

Fairness in Deceased Organ Matching

Nicholas Mattei
IBM Research AI
n.mattei@ibm.com

Abdallah Saffidine
UNSW Sydney
abdallah.saffidine@gmail.com

Toby Walsh
UNSW Sydney, Data61, TU Berlin
tw@cse.unsw.edu.au

Abstract

As algorithms are given responsibility to make decisions that impact our lives, there is increasing awareness of the need to ensure the fairness of these decisions. One of the first challenges then is to decide what fairness means in a particular context. We consider here fairness in deciding how to match organs donated by deceased donors to patients. Due to the increasing age of patients on the waiting list, and of organs being donated, the current “first come, first served” mechanism used in Australia is under review to take account of age of patients and of organs. We consider how to revise the mechanism to take account of age fairly. We identify a number of different types of fairness, such as to patients, to regions and to blood types and consider how they can be achieved.

Introduction

Kidney disease costs the Australian economy billions of dollars per year. Over ten thousand people in Australia are on dialysis, each costing hundreds of thousands of dollars in medical and welfare costs. Australia is especially challenged in this area as kidney disease is a major problem within the indigenous population. The incidence of end stage kidney disease in the indigenous population in remote areas of Australia is 18 to 20 times higher than that of comparable non-indigenous peoples.¹

A significant trend in Australia (as in other developed countries) is that age is now starting to play a major role in kidney disease. It is impacting both the demand and supply side of the kidney transplant market. On the demand side, the age of patients in Australia waiting to receive a kidney has increased significantly in recent years. In 2010, for example, just 11% of the waiting list were 65 years or older. In 2015, this had increased to 15%. Over the next 30 years, the proportion of the population of Australia aged over 65 years is predicted to double to around 25 per cent. This ageing demographic will likely further increase the age of people on the waiting list for a kidney transplant.

On the supply side of the market, the age of donated kidneys has also increased significantly. In 1989, the mean age

of donated kidneys in Australia was just 32 years old. In 2014, this had increased dramatically to 46 years old. Surgeons are now able and willing to transplant older kidneys into older patients. In 1989, the oldest transplanted kidney came from a donor aged 69 years. In 2014, this has increased to an 80 year old donated organ. A number of factors including increasing life expectancy, medical advances, and improved road safety have been driving these changes on both sides of the market.

Matching in Australia currently uses a mechanism based on first-come, first-served. Simply put, the longer one waits, the closer one moves to the top of the waiting list. An arriving organ is offered to the highest person on the list who is a compatible match. Given the limited supply of organs, whilst this mechanism is “fair” from a procedural perspective, it is now no longer viewed by many in the medical profession as “efficient” in terms of best use of the limited supply of organs. A 70-year-old patient may receive a kidney from a 30-year-old donor, and live 10 or even 15 years. But the organ might have lasted decades longer in a younger body according to UNOS statistics. Critics of the current system have argued that the organ’s full potential for giving life is wasted on an older person. In contrast, an organ from a 60-year-old donor transplanted into a 30-year-old patient may fail before the patient reaches old age, thereby creating the need for an additional organ. Worse still, the patient may be sensitized by the immune suppression drugs and so require an even closer match. Transplanting young organs into old patients, and old organs into young patients might therefore be considered less than optimal (Kaufman 2013).

There is thus a desire for a new mechanism that matches the age of patients to that of the organs and does so in a “fair” manner. In this paper, we consider what it means for a deceased organ matching mechanism to be fair when it takes account of features like age. When reducing the pool of candidate matches by age, other concerns come into focus like geography and blood type. For instance, matching nationally rather than at the state or even hospital level improves the quality of matches possible. This is especially important if we are decreasing the pool of possible candidates by ruling out certain matches based on age. However, there are concerns that organs will flow out of the less populated states and territories to the larger states where demand is highest. We suggest that fairness needs to be considered on multi-

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹These (and subsequent) statistics about kidney disease and transplant in Australia are taken from (Australia & New Zealand Dialysis & Transplant Registry 2015).

ple levels: patient, region, and blood type to name just three. We argue that regional and blood type fairness ultimately requires improving donation rates. On the other hand, we argue that a simple mechanism that matches the age of organs to age of patients gives an allocation that is stable under some rather natural preferences for the two sides of the market. There is thus no “incentive” to deviate from the outcome it returns. This simple mechanism also offers patients an uniform time on the waiting list so is procedurally “fair” like the current mechanism which ignores ages.

