

The Refugee Match:

A system that respects refugees' preferences and the priorities of states

Will Jones¹ and Alexander Teytelboym²

Abstract³

There is an urgent need to overcome the political deadlock preventing states substantively participating in burden-sharing in the international refugee regime, and, in particular, finding solutions to the European refugee crisis. We propose a centralised clearinghouse – a ‘two-sided matching system’ – to match refugees with states. Drawing on the success of matching in domains, such as education and healthcare, we outline the principles underlying matching system design, and illustrate in general terms how they could be applied to refugee protection. This matching system respects the priorities of states and gives agency to refugees. We then move to consider two specific empirical cases: international resettlement and the European migrant crisis.

Keywords: refugees, forced migration, market design, matching, resettlement, Europe, EU, burden-sharing.

Introduction

As the number of refugees reaches its highest level globally since the Second World War, and Europe faces the largest mass influx of refugees in its history (Alfred 2015), the international refugee regime is looking increasingly shaky. Faced with 200,000 refugees fleeing the ongoing slow-motion collapse of Syria and the rise of so-called Islamic State, the European Union response has been inconsistent, poorly coordinated, and (to date) inadequate to the scale of the problem. In this context, there have been

¹ Will Jones (william.jones@qeh.ox.ac.uk) is Departmental Lecturer in the Politics of Forced Migration, Refugee Studies Centre, University of Oxford

² Alexander Teytelboym (alexander.teytelboym@inet.ox.ac.uk) is the Otto Poon Research Fellow at the Institute for New Economic Thinking, University of Oxford

³ We would like to thank Alexander Betts, Cathryn Costello, Sarah Glatte, Olaf Kleist, Scott Kominers, Shengwu Li and Tom Scott-Smith for their insightful comments on earlier versions of this paper.

repeated calls for the revision or outright abolition of the Dublin II Regulation,⁴ particularly its requirement that the first country in the EU that asylum seekers enter be responsible for processing their claim (Costello and Guiffre 2015; Koser 2015b). As Constanza Hermanin, a senior analyst in migration and equality for the Open Society Foundation has noted, ‘after 25 years, it doesn’t work anymore.’ Across the world, the rise of far-right anti-migrant political parties and security rhetoric have led to governments’ implementing ever more restrictive asylum and immigration policies (Chimni 1998; Gibney 2010; Huysmans 2006; Weiner 1995).

At the time of writing, there are renewed calls to make the largely aspirational Common European Asylum System (CEAS) a reality. As Alexander Betts, Director of Oxford’s Refugee Studies Centre, noted in the *New York Times*, ‘Europe needs a comprehensive global refugee policy’ (Erlanger and Smale 2015). This is, therefore, a useful juncture to consider innovative suggestions as to how the mechanisms, which might underpin such a policy, could be made fairer, safer, and more effective. In this paper, we consider how insights from over thirty years of work in economics, operations research, and computer science on matching theory and market design can provide guidance on how to implement a desirable refugee policy. First, we briefly review the academic literature and recent applications of matching theory in different real world contexts. Second, we explain in general terms how matching theory could be productively applied to help states cooperate in refugee protection and burden-sharing, and how it makes that cooperation more likely. Then we consider two specific contexts: the global refugee resettlement mechanism and the current European refugee crisis. Finally, we outline and respond to some potential objections.

There are effectively two major problems we believe matching theory can address. First, the processes whereby refugees do or not do not get protected by particular countries are arbitrary, unfair, unsafe, and ineffective. Refugees play an extremely dangerous lottery when they select which country to apply for asylum in. In the case of global resettlement, refugees have little to no freedom over where in the world they will be resettled to. In the case of Europe, refugees must gamble by applying for asylum in only *one* European country. In both cases, this denudes refugees of agency and puts them through processes which are wracked with arbitrariness, ‘calculated kindness’ (Loescher and Scanlan 1998), and bad faith. Second, from the perspective of states, refugee flows are chaotic, unpredictable, and widely regarded as socially disruptive and destabilising. *Inter alia*, this foments anti-migrant politics, undermines sympathy with refugee groups, and limits the effectiveness of refugee reception programmes. States therefore seek to restrict and regulate migration flows, often in ways incompatible with their international legal obligations, to say nothing of the rights of refugees themselves. To the extent that states try to meet rapidly changing international movements of people with inflexible ad hoc arrangements, refugee protection becomes more expensive, wasteful, ineffective, and – for those reasons – less likely to happen at all.

⁴ The Dublin II Regulation (No. 604/2013) is the European Law ‘establishing the criteria and mechanisms for determining the Member State responsible for examining an asylum application lodged in one of the Member States by a third-country national.’

The two aims of expanding the agency of refugees and enabling states to control and manage migration are usually seen as in tension. It is our belief that both can be met more successfully through the use of a matching system – *the Refugee Match* – which uses no monetary transfers, better respects the freedom and choices of refugees, and serves the core interests of states in regulating their borders.

The proposed mechanism would require states to coordinate minimally with each other insofar as they would have to opt into it, but it need not force states to abandon any of their current agendas, be they liberal, or restrictive (for example, it is compatible with minimum or maximum refugee quotas, or neither). In consequence, a market-like system of the form we propose would give states additional incentives to engage in cooperative behaviour, increase quotas, and therefore promote interstate arrangements on burden-sharing, such as the CEAS.

Concretely, what we propose is a system that *matches* states to refugees and refugees to states. States and refugees give their preferences – over which refugees they most wish to host or which state they most wish to be protected in – to a centralised clearinghouse which matches them according to those preferences. Refugees, in principle, could submit their preferences from anywhere. The outcome of this mechanism satisfies a host of desirable properties. This system involves no money, works within quotas or under other distributional constraints, can be made to guarantee that all refugees within the system are hosted somewhere, can ensure that no refugee can be made better off (given their stated preferences) without making at least one other refugee worse off, and can be made ‘safe’ for states and refugees to honestly reveal their true preferences.

The use of economic theory in refugee protection is not new. Most prominently, Peter Schuck proposed a system in which states would be permitted to purchase and sell quota compliance obligations (Schuck 1997). It is important to recognise how different our proposal is. A matching system need not involve money, and ours is one such example: it merely sets out a way to find the optimal allocation of a given population refugees amongst a given population of states, given their preferences. No buying or selling is involved. We agree with those who regard the buying and selling of refugees in such a manner as morally repugnant, and likely to lead to worse protection in practical terms (Anker, Fitzpatrick, and Shacknové 1998). However, our proposal is compatible with this system, or any other burden-sharing agreement. It would work with Schuck’s tradable quotas as much as it would with the planned economy of Hathaway and Neve (Hathaway and Neve 1997), or the system of humanitarian visas bestowed at consular outposts outside the European Union recently advocated by Alexander Betts (Betts 2015). We are agnostic as to how burdens should be optimally shared, but do argue that any burden-sharing agreement which then uses a matching system is better for refugees and more likely to be assented to by states.

The idea of quota-trading has been recently revived by economists (J. Fernández-Huertas Moraga and Rapoport 2014; 2015; forthcoming). In their proposal, they also suggest that a matching system can play a (fairly minor) role in the allocation system. We focus explicitly on the design of the Refugee Match under various institutional constraints that takes quotas as given. The matching system is not merely an ancillary

component of a tradable quotas system. The details of its design are very important to its functioning, with or without tradable quotas.

In practice, the absence of a formalised system, which enables states to trade quotas in refugees, has not prevented them from engaging in actions with precisely the same effects. For example, many EU states are currently transferring resources to protect Syrian refugees in Turkey in the hope this will diminish the numbers attempting the onward journey to continental Europe ('Migrant Crisis: EU Plan Offers More Money for Turkey Camps' 2015). This is, in practical terms, no different from paying Turkey in order to transfer protection obligations from European states to Turkey, except insofar as it is less transparent than a formalised system of tradable quotas.

