

Is Policy Too Important to be Left to Empiricists? Lessons of the 2012 Nobel Prize in Economics

Richard Cornes and José A. Rodrigues-Neto¹

Abstract

Fifty years ago, a paper entitled ‘College Admissions and the Stability of Marriage’ was published in a somewhat obscure journal, the American Mathematical Monthly (currently a ‘B’ journal, according to the Australian Business Deans Council). The research program and policy developments that have flowed from that abstract and apparently slight seven-page paper recently led to the award of the 2012 Nobel Prize for Economics to one of its authors, Lloyd Shapley. (Shapley’s co-author, David Gale, died in 2008.) Shapley shared the Nobel Prize ‘for the theory of stable allocations and the practice of market design’ with US economist Alvin Roth, who has been responsible for much of the applied work that has built on Gale and Shapley’s insights. The history of the path leading from the abstract Gale/Shapley insights to the design of resource allocation mechanisms in 2012 is a fascinating and instructive one for many reasons. This article tries to give the reader an idea of what this literature is about, and of the many ways in which Matching Theory has led to real improvements in the design of operational resource-allocation mechanisms.

Historical background

Gale and Shapley were interested in situations in which a group of agents want to be matched with one another. The agents may be prospective marriage partners, firms and workers, schools and schoolchildren, universities and students, hospitals and junior doctors, and so on. Of course, institutions and individuals faced the need to be matched to one another long before the Gale/Shapley contribution. Decentralised procedures for college admission evolved in which colleges would independently solicit applications from students, and make offers to those considered most acceptable. This arrangement has several

¹ The Research School of Economics, The Australian National University; Richard.Cornes@anu.edu.au. We would like to thank the editor and referees for detailed comments. All remaining errors are our own responsibility.

drawbacks, of which one in particular regularly arose. If one institution wants to 'steal a march' on its competitors and attract the best students, it may be tempted to do so by getting in with an earlier offer, and trying to obtain binding acceptances from students. Thus, over time, one might expect to see institutions gradually making earlier and earlier offers. This is indeed what happened — by 1945, some hospitals in the US were making offers to medical students up to two years before they had graduated from medical school! Such unravelling of the market, as this process has come to be called, is a common form of market failure, and generates substantial inefficiency. Although each institution may correctly perceive it to be in its interest to move earlier than its competitors, the outcome is that each thereby forgoes any useful information about that student's quality that may emerge during those subsequent two years — information which, had it been available, might have led to an outcome in which some students and/or institutions obtain more desired matches. Ignoring such information may well leave all agents — both students and hospitals — worse-off.

In response to such inefficiencies, in time, clearinghouses evolved. Under such an arrangement, each institution may give up the right to make independent offers, instead committing itself to working through a central agency — the clearinghouse — which operates under prescribed rules and procedures concerning the submission of preferences over institutions by prospective students and the identification of acceptable students by institutions. Eventually a matching of students and colleges is achieved. Central clearinghouses evolved in various settings in a somewhat ad hoc way — one of the authors of the present discussion once applied for an undergraduate place in the UK through precisely such a clearing house, then known as UCCA (Universities Central Council on Admissions). Some seemed to work quite well, others less well, and some were quite short-lived. But at that time there was no clear understanding of what, precisely, might constitute a satisfactory clearinghouse, why some seemed to work well while others did not, and how one might go about analysing, let alone designing, effective clearinghouses.

Gale and Shapley changed that. They established an approach that has provided the accepted framework for the analysis of clearinghouses that involve no explicit money transfers — in this sense they are quite distinct from the typical commodity market — but which seek simply to attain a satisfactory matching of institutions and individuals. First, they stated clearly what one might mean by a 'satisfactory' matching. To this end, they introduced the idea of a 'stable matching'. Second, they showed that under certain conditions a satisfactory matching would certainly exist. Finally, and perhaps most importantly, their demonstration of the existence of such satisfactory mechanisms took the form of a simple algorithm that actually found such an outcome. Their proposed algorithm has many attractive properties.