Kidney Matching in Australia

The Organ and Tissue Authority of Australia is the formal body allocating donated kidneys to patients. Their charter requires the allocation of organs to be fair and efficient, though it is not formally defined what fairness and efficiency mean. In designing a new matching mechanism taking account of the age of patients and of organs, the Organ and Tissue Authority has decided to follow the US lead who already have revised their mechanism to take account of these ages using two measures: the Kidney Donor Patient Index (KDPI) and a patient’s Expected Post-Transplant Survival (EPTS) score.

The quality of a donated organ is measured by the Kidney Donor Patient Index (KDPI). This is an integer from 0 to 100 that is calculated from the age of the donor, their diabetic status, cause of death and other factors. A donated kidney with a KDPI of X has an expected risk of graft failure greater than $X\%$ of all donated kidneys. The quality of a patient on the waiting list is measured by their Expected Post-Transplant Survival (EPTS) score. This is also an integer from 0 to 100 that is calculated from the age of the recipient, their diabetic status, the number of prior organ transplants, their time on dialysis and a number of other factors. A patient with an EPTS of Y receiving a high quality donated kidney has an expected survival time that is greater than $100-Y\%$ of all patients. For a more detailed discussion of these measures see the data and simulators from www.srtr.org as well as the work of Bertsimas, Farias, and Trichakis (2013).

Whilst the Organ and Tissue Authority have decided to use KDPI and EPTS in their new mechanism, they have yet to decide on the precise details. Currently they collect KDPI and EPTS but do not use it when proposing a match. One candidate under consideration by the Organ and Tissue Authority is the BOX mechanism. This favors those matches that fit in the box bounded by $KDPI \leq 50$ and $EPTS \leq 25$. In other words, the top quartile of patients is offered the top half of organs. More precisely, the BOX mechanism ranks patients according to a lexicographical scoring function. The most important terms in the scoring function ensure a match between compatible types. The least important terms tie-break according to features like time on the waiting list. The middle term orders matches in the following order: (1) $KDPI \leq 50$ and $EPTS \leq 25$; (2) $EPTS < 25 < KDPI$; (3) $EPTS < 25 < KDPI < EPTS < 50$; (4) $EPTS < 50 < KDPI < EPTS < 75$; (5) $EPTS < 75 < KDPI$. The ordering can be viewed graphically (Figure 1). It favours matches in a rectangular box to the bottom left, then above the upper left diagonal, and then towards the lower right diagonal. The BOX mechanism has

some similarity to the current US mechanism that preferentially offers organs with $KDPI \leq 20$ to patients with an $EPTS \leq 20$. We have instead proposed the simple MIN mechanism (Mattei, Saffidine, and Walsh 2017). This picks a compatible patient for an arriving organ that minimizes $|KDPI-EPTS|$, tie-breaking by time on waiting list.

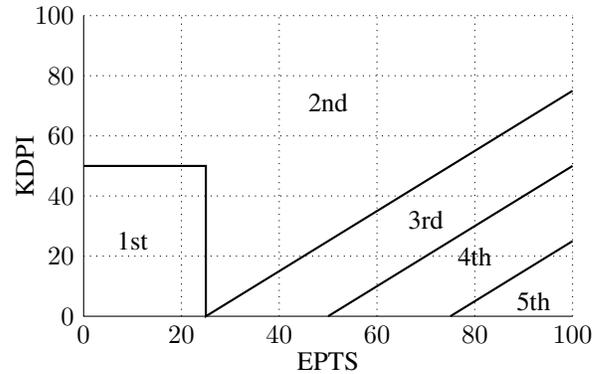


Figure 1: The BOX mechanism: the ordering is induced by a lexicographical scoring function with five key regions.

