⁰ The unit of money can be hundred thousands or million, if such units make the numbers seem more realistic.

\$5 (or equivalently forces the private developers to sell it at that price) *only to the poor individuals* (i.e., ones with wealth \$5) who do not own the apartment (with valuation greater than \$0). Suppose there is no regulation on the resale market.¹¹ As will be seen, the resale price will be strictly greater than \$5, so it pays to participate if you can. Hence, all those low wealth individuals would participate, the 100 of them, and they will receive the apartment. Among those, some will keep the apartment and some others will choose to resell the apartment, depending on the equilibrium resale price of the apartment. It can be seen that in equilibrium only those 50 individuals with valuation \$10 must sell to the rich people with valuation \$20.¹² Hence, in equilibrium only the high valuation individuals end up with the good, which is socially efficient.

Perhaps, it is unrealistic that the government can observe the individuals' wealth levels or disqualify the ones who already own an apartment. It is often not difficult to hide one's asset. It is also not difficult for one to beat the no-prior-ownership eligibility rule; people who own the house, even wealthy ones, can apply for the apartments in the name of their relatives who are not rich and do not own any apartments. Suppose it is difficult to exclude people. In fact, consider the worst possible scenario in which no one can be excluded from applying for the apartment. Clearly, whoever is assigned the apartment at \$5 can at least resell it at a higher price (a fact to be checked later), so everybody will apply for the apartment, and lotteries will have to be used to determine the winners. Since there are 100 apartments assigned to 300 applicants at random, each has $1/3^{\text{rd}}$ chance of receiving the apartment.

The eventual outcome, i.e., who eventually end up with the apartments, will depend on the realization of the lottery as well as the resulting resale process. Without knowing the details, we can establish the following facts. First, every rich individual with valuation \$20 must end up with the apartment. Either such an individual is lucky enough to be assigned in which case he or she will keep the apartment (because the resale price will not exceed \$20). Or else, such an individual will purchase from the resale market.¹³ Next, no individual with zero valuation will end up with the apartment. If such an individual ended up with the apartment, then the resale price must be zero, which cannot occur in equilibrium. Third, some low wealth individuals with valuation \$10 will get it, in fact each with probability $1/3$. Combining these facts, we conclude that the resulting allocation is at least as good as, but more than often better than, the market allocation.

Notice this assignment method (i.e., random assignment with resale) entails the problem of speculation, which has been much observed in the Korean housing market. Individuals with low valuation participate in assignment and succeed with some chance, and when they do, they will end up reselling the

¹¹ This is roughly descriptive of the regulation in place for new apartments in Korea until late 1990's. Random assignment with resale has been employed in other contexts as well, including the Oklahoma Land Opening and the resale of draft exemptions that occurred in the American Civil War.

¹² If not all of the former sells, then the price must be no greater than \$10, and all 50 rich people with valuation \$20 must demand the apartment, so the demand exceeds the supply. If some of those with valuation \$20 sell their assignment, then the price must be at least \$20, so the demand cannot exceed 50, resulting in excess supply. Hence, in equilibrium, precisely of those assigned, precisely the ones with valuation \$10 sell to the rich individuals with valuation \$20.

¹³ For this not to be the case, the resale price must be at least \$20. This means that all those assigned must have valuation \$20, which contradicts the hypothesis that an individual with valuation \$20 did not get assigned.

assignment at profit. Speculative activities are a natural consequence of a system that assigns a valuable good at a below-market price and does not restrict resale. These activities clearly reduce the odds that the individuals with high valuations get the assignment initially, but the problem of speculation does not vitiate the desirability of this mechanism. As noticed, this mechanism entails a better allocation than the competitive market, despite engendering speculation. Further, a simple way to eliminate speculation may not be desirable. For instance, relaxing the initial price cap below the market eliminates the speculative activities, but it will result in the competitive market outcome, which is even less efficient. Likewise, banning resale will eliminate speculation, but will only produce an inferior outcome. In practice, speculation will have other harms that are not reflected in the current analysis, and these issues must be carefully weighed. Nevertheless, switching to the market method need not be a better solution. A better approach will be to restrict the resale (but not prohibit), which will contain speculation problem but generate a better outcome than the market. At the minimum, this simple analysis suggests that the market approach need not be always the best way to assign resources if liquidity constraints are important.

