

The Local Refugee Match:

Aligning refugees' preferences with the capacities and priorities of localities

PRELIMINARY DRAFT

Will Jones¹ and Alexander Teytelboym²

Abstract³

Once refugees are granted protection in a particular host country, there is little concern about where in that country they are settled. Yet this matters enormously for refugees' chances to prosper in the new country and for the willingness of the local community to welcome them. We propose a centralised clearinghouse – a 'two-sided matching system' – to match refugees with localities. Drawing on the success of matching in domains such as public school choice, we outline principles underlying matching system design, and illustrate in general terms how they could be applied to refugee protection. This matching system respects the priorities and capacities of localities and gives agency to refugees. As an example, we describe in detail how such a system could work to meet the British government's commitment to resettle 20,000 Syrian refugees by 2020.

Keywords: refugees, forced migration, market design, matching, resettlement, United Kingdom, Syria.

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, Elizabeth Collett, Cathryn Costello, Scott Kominers, Katy Long, Gregory Maniatis, and Alvin Roth for their invaluable comments on earlier versions of this paper.

Introduction: the Problem of Who Goes Where and When

Imagine a world, similar to this one, where it is agreed that it is important that we educate every child. Imagine also that, once universal education has been decided by the demos, the state simply allocated each child to a school somewhere in the country. Thereafter, children were not permitted to switch schools, but have to make do. Little effort would be made to take into account anything about any child: where she and her parents might live, what skills she has, where her friends already attend, what her faith is, or what her interests are. The preferences of the children or of their parents would be completely ignored. In this world, there would be a lot of unhappy children in a lot of schools that would struggle to educate them.

This is probably still a better world than one in which children do not receive any education. But it is, clearly, vastly below the optimum. Without any additional resources, it would be possible to make everyone better off: some children and schools would be happy to swap places. Sometimes there are hard choices, where we can only make some schools better at the expense of others. But before we reach those choices, there are wins which come for free. These gains are achieved by successfully ‘matching’ children to schools.

This world is not so distant from how the protection of refugees is often conceptualised. The protection of refugees is often conceptualised as a numbers game: a simple matter of ‘how many?’ What this omits is the question of which particular refugee gets protection in which particular *locality* (any relevant form of local government responsible for the provision of services to refugees, such as towns, cities, borough, regions, federal states, councils or Local Authorities). People worry, rightly, about getting refugees out of dangerous zones to *somewhere* where they will be protected. But their final destination remains substantially under the radar of academic research and policy makers. After it has been determined that a particular refugee is being granted protection by a given hosting state, it remains necessary to find some way to determine which refugees end up where. In this paper, we argue that this can be done better than it is done currently.

When we do consider the obligations of local communities, most of the current debate focuses on *proportionality*, i.e. how we are to calculate what *number* of refugees a given locality ought to host. Much has been written about this already, with most suggesting some combination of (horizontal) fairness across localities modified by ability to host (e.g. Gibney 2015). This is encapsulated by the German *Königsteiner Schlüssel* system, which determines

quotas in a transparent way using a combination of each *Bundesland's* population and tax receipts. The second component of this debate could be thought of as *aptness*: when a large and diverse influx of refugees is resettled to a country, it is likely that many of them will be particularly suited to living in one place or another. They may have particular skills, which are locally in demand, or - conversely - they may have particular needs which a given locality is ideally positioned to meet. The willingness of localities to provide, the likelihood that refugees themselves will be in areas with the things they need, and the willingness of local communities to support refugees all depend on finding an apt match between refugees and hosts.

In this context, we propose a centralised *matching system* which would enable both refugees and localities to express, for themselves, their *preferences* as to where they would like to go, or which refugees they feel most capable of hosting. This system is not far removed from what is currently used in many cities such as Boston and New York, and countries such as the United Kingdom, to allocate children to schools. In these cases, the system does not decide which children go to school – they all do – and the system we propose here would not be appropriate for deciding which refugees get protected at all. Rather, once the questions of ‘how many and who?’ have been determined, it provides a fast and effective way of guiding governments in deciding which refugee should be hosted where. Not every refugee will be able to get their top choice locality, but the matching system finds an allocation that respects the second, third, fourth etc. preferences of all refugees without overstressing localities.

Firstly, we explain what aspects of status quo systems are sub-optimal in a manner we believe this proposal can address. Secondly, we briefly review extant contributions applying economic theory to institutional design for refugees, distinguishing this proposal from earlier suggestions.

How Resettlement Works

In the status quo, resettlement agencies and states do try to match refugees to local areas appropriately. For example, IOM and state bureaucracies conduct interviews in camps, and phone up locality representatives to enquire about hosting capacities. In private sponsorship systems, such as the Group of five (G5) system practised in Canada, refugees are relocated near their sponsors. In some countries, including Canada, refugees have unrestricted freedom

of movement once they have been relocated, and access to public funds wherever they decide to live. But this does not happen in the majority of cases. In most countries, states adopt policies which attempt to keep refugees in the areas to which they have been initially sent. For example, in the United Kingdom, there is extra funding made available to Local Authorities to help with the resettlement of Syrians under the Vulnerable Persons Resettlement Scheme. This funding is allocated to the Local Authority that first hosts the refugee: it does not automatically move with the refugee, and can only be transferred to another area with the bilateral agreement of both Local Authorities. Even in the first case, where refugees can move on from their first destination, there are still substantial costs and difficulties associated with moving (for example, refugees find themselves at the bottom of the social housing list). In the latter case, it is even more crucial that refugees and Local Authorities are matched to the best possible place first time.

Similarly, since the Immigration and Asylum Act of 1999, the UK has followed a policy that relocates asylum seekers away from the South East of England to a series of ‘dispersal zones’ within urban centres, largely in the North of England, Scotland, and Wales. This process was administered by the National Asylum Support Service (NASS, until its absorption into the Home Office). Although NASS entered into contracts with Local Authorities to house asylum seekers, there was no systematic process to attempt to make sure that asylum seekers and local areas were particularly suited to each other. This came under much criticism as being unacceptably punitive (Bloch and Schuster 2005), as removing refugees from precisely those social networks and community structures that could have fostered their integration (Darling 2011; Gill 2009), and leading to community tensions (things got bad enough for NASS to suspend relocations to six areas in 2004 (Casciani 2004)). Surveys conducted in these areas found that *both* those granted asylum and hosting communities expressed deep anxiety over the policy, and that existing residents felt particularly frustrated at a lack of consultation (D’Onofrio and Munk 2004). As many local communities announced that they would not participate in further rounds of contracts, NASS shifted to private service companies, who have come under considerable criticism for their expense (estimated in the region of £620 million by 2012) and the high volume of cases of abuse, harassment, and disproportionate force during deportation conducted by some of these companies (Grayson 2012). Obviously, the local communities where those granted asylum end up being hosted feel even more excluded from the process and thereby frustrated, and it has been argued that the privatisation

of provision in these ‘post-democratic’ cities has fostered further anti-refugee sentiment (MacLeod 2011).