Matching Theory

Two-sided matching theory is a mathematical framework for allocating of resources where both parties to the transaction need to agree to the match in order for a match to take place. One popular (among matching theorists) example of 'one-to-one' two-sided matching is (non-polygamous, consensual, heterosexual) marriage. In this case, the two 'sides' are men and women. A 'match' here is a marriage. No man can get married without the agreement of a woman, and vice versa. Men and women have preferences over partners on the other side. For example, Arabella may prefer Boris to Ciarán, whereas Davina may prefer Ciarán to Boris. Boris and Ciarán would also have preferences concerning Arabella and Davina.⁵ Economists often describe this as a 'matching market' in the sense only that it is a scenario in which agents make choices seeking to maximise their preferences within given constraints – there need not be money involved, let alone prices. Gale and Shapley (1962) showed that any such one-to-one two-sided matching market (such as renting apartments or getting taxi rides) has a stable outcome. 'Stable' here denotes a situation where, first, no individual prefers to be alone to the match that they have ended up with, and second, no pair would like to drop their respective partners and rematch with each other.

Gale and Shapley described a procedure to find such a stable matching. It works in the following way: men and women submit their (fixed) preferences (a ranking of all men or women) to the clearinghouse (normally, electronically to a computer). The rules of the clearinghouse specify which side gets to 'propose'. Suppose that women propose. In the first step women propose to their first-choice men. Some men may receive multiple offers. Those men then conditionally (i.e. non-bindingly) accept their most preferred woman, and reject all others. Whilst any final acceptance is, in the words of Gale and Shapley, 'deferred' until the very end of the algorithm, *all rejections are permanent*. The women whose proposals have been rejected then make a further proposal to their second most-preferred men. The men then consider both the new offers and the standing old offers, and once again select their most preferred option and reject all options. This procedure eventually terminates. At that point, the allocation (of marriages) is

⁵ They may also prefer to be single to marrying a particular partner.

stable. Stability of an outcome in the matching process turns out to be a key feature that prevents such a matching from unravelling (of course, real marriage markets are hardly stable, but we speculate that this is because people's preferences change). In effect, a stable outcome means that the matching is in equilibrium.

Gale and Shapley further showed that a stable outcome also exists in any given 'many-to-one' system, such as college admissions, where a single university can be matched with many students, but any given student only goes to one university. One well studied and carefully designed many-to-one matching system is the National Resident Matching Program (NRMP) in the United States, which allocates junior medical interns to hospitals (A. E. Roth 1984). The NRMP has, in fact, been using the Gale-Shapley algorithm to allocate since the 1950s. Besides producing a stable outcome, the ongoing success of the NRMP relies on two further important properties of this algorithm (A. E. Roth 1991). First, there is an outcome in this many-to-one matching system is Pareto-efficient, meaning that nobody can be made any better off without making any other party worse off. Work in refugee studies has rightfully raised concerns over the loaded and misleading ways in which the language of 'efficiency' is deployed in political discourse (Gibney and Hansen 2005). Alexander Betts has argued, however, that there are legitimate uses of the language of efficiency in refugee studies (Betts 2006). When we refer to efficiency in this paper, we simply mean Pareto efficiency (or 'allocative efficiency', in Betts' language).⁶ Second, reporting your ranking to a system that produces a stable outcome is 'safe' for one side of the match, meaning that that one side of the match has no incentive to misrepresent their preferences to the clearinghouse (e.g. their ranking of hospitals, universities, marriage partners, or whatever) in order to game the system (this is sometimes called 'incentive-compatibility', or 'strategy-proofness'). Moreover, Roth showed that if the doctors do the proposing, then they have no incentive to misrepresent their preferences and game the system because the Gale-Shapley procedure selects the outcome that is most preferred for every doctor, of all stable allocations (A. E. Roth 1986).⁷ For this reason, Roth suggested that the NRMP should switch from a hospital-proposing Gale-Shapley algorithm to a doctor-proposing one, in order to produce a system where doctors felt safe revealing their real preferences about which hospitals they wished to go to, and ended up with the best possible stable outcome for the doctors (rather than the hospitals).

Stability has been key to designs of many matching systems. A. E. Roth (1991) examined medical job application procedures in the United Kingdom. At the time, the UK operated with many different systems across different parts of the country, some of which used methods, which produced stable outcomes, some of which did not. He found that in almost all instances whenever a matching mechanism, is unstable in the technical sense it is also unstable in the vernacular: those systems unravelled as doctors and hospitals sought to 're-match' outside of the system (i.e. under the table).

⁶There is a more involved discussion of what efficiency can mean in the context of a proposed matching system for refugee below.

⁷Hospitals may be able to manipulate the Gale-Shapley algorithm even if they propose. In one-to-one matching whichever side proposes in a Gale-Shapley algorithm can safely reveal their true preferences.

Since the late 1990s, there has been an explosion in the number of applications of matching systems. One famous application of two-sided matching systems has been the allocation of students to public schools in Boston and New York (Abdulkadiroğlu and Sönmez 2003; Abdulkadiroğlu, Pathak, and Roth 2005). What Abdulkadiroğlu and Sönmez showed is that the procedure used at the time by the Boston school authority was neither stable nor strategy-proof. This meant that parents had incentives to try and game the system, and richer parents tended to be disproportionately good at doing so. Not only was there ‘re-matching’ outside the system (as students sought places in schools that they were not matched to), but the Boston school authorities also had no reliable information about which schools were genuinely preferred by students and their parents. The designers of the match managed to convince the school authorities to change their mechanism to Gale-Shapley, which they have used ever since. Other applications of two-sided matching have included the match of cadets to branches of the US Army (Sönmez and Switzer 2013) and admission of students to universities in Germany (Westkamp 2013).

It is also worth noting that matching does not necessarily need to be two-sided. Agents may not be neatly divided into sides, but may still wish to match. Just think of finding your room-mates. Famous applications of one-sided matching systems include pairwise exchange of kidneys (A. E. Roth, Sonmez, and Unver 2004) and the allocation of public housing (Abdulkadiroğlu and Sönmez 1999). Most importantly, most applications of matching systems, including the one we will describe in this paper, *do not involve the exchange of money*. Many systems, such as a market for organs, may be justifiably considered (what Roth calls) ‘repugnant’ if money were introduced in transactions (A. E. Roth 2007).⁸

In all these cases, matching systems were designed to suit the particular constraints of the problem in question. Economists are usually guided by attempting to design a matching system that produces outcomes that are stable, efficient, and strategy-proof, but what that means in particular cases changes. The point of designing the system is to ensure that students are matched to the schools they wish to go to, or tenants to the houses they actually wish to live in, or doctors to the hospitals they wish to work at. If the systems described above were not designed, but instead decentralised, they would probably be unable to overcome the informational and temporal frictions that plague them. The consequences of this are students ending up in schools they do not wish to attend, doctors mismatched to hospitals, and patients being unable to find the kidney that could save their life.⁹

A Many-to-One Match for Refugees: Why

⁸ A magisterial summary of the successes of matching systems and market design is by A. E. Roth (2015).

⁹ One example of a system which has repeatedly failed is the matching of clerks to judges in the US judiciary (Avery et al. 2007). The reason for the failure is the resistance of judges to playing by any centralised rules of the game, even though many judges recognise that the current system is broken. For example, judges are incentivised to give early ‘exploding’ offers (i.e. which expire very quickly if they are not accepted) to clerks, thereby often ending up hiring clerks they do not want, and clerks working for judges they do not particularly want to work for.

The situation of refugees is already a matching market. It just happens to be a ‘decentralised’ one: there is no clearinghouse, but states and refugees nonetheless match. In order to become a refugee in a particular country, they must be accepted by that country, and voluntarily choose to live there. A refugee can only avail themselves of the protection of one country, but one country can take many refugees, making this a many-to-one matching market, similar to those of the NRMP or college admissions.

In most general terms, a *centralised* matching system for refugees would allow refugees to apply for protection in several countries and allow countries to compete to protect different refugees. In one potential system, refugees could make one claim for asylum to a single centralised body, simultaneously specifying where, if successful, they would wish to be relocated. A different system would be one where states have identified in advance a population to be resettled (either through a group determination of status, or some other special dispensation). The match would then match refugees to countries. States come to the clearinghouse with a quota of refugees they are willing to accept (we discuss minimum and maximum quotas in the next section) and a ranking of the refugees. They are then matched.