The idea of using a non-market method with resale also raises some intriguing implications for ways in which many social programs may be improved. Many social and welfare programs do not allow the recipients of the benefits to transfer them for money. The logic above suggests that relaxing the transfer ban may improve the overall efficiency of the allocation. For instance, let's go back to the case of immigration visas being randomly assigned. Currently, the visas thus assigned cannot be transferred. It may not be a bad idea to allow for a (supervised) voluntary reallocation within the set of individuals who were considered for the assignment.¹⁴ In fact, the random assignment with resale will likely dominate the Becker proposal of auctioning off immigration visas. The candidates for immigrants, even though high skilled and having much to contribute, may be highly liquidity constrained. In such a case, markets will likely select the wealthy rather than the worthy immigrants. In fact, we have seen that the outcome would be worse than the one resulting from randomly assigning the visas to qualified individuals and allowing them to retrade.

We next consider several assignment problems for which money cannot be exchanged for some reasons. The market design challenge then will be designing in a sense a "market" in the absence of a usual market we all are familiar with. We will emphasize here that, despite the lack of standard markets, understanding of how they work will help us to design an effective and desirable mechanism even in this situation.

4.2 School Choice Design

¹⁴ For this reason, the reallocation procedure need not raise the problem of the visas getting in the hands of the "wrong" people. As long as the initial set of candidates are carefully chosen, there is no difference in this regard between the current system of random assignment without transfer and random assignment *with* one time supervised transfer.

Many countries, including Korea and the US, assign public school seats to students based on their residence; that is, students are assigned to public schools in their neighborhoods. While this residence-based assignment system may be natural, it clearly limits students' choice of public schools. Even the limited choices are not equally available to all families. Wealthy families can buy into good school districts and send their children to good schools, but poor families are often unable to do the same. The school choice initiative aims at expanding the choice of public schools beyond the boundary of a student's neighborhood. Starting with Minnesota in 1987, several school districts (New York City, Boston, Cambridge, Charlotte, Columbus, Denver, Seattle and St. Petersburg-Tampa) have adopted the initiative. In Korea, Seoul began its choice program in 2010.

While there is a broad public support for the school choice initiative, the precise method for implementing choice, i.e., how to assign school seats to students, remains actively debated. A recent debate surrounding the redesign of school choice programs in NY and Boston focused on the incentive issue: *how to elicit truthful reporting of preferences from students*.¹⁵

Before the redesign, the choice program used followed the so-called **Boston mechanism**. In the Boston mechanism, families submit rankings over schools and apply to the schools according to the rankings, starting with one's top choice school. Each school then considers those who chose it as their top choice, and assigns its seats according to its own priority ordering of students.¹⁶ If all seats are assigned, then stop; otherwise, the school, along with other schools with remaining seats, move down to the next group of students who have been turned down by their top choice schools, and assign seats again according to its priority order, and so on. This process continues until either all seats or all students are assigned.

Notice here how a student lists a given school in her ranking matters for her chance of assignment at that school. But this feature destroys the students' incentives for truthful reporting. To see this, suppose a school is a top choice of a large number of students, in fact much larger than the number of seats of that school. Some of these students may not choose to list that school as their top choices. Suppose a student indeed lists that school as her top choice. Even with the school listed as the top choice, its popularity among others means that her chance of getting assigned to that school would be still very small. But when failing to get in to that school, her chance of getting assigned to the second best school will be small as well since she will have a lower priority than the students who may list that school as the top choice. If she were to list the second-best school as the top choice, she may at least secure her placement at that school. Hence, a student in this case may find it in her best interest to manipulate her preference rankings. Such an incentive to "game the system" presents serious practical problems. First of all, it becomes very difficult for school administrators to advise the parents. In fact, it may not be appropriate for the administrators to advise

¹⁵ There are extensive discussions on this issue now in the context of school choice design. See Abdulkadiroglu and Sonmez (2003) and Abdulkadiroglu et. al (2005a,b).