A statement of the college admission problem

Suppose there are m students applying for college admission and n colleges available, each one of them offering some positions/places for new students. (Typically, m and n are different positive integer numbers.) Each student has a complete ranking over the colleges, based on her individual views and goals. Similarly, each college has a complete ranking over the new students. This may be based on information provided by students about academic attainments, but may well be influenced by other considerations — different colleges usually come up with different rankings over the students. What matters is that this ‘market’ has two sides made of indivisible units: students and positions. Each member, on each side of the market, has her individual complete ranking of agents on the other side. Ties are not allowed.

A matching is a list of which college each student will attend. This list may say that a particular student does not have a college; or it may say that a college does not have all (or any) of its positions filled. In other words, a matching is a list of matches between students and colleges. We permit each college’s list of students to exclude some students. That is, each college is allowed to deem some students unacceptable, meaning that the college prefers to stay with some unfilled posts rather than admitting what it believes to be unacceptable students. Similarly, each student may find some colleges unacceptable. We say that an ‘agent’ (a student or a college) *blocks* a proposed allocation if in the proposed matching the agent is to be matched to an unacceptable pair, according to her preferences.

In this situation, is there a natural matching/allocation? That is, what would be natural or efficient or fair outcomes? How should we match/assign students to academic places? Some outcomes could be very favourable to some or even to all students, while other outcomes would be better for colleges. How should we combine the preferences of all agents in both sides of the market? Depending on how a mechanism does this, there could be incentives for misrepresentations of preferences; that is, for strategic manipulations. How could we prevent agents from ‘gaming the system’?

An interesting consideration is that, on each side of the market, agents have idiosyncratic preferences. There may be some correlation between preferences or it may be the case that each agent ranks the other side of the market in a completely different way. There is not a unique ‘correct’ ranking.

Gale and Shapley argued that a desirable property for an allocation is that it would be not possible to find a student i and a college j such that i prefers j to her prescribed college and, at the same time, college j prefers i to at least one of its prescribed students. In other words, no pair of agents blocks the proposed/

prescribed matching. They defined a matching to be ‘stable’ if it is not blocked by any pair of agents and, at the same time, the individual rationality condition of every agent is satisfied; that is, no single agent blocks the proposed allocation.

A special case: The Stable Marriage Problem

The easiest of the original problems posed by Gale and Shapley is the ‘Stable Marriage Problem’. This is a particular case of the College Admissions Problem where each college can admit only a unique student. In other words, the matching allocation is one-to-one. Each agent on each side of the market matches with at most one agent on the opposite side of the market.

The original paper refers to one side of the market as ‘females’ while the other side would be composed of ‘males’. As before, each female ranks all males up to a point, considering the remaining males as unacceptable (she would rather be alone rather than matched to any unacceptable male), and, likewise, each male identifies and ranks the females that he considers acceptable.

Gale and Shapley’s contribution

We have already drawn attention to the first of Gale and Shapley’s contributions — their statement of stability as a desirable characteristic of a matching outcome. In addition they suggested an algorithm that allows us to find stable outcomes in the Stable Marriage Problem. This is the famous Deferred Acceptance Algorithm. A variation of this algorithm generates a stable matching in the more complex Stable College Admissions Problem.

They proposed the use of a centralised mechanism for discrete resource allocations. This might explain the lack of interest by US economists in the 1960s and 70s. The efficiency of free markets was a cornerstone of economics and the Arrow-Debreu Theorem had just made great academic impact.