Centralising this system would require harmonising eligibility procedures and refugee determination of status decisions. This is an important precondition for the system to work well. In any such system, it would need to be a certainty that, once the state and the refugee have been matched, this match will actually happen (which is to say, that individual will be granted refugee status in whichever country they have been matched to). States are sometimes thought to be ‘jealous’ guardians of their sovereignty, and unwilling to cede it in any circumstances. However, it has been a matter of broad consensus in International Relations across many disciplinary divides since at least the 1980s that there is a wide set of circumstances in which states can and do cooperate and surrender sovereignty to interstate regulatory frameworks (Keohane 1984; Milner 1992; Goldstein 2001; Koremenos, Lipson, and Snidal 2003; Hurrell 2007; Abbott et al. 2015). There is much disagreement on precisely why states cooperate, how much cooperation happens, and how durable that cooperation is, but it nonetheless is clear that states *at least* cooperate when they believe it to be in their clear interests to do so. In the context of refugee protection, this argument is most trenchantly made by Betts (2011). In the circumstances where states cannot be persuaded to surrender control of eligibility determination, this proposal is a non-starter. However, we would argue that what states would gain in return for this is much more *control* over who is ultimately settled in their country than they currently possess. States would be empowered to give their preferences as to which refugees they wish to accept to the system, just as refugees give theirs. The likely reason why some states currently are reluctant to sacrifice control over eligibility determination is that they have no control over which refugees come to their country (and, as a consequence, the outcome of the decentralised matching market is unstable).

Moreover, there is *prima facie* evidence that the circumstances we examine in this paper are cases where states are willing to sacrifice this power. In the case of refugee resettlement,¹⁰ refugee status determination is already not always conducted by the states which agree to resettle refugees (for example, the US government resettled 60,000 Bhutanese refugees from Nepal in 2006). A coalition of Western states successfully resettled two and a half million Indochinese refugees between 1975 and 1997 (Robinson 1998). There are therefore numerous circumstances in which states already surrender refugee status determination and a centralised matching system would merely require formalising this process. In the European case, there is a statutory commitment on the part of the member states of the European Union to harmonise eligibility procedures, the 2011 Qualification Directive, which was signed by all members of the EU except Denmark, Ireland, and the United Kingdom.

Even if states cannot be persuaded to use matching in burden-sharing, an individual state could set up a matching system across its different regions (e.g. British local authorities or German *Bundesländer*). This would have many of the advantages, without the issue of having to get sovereign states to opt in collectively. For example, British local authorities will have granular information about their particular advantages vis-à-vis other councils (e.g. some councils may have more capacity in child services than care for the elderly, and *vice versa*) and they should be encouraged to convey this information to refugees. Indeed, a refugee matching system could generate dramatic improvement in the welfare and integration of refugees within a given country's national borders and significantly reduce ghettoization and restrictionist arguments about capacity gaps.

If states opt into such a system, rather than determining the status of those refugees that have applied to them specifically, they would submit their ranked preferences to the clearinghouse. One disadvantage of this is that now states have to look at *all* the refugees in the system, rather than only those who happen to have applied to their respective national bureaucracies. Of course, this may not be particularly burdensome, as the ranking could be done using very straightforward principles (for example, a state could prioritise refugees that have been waiting longest, or face the immediate threat, or that they are best placed to assist).

Once states have committed to taking a given number of refugees, there is a question of how they prioritise which refugees they wish to receive. There are a variety of principles, which academics and policy-makers could seek to use in trying to determine who will be prioritised, given practical and political limits on how many refugees can be taken in. For example, the United Kingdom has recently stated it wishes to expand the Vulnerable Persons Relocation Scheme, which works to resettle Syrian refugees in the UK. The government has stated that their priorities are determined by greatest need, and an assessment of where the UK can singly make the greatest difference (McDermott and Vasagar 2015). On the other hand, the governments of Slovakia, Poland, and the Czech Republic have all signalled a

¹⁰ Resettlement here refers to 'the organized movement of pre-selected refugees to a destination country in which their settlement is expected to be permanent' (Van Selm 2014)

willingness to take more refugees, but only if they are Christian (Wasik and Foy 2015). Many would regard the latter as morally unacceptable and discriminatory, but the former as largely uncontroversial. The point is just, for the moment, that different states are free to rank these principles differently in the status quo, and reasonable, non-reprehensible states might rank the same refugees differently (for example, it would be eminently reasonable for Brazil and France to prefer Lusophone and Francophone refugees respectively).¹¹

Using a matching system does not dictate which principles states are allowed to use in ranking refugees, and the clearinghouse could permit or forbid the use of any criteria. Just as matching for doctors does not allow hospitals to engage in racist hiring practices, the clearinghouse would only allow states to rank refugees based on criteria which are compatible with the principles and goals of the 1951 Convention, or any other set of principles. Even if all states ended up ranking refugees in the same way (which is unlikely), the clearinghouse would still be an improvement on the status quo, as it would give refugees themselves more control over their lives.

The Refugee Match is simply the tool, which finds the best allocation, given whichever set of principles are embedded in the processes. Irrespective of those principles, what will be true of *any* matching system is that it will expand the agency of refugees, by giving them greater choice over where they are ultimately protected, without necessitating dangerous and uncertain journeys, and will give states greater control over which refugees they take in. Right now, the problem is that refugees have no choices, and states have no control. The Refugee Match can help both.

A Many-to-One Match for Refugees: How

What will the matching system for refugees practically look like? The first step is to agree what the two sides of the system are (i.e. which refugees, and which states). Any subset of states could agree to set up a matching system, and allocate any group of forced migrants. If, as we suggest here, the system is for asylum seekers, the second step is to create a centralised body that will determine the status of applications by individuals to the system. If the individuals are already refugees under the 1951 Convention, or for any other reason are deemed eligible for the system in question, then that is unnecessary. In order to create a system, which ensures that families, which do not wish to be separated, are not, the clearinghouse needs to accept preference submissions from households. The households would be treated as a single unit,¹² and states would then evaluate those households in the same way.¹³

¹¹ Note that countries can also, in principle, *agree* on what refugees have the highest priority.

¹² Which means participating states would need to adopt a common understanding of the family unit in question. Currently the definitions of family and dependent-relative status are specific to the jurisdictions of particular states, but this is no more than committed to in many domains of EU law, where a common definition of the family is required.

¹³ Families are not ‘unsplittable’ once their refugee status has been determined for a variety of reasons (for example, one member of the family can commit a crime incompatible with their refugee status without undermining the

One important consideration is whether to take into account interdependent preferences of extended families. For example, parents A and B together with their only child C should be treated as an independent household as should parent D with his children E and F. But what if B and D are cousins? A-B-C may have different preferences from D-E-F, but other things being equal, they would prefer to be resettled in the same country. This could undermine stability, efficiency, and strategy-proofness. This issue has arisen in matching couples in the doctor-hospital match (A. E. Roth and Peranson 1999). Theoretically, these sorts of preferences can create problems for the matching mechanism, but, in practice,¹⁴ when the number of people participating in the matching system is sufficiently large (and any meaningful system for refugees certainly would be), the system has often functioned very well. In the medical match, doctors who were couple had their preferences satisfactorily met, as could be the case for refugee families.

The next step is to determine the frequency at which the system is cleared (i.e. how often the matching is performed). For example, the NRMP happens annually, whereas the matches formed after speed-dating events happen every day. Designers of the match can adjust the frequency of clearing depending on the urgency and scale of the refugee flows.

Next, states need to determine in practical terms how they are going to rank the refugees that are in the system. Broadly speaking, they have two options: strict or weak ranking. Strict ranking, as in the NRMP, would require the states to assign a position to every single refugee they were willing to consider offering protection to. Realistically, given the numbers of refugees and the capacities of state bureaucracies, this is likely to only be achieved by using a points-based system. An alternative is to use weak ranking, as used in the Boston and New York public school system, where the refugees are ranked according to blocs of priority classes. For example, the state could deem those facing immediate personal danger to be the group with the highest priority. Within the classes, whatever they are, the ranking is done according to a lottery (this will only affect the refugees in a country's lowest-ranked priority class that has any but not all members accepted). From the point of view of refugees, the ranking is more straightforward: they simply submit a list of the countries they would be willing to be protected by in rank order (this also works if a given country offers more than one form of protection, such as full refugee status alongside some lesser form of temporary protection. This is explored in more depth below). If one wants to design a system, which guarantees that every refugee is granted protection somewhere, then refugees must rank every single country (otherwise, they are effectively saying that they would rather not be granted protection than go to a particular country).