Fairness to Age

We first consider the fairness of the two proposed matching mechanisms. We argue that, unlike the BOX mechanism, the MIN mechanism is procedurally fair to patients of different ages. In Figure 2 and 3, we plot the waiting time for patients using the two mechanisms in a simulation based on historical data for Australia for 2010 to 2014. When an organ arrives, it is matched nationally with a compatible patient using either the BOX or MIN mechanisms. We see that the BOX mechanism is not procedurally fair. Those patients on the waiting list with an EPTS of 25 or less spend much less time waiting than those with an EPTS of greater than 25. This is to be expected as the BOX mechanism preferentially favours patients with an EPTS of 25 or less. By comparison, the MIN mechanism is much more procedurally fair. The time on the waiting list is almost independent of EPTS. There is a small boundary effect against patients with an EPTS close to 0 (or close to 100) since they can only receive an older (or younger) organ. However, besides this small boundary effect, waiting time with the MIN mechanism is almost constant. This is to be expected given that EPTS and KDPI are population percentiles.

We can view the allocation of organs from deceased donors as a two-sided matching problem. On one side of the market, we have the patients on the waiting list. To maximize their post transplant survival time, each patient simply wants to receive the best quality organ, i.e., the organ with the lowest KDPI possible. Hence, the patients have identical preferences over the organs. On the other side of the market, we have the organs with preferences over the patients. Of course, organs don’t *actually* have preferences. They are just organs. And the donors are deceased so also arguably don’t have preferences for their organs at this point. This prefer-

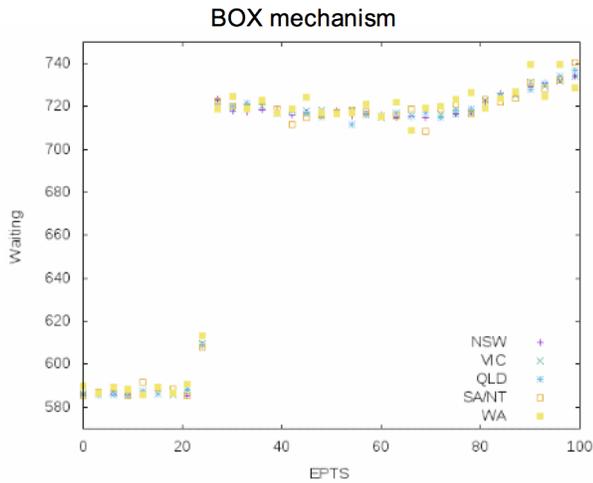


Figure 2: Distribution of waiting times according to EPTS of patient receiving organ using the BOX mechanism.

ence of the organs could be seen as a societal or medical preference, i.e., to ensure the maximum good from limited supply. We suppose then that the preference of the organ side of the market is to minimize $|KDPI - EPTS|$. Roughly speaking, the preference of this side of the market is to match age of organ to age of patient.

A fundamental notion in a two-sided matching market is stability (Roth and Sotomayor 1992). We do not want an allocation where a patient and organ not currently matched to each other would both prefer to be matched to each other rather than their current matching. That is, we want an allocation where there is no incentive to deviate. A patient (or organ) might prefer some other matching but that organ (or patient) does prefer this compared to their current matching. In the case of deceased organ matching, both sides of the market have identical preferences, so rather than a possible lattice of solutions (Gale and Shapley 1962), it turns out that under modest assumptions there is a unique stable allocation, ignoring the permutation of patients with identical EPTS and of organs with identical KDPI. This stable allocation is the one in which $KDPI = EPTS$, and the allocation that the MIN mechanism is likely to return in a large online market such as in Australia.

We suppose that the market is large so that there are enough compatible patients to ensure that $KDPI$ can be made equal to $EPTS$ in each match. There is then a unique stable allocation. In this allocation, all organs with a $KDPI$ of 0 are matched with patients with an $EPTS$ of 0. The patients receive organs with lowest possible $KDPI$ so cannot be happier. And the organs get matched to patients so that $KDPI = EPTS$ so again could not be happier. All the organs with a $KDPI$ of 1 are matched with patients with an $EPTS$ of 1. The organs get matched to patients so that $KDPI = EPTS$ so could not be happier. And the patients could only be happier if they were matched with an organ with a $KDPI$ of 0.