¹⁶ Schools' priority ordering of students may come from their test scores in the entrance examinations, SAT scores, or considerations such as whether a given student lives close to the school and/or whether the student has a sibling who is already attending the same school.

them to report truthfully, knowing well that that may not be in the best interest of the students. Even worse, not all families may be equally strategic or equally sophisticated in strategic behavior, which raises a fairness issue (see Abdulkadiroglu et al, 2007).

The incentive problem was an important motivation when the two school systems switched to the **Gale and Shapley's Deferred Acceptance** (henceforth, **DA**) algorithm. The DA algorithm works as follows: Just as with the Boston mechanism, students submit their preference rankings. Again, schools first consider students who listed them as their first choice, and select among them according to their priorities up to their capacity, and reject the others. So far, the procedure is the same as the Boston mechanism, except that the assignment at this point is only *tentative*, whereas the rejection is *final*. In the next round, schools consider the students who were turned down by their top choice schools and listed them as their second best choice. The schools then reselect students from those held from the previous round and the new applicants, again according to their priorities (i.e., without any discrimination between the two groups). This process continues until no students are rejected or there are no more schools to apply to. In one of the two latter cases, the process is complete, and the tentative assignment at that point becomes final.

Notice here, unlike the Boston mechanism, how a student lists a given school in her ranking doesn't affect her chance of assignment at that school, once she gets to that school in the process. This means that listing the best school as top choice doesn't sacrifice her shot at less preferred schools. Hence, this mechanism generates the incentive for the participants to report their (ordinal) preferences truthfully. Along with the incentive for truthful reporting, the DA algorithm also ensures fairness of allocation in the sense of eliminating justified envy: that is, under DA, no student can lose a seat of a school to another student who enjoys a lower priority. Further, the DA assigns to the best possible schools in any allocation that eliminates justified envy.

4.3 Kidney Exchange

Human organs such as kidneys are in constant short supply. As of 2005, there were more than 70,000 patients waiting for a kidney transplant in the US. Among them, only 16,500 received transplants, 9,800 from deceased donors and 6,570 from living donors, while 29,160 new patients joined the deceased donor waiting list and 4,200 patients died while waiting for a kidney.¹⁷ The primary reason for short supply is the lack of markets and monetary rewards for the donors. But there is also another subtle problem.

Patients in need of kidney transplants often have families or relative willing to donate their kidneys. Such donors are unable to donate their kidneys, however, if their blood-types or tissue-types are incompatible with those of their intended recipients. In this case, the patient must look for either a new

¹⁷ These figures are from SRTR/OPTN national data retrieved at <http://www.optn.org> on 2/27/2007.

donor or go back to the waiting list for a deceased kidney. Equally important, the donor's good will is wasted. Isn't there any way to save it?

There is a market-inspired solution for this problem. Suppose patient A_1 has a willing but incompatible donor A_1' . Suppose there is another incompatible patient-donor pair (A_2, A_2') in a similar situation. Suppose further that donor A_2' is compatible with patient A_1 , and donor A_1' is compatible with patient A_2 . Then, the two patients can simply exchange their donors. Both donors now donate their kidneys, and they do so, albeit not directly, to save the lives of their intended recipients. Exchanging incompatible donors can be seen as a form of a market, yet it can sidestep the moral objections since it does not involve monetary transfers. The value of such exchange is clear: it enables one to realize many donations that would not otherwise materialize.

Once kidney exchange of this simple form is allowed, one can imagine a much broader class of exchanges. For instance, suppose one finds a cycle of incompatible patient-donor pairs, $(A_1, A_1')-(A_2, A_2')-\dots-(A_k, A_k')$, such that A_i' is compatible with A_{i+1} , for $i=1, \dots, k-1$, and A_k' is compatible with A_1 . Then the exchange can take place along the cycle, benefiting all patients involved in the cycle. In practice, the length of the cycle required for efficiency need not exceed 4 (Roth, Sonmez, and Unver, 2007). The idea of exchange can be applied even when no such cycle can be found. For instance, suppose there is a chain of incompatible patient-donor pairs, $(A_1, A_1')-(A_2, A_2')-\dots-(A_k, A_k')$, such that A_i' is compatible with A_{i+1} , for $i=1, \dots, k-1$, but that A_k' is incompatible with A_1 . One can in such a situation motivate A_1' to donate to A_2 ---thus starting a long chain of transplants---by upgrading A_1' 's position in the queue for deceased kidney. Kidney exchanges of these forms have been studied recently; and the ideas have been put into practice at least in a limited scale.¹⁸