The set of all stable allocations — that is, allocations such that no subset of agents can collude to make all agents in this subset better off — corresponds with the notion of the core, a concept that was systematically developed during the 1950s and 60s within the research program on General Equilibrium. In any coalitional game, an allocation is in the core if there is no blocking coalition of any size. In the matching context, for a matching to be stable, there can be no unilateral deviation or any blocking coalition of two players. If an allocation is in the core, then, trivially, it is stable. For the Stable Marriage Problem, the reciprocal statement is also true. Why is this true? Well, if there is an improving

coalition with more than two agents, there must be a subset of agents in this coalition with one or two players that would prefer to deviate from the original allocation. If an allocation is stable, there is no such subset. Hence, this allocation must be in the core.

The Deferred Acceptance Algorithm (DAA) is remarkable in many ways. First, it proves the existence of at least one stable matching.

Second, the algorithm is reasonably simple and intuitive. While current economists may find it a very desirable and convenient feature, this is somewhat surprising for a paper written by two mathematicians 50 years ago. Most papers in mathematics at that time, and even nowadays, propose very complex structures and algorithms. By contrast, Gale and Shapley's rigorous demonstration of the existence of a stable outcome — to which we will return later — contains, as they point out (p.15), '...no obscure or technical terms. Knowledge of calculus is not presupposed. In fact, one hardly needs to know how to count.'

Third, the DAA runs in finite time and is relatively easy to a program in a computer. Gale and Shapley proved that their procedure has at most $n^2 - 2n + 2$ stages, where n is the number of agents on each side of the market.

Fourth, usually there are *multiple* stable allocations for any given instance of the Stable Marriage Problem. Depending on how we set the algorithm to run, it can generate the allocation that is best, among all stable ones, for males or the best allocation for all females. In other words, if we are restricted to the collection of stable matching allocations only, then all males would never disagree in their ranking of any two proposals. The same is true for females. Moreover, in the set of all stable matching allocations, the 'male best' stable matching is exactly the 'female worst' matching and vice-versa. According to mathematicians, the set of all stable allocations form what they call a 'distributive lattice', a very useful but somewhat unknown mathematical concept. Analogous considerations hold for the Stable College Admissions Problem. Topkis (1978) has more on lattices and optimisation on such domains.

Fifth, even though the algorithm itself may involve arbitrary choices, the set of all stable matching allocations does not depend on chance. Although chance does play a role sometimes, nobody can complain too much of 'bad luck'. If an agent is not matched by the Deferred Acceptance Algorithm to some other agent ranking relatively high in her preference list, it is likely that she is simply not ranked highly by any one of her most favoured matches. This property may be an important feature in persuading the public to accept the implementation of a central mechanism.

The Deferred Acceptance Algorithm

Consider an instance of the Stable Marriage Problem. We will describe a version of the Deferred Acceptance Algorithm (DAA) in which men propose to women. The algorithm has a sequence of such proposals. All agents start 'free'. Men can alternate between being free or 'engaged' several times during the process, but women cannot become free after they are engaged. Women who are free always accept an engagement proposal of any acceptable man, but will turn this man down if she ever receives a proposal from a more desirable agent. Engaged women reject proposals from men that are less desirable than her current engaged partner. Men who are free keep making proposals, but engaged men do not.

The version of the DAA in which men are the proposers works as follows:

- All agents start free.
- The following procedure is repeated as long as any man is free; when no free men remain, the algorithm terminates.
- The men in a randomly chosen order make proposals.
- The current man makes a proposal to the most desirable woman in his list who has not already rejected him.
- If this woman is free, then she and her proposer become engaged.
- Otherwise, if she prefers the proposer to her current partner, she becomes disengaged from her existing partner — who thereby becomes free — and becomes engaged to the current proposer; otherwise, she rejects the proposer.
- This procedure is repeated for the next man.
- As already stated, the algorithm ends when no free men remain.

When the algorithm ends, the couples who are engaged at that stage become matched.

It turns out that the order in which men make their proposals does not affect the outcome of the algorithm. In the process, each man starts with highly desirable partners, but as the algorithm proceeds, he can only go down in his preference list. However, each woman can only move up in her preference list, as she receives proposals from a range of different men, but switches partners only if it is in her advantage to do so (she rejects all proposals from men who are less desirable than her current engagement). Because both the set of men and the set of women are finite, the algorithm always ends in a finite number of steps.