There are three problems with systems for matching refugees to localities in the status quo, which our proposal can address. Firstly, these processes are invariably manual or ‘bespoke’: they rely on long interviews, contacting localities one by one, and an individual manually allocating refugees to particular areas. In cases of rapid relocation, even such bespoke matching becomes impossible. Trapp and Teytelboym (2017) show even with the data currently available to refugee agencies, it is possible to resettle refugees more effectively (e.g. resettle more refugees or achieve better measurable outcomes for refugees, see below) if the process were mathematically optimized. Secondly, this process is generally subject to minimal public scrutiny. This is a product of the focus on ‘how many’ rather than ‘where’. Once a state has committed to a headline resettlement figure, there is comparatively little discourse around how it was determined that a particular refugee ended up in a particular place. Thirdly, insofar as refugees are consulted about where they would like to go, their preferences are *inferred* and acted on by agencies, rather than directly and honestly stated and implemented. Often the resettlement agency decides on a particular measurable outcome, such as the probability of employment, and tries to match refugees in a way that maximize that outcome arguing that this outcome is in the best interest of both the agency and the refugee. The analogy is indeed, as we alluded to, above an educational system where the state simply decided which school your child should go to, based on what they *felt* was best for you, rather than actually asking you, and letting you decide, even where the state would have decided otherwise. Such systems are disempowering for the participants: they do not permit refugees to not only say what their preferences are, but to decide for themselves what is important to them.

This does not happen through any malice or incompetence on the part of those currently managing resettlement. Rather, there is currently no systematic way to connect the preferences of refugees with the capacities of local areas in a way which is rapid and comprehensive. For example, in 1999 it was decided to evacuate tens of thousands of Kosovar refugees, who were then relocated across Europe, Canada, United States, Australia and beyond (Fitzpatrick 2000; UNHCR 1999; Migration News 1999). Getting on to planes in Macedonia, some were unaware even of which country they were going to, let alone where in

that country.⁴ At the other end, communities and localities prepared to welcome and assist refugees, but frequently with very little information about the specific needs of the particular refugees they were receiving. This is not an isolated case, indeed, a 2015 report by Costello et al. (2014) to the European Parliament says the following:

Opting for coercion in the geographical distribution of asylum seekers and refugees may be costly for Member States in the longer term [...]. Policies of enforced dispersal outside main urban areas, adopted in Germany, Austria, Belgium, Denmark, the Netherlands, Sweden and the UK, have had negative impact on refugees' economic activity. The Swedish example is illuminating (Hatton, Constant, and Zimmerman 2013). Under a policy introduced in 1985, for instance, Sweden dispersed resettled refugees throughout the country, often to remote locations. A study conducted by Åslund and Rooth (2007) found that welfare dependency and non-employment increased as a result, particularly for those dispersed to remote areas with poor employment prospects.

There is, in fact, in a wealth of evidence from the Swedish case (where assignment of refugee to localities was almost random, allowing us to infer causality) that the initial locality has durable consequences for health, education, jobs, incomes and livelihoods of refugees (Åslund and Rooth 2007; Åslund and Fredriksson 2009a, 2009b; Åslund et al. 2011). Moreover, there is evidence that refugees actually spur mobility and increase labour market specialisation, benefitting the host population (Foged and Peri 2016).

Misplacement also fosters a xenophobic and restrictionist populism. Out of touch metropolitan elites, it is alleged, foist the costs and difficulties of refugee reception on communities with scant concern for the priorities and aspirations of the hosts (Myers 2015; *The Economist* 2016). For example, many trace the rising popularity of PEGIDA and similar xenophobic movements in the former East Germany to the Federal Government's insistence that they host large numbers of asylum seekers immediately following reunification (Joppke 1999; Gibney 2004). The argument made is that national politicians ignore the massive challenges and traumas of adjustment they faced at that time. This has seriously damaging consequences for the long-term willingness of communities to host refugees. Today, the British government is undertaking to resettle twenty thousand Syrian refugees by the end of the current parliament (in 2020). However, resettlement has proceeded slower than it could,

⁴ We are indebted to an anonymous reviewer for this observation.

due to the bespoke nature of the process. That 1,337 Syrians resettled in the UK through the Syrian Vulnerable Persons Resettlement Scheme between its launch in January 2014 and the end of December 2015 (Hawkins 2016) suggests that the scheme will have to become much faster in its later stages in order to meet its target. Further, the media are already raising concerns about possibly inapt matches. For example, the relocation of a Syrian family to the remote Isle of Bute in West Scotland received considerable media attention (Ackland 2016; Diamond 2015). This is not to single out Argyll and Bute Council for opprobrium: what is at stake is not their will or their good intentions, rather the possibility that a *mismatch* between what refugees need and what a community can provide risks undermining the success of resettlement in the short term, and support for the policy in the long term. Although participation in the scheme is voluntary, the absence of a process which transparently seeks to aptly match refugees to the most appropriate area fosters the widespread perception that refugees have been ‘foisted’ on communities ill-placed to deal with them, and having ‘perks that locals don’t enjoy’ (Reid 2015). These issues are sufficiently important that the European Court of Justice recently reaffirmed that although in general those granted international protection should be free to live where they like in a country, states did have the right to impose residence requirements on refugees “for the purpose of promoting their integration” (Robinson 2016).

In sum, current systems serve *nobody* as well as they could: hosting communities feel like their generosity is taken for granted, bureaucrats are saddled with an impossible task (where they attempt to make apt matches at all), and refugees themselves continue to know their lives are determined by factors outside their control. Finding a way to connect refugees with communities better is important to overcome the current bureaucratic bottlenecks, and to try and mollify communities worried about hosting refugees, but it is also crucial for the rights and livelihoods of refugees themselves. Many forms of protection and welfare support are not provided in every local area: the rights of disabled refugees, LGBT refugees, those suffering from PTSD or other forms of mental illness, all require particular services in order to properly vitiate their rights. In other contexts, many rights require appropriate community resources for their practical actualisation: not all local authorities have the same civil society capacities to call upon in making these rights real. Crucially, different areas may be able to provide differently: one community may have a well-organised network of Kinyarwanda-speaking churches, another will have Tigrinya-speaking Mosques. The practical ability of refugees to access such community resources, which could be crucial to their realisation of

their conception of the good, requires they be in particular parts of the country. To precisely that extent, the provision of these resources is also likely to relate directly to whether refugees and communities are able to integrate rapidly and durably.

Economics and Refugee Studies

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.⁵ 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,⁶ 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,⁷ or the system of humanitarian visas bestowed at consular outposts outside the European Union,⁸ recently advocated by Alexander Betts.⁹ We are agnostic as to how burdens should be optimally shared, but we 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,¹⁰ endorsed by some political theorists,¹¹ and criticised by others.¹² Fernández-Huertas Moraga and Rapoport are

⁵ P H Schuck, 'Refugee Burden-Sharing: A Modest Proposal', *Yale J. Int'l L.* 22, 1997: 243.

⁶ D Anker, J Fitzpatrick, and A Shacknove, 'Crisis and Cure: A Reply to Hathaway/Neve and Schuck', *Harvard Human Rights Journal* 11, 1998: 295.

