Summary: Inaugural Meeting of the Measuring and Interpreting Inequality Working Group

February 18, 2012

Summary prepared by Xan Vongsathorn

I. Overview

The conference initiated a conversation between market designers and experts on inequality. Simply allowing prices to clear a market may be deemed unfair by policymakers hoping to reduce inequality; in such cases, the tools of market design provide a natural way of *targeting* alternative social objectives. By the same token, the inequality literature's empirical approach to evaluating outcomes can be of great benefit in *assessing* the effectiveness of mechanisms developed by market designers.

Taken together, the conference papers contributed to a more complete picture of the current state of inequality, and of the tools available to influence inequality as it evolves over time. At the same time, the conference illuminated avenues for further research. It is hoped that the MIE meeting— and the work of HCE more generally—will lead to new work crossing field boundaries to understand and address socioeconomic inequality.

II. Connection between Market Design and Inequality

When inequality is a concern, free market solutions to allocating scarce resources may be considered unfair. How should resources be allocated in these settings, taking policy objectives into account? Market design provides a natural framework for carefully studying this problem. The tools of market design can be used to evaluate many of the *ad hoc* mechanisms that have arisen independently throughout the world. Market design also suggests alternative mechanisms which may better target policy goals.

A leading example–and a topic that recurred throughout the conference–is the role of market design in *school choice*. How should students be matched to schools? School choice programs must account for the preferences of students while respecting policy goals.

Many of the conference presentations considered how to incorporate *affirmative action* into school choice. Affirmative action addresses existing inequality by favoring disadvantaged groups; however, the best way to implement this goal is not obvious. Poorly chosen mechanisms may be needlessly inefficient, and may hurt many (or even all!) of the intended beneficiaries.

Manipulability

As a benchmark, we consider a basic design change which can help "level the playing field" in school choice.

In the *Boston mechanism*, which has been adopted in many school districts, students submit rank-order lists over schools, and seats at each school are assigned giving priority to students who ranked that school higher on their lists. That mechanism invites strategic behavior on the part of the students. By contrast, the *student-optimal stable mechanism*, in which student priority does not depend on student rank-order lists, is *strategyproof* for students, in the sense that truthful reporting of preferences is always optimal.

Strategyproofness promotes equality: Mechanisms which encourage strategic behavior will tend to favor students with the resources and ability to strategize. By contrast, a strategyproof mechanism does not favor sophisticated players.

Sonmez built upon and extended this point in his presentation of joint work with Parag Pathak, "School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation." He defined a natural measure of relative *manipulability* of mechanisms; under this metric, reductions in manipulability naturally promote equality.

To date, many school districts have not strategyproof mechanisms; nevertheless, a number of them have recently modified their mechanisms, often appealing to concerns about inequality and fairness. Sonmez presented data showing a strong trend toward *less manipulable* mechanisms in the United States and England, even when market designers were not involved. This trend is in line with the prescriptions of market design, and perhaps validates the goal of moving even further in the direction of strategyproofness.

Affirmative Action in Theory…

While strategyproofness and the student-optimal stable mechanism are attractive to market designers even absent inequality concerns, mechanisms may also be designed for the explicit purpose of reducing inequality.

Many school districts wish to match students to schools in a way that promotes diversity and benefits disadvantaged groups. A natural approach is to constrain the number of students of a given type who may attend a given school; some form of affirmative action approach is common practice in many cities. But what sorts of affirmative action constraints are most effective? What precise mechanism should we use? The answer is not immediately clear.

Towards an answer to this question, Scott Kominers presented joint work with Tayfun Sonmez on "Matching with Slot-Specific Priorities." In Chicago, students are matched to a set of selective public high schools in a simple way that incorporates affirmative action constraints. Students are ranked based on academic achievement and separated into four "tiers." There is open competition for 40% of the seats at each school, and 15% of the seats are reserved for each of the four tiers. In effect, the student-optimal stable mechanism is run, with the open seats filled before the reserved seats.

Kominers and Sonmez observed that the order in which seats are filled matters: If open seats are filled first, top minorities compete with majority students, whereas if open seats are filled second, lower-ranked minorities compete with majority students. We might therefore expect that the method used in Chicago is especially effective at promoting minority opportunity. Confirming this intuition, Kominers and Sonmez showed empirically that when open seats are instead filled last, about 18% of the seats change hands, with the most disadvantaged students being the most impacted.

… and Affirmative Action in Practice

Evaluation based directly on outcomes has thus far been largely *absent* from market design. If a mechanism is designed to boost the academic achievement of disadvantaged students, might we not want to consider the empirical effect of the mechanism on academic achievement?