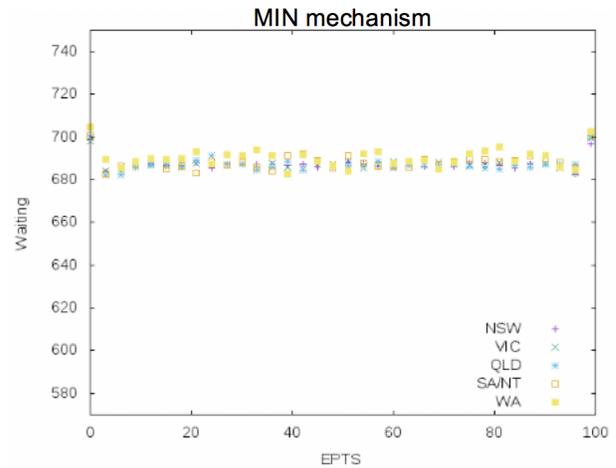


Figure 3: Distribution of waiting times according to EPTS of patient receiving organ using the MIN mechanism.

But none of these organ would prefer such a match. Similarly, the organs with a $KDPI$ of 2 are matched with patients with an $EPTS$ of 2, and so on. The unique stable allocation has $KDPI = EPTS$ in every match.

Fairness to Geography

Moving to a mechanism that takes account of $KDPI$ and $EPTS$ will decrease the pool of patients from which we match. We may desire therefore to consider matching at the national, rather than the state or hospital level to compensate. Especially for rarer blood types, matching that takes account of $KDPI$ and $EPTS$ will be better if we can use the larger national pool. Trivially the quality of the matching strictly increases when we match at the national level. We can perform all state level matches as well as those now possible nationally. Of course, there is a time penalty to matching nationally. However, it only takes about five hours to fly an organ across the country, and transplant surgeons assure us that transplant success depends on the quality of matching and not the amount of time the organ spends on ice. There is thus little to be lost, and much to be gained, if the states and territories can be persuaded to match nationally.

In Figure 4, we plot the distribution of donors and patients waiting for transplant according to state and territory. We compare this with the distribution of the population within Australia. Donation largely tracks population, as might be expected. The major exception is New South Wales (NSW) and the Australian Capital Territory (ACT)². Donation rates in NSW/ACT are much lower than in the rest of Australia. To compound this issue, the waiting list in NSW/ACT is proportionally much longer than in the other states and territories. In part, this may reflect that doctors are more likely to list patients on the waiting list in New South Wales than

²There are no facilities for kidney transplant in ACT so all ACT patients are dealt with in NSW.

in a state like Queensland (QLD) where the waiting list is smaller proportional to the population. The data may also reflect that patients gravitate towards the more sophisticated medical facilities available in a populous state like New South Wales.

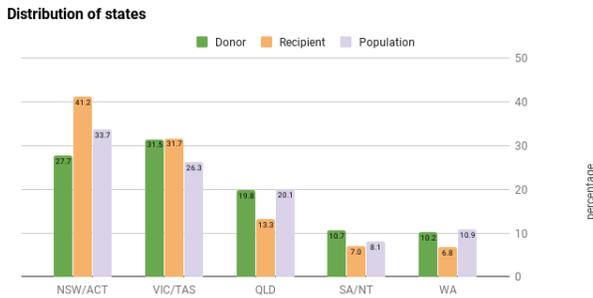


Figure 4: Distribution of donated organs, patients waiting transplant, and the wider population of Australia from 2010 to 2014.

