



The International Refugee Match: A System that Respects Refugees' Preferences and the Priorities of States

Choice of Asylum Country

Will Jones* and Alexander Teytelboym**

ABSTRACT

There is an urgent need to overcome the political deadlock preventing States from 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. Matching systems can operate independently or alongside other burden-sharing mechanisms, such as tradable refugee quotas (as suggested by Fernández-Huertas Moraga and Rapoport). 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

1. 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,¹ 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 the so-called Islamic State, the European Union (EU) response has been inconsistent, poorly coordinated, and (to date) inadequate to the scale of the problem. In this context, there have been repeated calls for the revision or outright abolition of the Dublin

* Will Jones is Lecturer in International Relations at Royal Holloway, University of London, UK. Email: will.jones@rhul.ac.uk

** Alexander Teytelboym is the Otto Poon Research Fellow at the Institute for New Economic Thinking, Oxford, UK. Email: alexander.teytelboym@inet.ox.ac.uk

¹ C. Alfred, “What History Can Teach Us About the Worst Refugee Crisis Since WWII”, *The Huffington Post*, 12 Sep. 2015, available at: http://www.huffingtonpost.com/entry/alexander-betts-refugees-wwii_55f30f7ce4b077ca094edaec (last visited 14 Mar. 2017).

III Regulation,² particularly its requirement that the first country in the EU that asylum-seekers enter be responsible for processing their claim,³ and it looks likely to be scrapped.⁴ As Costanza 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 governments to implement ever more restrictive asylum and immigration policies.⁵

At the time of writing, there are renewed calls to make the largely aspirational Common European Asylum System (CEAS) a reality.⁶ As Alexander Betts recently noted in the *New York Times*, “Europe needs a comprehensive global refugee policy”.⁷ 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 article, we consider how insights from over 30 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 in different real-world contexts. Secondly, we explain in general terms how matching 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

- 2 The Dublin III Regulation is the EU law instrument “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.”: Regulation (EU) No 604/2013 of the European Parliament and of the Council of 26 June 2013 establishing the criteria and mechanisms for determining the Member State responsible for examining an application for international protection lodged in one of the Member States by a third-country national or a stateless person (recast) [2013] OJ L 180/31.
- 3 C. Costello & M.-G. Guiffé, “‘Tragedy’ and Responsibility in the Mediterranean”, *Open Democracy*, 27 Apr. 27, 2015, available at: <https://www.opendemocracy.net/can-europe-make-it/mariagiulia-giuffe%C3%A9-cathryn-costello/crocodile-tears-tragedy-and-responsibility-i> (last visited 14 Mar. 2017); K. Koser, “Time to Reform the International Refugee Regime”, *OUPblog*, 21 Sep. 2015, available at: <http://blog.oup.com/2015/09/international-refugee-regime-reform/> (last visited 14 Mar. 2017).
- 4 G. Parker & D. Robinson, “Cameron Opposes Change to EU ‘Dublin’ Rule on Hosting of Refugees”, *Financial Times*, 20 Jan. 2016, available at: <https://www.ft.com/content/5e103cb4-bf8b-11e5-846f-79b0e3d20eaf> (last visited 14 Mar. 2017).
- 5 B. S. Chimni, “The Global Refugee Problem in the 21st Century and the Emerging Security Paradigm: A Disturbing Trend”, in A. Anghie & G. Sturgess (eds.), *Legal Visions of the 21st Century: Essays in Honour of Judge Christopher Weeramantry*, The Hague, Kluwer, 1998; J. Huysmans, *The Politics of Insecurity: Fear, Migration and Asylum in the EU*, London/New York, Routledge, 2006.
- 6 V. Chetail, P. De Bruycker, & F. Maiani (eds.), *Reforming the Common Asylum System: The New European Refugee Law*, Leiden/Boston, Brill/Nijhoff, 2016; V. Chetail, “Looking Beyond the Rhetoric of the Refugee Crisis: The Failed Reform of the Common European Asylum System”, *European Journal of Human Rights*, 5, 2016, 584–601; M. den Heijer, J. Rijpma & T. Spijkerboer, “Coercion, Prohibition, and Great Expectations: The Continuing Failure of the Common European Asylum System”, *Common Market Law Review*, 53(3), 2016, 607–42; C. Bauloz et al. (eds.), *Seeking Asylum in the European Union: Selected Protection Issues Raised by the Second Phase of the Common European Asylum System*, Leiden, Brill/Nijhoff, 2015; S. Velluti, *Reforming the Common European Asylum System – Legislative Developments and Judicial Activism of the European Courts*, New York, Springer, 2014; S. Peers, “The Second Phase of the Common European Asylum System – A Brave New World or Lipstick on a Pig?”, *Statewatch Analysis*, 2013.
- 7 S. Erlanger & A. Smale, “Europe’s Halting Response to Migrant Crisis Draws Criticism as Toll Mounts”, *The New York Times*, 28 Aug. 2015, available at: <https://www.nytimes.com/2015/08/29/world/europe/europe-migrant-refugee-crisis.html> (last visited 14 Mar. 2017).

current European refugee crisis. Finally, we outline and respond to some potential objections.