There is a potential concern about 'manipulability'. Put simply, states may worry that refugees will misrepresent themselves in order to get into higher priority categories (e.g. pretending to have a given educational qualification, or sexuality, if either of those are priority categories). This is a well-known

refugee status of the rest of the family), but this doesn't not undermine the stability of the system, as it does not create scenarios where states and refugees wish to re-match outside the system.

¹⁴ And, as it turns out, even in theory for large markets, see Kojima, Pathak, and Roth (2013).

problem, and sometimes leads to individuals making choices to enter into higher priority categories rather than misrepresenting themselves (as when individuals move house to get into a school catchment area to get neighbourhood priority). States must be able to trust the information they are receiving in order for them to consent to opt into such a burden-sharing mechanism at all. Unfortunately, many of the least manipulable priority categories are precisely those categories it would be unacceptable to use (e.g. it is very hard to manipulate one's race or sex). Some important categories are not: for example, some European states have already introduced DNA testing to verify family membership (UNHCR 2008). There are other categories of priority, which, in certain circumstances, would also be hard to manipulate: for example, who has waited the longest is often a matter of record. Assessing which refugees are in the most immediate danger could be similarly hard to manipulate if, for example, states wished to prioritise those who are registered in settlements immediately on the Turkish-Syrian border in the European case. Most obviously, prioritising children creates a relatively difficult status to manipulate, although it also means that the longer implementation of this system is delayed, the more people will be manipulated out of this category by growing up. Finally, the fewer priority classes there are, the easier it would be to verify applications and discourage manipulation.

Ultimately, this creates trade-offs, one between which priorities we pick, and one concerning how transparent those priorities are. Refugees are less likely to manipulate the system if they are uncertain about what the priorities of states are. Of course, this raises significant concerns about the transparency and accountability of the system. It would be for the creators of the Refugee Match to assess the relative importance of these goals. They could decide to create a fully transparent system (where the public and refugees know the priority categories and the ranking of priority classes in full), a semi-transparent system (where refugees know that all the categories on their application are being used, but not the ranking of priority classes), and a completely non-transparent system (where refugees neither know the priority ranking of priority classes, nor even which categories are being used). The first category is the most manipulable, and the last category is the least manipulable.

One of the key elements of the matching system for refugees is the determination of quotas. States can have both minimum and maximum quotas if they desire.¹⁵ The minimum quotas would need to be coordinated between the different states in order to ensure that a sufficient number of refugees are being allocated. In order to ensure that all refugees are allocated, it is important that no state ranks any refugee in the system as unacceptable, and that the minimum quotas of all participating states collectively add up to the total number of refugees. If this can be agreed, it is guaranteed that every refugee in the system will be offered protection somewhere, but if it not, then the refugees given the lowest priority by states would have to wait until those quotas are increased, and the system can clear again. States can unilaterally determine maximum quotas without affecting the functioning of the system. Matching systems with

¹⁵ States can also have subnational quotas in the same way as hospitals fill vacancies for different departments. For example, the German *Königsteiner Schlüssel* system allocates refugees to *Bundesländer* based on a system of fixed quotas (for example, Berlin has five per cent of the quota) without any regard for the preferences of refugees themselves.

maximum quotas are very well understood and are implemented in a variety of contexts, including the NRMP. Minimum quotas, however, present a challenge to designers of matching systems (Ehlers et al. 2014). In particular, it is difficult to ensure that the system is stable, efficient, and incentive-compatible. Economists have developed a variety of tools to overcome these difficulties and satisfy these three requirements to some extent in certain settings. For example, Kamada and Kojima develop a matching system for doctors in Japan that ensures that rural hospitals have a minimum complement of medical staff (Kamada and Kojima 2012).

Economists have also very recently developed tools to handle ‘distributional constraints’ (Kamada and Kojima 2015). One setting where this is often used is college admissions systems where affirmative action is a requirement (Hafalir, Yenmez, and Yildirim 2013; Kominers and Sönmez Forthcoming). Rather than just having minimum quotas, this system also insists that a certain distribution of students from particular backgrounds, ethnic identities, etc. are distributed in every college. In the context of refugees, there may be some reason why states would wish to decide in advance that the refugee populations they take must meet some distributional requirements. States could collectively pre-commit to taking a diverse population of refugees if, for example, one wished to prevent the scenario where some governments take in only Christian refugees. If this is something states wish to do or can be persuaded to accept, this feature can be built into the system.

One particularly attractive feature of our proposal is that the matching system can deal with different forms of protection. In the language of economics, this is called a ‘matching market with contracts’ (Hatfield and Milgrom 2005). In this setting, rather than just having a match, the agents have a choice about what type of match they form. For example, in the US Military Academy, cadets not only choose the branch of the military that they wish to serve in, but also have a choice over two types of contract that differ in terms of their length of service (Sönmez and Switzer 2013). In our context, the international refugee regime is often a confusing hodgepodge of different legal statuses stemming from various schemes, all offering slightly different forms of protection. Attempting to navigate this system often results in refugees applying for the ‘wrong’ status, losing their right to apply for a status they would have received, or applying for a sub-optimal status even when they might have been successful in securing a more substantive or durable form of protection. As with allowing refugees to apply to multiple countries, the solution here is to allow them to rank within the different ‘contracts’ potentially on offer. States, on the other hand, can have different maximum and minimum quotas for different types of protection status, building even more flexibility into the system.

The final decision that needs to be made is about precisely which matching mechanism to use. In the previous section, we described the famous Gale-Shapley algorithm, but there are many others. One interesting result from the economics literature is that there exists no matching mechanism that produces outcomes that are stable, efficient, and at the same time give all agents an incentive to truthfully report

their preferences (A. E. Roth 1982; 1984).¹⁶ Therefore, a choice needs to be made as to which of these properties are most important or to partially relax some of the constraints on what the system must guarantee.¹⁷ For example, some cities in the US use the Gale-Shapley algorithm to allocate students to public schools (which is strategy-proof and stable, but not necessarily efficient), whilst others use what is known as the Top-Trading Cycles algorithm which is strategy-proof and efficient, but not necessarily stable (Kesten 2006). In complex settings with minimum quotas and distributional constraints, more severe trade-offs between stability, efficiency, and strategy-proofness are present. Nevertheless, good matching systems have been designed and successfully implemented under these conditions (Biró et al. 2010).

Let us put the concepts of stability, efficiency and strategy-proofness in the context of refugee protection. A stable outcome in a matching system in which states actively rank all refugees would be one in which no country wants to exclude a refugee household and admit another instead, and no refugee household would rather leave the country they are already living in and be readily accepted by another country. Another way of looking at it, analogous to public school choice, would be that states (collectively or individually) establish an order of priorities, rather directly ranking refugees. In effect, a priority list is mathematically equivalent to a preference ranking, but the states do not make active choices about specific individuals, instead pre-committing themselves to a particular way of ranking refugees. Viewed this way, stability reflects a sort of ‘fairness’: no refugees’ priority for residence in a particular country will be violated (i.e. no refugee is assigned to a country when there is another refugee somewhere else who is higher up that country’s priority list, and wishes to go there).

Efficiency is trickier. Indeed, our definition of efficiency may change according to how we view the matching system. We can either include states in our considerations of efficiency or exclude them, to consider efficiency only from the point of view of refugees. If a country simply assigns priorities to refugees according to an international agreement, then we do not need to consider the interests of states in the efficiency calculation (country residence permits are simply objects to be assigned). Hence, assigning some refugee to a more preferred country (without hurting other refugees) may be more efficient even if it violates a country’s priorities. It may, however, be considered ‘unfair’ for receiving states because these refugees may have a lower priority for them than some other group of refugees. This is the efficiency criterion (i.e. one which only considers efficiency from the point of view of one side of the system) used in public school choice. If, on the other hand, countries make active preference decisions, an efficient outcome would be one in which it is not possible to improve the lot of a group of refugees *or of a group of countries* without hurting others. This would be how efficiency is judged in a hospital-doctor matching system.