There is concern, especially amongst the smaller states like South Australia (SA), Western Australia (WA) and the Northern Territory (NT), that donated organs will flow into New South Wales due to its comparatively much longer waiting list. We will compute the size of the potential flow. In 2014, there were 614 kidneys transplanted nationally. If we had matched nationally, rather than at the state level, we could expect on aggregate that 21 out of the 63 organs donated in Western Australia would flow out of the state (exactly one third), 23 out of the 66 donated organs would flow out of South Australia and the North Territory (slightly over one third), 40 out of the 122 organs donated in Queensland would flow out of state (just under one third). Almost all of the inflow of organs would be to New South Wales. Only one of the organs flowing out of SA/WA/NT/QLD would be expected on aggregate to end up in Victoria (VIC). The other 83 organs flowing between states and territories would end up on aggregate being given to patients in New South Wales. This inflow of 83 organs into New South Wales represents 13.5% of the total number of deceased organs donated nationwide.

Matching nationally ensures that everyone gets the same chance of a match irrespective of geography. This is very far from the case currently. In 2014, for example, the waiting list in New South Wales contained 474 patients at the start of the year, and increased to 500 by the end. Only 152 of the patients on the waiting list received a deceased organ. By comparison, in South Australia, the waiting list began and ended the year with 64 patients on it. During the course of the year, patients were added to and left the waiting list³ but a total of 67 patients received a deceased organ. Patients waitlisted in South Australia thus have a much greater chance of re-

³During 2014, one patient died waiting in South Australia, two received an organ from a living donor, 9 were taken off the list for medical and other reasons, and 79 new patients were added.

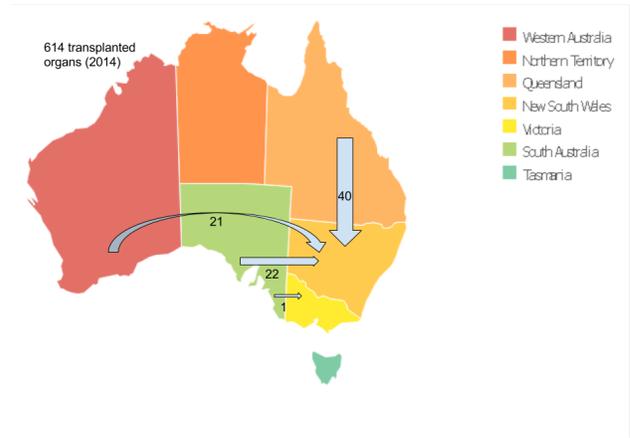


Figure 5: Aggregate flow of organs between states and territories in 2014 if matching had been nationally.

ceiving a transplant than in New South Wales. On average, patients waitlisted in South Australia wait about one year for transplant whilst those in New South Wales wait around three.

Matching nationally would result in greater fairness as waiting times across states would become more equal. On the other hand, matching nationally would result in a flux of organs out of the smaller states into New South Wales. The only way to prevent this is to harmonize donation rates within states. Our first recommendation to the Organ and Tissue Authority was therefore to mount a campaign to increase donation rates within New South Wales. We were told that such a campaign had already started.

Fairness to Blood Type

Reducing the pool of candidate matches by considering the age of organs and patients creates fresh challenges especially for rarer blood types. How do we treat different blood types fairly? We could, for example, permit organs coming from donors of blood type O to be transplanted into patients of blood type B since blood type O are universal donors. But this would disadvantage patients of blood type O waiting for transplant. Patients of blood type B would be matched out of the larger pool of organs of blood type O and B, whilst patients of blood type O would be matched out of the smaller pool of organs of blood type O only. This phenomenon has been observed in the European transplant market (Glander et al. 2010). On the other hand, there might be a greater demand for organs of blood type B relative to supply compared to blood type O. Transplanting a small number of organs from donors of blood type O into patients of blood type B could help to correct any imbalance.

Another concern when moving to a new mechanism based on age of organs and age of patients is that the quality of matching for rarer blood types will decrease as the pool size decreases since some matches will be ruled out based on age considerations. Historical data demonstrates where

challenges might arise in the Australian market. In Figure 6, we plot the distribution of donors and patients waiting for transplant according to blood type. We compare this with the distribution of blood types within the population. The exact blood type distribution in the population as a whole is not tracked, but we do have the distribution in the subset of the population donating blood.