There are effectively two major problems we believe matching can address. First, the processes whereby refugees do or do not get protected by particular countries are often arbitrary, unfair, unsafe, and ineffective. Some countries, such as Sweden and Germany, are considerably more assiduous than others in welcoming refugees, but which country a refugee ends up in is largely left to chance. 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 that are wracked with arbitrariness, “calculated kindness”,⁸ and bad faith. Secondly, 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 for 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 the 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.

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 a group of States to coordinate minimally with each other insofar as they would have to opt into the Refugee Match, 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. Obviously, it would be tremendously beneficial if every State participated in resettlement with such generosity that any refugee could have a literally unconstrained choice as to where they go. But such a world bears no resemblance to reality. However, somewhat less ambitiously and more realistically, a subset of countries could undertake to resettle a given population of refugees, as some States do now, and have historically. This could also be used in the instance of some other category of protection seekers, such as that outlined in the 2001 Temporary Protection

8 G. Loescher & J.A. Scanlan, *Calculated Kindness*, New York, Simon and Schuster, 1998.

Directive.⁹ Participating States and refugees would 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 that 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 changing hands between States or between States and refugees (similar to how a parent's choice of their children's primary school operates today), 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. Such a system could operate either in situations where a mass influx must be assisted rapidly or with systems that operate as a matter of routine (e.g. global resettlement). This system would only allocate the refugees to countries that are participating in the Refugee Match and therefore would have no effect on the right of refugees to apply for asylum by the conventional route of making it to a given State's territory, much as resettlement currently does not.

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.¹⁰ It is important to recognise how different our proposal is. The Refugee Match merely sets out a way to find the optimal allocation of a given population of refugees amongst a given population of States, given their preferences and given States' quotas. No buying or selling of quotas 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,¹¹ although in principle 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,¹² or the system of humanitarian visas bestowed at consular outposts outside the EU,¹³ recently advocated by Alexander Betts.¹⁴ We are agnostic as to how burdens should be optimally shared, but we do argue that any burden-sharing agreement that

9 Council Directive 2001/51/EC of 20 July 2001 on minimum standards for giving temporary protection in the event of a mass influx of displaced persons and on measures promoting a balance of efforts between Member States in receiving such persons and bearing the consequences thereof [2001] OJ L 212/12.

10 P. H. Schuck, "Refugee Burden-Sharing: A Modest Proposal", *Yale Journal of International Law*, 22, 1997, 243.

11 D. Anker, J. Fitzpatrick & A. Shacknove, "Crisis and Cure: A Reply to Hathaway/Neve and Schuck", *Harvard Human Rights Journal*, 11, 1998, 295.

12 J.C. Hathaway & R.A. Neve, "Making International Refugee Law Relevant Again: A Proposal for Collectivized and Solution-Oriented Protection", *Harvard Human Rights Journal* 10, 1997, 115.

13 U.I. Jensen, "Humanitarian Visas: Option or Obligation?", Study for the LIBE Committee, Directorate-General for Internal Policies, European Parliament, 2014.

14 A. Betts, "Let Refugees Fly to Europe", *New York Times*, 24 Sep. 2015, available at: <http://www.nytimes.com/2015/09/25/opinion/let-refugees-fly-to-europe.html?nytmobile=0> (last visited 14 Mar. 2017).

then uses a matching system is better for refugees and more likely to be assented to by States.

The idea of quota-trading has recently been revived by economists,¹⁵ endorsed by some political theorists,¹⁶ and criticised by others.¹⁷ Fernández-Huertas Moraga and Rapoport are the first to propose and mathematically develop how to apply matching models to the allocation of refugees across destinations, in combination with a quota-trading scheme in order to overcome the incentives of States to make themselves undesirable resettlement destinations. The main applications they consider are refugee resettlement in general and resettlement/relocation within the EU. In their papers, they use matching mechanisms to take into account the preferences of refugees over destinations, and possibly the priorities of destination countries over refugee types, as part of a broader allocation scheme where countries receive quotas that can be traded with others.

We would agree with Fernández-Huertas Moraga and Rapoport that the only way of making a tradeable quota system palatable is if it is implemented with a matching system, but in this article we seek to explore the value and the design details of a matching system in itself, with or without tradeable quotas. In their work, they work through the operation of matching mathematically for an audience of economists. The present article is not purporting to add to the stock of knowledge in mathematical economics. Rather, it develops the concept framed for an audience in refugee studies, adds considerable institution detail, considers fine-grained operational concerns, and evaluates the proposal with reference to common worries and normative frameworks within the field. Nevertheless, we expect that many insights in this article would be useful for building better mathematical models of refugee resettlement schemes.

In practice, the absence of a formalised system that 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.¹⁸ 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.