⁷ J C Hathaway and R A Neve, 'Making International Refugee Law Relevant Again: A Proposal for Collectivized and Solution-Oriented Protection', *Harv. Hum. Rts. J.* 10, 1999: 115.

⁸ U I Jensen, 'Humanitarian Visas: Option or Obligation?', Study for the LIBE Committee, Directorate-General for Internal Policies, European Parliament, 2014.

⁹ Alexander Betts, 'Let Refugees Fly to Europe', *New York Times*, 24 September 2015, Online Edition edition, <http://www.nytimes.com/2015/09/25/opinion/let-refugees-fly-to-europe.html?nytmobile=0>.

¹⁰ J Fernández-Huertas Moraga, and H Rapoport. 'Tradable Refugee-Admission Quotas and EU Asylum Policy'. *CESifo Economic Studies* 61, no. 3–4, 2015: 638–72. doi:10.1093/cesifo/ifu037. 'Tradable Immigration Quotas'. *Journal of Public Economics* 115, July 2014: 94–108. doi:10.1016/j.jpubeco.2014.04.002. 'Tradable Refugee-Admission Quotas (TRAQs), the Syrian Crisis and the New European Agenda on Migration'. *IZA Journal of European Labor Studies* 4, no. 1, 2015: 23.

¹¹ J Kuosmanen, 'What (If Anything) Is Wrong with Trading Refugee Quotas?', *Res Publica* 19, no. 2, 25 September 2012: 103–19, doi:10.1007/s11158-012-9198-y.

¹² M J Gibney, 'Refugees and Justice between States', *European Journal of Political Theory*, 25 May 2015, 1474885115585325, doi:10.1177/1474885115585325.

the first to propose and mathematically develop applying 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. Jones and Teytelboym (2017) elaborate on these ideas and discuss the tradeoffs between different mechanisms (similar to what we discuss below) that could be used for refugee relocation and resettlement.

Importantly, all these other schemes are *international*. Such schemes are particularly ambitious: they require the cooperation of sovereign states, there are a variety of thorny issues to be settled in making such proposals compatible with international refugee law, and are operationally more complex. Therefore, international refugee matching is unlikely to be implemented until the Local Refugee Match has been successfully shown to work at the domestic level. Nevertheless, an international refugee match is a laudable end goal which might eventually connect many successful local refugee matches.

Finally, what we propose is *not* a ‘market’ in the sense critiqued by Jonathan Darling (Darling 2016). In a refugee match, no money changes hands, there is no trading of any kind, and the match is made through a centralised government-operated clearinghouse. Although it is in principle compatible with the kind of privatised provision Darling critiques, that is because it is compatible with *any* set of implementing partners. It is a market in the technical sense of the economists: a situation where individuals (in this case refugees) are allocated a precious resource (resettlement), where it is therefore important to use that resource as effectively as is possible. Whilst we do therefore insist that resettlement *is* a matter of resource allocation in a manner Darling would resist, we hope to show that our system has the *opposite* consequence of his fears: rather than investing decisions in a depoliticised managerial class, it directly returns agency and dignity to the refugees and communities involved in the process.

How Two-sided Matching Systems Work

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. Matching theorists are primarily interested who matches with whom, or in the words of Alvin Roth “who gets what and why”. The theory has been developed over the past thirty years and found a number of applications, including allocation of junior doctors to hospitals (Roth 1984; Roth and Peranson 1999), children to American public schools (Abdulkadiroglu and Sönmez 2003; Abdulkadiroğlu, Pathak, and Roth 2005, 2005), exchange of kidneys among living donors (Roth, Sönmez, and Unver 2004), and allocation of social housing (Abdulkadiroğlu and Sönmez 1999).

One of the most widely-used applications of matching is school choice. Matching systems like this are used in many American cities and in most European countries in one form or another (for an up-to-date survey, see www.matching-in-practice.eu). The two ‘sides’ of the matching system are schools and children. Schools have ‘preferences’ over children expressed as ‘priorities’ of the ‘types’ of children the schools admit. For example, the highest priority children could be the one who live in the neighbourhood *and* have siblings in the school, the next priority could be children who *only* have siblings in the school, followed by those who live in the neighbourhood and so on. Children and their parents have explicit preferences over particular schools. One family may prefer Shermer High School to Middlesex Ridge School to Rydell High. Another family may have completely reversed preferences.

The school board would then invite parents (and schools) to submit their preferences to central clearinghouse.¹³ A computer algorithm would be run in order to produce an outcome – an assignment of children to schools. Economists argue that this outcome ought to have the following desirable properties:

- 1) *Feasible* – no school’s quota for the maximum number of children is violated and every child is assigned to at most one school.
- 2) *Stable* – no school and children can agree to ‘re-match’ (outside the matching system for example) such that a child ends up in a better school and the school does not

¹³ For example, the submissions to the Boston Public School matching system can be done via a website: <http://www.bostonpublicschools.org/Page/653>

violate another child's priority. Hence, no student's priority is violated.¹⁴

- 3) *Efficient* – no child can get into a more preferred school without another child ending up in a less preferred school.
- 4) *Safe* (or 'strategy-proof') – children cannot end up in a more preferred school by misreporting their preferences.¹⁵

Unfortunately, there exists no algorithm that can always satisfy all these properties (Balinski and Sönmez 1999) for any submitted preferences and schools' priorities over students. Therefore, school districts often choose which algorithm to use. The Gale-Shapley algorithm used in New York produces a feasible, stable and safe outcome that is also the best possible outcome for the students among all such outcomes. By contrast, the Top Trading Cycles Algorithm used in New Orleans produces feasible, efficient and safe outcomes that may not be stable. In the UK, for example, the use of certain algorithms is prohibited by law. The 2007 English Schools Admissions Code banned first-preference-first algorithms, such as the Boston mechanism, on the grounds that they are not safe (Pathak and Sönmez 2013; Pathak 2016).

Even though different mechanisms produce different outcomes, they are all likely to be a substantial improvement on a system whereby parents send individual application to each school and the admissions are done in a decentralized way. This sort of system would breed all sorts of perverse incentives: there might be renegeing on accepted offers, early 'exploding offers' (i.e. offers with rapid expiry dates), constant re-matching, inefficient early matching, unfilled school places and so on (Mongell and Roth 1991; Avery et al. 2007; Echenique and Pereyra 2016). While a completely decentralized approach may work reasonably well in some circumstances (e.g. PhD admissions), it has been long recognized that in the case of schools, substantial time and cost savings can be made by carefully centralizing the information about preferences. Although trade-offs can be made between different matching systems, most well-designed matching systems for school choice will be an improvement over a decentralized system in which thousands of children scramble to get places in hundreds of schools.

¹⁴ Stability *eliminates justified envy* among students so it can be viewed a *fairness* criterion in our context as it is in school choice. Stability derives its original connotation of the doctor-hospital matching system in which there was a genuine concern that unhappy doctors and hospital would wish to rematch.