Dennis Epple does exactly this, in joint work with Surendrakumar Bagde and Lowell J. Taylor, "Dismantling the Legacy of Caste: Affirmative Action in Indian Higher Education." The paper considers the impact of affirmative action in India on college enrollment rates and subsequent test scores. (Remarkably, the mechanism that matches students to colleges in this market is related to the Chicago mechanism studied by Kominers and Sonmez.)

The intended beneficiaries of affirmative action policies in India are the lower castes. But are they really helped, and if so, by how much? A market design evaluation would ask whether these students get into more preferred colleges as a result of affirmative action. However, students may not know what is good for them. Indeed, much of the rhetoric surrounding this controversial policy charges that affirmative action may inadvertently harm members of the lower castes, by inserting them into challenging environments for which they are not prepared. To settle the question, information on student rank-order lists is not sufficient. By appealing to outcomes directly, however, Epple and his coauthors argued that affirmative action does in fact benefit the lower castes. He presented evidence that the program increased both college attendance rates and subsequent test scores among the members of the lower castes. In addition, this approach can be used to quantify the magnitude of the gain.

Limitations of Affirmative Action and Priority Structure

Market design has more to say on the topic of affirmative action. Policymakers' true objectives are typically quite complex, and promoting diversity may come at the cost of other objectives such as nonwastefulness, fairness, and strategyproofness. Yet in many school districts, affirmative action constraints are simply mandated in the form of hard upper and lower bounds on the number of students of a given type at a given school. Two natural lines of inquiry follow from this. How can we implement the constraints imposed by law while best satisfying other policy goals? And might there be a mechanism that better achieves the true underlying policy objectives when legal constraints are relaxed in an appropriate way?

M. Bumin Yenmez addressed both of these questions in joint work with Lars Ehlers, Isa E. Hafalir and Muhammad A. Yildirim, "School Choice with Controlled Choice Constraints: Hard Bounds Versus Soft Bounds." Unfortunately it is not always possible to satisfy hard bounds as well as a natural notion of fairness and nonwastefulness. In light of this, Yenmez and his coauthors

presented a variant of the student-optimal stable mechanism that respects hard bounds and satisfies weakened notions of fairness and nonwastefulness. As this mechanism is not strategyproof for students, Yenmez and coauthors devised to a mechanism that achieves fairness, nonwastefulness, and strategyproofness when bounds are allowed to be *soft* in a certain sense. The latter mechanism encourages diversity, without forcing it to hold when student preferences are significantly out of line with policy-inspired constraints. In such cases, soft bounds may produce a result that is actually *more* in line with the policymaker's true objectives.

Controlled choice mechanisms ideally strike a balance between accommodating student preferences and achieving distributional goals. Indeed, school choice is not *choice* if the outcome is essentially determined by the priority structure imposed on the students! On a cautionary note, Caterina Calsamiglia presented joint work with Maia Guell and Antonio Miralles, "All About Priorities: On How There is No School Choice Under the Presence of Bad Schools." This work first considered school choice in Barcelona, where students are given priority at nearby schools. Under the Boston mechanism, the safe option is to rank highest a school at which the student has priority, and indeed most students do exactly this. Calsamiglia and coauthors used a change in the definition of "nearby schools" to show empirically that the outcome of the mechanism is highly sensitive to the priority structure. In this sense, student preferences play a small role relative to priorities; this provides a new reason why policymakers may wish to avoid the Boston mechanism.

Calsamiglia also discussed a theoretical model with a continuum of students and a binary priority structure, in which all students fear being placed in the same low-quality school. If the bad school is sufficiently undesirable, then under the Boston mechanism all students apply first to the school at which they have highest priority – irrespective of their preferences. Furthermore, under the student-optimal stable mechanism, students simply end up at the school for which they have highest priority. This model illustrates that priorities can potentially have significant impact on students' assignments.

Effects of School Choice on School Quality

John William Hatfield presented joint work with Fuhito Kojima and Yusuke Narita, "Promoting School Competition Through School Choice: A Market Design Approach." With student welfare as a top priority, the school choice literature has promoted the student-optimal stable mechanism (which gives the most preferred stable outcome for students) and the top trading cycles mechanism (which is Pareto efficient for students). However, while school choice has been hailed as a way of incentivizing higher school quality through increased competition for schools, past analysis has assumed school quality to be fixed.