In the rest of this article, we focus explicitly on the design of the Refugee Match under the institutional constraints of today. For example, we take States' interests and limited willingness to host refugees as given. We assume only minimal changes

15 J. Fernández-Huertas Moraga & H. Rapoport, "Tradable Refugee-Admission Quotas and EU Asylum Policy", *CESifo Economic Studies*, 61(3-4), 2015, 638-672; J. Fernández-Huertas Moraga & H. Rapoport, "Tradable Immigration Quotas", *Journal of Public Economics*, 115, 2014, 94-108; J. Fernández-Huertas Moraga & H. Rapoport, "Tradable Refugee-Admission Quotas (TRAQs), the Syrian Crisis and the New European Agenda on Migration", *IZA Journal of European Labor Studies*, 4(1), 2015, 23.

16 J. Kuosmanen, "What (If Anything) Is Wrong with Trading Refugee Quotas?", *Res-Publica*, 19(2), 25 Sep. 2012, 103-19.

17 M.J. Gibney, "Refugees and Justice between States", *European Journal of Political Theory*, 14(4), 2015, 448-63.

18 See, for example, the EU-Turkey joint statement of 18 March 2016, available at: <http://www.consilium.europa.eu/en/press/press-releases/2016/03/18-eu-turkey-statement/>. (last visited 14 Mar. 2017).

οικονομική επιβάρυνση
των Τουρκίας

to international refugee law are possible. We will argue that the Refugee Match can play an important complementary role within current international refugee assistance frameworks.

2. MATCHING THEORY

Two-sided matching theory is a mathematical framework for allocating 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.¹⁹ 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 Shapley²⁰ showed that any such one-to-one two-sided matching market (such as renting apartments or getting taxi rides) has a stable outcome and described an algorithm (which can be implemented quickly on a computer) that finds it. “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.

One well studied and carefully designed “many-to-one” matching system is the National Resident Matching Program (NRMP) in the US, which allocates junior medical interns to hospitals.²¹ 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.²² First, the 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.²³ Alexander Betts has argued, however, that there are legitimate uses of the language of efficiency in refugee studies.²⁴ When we refer to efficiency in this article, we simply mean Pareto efficiency (or “allocative efficiency”, in Betts’ language).²⁵

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

20 D. Gale & L. S. Shapley, “College Admissions and the Stability of Marriage”, *American Mathematical Monthly*, 69(1), 1962, 9–15.

21 A. E. Roth, “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory”, *The Journal of Political Economy*, 92(6), 1984, 991–1016.

22 A. E. Roth, “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*, 81(3), 1991, 415–440.

23 M. Gibney & R. Hansen, “Asylum Policy in the West: Past Trends, Future Possibilities”, in G. J. Borjas & J. Crisp (eds.), *Poverty, International Migration and Asylum*, Basingstoke, Palgrave Macmillan, 2005.

24 A. Betts, “What Does ‘Efficiency’ Mean in the Context of the Global Refugee Regime?”, *The British Journal of Politics & International Relations*, 8(2), 2006, 148–73.

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

Secondly, 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, Roth²⁶ 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.²⁷ 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 a key to the designs of many matching systems. Roth²⁸ examined medical job application procedures in the UK. At the time, the UK operated with many different systems across different parts of the country, some of which used methods that 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).

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.²⁹ 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 matching of cadets to branches of the US Army³⁰ and the admission of students to universities in Germany.³¹

26 A. E. Roth, “On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets”, *Econometrica*, 54(2), 1986, 425–27.

27 Hospitals 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.

28 Roth, “A Natural Experiment” above n 22.

29 A. Abdulkadiroğlu & T. Sönmez, “School Choice: A Mechanism Design Approach”, *The American Economic Review*, 93(3), 2003, 729–747; A. Abdulkadiroğlu, P. A. Pathak & A. E. Roth, “The New York City High School Match”, *American Economic Review*, 95(2), 2005, 364–367.

30 T. Sönmez & T. B. Switzer, “Matching With (Branch-of-Choice) Contracts at the United States Military Academy”, *Econometrica*, 81(2), 2013, 451–488.

31 A. Westkamp, “An Analysis of the German University Admissions System”, *Economic Theory*, 53(3), 2013, 561–589.

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 (as when students look for house-mates). Famous applications of one-sided matching systems include pairwise exchange of kidneys³² and the allocation of public housing.³³ Most importantly, most applications of matching systems, including the one we will describe in this article, *do not involve the exchange of money*. Many systems, such as a market for organs, may be justifiably considered (what Roth³⁴ calls) “repugnant” if money were introduced in transactions.³⁵

In all of 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 being mismatched to hospitals, and patients being unable to find the kidney that could save their life.³⁶

3. A MANY-TO-ONE MATCH FOR REFUGEES: WHY

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,

32 A.E. Roth, T Sönmez & M. U. Unver, “Kidney Exchange”, *Quarterly Journal of Economics*, 119(2), 2004, 457–488.

33 A. Abdulkadiroğlu & T. Sönmez, “House Allocation with Existing Tenants”, *Journal of Economic Theory*, 88(2), 1999, 233–260.