We urge curious readers to indulge themselves by reading the original proof of the existence of a stable matching in Gale and Shapley (1962: 11–12). The following argument captures its main insight. Suppose that man X_1 is matched to woman Y_1 , but prefers woman Y_2 to Y_1 . During the DAA, man X_1 proposes

to woman Y_2 before he proposes to woman Y_1 . The reason he is not matched to Y_2 is *either* because woman Y_2 has rejected X_1 's proposal (since she had a better partner) *or* because woman Y_2 accepted the proposal by man X_1 , but afterwards obtained a proposal from a more desirable man. In either case, Y_2 ended up with a man that she considers better than man X_1 .

Consider the following example of the men-proposing version of DAA: there are three men (X_1, X_2, X_3), and three women (Y_1, Y_2, Y_3). Preferences are as follows: the last option of each agent is to end up alone. Among members of the opposite sex, preferences are:

$X_1: Y_1, Y_2, Y_3$ (meaning that man X_1 prefers woman Y_1 to Y_2 , and prefers Y_2 to Y_3);

$X_2: Y_2, Y_3, Y_1$;

$X_3: Y_1, Y_2, Y_3$;

$Y_1: X_2, X_3, X_1$;

$Y_2: X_1, X_2, X_3$;

$Y_3: X_3, X_1, X_2$.

Suppose men make proposals in the order: X_1 first, then X_2 , and finally X_3 . In this case, man X_1 proposes to woman Y_1 , who accepts. Man X_2 proposes to woman Y_2 , who accepts. Man X_3 proposes to woman Y_1 , who accepts (because woman Y_1 prefers X_3 to X_1), making man X_1 free again. Man X_1 proposes to woman Y_2 who accepts (because woman Y_2 prefers X_1 to X_2), making man X_2 to become free again. Man X_2 proposes to woman Y_3 who accepts. The resulting matching is: (X_1, Y_2) , (X_2, Y_3) , and (X_3, Y_1) . The theory tells us that this matching is a stable matching in this particular example.

But there are multiple stable matching outcomes. The two versions — distinguished by whether men or women are the proposers — lead to different stable outcomes. For instance, if women make proposals in the order Y_1 first, then Y_2 , and finally Y_3 , the DAA proceeds as follows: woman Y_1 proposes to man X_2 , who accepts. Woman Y_2 proposes to man X_1 , who accepts. Woman Y_3 proposes to man X_3 , who accepts. The resulting matching is: (X_1, Y_2) , (X_2, Y_1) , and (X_3, Y_3) . This is also a stable matching.

Man X_1 and woman Y_2 are matched in both the man-proposing and woman-proposing versions. However, the man-proposing version of the DAA is strictly preferable to the woman-proposing version for men X_2 and X_3 . Indeed, man X_2 prefers to be matched with woman Y_3 (his second-best option) than with Y_1 , which is his third-best option. Man X_3 prefers Y_1 (his first option) to woman Y_3 , which is only his third-best option. On the other hand, women Y_1 and Y_3 strictly

prefer their partners under the woman-proposing version to those under the man-proposing version of the DAA. For instance, woman Y_1 prefers man X_2 rather than man X_3 .

Later developments by Roth and others

The National Resident Matching Program (NRMP) is a North American NGO that matches medical school students with residency positions that are available in any given year. The NRMP started to operate in 1952 and is sponsored by several medical associations. Before that, there was keen competition among hospitals for the best talent. Hospitals would make job offers independently of one another. Around 1944 individuals received offers up to two full years before their graduation. Agreement on the need for some kind of clearinghouse, or centralized market, was reached around 1950.²

Matches occurring two years before students complete their training tend to suggest that students lack incentive to study really hard during their last two years. As a consequence, they probably become less competent doctors. This is a clear source of inefficiency. Because many of the best positions are taken two years before they graduate, unmatched medical students also find fewer reasons to study hard in their last two years. Hence, matching a group of medical students two years before graduation may severely diminish the incentives of all students — both earlier matched and unmatched — to study hard.