In fact, the choice of mechanism may impact schools' incentives to improve themselves. To address this, Hatfield and coauthors identified circumstances under which a school is made better off when it becomes more attractive to students. Unfortunately, *no* stable or Pareto efficient mechanism always guarantees increases in school welfare following such quality improvements. On the other hand, stable mechanisms nearly have this property in large markets, while the top trading cycles mechanism does not. This gives a reason to prefer the student-optimal stable mechanism to top trading cycles.

Incentives for school quality are potentially quite relevant to inequality concerns. In particular, we should be particularly concerned if a mechanism enables schools to benefit by decreasing quality in the eyes of students they find *undesirable*, especially to the extent that these are disadvantaged

students. There is room for future empirical work considering the impact of school choice mechanisms on school quality, potentially giving special attention to changes in quality for different types of students.

III. Matching and the Evolution of Human Capital

In his invited address, Lones Smith discussed joint work with Axel Anderson, "Matching with Evolving Human Capital." Variation in human capital is both a form of inequality and a root cause of income inequality and other outcomes of concern. Thus the evolution of inequality over time is closely related to the underlying evolution of human capital.

Smith's joint work considers a dynamic one-to-one matching model in which agents match each period. Matched agents engage in joint production and augment their future human capital; both processes depend on current human capital levels. An agent's current match may be beneficial or detrimental to future human capital. And because future human capital matters, wages need not always increase with an agent's quality.

Whereas the school choice papers discussed above have focused on matching markets with non-transferable utility, the starting point for Smith and Anderson's work is Becker's 1973 theory of marriage, a static matching model with transferable utility. By adding in dynamic peer effects, Smith and Anderson endogenize the wage and human capital distributions. This allows them to consider different properties of the production and human capital transition functions, and their consequences for matching and the subsequent evolution of human capital.

After specializing to CES transition functions, Smith described four matching patterns and presented the results of simulations. Similar to Becker's 1973 model, positive (negative) assortative matching obtains when there is complementarity (substitutability) in the production and transition functions. In addition, two new patterns can arise when the production and transition functions are in tension with each other: individuals may separate into skill-based "clusters," with members of a given cluster rising to the top of the human capital distribution within the cluster—but no further. These patterns match a number of observed cross-sectional and dynamic wage patterns and suggest a reason for glass ceilings and poverty traps.

IV. Measurement Issues

The conference was bookended by two talks that highlighted issues of measurement.

Trending inequality is a matter of national importance, and sound measurement of inequality is a precondition for sound policy. At the same time, inequality is a multidimensional issue, and it is hardly obvious how to construct measurements that faithfully capture it. Addressing this, Bruce D. Meyer presented joint work with James X. Sullivan, "Consumption and Income Inequality in the U.S. since the 1960s." Meyer and Sullivan's work demonstrates that both the level and pattern of inequality are sensitive to the method of measurement, and suggests numerous shortcomings of previous measurements. Income measures have typically failed to account for taxes or transfers, while consumption measures emphasized nondurable consumption or looked only over short time

horizons. In addition, official statistics often focus on the extreme upper tail, even though these individuals make up a very small share of the population.

Meyer and Sullivan considered income inequality after taxes and transfers, and developed an improved measure of consumption that incorporates the flow value of durables. With these modifications, changes in inequality over time become much less pronounced. In particular, consumption inequality has been relatively flat since the mid-1980s. However, the large increase in inequality observed in the late 1970s and early 1980s appears robust to these measurement changes. In his talk, Meyer advised that more work is needed to understand what happened during that period.

Finally, the structure and evolution of inequality over time may be closely tied to the *social network* in which agents are embedded. In many settings it is natural to suppose that economic opportunities are a function of one's connections; moreover, these connections may endogenously influence future connections and opportunities. Thus it is valuable to understand the process by which the social network itself evolves.

In this vein, Bryan Graham presented "Homophily and Transitivity in Dynamic Network Formation." Graham introduced an empirical model of network formation that considers how the evolution depends on *homophily* and *transitivity*, as well as the current state of the network. Homophily is the desire to connect with similar agents ("birds of a feather flock together"), while transitivity describes the tendency to connect with nearby agents ("a friend of a friend is a friend"). In principle these can both have implications for inequality. Agents may be drawn to other agents of similar socioeconomic status (homophily), and it may then be difficult to expand beyond nearby agents (transitivity).

Networks are complex objects in which a single link can potentially affect the formation of many future links. However, with many observations of the network, we are effectively able to hold the link history fixed outside of a small neighborhood of a given link. By studying the frequency of link formation within these neighborhoods, we can learn about the determinants of network formation.