34 A. E. Roth, “Repugnance as a Constraint on Markets”, *The Journal of Economic Perspectives*, 21(3), 2007, 2.

35 A magisterial summary of the successes of matching systems and market design is A.E. Roth, *Who Gets What - and Why: The Hidden World of Matchmaking and Market Design*, London, Harper Collins, 2015.

36 One example of a system that has repeatedly failed is the matching of clerks to judges in the US judiciary. C. Avery et al., “The New Market for Federal Judicial Law Clerks”, *The University of Chicago Law Review*, 2007, 447–486. 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 and thereby they often end up hiring clerks they do not want and clerks end up working for judges they do not particularly want to work for.

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.³⁷ 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.³⁸ 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 are currently 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 article are cases where States are willing to sacrifice this power. In the case of refugee resettlement,³⁹ refugee status determination is already not always conducted by the States that agree to resettle refugees (for example, the US Government resettled 60,000 Bhutanese refugees from Nepal in 2006). A coalition of Western States

37 R. O. Keohane, *After Hegemony: Power and Discord in the World Political Economy*, Princeton, Princeton University Press, 1984; H. Milner, “International Theories of Cooperation Among Nations: Strengths and Weaknesses”, *World Politics*, 44(3), 1992, 466–496; J. Goldstein, *Legalization and World Politics*, Cambridge, MIT Press, 2001; B. Koremenos et al., *The Rational Design of International Institutions*, Cambridge, Cambridge University Press, 2003; A. Hurrell, *On Global Order: Power, Values, and the Constitution of International Society: Power, Values, and the Constitution of International Society*, Oxford, Oxford University Press, 2007; K. W. Abbott et al., *International Organizations as Orchestrators*, Cambridge, Cambridge University Press, 2015.

38 A. Betts (ed.), *Global Migration Governance*, Oxford, Oxford University Press, 2011.

39 Resettlement here refers to “the organized movement of pre-selected refugees to a destination country in which their settlement is expected to be permanent”: J. Van Selm, “Refugee Resettlement”, in E. Fiddian-Qasbiyeh et al. (ed.), *The Oxford Handbook of Refugee and Forced Migration Studies*, Oxford, Oxford University Press, 2014.

successfully resettled two and a half million Indochinese refugees between 1975 and 1997.⁴⁰ 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 EU to harmonise who qualifies for refugee status, the Qualification Directive, which is binding on all members of the EU except Denmark.⁴¹

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 the longest, or face an 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 UK 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.⁴² On the other hand, the Governments of Slovakia, Poland, and the Czech Republic

40 C. Robinson, *Terms of Refuge: The Indochinese Exodus & the International Response*, London, Zed Books, 1998.

41 See Directive 2011/95/EU of the European Parliament and of the Council of 13 December 2011 on standards for the qualification of third-country nationals or stateless persons as beneficiaries of international protection, for a uniform status for refugees or for persons eligible for subsidiary protection, and for the content of the protection granted (recast) [2011] OJ L 337/9; and Council Directive 2004/83/EC of 29 April 2004 on minimum standards for the qualification and status of third country nationals or stateless persons as refugees or as persons who otherwise need international protection and the content of the protection granted [2004] OJ L 304/12 binding upon the UK and Ireland.

42 J. McDermott & J. Vasagar, "Britain to Take 20,000 Syrian Refugees — but Not from Europe", *Financial Times*, 7 Sep. 2015, available at: <http://www.ft.com/cms/s/0/d709bb32-5581-11e5-9846-de406ccb37f2.html#axzz3mBzPbmKS> (last visited 14 Mar. 2017).

have all signalled a willingness to take more refugees, but only if they are Christian.⁴³ 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).⁴⁴

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 that are compatible with the principles and goals of the 1951 Convention relating to the Status of Refugees, or any other set of principles (e.g. international human rights law). 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 that 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.

4. 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. Matching is not about *who* is to be resettled: it is a system that can be used *after* a commitment has been made by a given set of States to protect a given set of refugees. For example, EU Member States have pledged to relocate 160,000 refugees from Greece and Italy in order to alleviate the chronic strain on those nations' resources.⁴⁵ Matching is not a tool for selecting which refugees should be included in that number, it is a tool for deciding who goes where after that decision has been made.

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

43 Z. Wasik & H. Foy, "Poland Favours Christian Refugees from Syria", *Financial Times*, 21 Aug. 2015, available at: <http://www.ft.com/cms/s/0/6edfd30-472a-11e5-b3b2-1672f710807b.html#axzz3mBzPhmKS> (last visited 14 Mar. 2017).