¹⁶ It is actually much worse than that. For any quotas, it is possible for some preferences that the doctor-proposing Gale-Shapley algorithm results in the doctors getting either their worst or second-worst choice (Kesten 2010).

¹⁷ One particularly promising direction has been to consider systems with a large number of participants where these properties are only approximately satisfied (for example, Azevedo and Budish (2013) show how a weaker version of strategy-proofness can square this impossibility circle in large markets).

A system, which is strategy-proof for refugees, would ensure that no refugee had any reason to misrepresent their genuine preferences over what countries they wished to live in. There would no longer be any reason for a refugee to gamble seeking protection in one country based on an inaccurate belief that they would be more likely to receive protection there, thereby forgoing protection in another country. On the other hand, if the system were designed in order to make it strategy-proof for states, no state would have any incentive to dishonestly misrepresent their position vis-à-vis the refugees they were willing to accept.¹⁸

It would be for the creators of this system to decide which of these properties to prioritise and what distributional constraints to impose - these desiderata will determine precisely what mechanism should be used (see, for example, Echenique and Yenmez (2015)). Efficiency seems like an obvious first-order priority. In addition, it may be necessary to ensure the system is strategy-proof for refugees, making it impossible for refugees to game the system, in order to achieve political buy-in. On the other hand, states may care more about creating a system in which it is safe for them to accurately report their preferences, and they can trust that other states are doing the same, thus requiring strategy-proofness for states.

If states do not actively rank refugees, but rather assign priorities to groups of refugees, then the role of stability as ‘fairness’ is important in order to ensure that the system is transparent and does not arbitrarily discriminate against groups that the state has publicly committed to prioritising.

If states actively rank refugees, then there is a good case for sacrificing stability so that the system can be more efficient and strategy-proof. Preventing re-matching in many settings, such as the doctor-hospital match, or the public school match, is extremely difficult: individuals frequently go outside the system to try and renegotiate their deals under the table, leading to the unravelling of the whole system. This is why stability is normally considered important (A. E. Roth 1991). However, in our case, it is unlikely that a refugee will be matched to a country, receive refugee status, and then go on to apply for asylum in another country. This is particularly unlikely if states collectively agree that all such applications must now go through the central mechanism. For example, in Europe, while refugees in the Schengen states may travel to other Schengen states, they may not stay for longer than 90 or work without a permit. Moreover, they need a visa to travel to non-Schengen states (UNHCR 2007). Therefore, sacrificing stability does not affect the robustness of the outcome of the matching system because most ‘re-matchings’ that would render an outcome unstable are prohibited, costly or time-intensive. In a well-designed matching system for refugees, stability comes for free.

One remaining issue is the *dynamic* nature of the matching system, which is a feature that has only recently been considered by economists (Dur 2011; Pereyra 2013; Thakral 2015; Akbarpour, Li, and Gharan 2015;

¹⁸ Indeed, there are also mechanisms that, under certain conditions, can even make it safe for the states to reveal their true *quotas*. These certain conditions, however, are extremely unlikely to obtain, and are therefore too strong for our purposes. In a matching system with a *large* number of countries and refugees, neither side would have any incentive to misreport their preferences (Immorlica and Mahdian 2005; Kojima and Pathak 2009). But there is no guarantee that there will be enough countries participating in any given matching system to ensure this.

Kadam and Kotowski 2015). Dynamic matching occurs when agents do not enter the matching system all at once, but rather over time. We might worry that if the clearing of the system happens at intervals that depend on the flow of refugees, some refugees might try to delay entry into the matching system because they expect the matching system to be less competitive in future and hope for a better outcome. However, this concern is mitigated by the inherent uncertainty about the clearing period that a refugee ends up in. This uncertainty is very natural: backlogs of applications often pile up making it hard to tell how long it takes to process an asylum request. But this uncertainty makes it hard for refugees to manipulate the mechanism, incentivising them to apply for asylum as soon as they feel is necessary in the knowledge that their application will be processed on a first-come-first-served basis (or a version thereof adjusted for vulnerability and immediacy of threat).

A Global Resettlement Mechanism

Globally, resettlement needs (i.e. refugees that have been deemed to be vulnerable in their current location) have been estimated to be as high as 800,000. However, only some 85,000 places are available on an annual basis (Van Selm 2014). Clearly, the key problem here is that there are nowhere near enough places. Our proposal does not purport to solve that problem. The scarcity of places in resettlement schemes, however, makes it particularly crucial that those places are allocated in the best way possible. Paradoxically, states currently have too much power but not enough choice. On the one hand, given the limited number of places beyond confirmation of refugee status and the determination of need, states are more or less unconstrained to cherry-pick the refugees they deem most desirable. Often, it will happen that refugees with particular skills or that better meet amorphous ‘integration criteria’ get resettled first. On the other hand, resettlement for states arise in the context of emergencies where they are presented with a binary choice between taking on an entire group *tout court* themselves or nothing. For example, European governments panicked in the face of large flows out of Libya following the civil war there and, in the end, resettlement was negligible (Koser 2015a).

In this context, we propose that the Refugee Match can help with the following problems.

First and most obviously, it can give refugees more choice than they currently possess about where they are resettled rather than being directed by the arbitrary whims of bureaucracies. Vulnerable refugees could state their preferences as to which of the twenty-odd states operating resettlement schemes they most prefer. Refugees themselves have better information about their lives, aspirations and the quality of life they would have in a particular country than hopelessly overstretched technocrats.

Second, a matching system can make the process less arbitrary and contingent on ad-hoc bilateral deals. In the status quo, states identify particular populations, which come to their attention for a variety of reasons, and simply resettle large numbers from those areas. Whether a refugee population becomes resettled thus is more to do with the competence and drive of particular civil servants than either a sober

assessment of who the vulnerable or the core interests of the state. If states are determined to form particular bilateral relationships, they can continue to do so alongside the matching system. American resettlement of Iraqis that worked with the US Army during the occupation of Iraq might fit this case. However, under a matching system, states can continuously adjust their quotas, assess priority using universal standards of vulnerability, and let the algorithm do the hard work.

Third, because the computer algorithm finds the matching outcome as soon as the preferences of both sides are submitted, the matching system can speed up resettlement. Refugees in extremely vulnerable situations often wait for years on referral lists. Much of what makes that process slow (such as vulnerability determination and the limited number of places) cannot be ameliorated by matching but it can short-circuit the oft-torturous process of negotiation between UNHCR, IOM and state bureaucracies.

Fourth, because matching enables states to effectively share burdens whilst also empowering them to receive refugees they prefer, this mechanism may encourage further states to participate in resettlement. If states realise that they are not under an obligation or expectation to resettle single-handedly an entire population of refugees, they may have more incentive to join the matching system. In addition, states may find it more palatable to settle a diverse set of smaller refugee groups than a larger and more homogeneous one.

The European Refugee Crisis

The substantial inflows of refugees into Europe have dramatically exposed the inadequacy of existing arrangements. The disintegration of Syria, where no less than nine million Syrians have been displaced, of whom three million are refugees, shows no sign of being resolved swiftly (UNHCR 2015). The overwhelming majority of refugees are in neighbouring countries: Jordan, Lebanon, and Turkey. There are now over a million refugees in Lebanon, in a total population of four million (*ibid.*). Jordan and Lebanon have closed their borders. Even for those that have reached refugee camps in the region, the situation continues to worsen. Left without the right to work, access to security, basic legal protections, and facing ongoing racism and violence, it is unsurprising that many refugees continue to attempt the long and dangerous journey to Europe. We should assume this situation will be enduring and that refugee flows will continue.