Of course, two full years before graduation the information available regarding future doctors' skills is less precise than what we would obtain if matches were made at graduation time. This was another significant source of inefficiency.

Up until 1995, the version of the Deferred Acceptance Algorithm used by the NRMP always generated the matching assignment that was the best for hospitals, among the subset of all stable allocations, and worst for doctors. Since 1998 the revised NRMP algorithm finds the doctor-optimal stable matching. Interestingly, only a small proportion of the assigned doctor-hospital pairs would change, whether we run the previous or the current algorithm. This is a positive feature of the procedure. Later, the NRMP started to allow for the possibility that couples would submit joint preference lists. This makes the assignment problem significantly harder.

The unravelling of the market has been particularly acute in the process by which US federal court judges hire federal judicial law clerks. Avery *et al.* (2001)

2 Roth and Peranson (1999: footnote 5) claim that the '...initial NRMP algorithm, dating from 1951, [was] the first we know of'. For more details on the history of the DAA and NRMP, see Roth (1984), Gusfield and Irving (1989) or Roth and Sotomayor (1992).

note that, at the time when many prominent law schools in 1999 abandoned their attempts to regulate the time at which letters of recommendation could be sent, the clerk-hiring process got under way two years before the clerkship positions themselves were due to begin. The behaviour of judges in their attempts to jump the gun on their rivals in making early offers and seeking binding acceptances in order to attract the best candidates, led the *New York Times* to refer to the process as a judicial ‘free-for-all’ in which judges ‘behave[ed] like 6-year-olds’. One judge likened the process to a ‘calf scramble ... the low point of many western rodeos’ — not an image that most of us have of the judiciary!

Some significant characteristics of this piece of intellectual history

In the original formulation, the Stable Marriage Problem assumes that both sides of the market have n agents. It turns out that this assumption is not terribly important. The Deferred Acceptance Algorithm runs unchanged when the numbers of agents on each side of the market are distinct/different from each other.

In reality, a large part of the work done by Alvin Roth and his co-authors deals with institutional particularities. Yet the central ideas of Matching Theory are always present and the researchers continue to create and develop new abstract ideas, concepts, algorithms and mathematical techniques that will, in the future, become the basic knowledge upon which concrete, applicable real-world solutions will be developed.

The lesson is two-fold: first we learn from this history that one must be willing to adapt abstract frameworks to tackle real-life problems — what Roth refers to as ‘social engineering’. Second, the search for pure, abstract algorithms and formal mathematical results in idealised environments is crucially important. Abstract models have provided critically important changes in paradigm. They can shed light on difficult problems and indicate directions where possible solutions lie and where there is need for more knowledge and more pure or applied research. Too frequently governments have a myopic focus. However, the power of very abstract ideas is enormous.

Investment in abstract modelling may save public funding because the same central principles can typically be used in a wide variety of applications. In short, a single investment may generate many fruits, if we are prepared to adapt it to the idiosyncratic features of each institutional framework.

The award of this particular Nobel Prize provides a good opportunity for reflection. Where are our tax dollars best spent — in very concrete yet somewhat limited research of clear immediate practical use, or in more abstract research with no immediately apparent connection to reality, but deep scientific significance? For instance, would it be better to invest one dollar with the certainty of collecting two, or to invest one dollar with a one per cent probability of collecting a thousand?