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

45 J. Kanter, "EU Nations Urged to Accept 160,00 Migrants", *New York Times*, 9 Sep. 2015, available at: <https://www.nytimes.com/2015/09/10/world/europe/europe-migrant-crisis-jean-claude-juncker.html?r=0> (last visited 14 Mar. 2017).

unnecessary. In the European context, there has been much progress towards achieving “joint processing” of applications according to a standardised set of procedures.⁴⁶ In order to create a system that ensures that families who 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,⁴⁷ and States would then evaluate those households in the same way.⁴⁸ 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.⁴⁹ Theoretically, these sorts of preferences can create problems for the matching mechanism, but, in practice,⁵⁰ 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 a 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. There is a trade-off here. Less frequent clearing “thickens” the system, removing incentives to cheat the system, and offering more choice (if there is a larger, and therefore more diverse set of participants on both sides, it is more likely particularly fitting matches will be made). However, it increases the amount refugees spend in potentially vulnerable situations.

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,

46 H. Urth et al., “Study on the Feasibility and Legal and Practical Implications of Establishing a Mechanism for the Joint Processing of Asylum Applications on the Territory of the EU”, European Commission, 2013.

47 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.

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

49 A. E. Roth & E. Peranson, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic D”, *American Economic Review*, 89(4), 1999, 748–782.

50 And, as it turns out, even in theory for large markets, see F. Kojima, P. A. Pathak & A. E. Roth, “Matching with Couples: Stability and Incentives in Large Markets”, *The Quarterly Journal of Economics*, 128(4), 2013, 1585–1632.

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 that 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 problem and sometimes leads to individuals making choices to enter into higher priority categories by misrepresenting themselves (as when individuals move houses in order 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.⁵¹ 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 the 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

⁵¹ UNHCR, *UNHCR Note on DNA Testing to Establish Family Relationships in the Refugee Context*, Geneva, UNHCR, 2008; available at: <http://www.refworld.org/docid/48620c2d2.html> (last visited 14 Mar. 2017).

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 of 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.⁵² 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 is 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 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.⁵³ 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 developed a matching system for doctors in Japan that ensures that rural hospitals have a minimum complement of medical staff.⁵⁴

Economists have also very recently developed tools to handle “distributional constraints”.⁵⁵ One setting where this is often used is college admissions systems where affirmative action is a requirement.⁵⁶ 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

52. 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.

53. L. Ehlers, I. E. Hafalir, M. B. Yenmez & M. A. Yildirim, “School Choice with Controlled Choice Constraints: Hard Bounds versus Soft Bounds”, *Journal of Economic Theory*, 153, 2014, 648–683.

54. Y. Kamada & F. Kojima, “Stability and Strategy-Proofness for Matching with Constraints: A Problem in the Japanese Medical Match and Its Solution”, *The American Economic Review*, 102(3), 2012, 366–370.

55. Y. Kamada & F. Kojima, “Efficient Matching under Distributional Constraints: Theory and Applications”, *American Economic Review*, 105(1), 2015, 67–99.

56. I. E. Hafalir, M. B. Yenmez & M. A. Yildirim, “Effective Affirmative Action in School Choice”, *Theoretical Economics*, 8(2), 2013, 325–363; S. D. Kominers & T. Sönmez, “Designing for Diversity in Matching”, *Theoretical Economics*, forthcoming, available at: http://scottkom.com/articles/Kominers_Sonmez_Designing_for_Diversity_in_Matching.pdf (last visited 14 Mar. 2017).

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.

Another factor that distinguishes this system from most applications of matching models is that it might be reasonably expected that some States would prefer *not* to be chosen. In such circumstances, countries could deliberately try and make themselves as unwelcoming to refugees as possible, in order to encourage refugees to try to avoid them. In the simplest version of the system, States could pretend to be generous by offering places, then act to guarantee that those places would be unfilled, and a State could thereby shirk their responsibilities to resettle with no penalty. This could result in some States taking disproportionate burdens, leading to collective action breaking down, and the collapse of the system.

This problem can be addressed in a variety of ways. Firstly, minimum quotas offer a partial mitigation of this problem (particularly if the minimum quotas are also maximum quotas that collectively sum to the total number of refugees to be resettled) as such a system would guarantee that some refugees were resettled to the least desired countries. If those countries know that they will be hosting refugees no matter how many others host, their incentives flip: a State that knows it will be welcoming some refugees only has incentives to be as desirable as possible to the refugees it most wishes to host. In fact, the State would have an incentive to invest in welcoming refugees no matter how much they may dislike the idea of hosting refugees in general. This solution gives rise to three key concerns. Firstly, this may be seen as ethically unacceptable, as it would involve forcing refugees to relocate to countries that they had not expressed a preference for. Secondly, it might be thought that this therefore undermines the benefits that accrue from finding apt matches. Thirdly, this may increase the likelihood that refugees will withdraw from one round of matching in the hope of getting one of their preferred matches in a subsequent round, further unravelling the system.