¹⁵ It may even be preferable that a matching system is not manipulable by *groups* of parents acting together to game it.

Our argument for the rest of this paper is that a similar logic applies in the case of refugees. Rather than who gets what and why we will focus on who goes where and when.

How the Local Refugee Match would work

The Local Refugee Match would function similarly to the matching systems used successfully to match children to schools. This system would come into effect after it is agreed that a given population of refugees are to be resettled, and that particular communities must host some proportion of that total number. At that point both parties submit a ranking to a centralised clearinghouse: refugees (as families) would submit their preferences over where they wish to go, and communities their priorities as to which categories of refugees they feel best able to help.

Localities would not rank refugees individually. This is not least for logistical reasons, but also opens a can of ethical worms, and might enable unscrupulous localities to use names as proxies for unethical criteria (e.g. second names which identify the refugee's faith or ethnicity). Instead, they would have 'priority categories' corresponding to their provision capacities, which they would rank. The full ranking of priority categories is the 'priority structure' of a particular locality. The provision capacities of localities are more diverse than is usually thought: for example, some hospitals specialise in providing for particular conditions. In a locality with a hospital treating unusual medical conditions (e.g. tropical medicine), the highest priority might be for refugees who have those conditions. Other priority categories might include: the suitability of accommodation, particular care services, the availability of particular forms of in-kind welfare, educational opportunities (e.g. spaces in schools), employment opportunities, the presence of particular civil society groups in a position to play support roles in refugee reception, and other integration services. Like the case of schools, higher priority will be given to refugees who satisfy several categories (e.g. those with a rare medical condition *and* family in the area will be prioritised over those who have either a rare medical condition or family).

The central state decides what the priority categories could be, but localities themselves could control their ranking of those categories. Deciding what categories it is permissible to rank on is important in order to prevent localities attempting to prioritise refugees in morally repugnant ways (e.g. were a locality to try and take only white refugees). One possible way to

do this would be to make the possible priority categories correspond to the categories of vulnerability and particular need already collected by UNHCR and other resettlement agencies. Then, it would be a simple matter to guarantee that refugees with those particular needs were matched to localities which actually possessed the capacity to meet them.

Some of these services may be more limited than the total hosting capacity of the locality, so a locality would be able to have both a headline quota (e.g. one hundred families) and a sub-quota for each priority category (e.g. for disability services). So the matching algorithm would find an outcome which satisfies all the priority categories (i.e. it doesn't overfill quotas in any category) without violating the overall priority structure or some of the properties of a 'good' outcome (see below) (Echenique and Yenmez 2015; Sönmez and Kominers 2015). Localities often know what these capacities are, but the task of the central state manually collecting, processing, and then making decisions with this information can be excruciatingly slow, and mistakes are understandably made. In contrast, a matching system processes all this information nigh instantly, and without a cumbersome bureaucratic processes.

One difference with school choice is that normally the priorities of schools are uniform: they are all imposed by whichever level of government regulates education (Local Education Authority in the UK). Whereas in the case of localities hosting refugees, it is clear that there is going to be huge differences in hosting capacities. Therefore it is sensible to allow localities some degree of discretion over their priority structure. For example, one locality might wish to prioritise rare medical conditions over family connections, whilst another might have a very different ranking.

Refugees have diverse preferences over where they most wish to go: they will have eclectic skills, needs, and life goals. For example, refugees with children have very different aspirations to refugees in or nearing retirement. Different refugee families will have friends or relatives in different parts of the country, will speak different languages, be of different faiths, and so on. Right now, centralised bureaucracies try to collate this information and infer preferences on behalf of refugees. While bureaucrats may have some idea about the top preference of a refugee family, they are extremely unlikely to identify correctly their second, third, fourth etc. preferences. The work of Emma Stewart and Gareth Mulvey amply demonstrates the imperfections of extant integration policies in the United Kingdom, many of which spring from the difficulties a government faces in taking particular individualised needs and aspirations into account (Stewart, Mulvey, and Council 2011; Stewart and Mulvey

2014). In the Local Refugee Match, refugees simply list their geographical preferences, much as the parents of children simply list their preferences over schools. The clearinghouse can restrict the list refugees pick from to a subset of localities which can meet their particular protection needs (e.g. all the areas with the appropriate medical provision). After that, it is for them to decide what by what criteria and how they wish to rank those areas. As such, the Local Refugee Match allows refugees and communities to decide *for themselves* what is most important to them.

In order for these preferences to be meaningful, an important role of the central government is to ensure that refugees receive accurate and comprehensive information about these areas in order to help them make a decision. In many cases this information already exists and simply needs to be put in a format refugees can take advantage of (e.g. translated into Kurdish). In some cases this information is already presented: for example, the Danish Refugee Council prior to resettling refugees in Denmark provide them with detailed information on their new hosting area. Refugees already face waiting periods between being told they are being relocated and actually making the journey, in which they have time to evaluate this information and find out as much more information as they like, using their own sources (refugees often communicate with their co-nationals who have already reached particular countries and have access to the internet). The burden of information gathering need not fall only on the local government and a lot of the information could be aggregated using crowd-sourcing on the part of an active citizenry.

Much as with school choice, there are trade-offs in the selection of algorithm. In the context of refugees, the desirable properties which could potentially be satisfied would be:

- (1) *Feasible* – no quota for priority categories either at the local or national are violated, which is to say no refugee is matched to a locality which had stated they lacked the capacity to effectively assist them. Any refugee is matched to at most one locality. No families are split.¹⁶

¹⁶ The presence of families may require some difficult tradeoffs to be made. For example, it is unclear whether a single individual satisfying priority categories A and B should be matched ahead of a two-person family with one member in priority category A and one member in B. A similar ethical issue has arisen previously in pairwise living-donor kidney exchange (Roth, Sönmez, and Ünver 2005, 2007). We thank Tayfun Sönmez for pointing this out.

- (2) *Stable* – refugees and localities do not end up dissatisfied with their choice and wanting to ‘re-match’ by moving to a different locality where they would have a higher priority than some other refugee.
- (3) *Efficient* – no refugee can be made better off given their preferences without making at least one other refugee worse off.
- (4) *Safe* – localities and refugees can honestly reveal their true preferences and priorities, and have no incentive to try and game the system.

The choice of the algorithm matters, because no system can fulfil all four of these conditions simultaneously. We would suggest that the most important properties in the case of the Local Refugee Match are feasibility, safety, and efficiency. Whilst efficiency and feasibility are clearly desirable, strategy-proofness is important because refugees have unequal levels of access to information about different localities, and if the system is gameable, those who are more informed may be at an advantage. In order to remove that advantage, and ‘level the playing field’, it would be important to convey to refugees that they cannot make themselves better off by lying about their preferences. Stability is not an especially important property in this case because re-matching (that leaves some refugees worse off) can either be legally prohibited or made exceptionally costly. For example, localities could face extremely serious legal penalties if they were to try and ‘drop’ a refugee in favour of another, or attempt to violate their own stated priority structures. Nevertheless, an unstable outcome might be perceived as ‘unfair’ if refugees can see that there more desirable localities in which they would have had a higher priority.