We would argue that research in economic theory has, on average, more of a ‘public-good nature’ than a good deal of applied research. Very applied research is designed — by its very nature — to deal with a very particular case. Lessons may sometimes be transported to some other situations, but typically on a rather small scale. By contrast, the ideas in Gale and Shapley’s analysis of matching are non-rival — their use in one setting does not limit their use in others. Cornes and Sandler (1996) provide an extended discussion of this non-rivalry property in the context of public-good theory, and Romer (1990) discusses its crucial importance in the context of technological change and welfare. A matching algorithm can be used time and time again — to match medical interns with hospitals, students with schools, law clerks with law courts or legal firms. Further, once it is developed, its extension and adaptation for other uses is a relatively low-cost matter, as is evidenced in the developments from Gale and Shapley’s original insights up to the more complicated algorithms being developed today by Roth and many others.

The example of combinatorial analysis and auctions

A new important research trend inside auction theory is the study of combinatorial auctions; that is, procedures to auction bundles of heterogeneous goods, instead of a single product. The idea is that these different products may complement or substitute each other and these degrees of complementarities or substitutability may be idiosyncratic. For instance, suppose you want to sell a table and a set of chairs that fit the table. Is it better to hold two auctions, one for the table and another for the chairs, or to sell the entire bundle in a single auction?

Matching Theory and Auction Theory are the two faces of what current researchers call “Market Design”. Rassenti *et al.* (1982) use experimental analysis to study the design for airport time-slot auctions. They propose the use of a combinatorial auction mechanism. De Vries and Vohra (2003) provide a classic survey of combinatorial auctions. Their survey reveals that over the last 10 years there have been substantial theoretical advances on this front and

that many new applications have emerged in a wide variety of contexts — sales of many products and services such as furniture, airport time slots for airlines, delivery routes, and network routing are among the huge variety of situations in which combinatorial analysis is proving useful. More recently, Vohra (2011) has shown how combinatorial reasoning, in the form of linear programming, can be used to develop and understand modern mechanism design.

Other significant areas within economics that use the ideas of combinatorial analysis are Voting Theory and Matching Theory, as well as all problems related to the allocation of indivisible goods. However, for some traditional economists educated to work only with calculus, optimisation and differential methods with continuous variables, the language of combinatorial analysis represents a huge challenge. It is so because, unfortunately, discrete mathematics, even its elementary ideas and concepts, is no longer part of the typical university training in the field of economics.

The basic ideas of permutations and combinations are extremely helpful in the education of university students. Few concepts are as efficient to help students learn how to think analytically, in an organised way, as the basic counting and ordering principles of combinatorial analysis. Their presence in the syllabus helps to make students learn how to think and how to organise complex structures and ideas in their minds. Ironically, these virtues may partly explain their gradual disappearance from university economics syllabuses since the 1960s — under ever-increasing time and resource pressures, it may have been easier and quicker to teach alternative approaches that rely more heavily on a relatively uncritical application of ‘cookbook recipes’.

Current applications of Matching Theory

In the last two decades ever more applications for Gale and Shapley’s Matching Theory have been found. Although they rarely make front-page news in the manner of the latest macroeconomic or financial ‘crisis’, these applications are beginning to have a significant beneficial impact on the lives of millions of people — as resident doctors and patients with transplanted kidneys in the USA or parents of children benefiting from a better high school allocation in New York or Boston public schools system would tell you — and to change the world we live in.

Matching Theory is used today to allocate secondary students to public high schools in New York and Boston, two very large educational systems. Before entering as incoming freshmen, all students submit a preferred high school list,

and high schools set criteria that effectively rank all new students in the city system. This work was done by Alvin Roth together with Atila Abdulkadiroglu, Parag A. Pathak and Tayfun Sonmez (Abdulkadiroglu *et al.* 2005).

Roth (2002) brings quantitative evidence about observed benefits of stability in matching mechanisms used in practice, focusing on different regional markets for new physicians and surgeons in the United Kingdom. Roth's paper describes quantitative evidence found in previous papers and displays (Table 1: 1354) the average costs of markets. Among markets with a clearinghouse, one group indicates similar costs between the DAA and the Newcastle Priority Algorithm — an alternative algorithm that does not take the Gale and Shapley concept of stability into account. However, in two other cases, the costs associated with the DAA are significantly lower than the cost associated to the Newcastle Priority Algorithm.