4.4 Course Assignment

At universities or even high schools, students are allowed to choose courses. The problem is that each course comes with a capacity constraint because the large class is either infeasible or undesirable. Popular courses are usually in high demand, so their capacity constraints are acutely binding. A typical procedure used in a situation like this is to follow "first-come-first-served" and close the class when the capacity is filled. But this is unlikely to be efficient; students must differ in their values from taking different courses and their ability to benefit from different classes. These preference differences are unlikely to be accounted for in the first-come-first-served method.

A few Business schools have adopted innovative procedures that take students' preferences into consideration. A simple such procedure is to use a variant of **random serial dictatorship** (henceforth, **RSD**). RSD orders the students uniformly randomly, and assigns the students in sequence of the realized order their

¹⁸ See Sonmez and Unver (2009) for an excellent survey in research on kidney exchange.

most preferred set of courses that are remaining at their turns. The resulting allocation is efficient relative to students' preferences in the sense that no reallocation of the courses can make all students uniformly better off. The assignment is also *strategy-proof*, meaning that the students have dominant strategies to report their ordinal preferences truthfully. But RSD is not fair ex post since the students with the favorable draws of serial turns enjoy the best possible courses, and those with unfavorable turns end up with the worse courses. A **random draft dictatorship (RDD)** appears to be much more fair. Like RSD, RDD orders the students uniform randomly but assigns each student only one course at his turn, and repeats this process several rounds.¹⁹ This method resembles the draft system that professional leagues employs to assign athletes their teams in the rookie markets, and has the benefit of limiting the ability by any team to monopolize several best players. RDD possesses the same fairness benefit in the course allocation context, although the strategy-proofness property is lost.²⁰

Another idea is to endow each student with a predetermined budget in artificial currency points, say 100 tokens, and allow the students to purchase probabilistic shares of alternative courses at prices that are computed to clear markets. This method---originally suggested by Hylland and Zeckhauser (1979) further developed by Budish, Che, Kojima and Milgrom (2011) into the course allocation context---applies the classical idea of competitive markets to implement a desirable outcome. By the first-welfare theorem, the allocation of courses is Pareto efficient ex ante (i.e., in terms of the probabilistic shares), and is fair in the sense of envy free.

4.5 Possession Assignment in the NFL Overtime Games²¹

In a National Football League overtime, a coin is tossed to determine which team will receive the kick off and thus the first-possession. In its current sudden-death format, whichever team scores the first, even with a field goal, wins the game.²² For this reason, acquiring starting-possession confers a team a significant advantage in odds of eventual winning. This feature makes the subsequent play after coin-flipping something of an anti-climax. It would be desirable for fans, and ultimately for the league, if the teams were not treated differently based on an ad-hoc procedure such as coin flipping, but rather if they were treated more evenly based on their natural strengths. But how do we ensure fairness in this latter sense? The difficulty is that one team will have start with the possession. The question is which team ought to have a possession and at what yard from his end line.

An ingenious solution to the problem is to introduce the idea of market. But how? One can use the yard from one's own end line as a currency to trade the possession. Here is how one can implement this

¹⁹ See Budish and Cantillon (2010) for examining this mechanism.

²⁰ In early rounds, a student may not claim her most-preferred course if it is not very popular, but rather claim the most popular course, knowing that such a course will be demanded by others and disappear in the early rounds.

²¹ The discussion here is largely based on Che and Hendershott (2006, 2009).

²² Due to the controversy discussed here, the NFL league has now instituted a new rule that requires scoring by a touchdown to win overtime in a playoff game. In a regular season game, however, scoring by a field goal wins overtime.

idea specifically. The opposing teams (say their coaches) can bid on the yard line from its one end line, namely offering to start its offense at that yard line. Possession of the football close to one's end line is disadvantageous for the team since it requires moving the ball a lot of distance to score, whereas possession of the ball far away from its end line is advantageous. The team that makes the low bid (in yards) is then declared as the winner and begins its offense at that yard line.