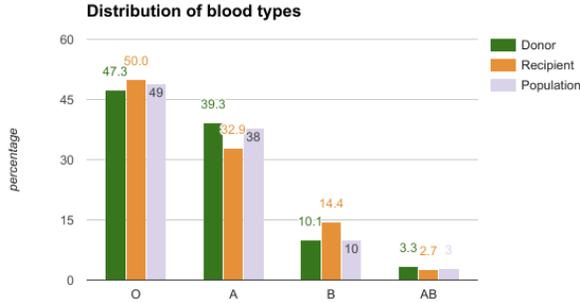


Figure 6: Distribution of blood types in donated organs, in patients waiting transplant, and in the wider population for Australia from 2010 to 2014.

Donation tracks population quite well. This should perhaps not be too surprising. Those ethnic and other groups willing to donate their blood are perhaps also likely to be those willing to donate their organs. However, demand is somewhat different. In particular, there is a greater percentage of patients of blood type B waiting transplant than donation. 14.4% of the waiting list have blood type B yet only 10.1% of organs donated are blood type B. On the other hand, patients of blood type A are better off. 39.3% of donated organs are of blood type A, yet only 32.9% of the waiting list have this blood type. There is some medical evidence that people of blood type A are less prone to kidney disease.

To see how we might improve fairness across blood types, we set up a simple linear model. Let p_t be the fraction of the patients with blood type t . For instance, p_A and p_{AB} are the fraction with blood type A and AB respectively. Let o_t be the fraction of the organs with blood type t . For instance, o_A and o_{AB} are the fraction with blood type A and AB respectively. We suppose matching takes into account other factors like HLA type but consider here just the size of the pool from which possible matches are drawn.

We suppose that within our matching procedure a fraction x_{t_1, t_2} of the donated organs are of type t_1 and are considered for transplant to patients of blood type t_2 subject to HLA match and other factors like age. Thus $0 \leq x_{t_1, t_2} \leq o_{t_1}$. We insist on blood type compatibility for transplant. Hence, $x_{A, O} = x_{A, B} = x_{B, O} = x_{B, A} = x_{AB, O} = x_{AB, A} = x_{AB, B} = 0$. That is, we cannot transplant organs of type A to patients of type O, organs of type A to patients of type B, etc.

We have a conservation law for organs of each blood type.

This requires:

$$\begin{aligned} x_{O, O} + x_{O, A} + x_{O, B} + x_{O, AB} &= o_O \\ x_{A, A} + x_{A, AB} &= o_A \\ x_{B, B} + x_{B, AB} &= o_B \\ x_{AB, AB} &= o_{AB}. \end{aligned}$$

We suppose that the mechanism is fairest when the fraction of organs available for a given blood type is as close as possible to the fraction of patients of this blood type. We introduce therefore some variables to measure this. Let z_t be the ratio of the fraction of organs available for blood type t and the fraction of patients of blood type t .

$$\begin{aligned} z_O &= \frac{x_{O, O}}{p_O} \\ z_A &= \frac{x_{O, A} + x_{A, A}}{p_A} \\ z_B &= \frac{x_{O, B} + x_{B, B}}{p_B} \\ z_{AB} &= \frac{x_{O, AB} + x_{A, AB} + x_{B, AB} + x_{AB, AB}}{p_{AB}}. \end{aligned}$$

To maximize fairness, we consider an egalitarian objective in which we maximize $z = \min(z_O, z_A, z_B, z_{AB})$. This can be solved in polynomial time using linear programming. Let's consider how this model fares on the historical Australian data. We have the following input data.

	O	A	B	AB
o_t	0.473	0.393	0.101	0.033
p_t	0.500	0.329	0.144	0.027

Maximizing fairness gives the following organ fractions.

x_{t_1, t_2}	O	A	B	AB
O	0.446	0.000	0.027	0.000
A		0.393	0.000	0.000
B			0.101	0.000
AB				0.033

And the corresponding z fractions.