The more diverse the preferences are, the bigger the gains from using a matching system. Preferences and priorities are likely to be diverse for both localities and refugees. Moreover, the ‘side’ that has the *more* diverse preferences is the one that is more likely to get what they want. Which is to say: if all localities have similar priority structures, then the system is likely to make the desires of refugees themselves decisive.

Similarities and differences between Local Refugee Match and the design of other matching systems

While our ideas undoubtedly draw on a vast theoretical and practical literature in matching system design, it is worth emphasising what makes the Local Refugee Match a special and important case.

First, as in school choice, we will work off the assumption that localities are not strategic. The easiest way for locality to strategise even under the current system is to misreport their capacities to host refugees. Clearly, many localities do so and in fact there is no matching system that can induce them to reveal this capacity truthfully while producing stable outcomes (Sönmez, 1997). A further question is whether localities will try to misreport their priorities over refugees. Once again, if they do so, there is no way to induce them to reveal priorities honestly if the mechanism is stable (even if when the capacities are fixed, see Roth (1982)). The best we can hope for is for non-manipulation by one side of the system: refugees or localities. Hence, an important question is: which side is *more* strategic, refugees or localities? We expect that since refugees have a lot more at stake in the process, they will be the ones with the greatest incentives to strategically manipulate their preferences. However, as in other matching system design contexts, this is something that we would need to observe and settle empirically once the systems are tested.

Second, there might be some concern that refugees would only express limited choices. In school choice, for example, families can typically express about a dozen choices, but since schools are local, many families get one of their top 5 or 6 choices. Note how different this is in the case of college or university admissions. If the number of choices is restricted (as it is, for example, to five in the UK), the students have an incentive to only apply to colleges that they are likely to get into. There is a very strong reason for the applicant to be strategic: she might end up without a university place. Most governments do not commit to offering every applicant a college or university place. But governments *do* commit to finding a new home for every refugee family that it plans to resettle. So if refugees were asked directly where they would like to be resettled, they might reasonably misreport their preferences. In particular, they might wish to say that they would only like to be resettled in one particular locality or nowhere else. We argue such preference elicitation design is unsuitable for our purposes. Since the resettlement agency has committed to resettling this particular refugee family, this means that provided that it cannot resettle the family to the most preferred locality (especially since this locality might be listed by many refugee families), it have to assume that all other localities are equivalent in preference for the family. This is likely to be far from the family's true preferences. Of course, we cannot expect refugees to carefully rank preferences over

hundreds of localities, but we can expect them to rank *properties* or *features* of localities (such as whether they are urban or rural, have low crime, good school, parks etc.). The resettlement agency can therefore deduce what the refugees' preferences would be over *all* localities by using locality-level data. Some efforts in the design of such information systems are already under way and we discuss in more detail how this could be done in the British context.

Third, unlike in many other matching systems, in refugee resettlement we are often dealing with families rather than individuals. Families have multiple members who have different needs. In contrast to school choice where one child takes one school seat, a single family might take up 5 or 6 resettlement places in one locality. This can create problems even for the existence of stable outcomes. Therefore new notions of fairness will need to be put forward (see Delacretaz et al. (2017)).

The British Local Refugee Match

The United Kingdom provides an ideal ground to implement the Local Refugee Match. Since the UK is an island and not part of the Schengen Agreement, it has not experienced the vast irregular migration flows of continental Europe. As the Syrian crisis unfolded, on 7 September 2015, the British government announced its commitment to hosting 20,000 refugees over the next five years: the Syrian Refugee Resettlement Programme – a joint unit between the Home Office and the departments for International Development and Communities and Local Government. Many of the costs associated with refugees are borne by centrally-funded services (e.g. the Department of Work and Pensions), but the government has also undertaken to cover the additional costs of Local Authorities of hosting refugees for at least the first year, which has been estimated to range from £8,500 (adults and children under three) to £14,000 (school-age children) (Dedman 2015). The £460 million committed by the government is intended to meet these costs in full for the first year. This funding is in addition to housing support and educational funding, which are separately funded. However, the subsequent costs of refugee resettlement borne by local authorities until 2019-2020 are supposed to be met with only an additional £130 million. Therefore, local authorities will inevitably have to share some of the costs with the central government. Whilst there is no indication at present that Local Authorities are concerned about the costs at present, Local

Authorities may justifiably worry about funding shortfalls in the future.¹⁷ Minimising secondary movement in the first year is therefore particularly important.

There are several peculiarities of the British system, which create additional complexities and opportunities for good matching system design, which we now discuss. It's not unreasonable to think that many of these peculiarities will arise in other Local Refugee Matches in the future.

Voluntary Participation

First, while the overall quota/target has been set, the participation of Local Authorities is entirely voluntary. In Scotland, all thirty two Local Authorities are participating (Brooks 2015), but such solidarity is extremely unusual. In general, Local Authorities are likely to need incentives to join the Local Refugee Match. The simplest way to resolve this would simply be for the government to force Local Authorities to participate, as is the case in Germany. This seems unlikely, even if these quotas could then be traded by Local Authorities, as has been suggested in the context of international resettlement by Hillel Rapoport and Jesús Fernández-Huertas Moraga (2014; 2015; forthcoming). Equally, the government could drastically raise the financial support for each refugee. This seems even less likely in the current British context.

In this context, we suggest two other levers which can square this circle, and do not involve coercion. Firstly, it is important to encourage early entry into the matching system to ensure that refugees could be processed early are not disadvantaged from having less choice. Local authorities have an incentive to participate in the scheme when they believe the fewest other authorities are participating, as when the system is less competitive on that 'side', they are likely to get particularly apt matches given their priority structures. What this means is that if it appears that not many Local Authorities are participating in early rounds of the Local Refugee Match, they have every incentive to leap in at that point.

Second, the central government could increase the transfer payment for each additional refugee. Instead of a flat payment scheme paying out a fixed sum for every refugee, the total

¹⁷ There is also a discretionary central government fund that Local Authorities can apply to cover additional costs of hosting, such as adaptation of housing for disabled refugees.

payment to a Local Authority could be determined by an increasing function in which the per refugee payment rises as the number of refugees hosted rises (up to a certain maximum payment).

The payment function could increase in two ways: it could be concave or convex (see Appendix). This has important implications for the functioning of the system. In the status quo, the payment per refugee is either above or below the actual cost of hosting. In the UK today, it currently appears that the marginal (and average) cost is roughly two thousand pounds above what the government has budgeted to pay Local Authorities, which means Local Authorities will be forced to make up the difference from their central budgets. In this scenario, the only councils that participate will be the exceptionally generous, such as those in Scotland, and they will be either financially penalized for their generosity or the government will have to agree to spend much more money than it originally intended. What this means is that the status quo neither guarantees that the total target number of refugees will be relocated, nor that the budget constraints will be met.