The current system, where the first arrival country has responsibility for refugees, is fundamentally unfair to countries such as Greece, Italy, or Hungary, and creates chaos and tragedy as European states play pass-the-parcel with human lives. It has never been clearer that a new deal on responsibility-sharing within Europe is needed. The Dublin II Regulation is unworkable. History suggests interstate cooperation to meet such a challenge is not impossible: from the Comprehensive Action Plan for Indochinese Refugees in 1989 to the European Evacuation of Kosovar refugees from Macedonia in 1999, a clear plan and astute political leadership have led to negotiated cooperation in the past. Crucially, a workable

solution requires cooperation across the EU. It is the hope of many that sufficient popular pressure of the sort that followed the publication of the tragic photographs of Aylan Kurdi, may be enough to get states on board, but there is no conflict between that and trying to persuade governments that a coordinated, managed response, is not only better for refugees, but ultimately for states as well.

Given the widespread recognition of the urgency of forming a comprehensive European asylum policy, the crucial role matching can play is in providing states with a framework which might persuade them to participate in burden-sharing arrangements to meet the crisis, thus overcoming the current deadlock.

A matching system can run either alongside or instead of a system modelled on Dublin II, where a refugee applies for protection in one and only one country. This, for example, could replace the current situation where Syrian refugees in the countries bordering Syria may apply for humanitarian visas from particular countries, but must pick a particular country to apply to. Refugees are then tied to a particular application to a particular country (embassies invariably keep the original documentation refugees provide, in part to prevent multiple applications). It advantages nobody that at present refugees must play the lottery in deciding where to apply to, and states cannot evaluate from the entire pool of applicants. Allowing any refugee to apply to a single system from *any* embassy is therefore in the interests of refugees and states.

Another potential use would be with something analogous to the 2001 Temporary Protection Directive, which created a set of exceptional procedures to deal with the crisis in the former Yugoslavia. A matching system could have been used to allocate refugees within that system without undermining any of its core aims, or affecting the rest of European refugee regime. In particular, the proposed system could exclude discriminatory categories such as religion and race, which are incompatible with European law, as well as being morally unacceptable, and focus on categories of vulnerability, suitability for integration, presence of family, and so on. The system would then match refugees and states. This, as before, necessitates that eligibility is centrally determined. Yet this is no more than what is already aspired to in the 2011 Qualification Directive. What is currently lacking is a way to persuade states to implement this in practice. But we argue that using matching has further advantages over and above simply persuading states to participate.

First, matching may save lives. The matching system allows refugees to submit an application for asylum to every participating state, and they could, in principle, submit it remotely. It could include, for example, regional processing centres in the Middle East and North Africa. The core advantage is that, to the extent refugees can have confidence that there is a fair and effective system that will grant them protection without attempting a dangerous crossing, they are less likely to take that risk. Unlike naval security operations in the Mediterranean, such as Operation Triton – which merely empower human traffickers and make the journey more dangerous rather than deal with the underlying demand – this is a system which makes refugees outside Europe no less likely to secure protection than those that have reached it (or at least one with a credible and safe alternative to smugglers). Of course, creating a system designed to

prevent deaths at sea and provide people with an alternative to dangerous crossings is principally and practically separate from the question of dealing with refugees that have already made it to Europe – these refugees must be addressed separately.¹⁹

Europe also offers a case study of how rapid matching may help manage a sudden chaotic influx. A matching system can be deployed as soon as the states agree on the quotas and the criteria for ranking refugees. To the extent that matching is successful, it reduces the chaos and volatility of rapid population movements as fewer people try and reach Europe, because getting to a particular destination has become substantially less important. Furthermore, the unmanaged chaos under the status quo does not make it likely that refugees will end up in the countries they genuinely wish to go to because they have family there, speak a relevant language, have contacts which might help them secure employment, and so on. This is all in the interests of states, because first, states wish to minimise the anarchy of sudden population flows, with the concomitant stress on infrastructure and scarce resources, and second, because it guarantees that they will also end up protecting those refugees they most prefer.

Notably, even if not all states can be persuaded to participate, it still makes sense for those states that are willing to use the matching system to do so.²⁰ Hopefully, it should become clear in time to any states staying out of the matching system that they thereby receive less-preferred refugees in comparison to states which do participate, thus encouraging wider participation. For this reason, the refugee match is precisely what can persuade states to opt into refugee burden-sharing, surrender control of eligibility determination, and ultimately protect more refugees.

Further considerations

It might be thought this system does not offer genuine choice to refugees, but merely creates a system in which states have all the power, and refugees have none. We argue the opposite. We note that it is states, in the current system, which have virtually *all* of the power. They may be constrained by international norms or legal obligations, but only weakly at best. Yet neither side has any choice. States must process only those who apply within their borders, refugees must accept the protection or resettlement that is offered to them, even where there is a country that they would have preferred to go to, that also would have preferred to take them. A matching system is *exactly* the situation where both sides need to agree. Refugees already exist within an international system of states, where it is necessary for a state to agree to render protection in order for a refugee to receive it. *Pace* total global open borders, the assent of states is

¹⁹ Right-wing politicians often invoke a distinction between deserving refugees ‘there’ in order to justify restrictive policies attacking the rights of alleged queue-jumpers ‘over here’. It is plausible that we undermine this discourse to the extent that we make it (even more) flagrantly untrue.

²⁰ Designing mechanisms that encourage agents to participate in it rather than contracting and transacting outside it is an active area of research in matching theory. The bottom line is that it may hard to create bullet-proof incentives for participation in many mechanisms especially if there is uncertainty about the outcome (see, for example, B. N. Roth and Shorrer (2015)).

a precondition for protection, and only agreement from both sides gives refugees any choices at all. Currently, much activism works assumes a strict dichotomy, where refugees can either have unconstrained choice, or not choice at all. For example, many European refugee activists advocate a system, which would give refugees to the EU total freedom as to where in the EU they decided to live while states are compensated for every refugee they take. This would be the best outcome for all refugees, but there is little-to-no chance states would assent to it. But this dichotomy also misses that in most scenarios justice does not require that we give individuals total choice, but choice from a reasonable set of options. Few would advocate that students must have the right to go to any university they wish, or doctors to work at any hospital. What is important is that they go to good universities and hospitals, and if *both* sides of the system have *some* choice, that is more likely to happen. Those choices, although not complete, are nonetheless real. If it really is the case that *almost every* Syrian refugee makes Germany their first preference, then of course they will not all be able to go there, but refugees would still have qualitatively more agency if they were to achieve their second, third, fourth, or fifth preference.

Were every refugee to submit the same preference ranking *from beginning to end*, states would merely pick the refugees they prefer (rather as if everyone in a given school district agreed completely on the desirability of the schools). Scenarios of ‘perfect correlation of preferences’, such as this, do indeed hand all of the power to the other side of the system (be that states, or schools). But this is extremely unlikely. In the Boston public school matching system, the image of a single hopelessly oversubscribed school is largely mythical.²¹ There are abundant reasons to believe that the preferences of refugees are as heterogeneous as they themselves are, as are those of countries. This heterogeneity is exactly what makes the matching system so appropriate and powerful.

Second, it might be thought that neither states nor refugees possess sufficient information for matching to be effective. In the status quo, states compete at telling refugees how unpleasantly they will treat them (as when Geert Wilders released YouTube videos telling refugees they had no chance of ever getting status in the Netherlands, and similar flyer campaigns of the Australian government in camps across Asia). Matching changes the game completely. In a matching system, where it is fixed in advance how many refugees will be matched to a given country, but not *which* refugees, it is suddenly important for states to compete at being *attractive*. This no longer affects the total number of refugees they will host (which was determined in advance), but it does affect *who* they will host. Because the match is two-sided, a state cannot merely take in the refugees it wants, if those refugees would prefer to go somewhere else.

²¹ This might be partly due to neighbourhood priorities. Between forty-three and sixty-six per cent of students were allocated to the first choice school and none of the students were allocated to their bottom three choices. In grade six, over eighty per cent of students get one of their first two choices and all students were allocated to their top five choice school (Dur et al. 2013).