A second solution is proposed by Fernández-Moraga and Rapoport.⁵⁷ In their work, they envisage a centrally imposed sanction mechanism that fines States that end up with unmatched places after refugees are allocated. Further, they point out that a tradeable quota system is a logical way to determine the level (or “price”) of that sanction. Such sanctions would be the stick that would incentivise States to become more attractive, rather than less, to refugees. The key barrier such a solution faces is that it requires a central authority capable of enforcing financial sanctions on States for non-compliance. In contrast to everything else we discuss here, which requires minimal infractions upon sovereignty, this would necessitate that States sacrifice considerable sovereignty to such a body. In the context of the EU, these powers already exist in theory, although historically their implementation has been patchy, to say the least. There is a long and dishonourable history of States flouting the demands of Brussels when they are seen as too onerous.

57 Fernández-Huertas Moraga & Rapoport, “Tradable Immigration Quotas”, above n 15; Fernández-Huertas Moraga & Rapoport, “Tradable Refugee-Admission Quotas and EU Asylum Policy”, above n 15; Fernández-Huertas Moraga & Rapoport, “Tradable Refugee-Admission Quotas (TRAQs)”, above n 15.

Even if this system could be enforced, there are potential consequences which many would regard as undesirable. In practical terms under the proposed initial quota allocation such a system implemented in the present context could lead to some of the poorest countries in the EU paying some of the richest to take their quotas. In one paper, Fernández-Moraga and Rapoport perform a number of simulations to illustrate the possible workings of their proposed system.⁵⁸ For example, in their proposal, Hungary receives an initial quota of 1,100, which it then pays to reduce to 1. In contrast, the UK, Ireland, and Denmark all begin on initial quotas of 0 (due to their ability to opt out of European Council decisions taken under Justice and Home Affairs competence) and are paid to raise their admission quotas. Hence, the distributional consequences of such a scheme are in practice dependent on the initial allocation of quotas. Given how unpopular the imposition of such quotas is, trading in practice only partially ameliorates this concern.

The EU quota proposal (used as a benchmark initial allocation by Fernández-Huertas Moraga-Rapoport) also gives Greece and Italy an initial allocation of zero since the EU relocation proposal only deals with a small fraction of the refugees, so Greece and Italy would continue to have to host a significant number of current and arriving refugees. Such a system would result in considerable financial reimbursement for such States, if they agree to host more than their initial quota, but if the constraints they face are not easily remedied by finance (e.g. physical limitations on infrastructure and space), or they are otherwise limited from using that money (if, let's just imagine, strictures imposed as part of a bailout prevented them from spending money), this could be of limited use. In Fernández-Huertas Moraga and Rapoport's terms, if there is limited fungibility between "physical solidarity" (e.g. resettling refugees) and "financial solidarity" (e.g. paying someone else to do it), then a system of fines and cash transfers will not provide a solution.

Our preferred solution is to require refugees to express preferences over the entire set of countries within the system, combined with minimum quotas that guarantee that each participating country will receive some refugees (and add up to the total number of refugees to be resettled). Such a system would guarantee that every refugee is matched, no spaces are left unfilled, and a State has no incentive to be unattractive. It requires no additional enforcement capacity (of course, States may refuse to participate in a system, but that is much less damaging than States participating then trying to become as unattractive as possible). Such a system faces two obstacles. Firstly, some refugees may still get matched to countries that are extremely low down their list of preferences, even their last, creating strong incentives to move across borders. Secondly, it might be thought that refugees lack the information to express meaningful preferences over such a large set of countries. The first problem can only be mitigated, not fully overcome. It is less of a problem to the extent that both refugees and States have heterogeneous preferences. For example, even if a particular country is undesirable for most, that is fine if it is desirable for even a small subset of refugees (say, those with the skills to fill their particular labour shortages). Indeed, it might be hoped that over the long term, such a system would give incentives to both refugees seeking resettlement and receiving States to engage in further

⁵⁸ Fernández-Huertas Moraga & Rapoport, "Tradable Refugee Admission Quotas (TRAQs)" above n 15.

such specialisation. More generally, there is some evidence from other contexts in which matching systems have been implemented that this is less of a problem than is commonly assumed. In the context of public school choice, the default assumption of many is that parents largely agree on what the good and bad schools are (i.e. they have almost perfectly correlated preferences). However, this concern turns out to be unfounded. In the context of the Boston school match, Dur et al. find that between 43 and 66 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 80 per cent of students get one of their first two choices and all students were allocated to their top five choice school.⁵⁹ Public school choice is different from refugee resettlement, not least because there are enough public school places in Boston for every child. But the point about incorrectly assuming homogeneity of preferences still applies in our case. Concerns about informational gaps on the part of refugees could be overcome either by providing better information to refugees prior to their expressing preferences (there is already a window between being selected for resettlement and being moved which could be used for this purpose), or through systems that “impute” preferences by asking refugees to express preferences over the *properties* of areas they prioritise in resettlement and using that information to construct a ranking of countries.⁶⁰

The penultimate decision that needs to be made is about precisely which matching mechanism to use. One interesting result from the economics literature⁶¹ 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.⁶² 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.⁶³

Finally, the clearinghouse has to be created and supervised. Creating the clearinghouse is a technical challenge analogous to the creation of other international market-based systems such as the EU Emissions Trading System. Once these decisions have been taken, and the clearinghouse created, the actual running of the clearinghouse can be left to technocrats.