Currently, it seems likely that the British government will either have to spend more money, or accept that it cannot get local councils to voluntarily accept hosting the target number of refugees. There are two options: the government can either insist that it is going to host as many refugees as is possible within their original spending plan, or commit to hosting a given number, even if the costs rise. In the former case (i.e. where there is a Spending Target), the priority is to find places for as many refugees as possible within that constraint. In the second case (where there is a Refugee Number Target), the challenge is find a way to host these refugees in the areas where they can be protected and provided for in the most affordable way.

The Spending Target system has payments determined by a concave function. Here, the central state can set an amount it is willing to spend, and the system determines the most refugees which can be allocated within that constraint. Local Authorities announce a maximum they would be willing to accept (their 'break-even' point), and would be willing to accept any number of refugees below that point (as the marginal payment is higher than the cost of hosting for all refugees below that point). The Spending Target system is likely to lead to a larger number of authorities are likely to host small number of refugees.

In the Refugee Number Target system, payments are determined by a convex function. Here,

the central state sets a number of refugees it wishes to resettle, and the system determines the cheapest allocation which delivers that. Local Authorities announce a minimum they would be willing to accept (their ‘break-even’ point), and are willing to accept further refugees up to their actual total capacity (it is unlikely that Local Authorities will actually be asked to host refugees at their maximum hosting capacity, where the marginal payments are astronomical). In this case, a smaller number of authorities are likely to host large numbers of refugees.

The intuitive preference of most in refugee studies would be for a system where a state commits to relocating whatever the cost, but the preference of government bureaucrats will undoubtedly be for a system where the total cost is under control. Within the context of a matching system with voluntary Local Authority participation, both may wish to think twice. It is not obvious whether the Spending Target or the Refugee Number Target would give refugees a better deal. The Spending Target would likely reduce the number of options refugees can select from, but since these Local Authorities’ quotas would be high, this might be more preferable if participating Local Authorities are popular destinations for refugees. The Refugee Number Target, on the other hand, would provide a more diverse choice for refugees and hence a possibility of finding very apt matches, but there would be more intense competition for the places in Local Authorities with small quotas. There is ongoing debate concerning the extent and manner in which refugees, like other migrants, should be concentrated or dispersed (Casciani 2004; Bloch and Schuster 2005). Depending on one’s view regarding these debates, the Spending or Refugee Number Target will be preferable. A knee-jerk preference one way or the other is unwarranted.

Too Much Information?

In the UK, there are roughly 350 Local Authorities that could host the refugees. Even if half of them enter the Local Refugee Match, it would be unreasonable to ask refugees to rank all of them in a preference list. It is much more important to provide high-quality information about each Local Authority to allow refugees to make an informed choice regarding those areas they do choose between. Refugees, unlike middle-class parents, are not in a position to pick up the *Good Schools Guide*,¹⁸ call on the Head Teacher, and spend hours on Mumsnet¹⁹

¹⁸ An independently published guide to all British schools, designed to inform parents.

¹⁹ Mumsnet is a British website for parents with a reputation for, amongst other things, particularly forensic dissections of the quality of schools.

to find the perfect school. Therefore, a system which genuinely allows refugees to express their preferences needs to ensure they are faced with manageable choices and good information.

To achieve this, the system can be simplified in two ways. First, the Home Office could pre-select only those Local Authorities that satisfy acute needs of the refugees, for example by only presenting them with Local Authorities that have the capacity to cater for a given medical condition, or welfare need (e.g. assisted housing). Second, the Home Office should only ask refugees to rank a reasonable number of Local Authorities: for example, parents do not need to rank all 120 state schools in Boston. In that matching system, they are asked to rank at most twelve.

In the context of the Local Refugee Match, what this would mean is that a given refugee would be handed a subset of the participating authorities to rank between. They could be confident that this subset was already of the only acceptable Local Authorities (which alone should make the process considerably less stressful). The government could provide, along with this list, a short précis of the local areas in question (this is information Local Authorities already have written: it would usually only be a matter of translation), and if they had time, refugees would be at liberty to acquire as much extra information as they liked through as many sources as they have access to. In the current British system, refugees often spend months ‘in the pipeline’ after they have been guaranteed relocation before they fly to the UK. This means there is a pre-existing window to allow refugees to make this choice as deliberately and reflectively as possible. In cases of more rapid relocation, having a couple of days to acquire information and think is better than having no choice at all. Even where refugees have little ability to acquire more information beyond that related to them by the central state, they have access to crucial information the government does not have: they know who they are. They know their aspirations, capacities, skills, hopes, and dreams. Trying to communicate all of this to Local Authorities is near-impossible. It is simpler to let refugees use this information for themselves.

Batching and Matching

Processing refugees can be a fairly slow process. It is therefore important to decide how frequently the matching system is run and a matching is found. In the case of schools, the

system is run once a year when all the applications are processed at once and all the children of a particular year group are matched. In the case of Syrian refugees coming to Britain, it would not be reasonable to wait until all 20,000 applications are processed and preferences are submitted before running the system. This would, however, ensure the most apt matches possible for the refugees and Local Authorities as a whole, but would involve unacceptably long waiting periods in oft-dangerous areas. Conversely, the system could be run once every three months or even more frequently to clear all the preferences of refugees collated in that period. The frequency of matching could be constrained by the speed of processing applications by the Home Office and various Local Authority characteristics, such as the speed of the housing market turnover. Regular matching intervals reduce operational costs, such as booking charter flights.

There are serious trade-offs in deciding the frequency of clearing (i.e. what batches refugees will be matched in) (Akbarpour, Li, and Gharan 2015; Dur 2011; Pereyra 2013).²⁰ If matching is done too frequently, some refugees who were deemed to have higher priority at an earlier clearing round may have lower priority (for a given Local Authority's priority structure) than a further set of refugees entering the system only slightly later. There might have been another refugee submitting their preferences a day later who would have had higher priority for the scarce protection capacity at that Local Authority, which has now been assigned. Avoiding this requires processing refugees less frequently, and in larger batches. However, refugees would then need to wait for longer rather than being protected right away. If the system matches too frequently and there are therefore fewer participating Local Authorities in any given round, refugees may end up damaging their own interests and that of other refugees in the Local Refugee Match. This happens by attempting to manipulate the system, because they decide to wait for a particularly desirable Local Authority (who they hope will be participating in a future round).

Frequent matching would not mean refugees were forced to make rapid decisions before they felt comfortable with their choices – they could enter any round they choose. This would only be problematic if the system was such that refugees thought they could game the system by holding off submitting in one round in favour of some later round. One way to avoid this problem is to allow refugees to submit preferences with plenty of time to think about their decision, but not inform them which particular batch they will enter.

²⁰ Trapp and Teytelboym (2017) quantify the costs of frequent batching in one resettlement context.