This market approach to setting the overtime rule has some desirable property. The odds of winning game when a team begins its offense at the equilibrium bid is equalized to its odds of winning when the team begins its offense at exactly the same yard line (from the team's own end line). As long as offense at a given yard line is more likely to result a victory than a defense against the opponent beginning at that yard line (from its own end zone), a team will lower its bid until that equality is restored. Hence, there is no regret and envy in equilibrium. In particular, if teams are evenly matched in skills and strength, the bidding will result in the equal chance of winning for each time at the start of the over time play.

5. Incorporating Psychology and Behavioral Science into Market Design

Market design has so far proceeded from both the mechanism design and social choice paradigms, both of which usually take as an accepted hypothesis the traditional perception about human behavior, namely that of pursuing self-interest with unlimited rationality. Yet, studies in psychology and recent behavioral economics have taught us how the actual human behavior may depart, sometimes dramatically, from this traditional perception and, more importantly, how the failure to appreciate that departure may result in misleading policy implications. In this regard, one important avenue of future research is to incorporate the lessons from psychology and behavioral sciences into market design. While still lagging in research, the practice of market design already recognizes the importance of behavioral and psychological aspects of human behavior. Let us provide two cases in point.

Online matchmaking is a dating service that recommends its participants whom they may wish to meet, make friends with, date, and possibly marry. Not surprisingly, identifying the right match is the key to its successful operation. Matchmaking models, such as match.com, eHarmony.com, and chemistry.com, have developed elaborate algorithms to accomplish this goal. Traditionally, the algorithms identify a match based on personal preferences and personality type data obtained via extensive questionnaires that users fill out. These models classify people in several dimensions of personality they developed and offer a list of people who are likely to be compatible with each other. One significant issue with this approach, however, is that the participants' survey reports are subjective and thus often involve biases, both intentional and unintentional. It is difficult to verify whether what users say about themselves are what they really are, and also sometimes users are not clearly aware of what they prefer.

To address these issues, some matchmaking models now turn to a more behavioral based approach, similar to a system used by Amazon and Netflix, namely collaborative filtering. For example, a system developed by IntroAnalytics identifies users' preferences based on users' behavior on the site without survey data (Swain and Devlin, 2011). The system may look at who an applicant clicked to check out detailed information about, who he/she actually sent out emails for contacts, how many email exchanges he/she had with the person he/she contacted before meeting in person, and so on. The idea is that individuals often can "act" more revealingly than they can "articulate" about their preferences. This method can filter out some of the biases that a survey method may overlook.

Even such a behavioral-based approach can be further informed by the psychology. Take the case of the highly-publicized Netflix contest. Netflix is a video-rental Internet company, well-known for its effective system of recommending movies to their users, *cinematch*. This system relies on a form of collaborative filtering, namely recommending users the movies that other users who share a similar behavioral history have rented. The idea is that the behavioral history is a good predictor of one's preferences, so those with a similar behavioral history will agree on what they consider to be good movies. In October 2, 2006, Netflix announced that it would run a contest to improve *cinematch*. Netflix would share the database of its user behavior and the contest would end with the model that improves the accuracy of *cinematch* by 10%. The winner would receive a prize of \$1million and the continued ownership of intellectual property of an improved recommendation system (Lohr, 2009).

One of the contestants was Gavin Potter. His model is based on collaborative filtering but it had one distinctive component that sets it apart from the other models. It incorporated human psychology into the mathematical model, in particular the psychological bias called "anchoring," the concept first theorized by Tversky and Kahneman (1974). Anchoring refers to the bias in human judgment that people tend to anchor the subsequent evaluation on the evaluation that they start with. In other words, a positive or negative score a user may give to the first movie that he or she rates will inflate or deflate the rating of movie the user will evaluate next. The anchoring bias also entails inertia in the sense that users tend to give rating similar to what they gave recently. Potter's model tried to find a way to account for these biases (Ellenberg, 2008).

Although Potter did not win, the idea of incorporating psychological factors helped the winning submission by BellKor. The final revision of BellKor's model accounted for similar kinds of judgment biases, so called temporal variability in the data (Bell, Koren, and Volinsky, 2010). This reflects the fact that movies' popularity may change over time and that users baseline rating can change over time (e.g., users used to give 3 stars to a movie that they rated average but now they give 4 stars for average). Integration of these two aspects was the major improvement in the BellKor's final submission that won the contest.