An unexpected benefit of using Matching Theory in practice was to decrease corruption in some cases. As the centralised mechanism was implemented in school districts, local authorities lost their power to accept or reject students to high schools. To the surprise of some, extra positions became available in many schools.

Matching Theory has had a dramatic effect on the allocation of high school students. Every year, in New York City School District alone, there are around 100 000 new high school students entering the system. All of their families provide a list of preferences and the schools provide their criteria to rank the entrant students. Obtaining a better allocation of school leaves thousands of families better off and their kids will probably learn more and become more productive workers. As economists know well, enhancing human capital leads to a permanent increment in productivity.

Organ transplantation is another major beneficiary of Matching Theory. Typically, the waiting list for deceased donors is very long. Many patients die before a compatible organ is available. A partial solution to this problem is to increase the incentives for live donors. Often a patient can find a living donor, usually a spouse or sibling. However, rarely is the donor organ compatible with the patient's biology. In this context, Matching Theory and Market Design have played a crucial role in the design of the institutional rules that generate the best incentives and the largest possible number of successful organ donations. Indeed, it seems to have been precisely this application that drew the Nobel committee's attention to the work of Alvin Roth.

A feature of the market for organ donation is that it is illegal to buy or sell human organs in most countries, including all developed countries. Even if both parties agree on the terms, the law does not allow for such transactions.

(For example, in the US, the National Organ Transplant Act forbids contracts for organ transplantations.) Moreover, organs are indivisible. This is exactly the framework of Matching Theory, and indeed much has been done recently.

Alvin Roth, together with Itai Ashlagi, David Gamarnik, Utku Unver and Tayfun Sonmez, and with the support from the US National Science Foundation, developed the ideas, procedures and software needed to match living kidney donors with recipients. Their research and ideas from idealised abstract models are achieving what had earlier seemed impossible: that a single extra donation can help to spark a long chain of donations, saving many lives directly and decreasing the waiting times of patients in the waiting lists, which in turn, greatly increases the probability of success in organ transplantations. Also, reduced waiting lists improve dramatically the welfare of thousands of patients. It is no longer the case that one donor helps a single sick person.

The literature on two-sided matching uses elements of economic theory as well as empirical and experimental analysis, as explained by Roth (2002).

Conclusions

The history of Matching Theory teaches us several valuable lessons.

First, it could be dangerous to dismiss basic research as necessarily less important because of its lack of obvious immediate impact on pressing problems of the day. This is a crucial observation these days as policy circles discuss the introduction of policies that try to measure research impact.

Second, some of the most applicable and useful ideas often come along first in the most abstract forms. They require time to mature and to prove their true value; in this case a few decades. This does not fit with the politics of the democratic electoral cycle. Hence, support for deep research must be a national priority above and beyond the politics of the day. Funding for basic research should be allocated in a way that it becomes untouchable to political actors, both within and outside universities. Matching Theory is a great example — perhaps the best — of a line of research in economics which, initially, may have seemed like a piece of academic self-indulgence, yet ultimately having — as its authors surmised — real significance to real allocation issues. It is not unfair to ask how the Gale/Shapley piece would have been shoehorned into the straitjacket of national priorities.

Third, the power of abstract mathematical ideas should never be underestimated. Looking to social sciences or policy studies through the lens of what seems to be of immediate national significance will probably lead to myopic incentive

schemes that are detrimental to mathematics and to the creation of powerful abstract new ideas and concepts. At the end of the day, such short-sightedness may generate an academic sector producing less significant outcomes, most of which use, at best, simplistic and conventional ideas. This would be a poor path for our society to follow.

Fourth, Matching Theory is a great example of the breadth of current economic theory. It explores the ideas behind allocation of resources, and their implications for human wellbeing. Of course, topics related to the allocation of goods and agents are always interesting to economists, whatever the particular circumstances.