	O	A	B	AB
z_t	0.89	1.19	0.89	1.22

Because of the drowning effect of the minimum function, there are multiple solutions with the same maximal minimum z value. To minimize transfer between blood types, we chose the solution with the maximal number of decision variables x_{t_1, t_2} for $t_1 \neq t_2$ set to zero. In this case, the only organs matched across blood type are organ of blood type O which are matched with patients of blood type B.

There are a number of ways we could translate this into practice. The simplest would simply be to toss a coin when a new organ arrives to decide which groups of patients is used in its matching. In this case, when an organ arrives of blood type O, with probability $\frac{x_{O, O}}{o_O}$, we consider patients of blood type O (that is, with probability 0.943), and otherwise (with probability $1 - \frac{x_{O, O}}{o_O} = \frac{x_{O, B}}{o_O} = 0.057$) we consider patients of blood type B. In short, with a 5.7% chance, we transfer an

organ of type O to a patient of type B, otherwise we match organs to patients of the same blood type.

Note that we have been unable to achieve complete fairness as the z values are not identical. We can equalize the treatment of patients of blood type O and blood type B (that is, we can equalize z_O and z_B). But these blood types are at a disadvantage compared to blood types A and AB (since z_A and z_{AB} are larger). We cannot use the relative excess of organs of blood type A and AB to help the relative excess of patients of blood type O and B. We simply need more organs of blood type O and B to give to patients of blood type O and B. Blood type A is also at a slight disadvantage compared to AB (since z_A is smaller than z_{AB}). We also cannot fix this problem by transferring organs between blood types. This illustrates a fundamental impossibility to be fair to the different blood types. As organs can only be transferred across blood type in one direction, there will be online organ matching problems, like the one in Australia, where we cannot treat patients of different blood type equivalently. Based on this analysis, we have advised the Organ and Tissue Authority to consider a publicity campaign to increase donation of organs from members of the public with blood types O and B.

Related Work

As mentioned earlier, the US has already adopted a matching mechanism that takes account of KDPI and EPTS. Other countries have also adopted mechanisms that take account of age. For example, the National Kidney Allocation Scheme (Johnson et al. 2010) introduced in the UK in 2006 has five hierarchical tiers, from highly sensitized pediatric patients for whom the kidney is a zero-mismatch down to not highly sensitized adult patients for whom the kidney is not a zero-mismatch. Within each tier, the recipient is chosen by a points system based on waiting time, HLA match, age difference between donor and patient, and blood group points which (as we have proposed in considering fairness across blood type in Australia) allocate a proportion of O type kidneys to B type recipients.

In Mattei, Saffidine, and Walsh (2017), we looked at deceased organ matching from an axiomatic perspective. We proved, for example, that no mechanism for deceased organ matching satisfies participation: patients may improve the quality of their match by strategically delaying their arrival. In addition, Mattei, Saffidine, and Walsh (2017) studied experimentally the efficiency of MIN and the more complex BOX mechanism under consideration by the Organ and Tissue Authority of Australia. In proposing a new mechanism, we cannot overlook the challenging issues of managing the transition (Ahlert and Kliemt 2013). It may be difficult to persuade patients to buy into a new mechanism if they will wait longer compared to the old mechanism. In this case, the current “first come, first served” mechanism gives patients a uniform waiting time on average, whilst the proposed new MIN mechanism also offers a uniform waiting time on average. Patients can expect to wait a similar wait time on average as now. However, this is only on average and there will be individual winners and losers.

Kidney exchange problems have been studied by economists (Roth, Sönmez, and Ünver 2005; 2004), medics (Montgomery et al. 2006), and computer scientists (Dickerson, Procaccia, and Sandholm 2012). Axiomatic and fairness concerns have been considered in kidney exchanges (Dickerson, Procaccia, and Sandholm 2014), as well as the online feature that such exchanges are repeated over time (Awasthi and Sandholm 2009; Dickerson and Sandholm 2015). There has also been research on strategy-proof mechanisms for kidney exchange at the patient and hospital level (Ashlagi et al. 2015), and on exchanges that combine kidneys with other organs such as lungs (Luo and Tang 2015) and livers (Dickerson and Sandholm 2016). Within Australia, only a few dozen of the more than one thousand kidneys transplanted each year come from paired exchange. Deceased donors provide the majority of transplanted organs, with paired exchange representing less than 3% of all transplants. Deceased donors are likely to provide the majority of transplanted organs until we have xeno-transplants or can grow new organs. The fairness of matching deceased organs to patients is thus of considerable importance.