Only an indication

One issue confronting any government contemplating implementing a new scheme is the simple worry that this may all go horribly wrong for some reason unanticipated by the ivory tower academics who dreamt it up. Given how many lives are at stake, we should not be sanguine about this risk. In some cases, matching systems must be designed and implemented with no possibility of error. For example, systems which allocate children to schools would have horrific legal, social, and human ramifications if they went wrong, and had to be reneged upon. Telling children they are going to particular school, only to need to change that allocation later, is disruptive and cruel. In this case, it is important to avoid a scenario where refugees are told they are going to a particular area only for that promise to subsequently crumble. Crucially, this system is *not* for determining who is to be relocated: that is still determined by the conventional processes. A refugee guaranteed that they will be resettled in the UK has no chance of being suddenly bumped out of the system through the use of matching. It would still be regrettable to announce the results of a match which placed a given refugee in, say, Leeds, enabling them to start planning, learning about their new area, and making contacts, only for them to learn they were really going to Bognor. For this reason, we would suggest that the early stages of the system be manually checked, and run *indicatively*. In this scenario, the matching algorithm could be run, yielding a match, which could then be subsequently checked by the central government, and adjusted if need be, using qualitative information available only to bureaucrats, and any information which only became known after the match. All this could be done before the results are released to either refugees or Local Authorities, reducing the likelihood of embarrassing reversals. Small improvements to the Local Refugee Match can be made over time without taking the whole system down.

Conclusion: what this does (and what it doesn't)

The matching system is not a silver bullet. It will not solve the problems which ultimately cannot be overcome without a dramatic increase in the numbers of refugees which states are willing to accept. Similarly, it is not applicable in all contexts, such as Canadian private sponsorship, where local sponsors (more or less freely) choose the refugees they wish to assist. However, where it is implemented, it has several important advantages.

Firstly, like systems that allow parents to express preferences about which school their child goes to, this system raises the likelihood of achieving more *apt* matches. Not every refugee will get their first choice because inevitably some localities will be vastly over-subscribed, but the refugees' preferences over lower-ranked localities will be respected. The localities refugees are matched to will be *more likely* to have the particular capacities they need, correspond to their vision of their future, and are more likely to be able to offer them employment and other opportunities suited to them. The quality of life which a hosting state can offer to its refugees is likely to be dramatically higher as a result.

Secondly, and to the extent that the system delivers *apt* matches, that alone may increase the willingness of communities to host refugees in the long term. This is possible because, on the one hand, local communities may be less nervous about the effect of refugee arrivals on their lives, if they can be confident the system has found the best matches possible. On the other hand, *apt* matches should also translate into better integration: refugees will be more likely to find jobs, not strain local resources, suffer unnecessarily from a lack of appropriate welfare provision, and so on. Better-integrated and happier refugee populations should foster communities which are more confident about their ability to host refugees in future.

The third key advantage of this system is its speed. The slow progress of the British government in reallocating Syrians to the UK is a case in point. At their current pace they cannot hope to meet their target on time. This is not an isolated case: the protection of Kosovars across Western Europe and elsewhere, Indo-Chinese refugees resettled in the West under the Comprehensive Plan of Action, or Germany's current mass influx of Syrian refugees are all situations where a bespoke system of manually trying to allocate refugees to appropriate areas is simply not possible. The Local Refugee Match instantly delivers an allocation once the preferences of refugees and the priority structures of localities are aggregated. Collecting preferences in this form is considerably more efficient than the current processes. Obviously, the system would work best if refugees could collect as much information about their potential destinations as they wished before expressing their preferences, but where these decisions have to be made relatively rapidly, a matching system is particularly invaluable.

The fourth and most important advantage of this system is that it gives refugees themselves some control over their lives. Systems which consult refugees, infer their preferences, and then deliver back a decision are – however solicitous – deciding *for* refugees, not allowing refugees themselves to be the decision-makers. Of course, this is not an unconstrained choice. Neither are the choices of parents over schools. The debates around refugee protection are often polarised between an activist demand that refugees be given *total* freedom (which, in most cases, is unlikely to be forthcoming), and a restrictionist retort to the effect that refugees must ‘understand’ that they cannot ‘pick and choose’ (Harding 2015). In most spheres of human life it is granted that being offered some choice, constrained by what is available, and the choices of others, is valuable even if that choice is constrained: the choices of students applying to schools are constrained by the availability of places as well as the legitimate demands of other students. They cannot simply pick and choose. Nonetheless, their choices are real and valuable. It is not too much to ask that the same freedoms be extended to refugees.

References:

- Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2005. 'The New York City High School Match'. *American Economic Review* 95 (2): 364–367.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 1999. 'House Allocation with Existing Tenants'. *Journal of Economic Theory* 88 (2): 233–260.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. 'School Choice: A Mechanism Design Approach'. *The American Economic Review* 93 (3): 729–747.
- Ackland, Olivia. 2016. 'Meet the Syrian Refugees Being Housed on a Remote Island off the Scottish Coast'. *The New Statesman*, January 20. <http://www.newstatesman.com/world/middle-east/2016/01/meet-syrian-refugees-being-housed-remote-island-scottish-coast>.
- Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan. 2015. 'Dynamic Matching Market Design'. Mimeo. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2394319.
- Åslund, Olof, Per-Anders Edin, Peter Fredriksson, and Hans Grönqvist. 2011. 'Peers, Neighborhoods, and Immigrant Student Achievement: Evidence from a Placement Policy'. *American Economic Journal: Applied Economics* 3 (2): 67–95.
- Åslund, Olof, and Peter Fredriksson. 2009. 'Peer Effects in Welfare Dependence Quasi-Experimental Evidence'. *Journal of Human Resources* 44 (3): 798–825. doi:10.3368/jhr.44.3.798.
- Åslund, Olof, and Dan-Olof Rooth. 2007. 'Do When and Where Matter? Initial Labour Market Conditions and Immigrant Earnings'. *The Economic Journal* 117 (518): 422–48. doi:10.1111/j.1468-0297.2007.02024.x.
- 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*, 447–486.
- Balinski, Michel, and Tayfun Sönmez. 1999. 'A Tale of Two Mechanisms: Student Placement'. *Journal of Economic Theory* 84 (1): 73–94.
- Bloch, Alice, and Liza Schuster. 2005. 'At the Extremes of Exclusion: Deportation, Detention and Dispersal'. *Ethnic and Racial Studies* 28 (3): 491–512.
- Brooks, Libby. 2015. 'Scotland to Take One-Third of Syrian Refugees Arriving before Christmas'. *The Guardian*, November 5, sec. World news. <http://www.theguardian.com/world/2015/nov/05/scotland-to-take-one-third-of-syrian-refugees-arriving-before-christmas>.
- Casciani, D. 2004. 'Asylum City Dispersals Suspended'. *BBC News*, Online edition. <http://news.bbc.co.uk/1/hi/uk/4013431.stm>.
- Costello, Cathryn, Madeline Garlick, Elspeth Guild, Violeta Moreno-Lax, Minos Mouzourakis, and Sergio Carrera. 2014. 'New Approaches, Alternative Avenues and Means of Access to Asylum Procedures for Persons Seeking International Protection'. PE509.989. European Parliament.
- Darling, Jonathan. 2011. 'Domopolitics, Governmentality and the Regulation of Asylum Accommodation'. *Political Geography* 30 (5): 263–271.
- . 2016. 'Privatising Asylum: Neoliberalisation, Depoliticisation and the Governance of Forced Migration'. *Transactions of the Institute of British Geographers* 41 (3): 230–43. doi:10.1111/tran.12118.
- Dedman, Simon. 2015. 'Syria Refugees to Cost "up to £23k Each" in First Year in UK'. *BBC News*. October 19. <http://www.bbc.co.uk/news/uk-34567209>.
- Diamond, Claire. 2015. 'How Will Syrian Refugees Fare on a Scottish Island?' *BBC News*. November 18. <http://www.bbc.co.uk/news/uk-scotland-34855250>.
- D'Onofrio, L., and K Munk. 2004. 'Understanding the Stranger'. London: ICAR.
- Dur, Umut. 2011. 'Dynamic School Choice Problem'. Mimeo. <http://www.gtcenter.org/Archive/2011/Conf/Dur1312.pdf>.
- Echenique, Federico, and Juan Sebastián Pereyra. 2016. 'Strategic Complementarities and Unraveling in Matching Markets'. *Theoretical Economics* 11 (1): 1–39.
- Echenique, Federico, and M. Bumin Yenmez. 2015. 'How to Control Controlled School Choice'. *American Economic Review* 105 (8).