These cases illustrate that incorporating human psychology and behavioral sciences can fruitfully inform and guide the design of some markets. It remains unclear, and thus a promising future research

subject, to what degree the psychological factors matter for different problems of market design and how they can be effectively incorporated in the design of markets.

References

- Abdulkadiroglu, A., Pathak, P., Roth, E., and Sonmez, T. (2005a), "The New York City High School Match," *American Economic Review, Papers and Proceedings*, 95, 364-367.
- Abdulkadiroglu, A., Pathak, P., Roth, E., and Sonmez, T. (2005b), "The Boston Public School Match," *American Economic Review, Papers and Proceedings*, 95, 368-371.
- Abdulkadiroglu, A., Pathak, P., Roth, E., and Sonmez, T. (2007), "Changing the Boston Mechanism: Strategyproofness as Equal Access," Unpublished manuscript, Harvard University.
- Abdulkadiroglu, A. and Sonmez, T. (2003), "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729-747.
- Bell, R. Koren, Y., and Volinsky, C. (2010), "AT&T Labs Research — Leading Invention, Driving Innovation," May 19, http://www.research.att.com/articles/featured_stories/2010_05/201005_netflix2_article.html.
- Budish, E., and Cantillon, E. (2011), "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," unpublished working paper.
- Budish, E., Che, Y.-K., Kojima, F., and Milgrom, P. (2011), "Designing Random Allocation Mechanisms: Theory and Applications, unpublished working paper, Columbia University.
- Che, Y.-K. (1993), "Design Competition through Multidimensional Auctions," *RAND Journal of Economics*, 24, 668-680.
- Che, Y.-K., Choi, S., and Kim, J. (2011), "An Experimental Study of Sponsored-Search Auctions," unpublished working paper, Columbia University
- Che, Y.-K., Kim, J., and Gale, I. (2011), "Market versus Non-Market Assignment of Ownership," unpublished working paper, Columbia University.
- Che, Y.-K. and Hendershott, T. (2008), "How to Divide the Possession of a Football?" *Economics Letters*, 99, 561–565.
- Che, Y.-K. and Hendershott, T. (2009), "The NFL Should Auction Possession in Overtime Games," *The Economists' Voice*, October.
- Edelman, B. and Ostrovsky, M. (2007), "[Strategic Bidder Behavior in Sponsored Search Auctions](#)," *Decision Support Systems*, 43, 192-198.
- Edelman, B., Ostrovsky, M. and Schwarz, M. (2007), "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97, 242-259.
- Ellenberg, J. (2008). This Psychologist Might Outsmart the Math Brains Competing for the Netflix Prize. February 25, http://www.wired.com/techbiz/media/magazine/16-03/mf_netflix.

- Hylland, A., and Zeckhauser, R. (1979), "The Efficient Allocation of Individuals to Positions." *Journal of Political Economy*, 87, 293–314.
- Kwerel, E. and Williams, J. (1993), "Moving toward a Market for Spectrum," *Regulation*, 2, 53-62.
- Lohr, S. (2009), "Netflix Awards \$1 Million Prize and Starts a New Contest," NY Times, September 21, <http://bits.blogs.nytimes.com/2009/09/21/netflix-awards-1-million-prize-and-starts-a-new-contest/>.
- Milgrom, P. (2004), *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.
- The Economist, 2010, "The Price of Entry: a new proposal from Gary Becker to make a market in immigration," http://www.economist.com/node/16424085?story_id=16424085&CFID=168980692&CFTOKEN=59876854.
- Roth, A. (2002), "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics," Fisher-Schultz Lecture, *Econometrica*, 70, 1341-1378.
- Roth, A., Sonmez, T., and Unver, U., (2007), "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences," *American Economic Review*, 97, 828-851.
- Sonmez, T. and Unver, U. (2009), "Market Design for Kidney Exchange," unpublished working paper, Boston College.
- Swain, F. and Devlin, H. (2011), "Connected: How maths is helping us to get it on.. online," *The Times*, London, February 3, p. 46
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1130.
- Varian, H. (2007), "Position Auctions," *International Journal of Industrial Organization*, 25, 1163-1178.