59 Dur et al., “The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice”, *Working Paper*, 2013.

60 We discuss this possibility in considerably more depth in a forthcoming paper, available from the authors on request.

61 A. E. Roth, “The Economics of Matching: Stability and Incentives”, *Mathematics of Operations Research*, 7(4), 1982, 617–628.

62 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. See K. Onur, “School Choice with Consent”, *The Quarterly Journal of Economics*, 125(3), 2010, 1297–1348.

63 One particularly promising direction has been to consider systems with a large number of participants where these properties are only approximately satisfied (for example, E. M. Azevedo & E. B. Budish, “Strategy-Proofness in the Large”, *Chicago Booth Research Paper*, no. 13, 2013, available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2237425 (last visited 14 Mar. 2017). This paper shows how a weaker version of strategy-proofness can square this impossibility circle in large markets).

5. STABILITY, EFFICIENCY, AND FAIRNESS IN RESETTLEMENT

Let us put the concepts of stability, efficiency, and strategy-proofness in the context of refugee protection. Stability in our context 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). As in public school choice, States (collectively or individually) establish an order of priorities, rather than directly ranking refugees. In effect, States do not make active choices about specific individuals, instead pre-committing themselves to a particular way of ranking refugees.

Efficiency is straightforward. Since a country simply assigns priorities to refugees according to an international agreement, 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 also 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.

A system that 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.⁶⁴

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).⁶⁵ 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.

64 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 exist, 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. See N. Immorlica & M. Mahdian, “Marriage, Honesty, and Stability”, in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, Philadelphia, Society for Industrial and Applied Mathematics, 2005, 53–62, available at: <http://dl.acm.org/citation.cfm?id=1070432.1070441> (last visited 14 Mar. 2017); F. Kojima & P. A. Pathak, “Incentives and Stability in Large Two-Sided Matching Markets”, *The American Economic Review*, 99(3), 2009, 608–627. But there is no guarantee that there will be enough countries participating in any given matching system to ensure this.

65 F. Echenique & M. B. Yenmez, “How to Control Controlled School Choice”, *American Economic Review*, 105(8), 2015, 2679–2694.

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.⁶⁶ 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 days or work without a permit. Moreover, they need a visa to travel to non-Schengen States.⁶⁷ 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.

6. A GLOBAL RESETTLEMENT MECHANISM

In mid-2016, the United Nations High Commissioner for Refugees (UNHCR) estimated that 2017’s global resettlement needs will be over 1,190,000 persons.⁶⁸ However, only some 85,000 places are available on an annual basis.⁶⁹ 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. As many have argued, international emergencies and crises, such as the current Syrian situation, have unjustly distracted attention from the plight of those in “protracted refugee situations” (i.e. five years or more), where resettlement places are extremely limited.⁷⁰ In particular, Milner and Loescher recommend “more effective and strategic use of resettlement as a tool of burden and responsibility sharing, including through the application of a group resettlement referral methodology”.⁷¹ In this context, matching can make resettlement a more effective durable solution and increase the likelihood that the limited places at our disposal are used well.

66 Roth, “A Natural Experiment in the Organization of Entry-Level Labour Markets” above n 22.

67 UNHCR, “Many Refugees Misunderstand Schengen Expansion”, UNHCR News Stories, 21 Dec. 2007, available at: <http://www.unhcr.org/476bb7e24.html> (last visited 14 Mar. 2017).

68 UNHCR, *UNHCR Projected Global Resettlement Needs, 22nd Annual Tripartite Consultation on Resettlement, Geneva, 2016*, available at: <http://www.unhcr.org/protection/resettlement/575836267/unhcr-projected-global-resettlement-needs-2017.html> (last visited 14 Mar. 2017).

69 Van Selm, “Refugee Resettlement” above n 39.

70 J. Milner & G. Loescher, “Responding to Protracted Refugee Situations: Lessons from a Decade of Discussion”, *Forced Migration Policy Briefing*, Oxford, Refugee Studies Centre, 2011.

71 *Ibid.*

In some cases, resettlement functions according to criteria set by UNHCR and tries to encompass priority needs. When resettlement works well, UNHCR acquires much information about refugees through interviews and background checks. Generally speaking, a refugee will be referred to one resettlement country after UNHCR has inferred the refugee's preferences using interviews and so on. This is a system where UNHCR manually processes this information and ultimately makes the decision for the refugee. As such, matching can be distinguished in two ways. Firstly, even if UNHCR wished to continue with a system where they inferred the preferences of refugees, it would remain possible to automate the actual matching, which would have several advantages we list below. In identifying one country for resettlement, UNHCR is effectively inferring the refugee's "top" choice alone. If UNHCR could infer a preference list across participating countries, an automated matching system could still indicate a more desirable allocation within their limited resources. Secondly, this bespoke inferring of preferences through interviews is extremely time- and resource-consuming and could be substituted for simply asking the refugees themselves what countries, out of the participating States, they would most wish to resettle to. What does not happen currently is that *after* it has been determined that a refugee will be resettled, they *themselves* are given agency as to where they go.