As John Maynard Keynes famously wrote in *Essays in Persuasion*: ‘If economists could manage to get themselves thought of as humble, competent people on a level with dentists that would be splendid.’

In this piece, we have drawn attention to economists who, through their association not with dentists but with school administrators, renal transplant surgeons and others, have greatly enhanced the wellbeing of many individuals. We suspect that many contemporary critics of economics, and indeed some of our colleagues within the discipline, are simply unaware of its involvement in the design of resource allocation mechanisms of the kind that we have described, and of their potential benefits.

Finally, Matching Theory is an area with proven positive results in terms of public policy. However, in Australia, high school graduates are not allocated to universities with the use of a centralised mechanism that uses the DAA. The implementation of such a scheme would improve access of students to universities and would also lead to better matches between the two sides of the market. There are plenty of other areas that could benefit if Australian policymakers decided to use Matching Theory to formulate and implement better public policies.

We leave the last word to the writer of the piece in *The Economist*³ which announced the award of the 2012 Nobel Prize: ‘The recognition of Mr Shapley’s and Mr Roth’s work is also a reminder: that for all the bad press economics has received since the [financial] crisis, the discipline still brims over with insights that can solve real-life problems.’

3 20 October 2012.

References

- Abdulkadiroglu, A., Pathak, P., Roth, A. E. and Sonmez, T. 2005, 'The Boston Public School Match', *The American Economic Review Papers and Proceedings* 95: 368–71.
- Avery, C., Jolls, C., Posner, R. A. and Roth, A. E. 2001, 'The market for federal Judicial law clerks', *University of Chicago Law Review* 68: 793–902.
- Cornes, R. C. and Sandler, T. 1996, *The Theory of Externalities, Public Goods and Club Goods* (2nd edition), Cambridge University Press.
- de Vries, S. and Vohra, R. 2003, 'Combinatorial auctions: A survey', *INFORMS Journal on Computing* 15(3): 284–309.
- Gale, D. and Shapley, L. S. 1962, 'College admissions and the stability of marriage', *American Mathematical Monthly* 69: 9–15.
- Gusfield, D. and Irving, R. W. 1989, *The Stable Marriage Problem: Structure and Algorithms*, Foundations of Computing Series, The MIT Press.
- Hajduková, J. 2004, *On Coalition formation games*, IM Preprints series A, No. 5/2004.
- Knuth, D. E. 1976, *Mariages Stables*. Le Presses de l'Université de Montréal, Montreal.
- O'Brien, M. and Kellan, A. 2012, *New Software Matches More Kidney Donations, Faster*, National Science Foundation. Available at: http://www.nsf.gov/news/special_reports/science_nation/kidneyexchange.jsp
- National Resident Matching Program. Information available at: <http://www.nrmp.org> and http://en.wikipedia.org/wiki/National_Resident_Matching_Program
- Rassenti, S. J., Smith, V. L. and Bulfin, R. L. 1982, 'A combinatorial auction mechanism for airport time slot allocation', *The Bell Journal of Economics* 13: 402–17.
- Romer, P. M. 1990, 'Endogenous Technical Change', *Journal of Political Economy* 98: S71–102.
- Roth, A. E. 1984, 'The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory', *Journal of Political Economy* 92: 991–1016.

Roth, A. E. 2002, 'The economist as engineer: Game theory, experimentation, and computation as tools for design economics', *Econometrica* 70: 1341–78.

Roth, A. E. and Sotomayor, M. A. O. 1992, *Two-sided Matching: A Study in Game-theoretic Modeling and Analysis*, Cambridge University Press.

Roth, A. E. and Peranson, E. 1999, 'The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design', *The American Economic Review*: 748–80.

Topkis, D. M. 1978, 'Minimizing a submodular function on a lattice', *Operations Research* 26: 305–21.

Vohra, R. V. 2011, *Mechanism Design: A Linear Programming Approach*, Cambridge University Press.