Hence, states will have strong incentives to invest in refugee reception and inform the refugees about how welcoming they will be.²²

Third, it might be thought that a system of this form commodifies refugees, turning them into objects to be traded, exchanged, and competed over by states. Again, the opposite is the case. A two-sided system respects the preferences of agents on both sides. Rather than turning refugees into numbers, or ‘human cargo’ where refugees are passive recipients of the largesse of states, refugees become agents with preferences, choices, and agendas, and those preferences are now taken into account. A matching system is distinct from a traditional market for a commodity, as it requires consent of both parties, and does not involve money. It is thus no more commodifying than marriage, friendships, or public education.

Fourth, some will undoubtedly think it unacceptable to promote or give any legitimation to a discourse of ‘desirable’ and ‘undesirable’ refugees. We have several responses. Merely by centralising matching in a clearinghouse, we establish shared norms about what is acceptable and what is not. It prevents states from unilaterally creating arbitrary policies for political purposes.²³ Secondly, a system where states must submit these preferences to a centralised system makes them substantially more accountable to public scrutiny and judicial review than a system where states identify and take desirable refugees on the quiet. More importantly, a preference ranking is not an index of desirability, but rather of priority. For example, a university admissions system can accord priority rank to those from untraditional backgrounds, or who face additional barriers to attending university. In the same way, a preference ranking of states can reflect an assessment of who faces the greatest need, vulnerability, or any number of other variables to do with a refugee’s suitability for a given country that have nothing to do with evaluations of desirability. All that the matching system does is aggregate the priorities of states and refugees, and arrive at the optimal outcome within the constraints. It is undeniable that some states currently behave reprehensibly and others do not. But whether or not they do, both selfish and altruistic preferences would be better satisfied within the Refugee Match.

Conclusion

Of course, in order to *actually* solve refugee crises of this form, states would have to accept enormous inflows of refugees, and find a way to resolve the ongoing conflagration in Syria and elsewhere. Matching systems, such as the one we propose here, are never the total solution to the various issues they seek to address. They are merely a substantial improvement on the status quo within the constraints of the politically palatable, and may give states incentives to relax these constraints. Although matching mechanisms cannot make states behave morally, they will nonetheless improve the situation for refugees

²² By the same token, states might have an incentive to discourage certain refugees from applying because they cannot legally discriminate against them within the priorities defined by the matching system (although odious states that prioritise their Islamophobia over filling their skills gaps will be punished economically in the long term).

²³ Such as when states announce a willingness to cherry-pick telegenic orphans plucked straight from the Syrian border

whether or not states can be made to act in accordance with their legal and moral obligations. Our work, therefore, provides a pragmatic proposal in the spirit of those who argue that states will contribute towards efforts to protect refugees when they recognise a relationship between the rights of refugees and their own interests (Betts 2009).

The Refugee Match is a realistic, pragmatic, quickly implementable, and just improvement on much of the current international refugee regime. In this paper, we outlined how a matching system, which respects the preferences and choices of refugees *and* the priorities of states, can better protect the human rights of the vulnerable, and increase the likelihood that states will participate in burden-sharing.

This exposes a particularly unhelpful and misleading platitude in the current rhetoric of Europe's politicians. To take one example, in September of 2015, Thomas de Maizière, Interior Minister of Germany, announced the closure of the German border to Syrian refugees. He told a press conference that "asylum seekers must understand that they cannot choose the states where they are seeking protection" (Harding 2015).

Nonsense. Any system, which genuinely upheld the rights of refugees, would have to start by respecting their choices. Asylum seekers *must* be able to choose the states where they want to spend their lives. The Refugee Match is a good start.

Works Cited:

- Abbott, Kenneth W., Philipp Genschel, Duncan Snidal, and Bernhard Zangl. 2015. *International Organizations as Orchestrators*. Cambridge University Press.
- Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2005. 'The New York City High School Match'. *American Economic Review* 95 (2): 364–67.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 1999. 'House Allocation with Existing Tenants'. *Journal of Economic Theory* 88 (2): 233–60.
- . 2003. 'School Choice: A Mechanism Design Approach'. *The American Economic Review* 93 (3): 729–47.
- Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan. 2015. 'Dynamic Matching Market Design'. Mimeo. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2394319.
- Alfred, Charlotte. 2015. 'What History Can Teach Us About The Worst Refugee Crisis Since WWII'. *The Huffington Post*, September 12. http://www.huffingtonpost.com/entry/alexander-betts-refugees-wwii_55f30f7ce4b077ca094edaec.
- Anker, Deborah, Joan Fitzpatrick, and Andrew Shacknove. 1998. 'Crisis and Cure: A Reply to Hathaway/Neve and Schuck'. *Harvard Human Rights Journal* 11: 295.
- Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin E. Roth. 2007. 'The New Market for Federal Judicial Law Clerks'. *The University of Chicago Law Review*, 447–86.
- Azevedo, Eduardo M., and Eric B. Budish. 2013. 'Strategy-Proofness in the Large'. *Chicago Booth Research Paper*, no. 13-35. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2237425.
- Betts, Alexander. 2006. 'What Does "Efficiency" Mean in the Context of the Global Refugee Regime?'. *The British Journal of Politics & International Relations* 8 (2): 148–73.
- . 2009. *Protection by Persuasion: International Cooperation in the Refugee Regime*. Cornell University Press.
- . , ed. 2011. *Global Migration Governance*. Oxford University Press.
- . 2015. 'Let Refugees Fly to Europe'. *New York Times*, September 24, Online Edition edition. <http://www.nytimes.com/2015/09/25/opinion/let-refugees-fly-to-europe.html?nytmobile=0>.
- Biró, Péter, Tamás Fleiner, Robert W. Irving, and David F. Manlove. 2010. 'The College Admissions Problem with Lower and Common Quotas'. *Theoretical Computer Science* 411 (34): 3136–53.
- Chimni, B.S. 1998. 'The Global Refugee Problem in the 21st Century and the Emerging Security Paradigm: A Disturbing Trend'. In *Legal Visions of the 21st Century: Essays in Honour of Judge Christopher Weeramantry*, edited by Antony Anghie and Garry Sturgess. Martinus Nijhoff Publishers.
- Costello, Cathryn, and Maria-Guilia Guiffré. 2015. "'Tragedy" and Responsibility in the Mediterranean'. *Open Democracy*. April 27. <https://www.opendemocracy.net/can-europe-make-it/mariagiulia-giuffre%C3%A9-cathryn-costello/crocodile-tears-tragedy-and-responsibility-i>.
- Dur, Umut. 2011. 'Dynamic School Choice Problem'. Mimeo. <http://www.gtcenter.org/Archive/2011/Conf/Dur1312.pdf>.
- Echenique, Federico, and M. Bumin Yenmez. 2015. 'How to Control Controlled School Choice'. *American Economic Review* 105 (8).
- Ehlers, Lars, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim. 2014. 'School Choice with Controlled Choice Constraints: Hard Bounds versus Soft Bounds'. *Journal of Economic Theory* 153: 648–83.
- Erlanger, Steven, and Alison Smale. 2015. 'Europe's Halting Response to Migrant Crisis Draws Criticism as Toll Mounts'. *The New York Times*, August 28.
- Fernández-Huertas Moraga, Jesús Fernández-Huertas, and Hillel Rapoport. 2015. 'Tradable Refugee-Admission Quotas and EU Asylum Policy'. *CEifo Economic Studies* 61 (3-4): 638–72. doi:10.1093/cesifo/ifu037.