In the worst cases, 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 arises 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.⁷² 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 internal procedures of bureaucracies. Vulnerable refugees could state their preferences as to which of the twenty-odd States operating resettlement schemes they most prefer. We think it likely that refugees themselves have better information about their lives, aspirations, and the quality of life they would have in a particular country than an overstretched bureaucracy trying to process this information manually.

Secondly, 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 or categories (e.g. women at risk, or lacking safety *in situ*). Whether a refugee population becomes resettled thus frequently has more to do with the competence and drive of particular civil servants than either a sober

72. K. Koser, *Responding to Migration from Complex Humanitarian Emergencies: Lessons Learned from Libya*, Chatham House, Briefing Paper, Nov. 2011, available at: <http://lastradainternational.org/lsidocs/Migration%20lessons%20learned%20from%20Libya.pdf> (last visited 14 Mar. 2017).

assessment of who the vulnerable are 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.

Thirdly, 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 the UNHCR, the International Organization for Migration, and state bureaucracies.

Fourthly, because matching enables States to effectively share burdens while 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.

7. 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.⁷³ 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.⁷⁴ 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, or basic legal protections and facing on-going 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.

Prior to the emergence of this crisis, the European Commission was already considering joint processing, relocation, long-term residence, protected entry procedures, and resettlement.⁷⁵ These reports amass a body of policy knowledge and ideas

73 UNHCR, "Syrian Refugees Inter-Agency Regional Update", webpage, 20 Sep. 2015, available at: <http://data.unhcr.org/syrianrefugees/regional.php> (last visited 14 Mar. 2017).

74 *Ibid.*

75 J. Van Selm et al., "Study on The Feasibility of Setting up Resettlement Schemes in EU Member States or at EU Level, against the Background of the Common European Asylum System and the Goal of a Common Asylum Procedure", European Commission, 2003; J. Van Selm, N. Lassen & J. Doomernik, "Study on The Transfer of Protection Status in the EU, against the Background of the Common European Asylum System and the Goal of a Uniform Status, Valid throughout the Union, for Those Granted Asylum", DG JAI/A2/2003/001, European Commission, 2004; European Commission,

as to how resettlement and relocation can be implemented. What they do not consider is the possibility of using matching to determine which refugees go to which participating States within the schemes they outline.

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 to replace the Dublin III Regulation. History suggests interstate cooperation to meet such a challenge is not impossible: from the Comprehensive Plan of Action 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 that 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 III, 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, 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 mass relocation of Kosovans across Europe in the late 1990s and early 2000s.⁷⁶ In this case, almost no attempt was made to match refugees and hosting locations. Refugees were sometimes unaware of which country the plane they boarded was taking them to. A matching system could have been used to allocate refugees within that system quickly to a country where they were most likely to flourish for the duration of their stay without affecting the rest of the 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,

Directorate-General Home Affairs, "Study on the Feasibility of Establishing a Mechanism for the Relocation of Beneficiaries of International Protection", European Commission, 2010.

⁷⁶ K. Kerber, "The Temporary Protection Directive", *European Journal of Migration and Law*, 4(2), 2002, 193-214.

over Karabagh
Don 31/1/00

necessitates that eligibility is centrally agreed upon. 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. We argue that using matching has further advantages over and above simply persuading States to participate.

Most simply, 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 that 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.⁷⁷

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 (or move across) 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 that 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.⁷⁸ 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

77. 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.

78. 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 & R. I. Shorrer, “A Model of Mechanism Design in the Presence of a Pre-Existing Game”, *Mimeo*, Feb. 2015, available at: http://scholar.harvard.edu/files/ran/files/media_tors_0.pdf?m=1426648159 (last visited 14 Mar. 2017)).

precisely what can persuade States to opt into refugee burden-sharing, surrender control of eligibility determination, and ultimately protect more refugees.

8. FURTHER CONSIDERATIONS

It might be thought that 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, that 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, 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 a precondition for protection, and only agreement from both sides gives refugees any choices at all. Currently, many activism works assume a strict dichotomy, where refugees can either have unconstrained choice or no choice at all. For example, some European refugee activists advocate a system that 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.⁷⁹ 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.

Where 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, as Dur et al. point out.⁸⁰ 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.

79 E.g. G. Granjon, “Freedom of Movement under European Conditions”, *Euractiv*, 28 Jul. 2014, available at: <http://www.euractiv.com/section/social-europe-jobs/opinion/freedom-of-movement-under-european-conditions/> (last visited 14 Mar. 2017).

80 Dur et al., “The Demise of Walk Zones in Boston” above n 59.