Conclusions

We have considered fairness in organ matching from deceased donors. We suggested that fairness needs to be considered on multiple levels: patient, region, and blood type to name just three. We argued that regional and blood type fairness ultimately requires improving donation rates. On the other hand, we argued that a simple mechanism that matches the age of organs to age of patients gives an allocation that is stable under a simple preference model. Experiments demonstrate that this mechanism is also procedurally “fair” as waiting times for patients are largely independent of their age. We compared this with the other candidate mechanism proposed for allocating organs in Australia where wait times depend on the age of patients. With this mechanism, the youngest quartile of patients in Australia would have wait times around 20% less than the other quartiles.

Acknowledgements

Data61 is supported by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. Abdallah Saffidine is the recipient of an ARC DECRA Fellowship (DE 150101351). Toby Walsh is supported by the European Research Council and by AOARD Grant FA2386-12-1-4056.

References

- Ahlert, M., and Kliemt, H. 2013. Problems of priority change in kidney allocation and beyond. *The European Journal of Health Economics* 14(3):383–390.
- Ashlagi, I.; Fischer, F.; Kash, I. A.; and Procaccia, A. D. 2015. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91:284–296.

- Australia & New Zealand Dialysis & Transplant Registry. 2015. 38th annual ANZDATA report. Available from <http://www.anzdata.org.au/>.
- Awasthi, P., and Sandholm, T. 2009. Online stochastic optimization in the large: Application to kidney exchange. In *Proc. 21st IJCAI*.
- Bertsimas, D.; Farias, V. F.; and Trichakis, N. 2013. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.
- Dickerson, J. P., and Sandholm, T. 2015. Futurematch: Combining human value judgments and machine learning to match in dynamic environments. In *Proc. 29th AAAI*, 622–628.
- Dickerson, J. P., and Sandholm, T. 2016. Multi-organ exchange: The whole is greater than the sum of its parts. *JAIR*.
- Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2012. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proc. 11th AAMAS*, 711–718.
- Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2014. Price of fairness in kidney exchange. In *Proc. 13th AAMAS*, 1013–1020.
- Gale, D., and Shapley, L. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69:9–15.
- Glander, P.; Budde, K.; Schmidt, D.; Fuller, T.; Giessing, M.; Neumayer, H.-H.; and Liefeldt, L. 2010. The blood group O problem in kidney transplantation: time to change? *Nephrology Dialysis Transplantation* 25(6):1998–2004.
- Johnson, R.; Fuggle, S.; Mumford, L.; Bradley, J.; Forsythe, J.; and Rudge, C. 2010. A new UK 2006 national kidney allocation scheme for deceased heart-beating donor kidneys. *Transplantation* 89(4):387–394.
- Kaufman, S. 2013. Fairness and the tyranny of potential in kidney transplantation. *Current Anthropology* 54(S7):S56–S66.
- Luo, S., and Tang, P. 2015. Mechanism design and implementation for lung exchange. In *Proc. 24th IJCAI*, 209–215.
- Mattei, N.; Saffidine, A.; and Walsh, T. 2017. Mechanisms for online organ matching. In *Proc. 26th IJCAI*.
- Montgomery, R. A.; Gentry, S. E.; Marks, W. H.; Warren, D. S.; Hiller, J.; Houp, J.; Zachary, A. A.; Melancon, J. K.; Maley, W. R.; Rabb, H.; Simpkins, C.; and Segev, D. L. 2006. Domino paired kidney donation: a strategy to make best use of live non-directed donation. *The Lancet* 368(9533):419.
- 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.; Sönmez, T.; and Ünver, M. U. 2004. Kidney exchange. *The Quarterly Journal of Economics* 119(2):457–488.
- Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2005. Pairwise kidney exchange. *Journal of Economic theory* 125(2):151–188.