- Fernández-Huertas Moraga, Jesús, and Hillel Rapoport. 2014. 'Tradable Immigration Quotas'. *Journal of Public Economics* 115 (July): 94–108. doi:10.1016/j.jpubeco.2014.04.002.
- Fernández-Huertas Moraga, and Hillel Rapoport. forthcoming. 'Tradable Refugee-Admission Quotas (TRAQs), the Syrian Crisis and the New European Agenda on Migration'. *IZA Journal for European Labor Studies*
- Gale, David, and Lloyd S. Shapley. 1962. 'College Admissions and the Stability of Marriage'. *American Mathematical Monthly*, 9–15.
- Gibney, Matthew. 2010. 'The Securitization of Humanitarian Migration: Digging Moats and Sinking Boats. By Scott D. Watson.' *Journal of Refugee Studies* 23 (3): 404–6.
- Gibney, Matthew, and Randall Hansen. 2005. 'Asylum Policy in the West: Past Trends, Future Possibilities'. In *Poverty, International Migration and Asylum*, edited by George J. Borjas and Jeff Crisp. Palgrave Macmillan.
- Goldstein, Judith. 2001. *Legalization and World Politics*. MIT Press.
- Hafalir, Isa E., M. Bumin Yenmez, and Muhammed A. Yildirim. 2013. 'Effective Affirmative Action in School Choice'. *Theoretical Economics* 8 (2): 325–63.
- Harding, Luke. 2015. 'Refugee Crisis: Germany Reinstates Controls at Austrian Border'. *The Guardian*, September 13, sec. World news.
http://www.theguardian.com/world/2015/sep/13/germany-to-close-borders-exit-schengen-emergency-measures?CMP=Share_iOSApp_Other.
- Hatfield, John William, and Paul R. Milgrom. 2005. 'Matching with Contracts'. *American Economic Review*, 913–35.
- Hathaway, James C., and R. Alexander Neve. 1997. 'Making International Refugee Law Relevant Again: A Proposal for Collectivized and Solution-Oriented Protection'. *Harv. Hum. Rts. J.* 10: 115.
- Hurrell, Andrew. 2007. *On Global Order : Power, Values, and the Constitution of International Society: Power, Values, and the Constitution of International Society*. OUP Oxford.
- Huysmans, Jef. 2006. *The Politics of Insecurity: Fear, Migration and Asylum in the EU*. Routledge.
- Immorlica, Nicole, and Mohammad Mahdian. 2005. 'Marriage, Honesty, and Stability'. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 53–62. SODA '05. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
<http://dl.acm.org/citation.cfm?id=1070432.1070441>.
- Kadam, Sangram V., and Maciej H. Kotowski. 2015. 'Multi-Period Matching'. Mimeo.
<http://www.hks.harvard.edu/fs/mkotows/1305-multiperiod-wpWEB.pdf>.
- Kamada, Yuichiro, and Fuhito Kojima. 2012. 'Stability and Strategy-Proofness for Matching with Constraints: A Problem in the Japanese Medical Match and Its Solution'. *The American Economic Review*, 366–70.
- . 2015. 'Efficient Matching under Distributional Constraints: Theory and Applications'. *American Economic Review* 105 (1): 67–99.
- Keohane, Robert O. 1984. *After Hegemony: Power and Discord in the World Political Economy*. Princeton NJ: Princeton University Press.
- Kesten, Onur. 2006. 'On Two Competing Mechanisms for Priority-Based Allocation Problems'. *Journal of Economic Theory* 127 (1): 155–71. doi:10.1016/j.jet.2004.11.001.
- . 2010. 'School Choice with Consent'. *The Quarterly Journal of Economics* 125 (3): 1297–1348.
- Kojima, Fuhito, and Parag A. Pathak. 2009. 'Incentives and Stability in Large Two-Sided Matching Markets'. *The American Economic Review* 99 (3): 608–27.
- Kojima, Fuhito, Parag A. Pathak, and Alvin E. Roth. 2013. 'Matching with Couples: Stability and Incentives in Large Markets*'. *The Quarterly Journal of Economics* 128 (4): 1585–1632. doi:10.1093/qje/qjt019.
- Kominers, Scott Duke, and Tayfun Sönmez. Forthcoming. 'Designing for Diversity in Matching'. *Theoretical Economics*.
http://scottkom.com/articles/Kominers_Sonmez_Designing_for_Diversity_in_Matching.pdf.

- Koremenos, Barbara, Charles Lipson, and Duncan Snidal. 2003. *The Rational Design of International Institutions*. Cambridge University Press.
- Koser, Khalid. 2015a. *Responding to Migration from Complex Humanitarian Emergencies: Lessons Learned from Libya*. Accessed September 23. <http://lastradainternational.org/lsidocs/Migration%20lessons%20learned%20from%20Libya.pdf>.
- . 2015b. 'Time to Reform the International Refugee Regime'. *OUPblog*. Accessed September 21. <http://blog.oup.com/2015/09/international-refugee-regime-reform/>.
- Loescher, Gil, and John A. Scanlan. 1998. *Calculated Kindness*. Simon and Schuster.
- McDermott, John, and Jeevan Vasagar. 2015. 'Britain to Take 20,000 Syrian Refugees — but Not from Europe'. *Financial Times*, September 7. <http://www.ft.com/cms/s/0/d709bb32-5581-11e5-9846-de406ccb37f2.html#axzz3mBzPhmK5>.
- 'Migrant Crisis: EU Plan Offers More Money for Turkey Camps'. 2015. *BBC News*. Accessed October 9. <http://www.bbc.co.uk/news/world-europe-34451660>.
- Milner, Helen. 1992. 'International Theories of Cooperation Among Nations: Strengths and Weaknesses'. *World Politics* 44 (03): 466–96. doi:10.2307/2010546.
- Pereyra, Juan Sebastián. 2013. 'A Dynamic School Choice Model'. *Games and Economic Behavior* 80: 100–114.
- Robinson, Court. 1998. *Terms of Refuge: The Indochinese Exodus & the International Response*. Zed Books.
- Roth, Alvin E. 1982. 'The Economics of Matching: Stability and Incentives'. *Mathematics of Operations Research* 7 (4): 617–28.
- . 1984. 'The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory'. *The Journal of Political Economy*, 991–1016.
- . 1986. 'On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets'. *Econometrica* 54 (2): 425–27.
- . 1991. 'A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom'. *The American Economic Review*, 415–40.
- . 2007. 'Repugnance as a Constraint on Markets'. *The Journal of Economic Perspectives* 21 (3): 2.
- . 2015. *Who Gets What - and Why: The Hidden World of Matchmaking and Market Design*. Harper Collins Publ. UK.
- Roth, Alvin E., and Elliott Peranson. 1999. 'The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic D'. *American Economic Review* 89 (4): 748–82.
- Roth, Alvin E., Tayfun Sonmez, and M. Utku Unver. 2004. 'Kidney Exchange'. *Quarterly Journal of Economics* 119 (2): 457–88.
- Roth, Benjamin N., and Ran I. Shorrer. 2015. 'A Model of Mechanism Design in the Presence of a Pre-Existing Game'. Mimeo. http://scholar.harvard.edu/files/ran/files/mediators_0.pdf?m=1426648159.
- Schuck, Peter H. 1997. 'Refugee Burden-Sharing: A Modest Proposal'. *Yale J. Int'l L.* 22: 243.
- Sönmez, Tayfun, and Tobias B. Switzer. 2013. 'Matching With (Branch-of-Choice) Contracts at the United States Military Academy'. *Econometrica* 81 (2): 451–88.
- Thakral, Neil. 2015. 'Matching with Stochastic Arrival'. Mimeo. https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=ESWC2015&paper_id=1812.
- UNHCR. 2007. 'Many Refugees Misunderstand Schengen Expansion'. <http://www.unhcr.org/476bb7e24.html>.
- . 2008. 'UNHCR Note on DNA Testing to Establish Family Relationships in the Refugee Context'. <http://www.refworld.org/docid/48620c2d2.html>.

- . 2015. 'Syrian Refugees Inter-Agency Regional Update'.
<http://data.unhcr.org/syrianrefugees/regional.php>.
- Van Selm, Joanne. 2014. 'Refugee Resettlement'. In *The Oxford Handbook of Refugee and Forced Migration Studies*, edited by Elena Fiddian-Qasmiyeh, Gil Loescher, Katy Long, and Nando Sigona. Oxford University Press.
<http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199652433.001.0001/oxfordhb-9780199652433-e-014>.
- Wasik, Zosia, and Henry Foy. 2015. 'Poland Favours Christian Refugees from Syria'. *Financial Times*, August 21. <http://www.ft.com/cms/s/0/6edfdd30-472a-11e5-b3b2-1672f710807b.html#axzz3mBzPhmK5>.
- Weiner, Myron. 1995. *The Global Migration Crisis: Challenge to States and to Human Rights*. HarperCollins College Publishers.
- Westkamp, Alexander. 2013. 'An Analysis of the German University Admissions System'. *Economic Theory* 53 (3): 561–89.