- Fernández-Huertas Moraga, Jesús Fernández-Huertas, and Hillel Rapoport. 2015. 'Tradable Refugee-Admission Quotas and EU Asylum Policy'. *CESifo 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*.
- Fitzpatrick, Joan. 2000. 'Temporary Protection of Refugees: Elements of a Formalised Regime'. *American Journal of International Law* 94 (2): 279–306.
- Foged, Mette, and Giovanni Peri. 2016. 'Immigrants' Effect on Native Workers: New Analysis on Longitudinal Data'. *American Economic Journal: Applied Economics* 8 (2): 1–34.
- Gibney, Matthew. 2004. *The Ethics and Politics of Asylum: Liberal Democracy and the Response to Refugees*. Cambridge, UK ; New York: Cambridge University Press.
- . 2015. 'Refugees and Justice between States'. *European Journal of Political Theory*, May, 1474885115585325. doi:10.1177/1474885115585325.
- Gill, Nicholas. 2009. 'Governmental Mobility: The Power Effects of the Movement of Detained Asylum Seekers around Britain's Detention Estate'. *Political Geography* 28 (3): 186–196.
- Grayson, J. 2012. 'Asylum Seeker Housing Managed by for-Profit Prison Guards? Why Not?' *Open Democracy*. <http://www.opendemocracy.net/ourkingdom/john-grayson/asylum-seeker-housing-managed-by-for-profit-prison-guards-why-not>.
- 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.
- Hatton, T. J., A. F. Constant, and K. F. Zimmerman. 2013. 'Refugee and Asylum Migration'. In *International Handbook on the Economics of Migration*. Northampton: Edward Elgar Publishing.
- Jones, Will and Alexander Teytelboym. 2017. 'The International Refugee Match: A System That Respects Refugees' Preferences and the Priorities of States'. *Refugee Survey Quarterly* 36 (2): 84–109.
- Joppke, Christian. 1999. *Immigration and the Nation-State: The United States, Germany, and Great Britain*. Clarendon Press. https://books.google.co.uk/books?hl=en&lr=&id=ih-WRR_EHkUC&oi=fnd&pg=PA1&dq=joppke+not+a+country+of+immigration&ots=mtThqU2EvD&sig=_erQnPtPHNuVliLhouN8dt8wmi0.
- MacLeod, Gordon. 2011. 'Urban Politics Reconsidered Growth Machine to Post-Democratic City?'. *Urban Studies* 48 (12): 2629–2660.
- Migration News. 1999. 'Kosovar Refugees' 6 (5). <https://migration.ucdavis.edu/mn/more.php?id=1801>.
- Mongell, Susan, and Alvin E. Roth. 1991. 'Sorority Rush as a Two-Sided Matching Mechanism'. *American Economic Review* 81 (3): 441–464.
- Myers, Rupert. 2015. 'This Fragile pro-Refugee Coalition Is Leaving Voters Behind. We Need to Take Them with Us'. *The Telegraph*. September 8. <http://www.telegraph.co.uk/news/worldnews/europe/11848519/This-fragile-pro-refugee-coalition-is-leaving-voters-behind.-We-need-to-take-them-with-us.html>.
- Pathak, Parag A. 2016. 'What Really Matters in Designing School Choice Mechanisms'. Mimeo.
- Pathak, Parag A., and Tayfun Sönmez. 2013. 'School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation'. *American Economic Review* 103 (1): 80–106.
- Pereyra, Juan Sebastián. 2013. 'A Dynamic School Choice Model'. *Games and Economic Behavior* 80: 100–114.
- Reid, Sue. 2015. 'Senseless! Why Are Syrian Refugees Being Foisted on a Remote Scottish Island with High Unemployment and Poverty - Then given Perks Some Locals Don't Enjoy?' *The Daily Mail*, December 19, Online edition. <http://www.dailymail.co.uk/news/article->

- 3366575/Senseless-Syrian-refugees-foisted-remote-Scottish-island-high-unemployment-poverty-given-perks-locals-don-t-enjoy.html.
- Robinson, Duncan. 2016. 'Europe's Top Court Places Curbs on Germany's Migration Policy'. *Financial Times*. March 1. <http://www.ft.com/cms/s/0/624a729a-df98-11e5-b67f-a61732c1d025.html#axzz428ya822z>.
- Roth, Alvin E. 1982. 'The Economics of Matching: Stability and Incentives'. *Mathematics of Operations Research* 7 (4): 616-628.
- . 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.
- Roth, Alvin E., and Elliott Peranson. 1999. 'The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design'. *American Economic Review* 89 (4): 748-782.
- Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver. 2004. 'Kidney Exchange'. *Quarterly Journal of Economics* 119 (2): 457-88.
- . 2005. 'Pairwise Kidney Exchange'. *Journal of Economic Theory* 125 (2): 151-188.
- . 2007. 'Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences'. *The American Economic Review*, 828-851.
- Sönmez, Tayfun. 1997. 'Manipulation Via Capacities in Two-Sided Matching Markets'. *Journal of Economic Theory* 77 (1), 197-204.
- Sönmez, Tayfun, and Scott Duke Kominers. 2015. 'Matching with Slot-Specific Priorities'. *Theoretical Economics*.
- Stewart, Emma, and Gareth Mulvey. 2014. 'Seeking Safety beyond Refuge: The Impact of Immigration and Citizenship Policy upon Refugees in the UK'. *Journal of Ethnic and Migration Studies* 40 (7): 1023-1039.
- Stewart, Emma, Gareth Mulvey, and Scottish Refugee Council. 2011. *Becoming British Citizens?: Experiences and Opinions of Refugees Living in Scotland*. Scottish Refugee Council. <http://core.ac.uk/download/pdf/9558191.pdf>.
- The Economist*. 2016. 'Radikale Rechte', February 27. <http://www.economist.com/news/europe/21693636-after-falling-years-xenophobic-extremism-germany-rise-radikale-rechte?frsc=dg%7Cc>.
- Trapp, Andrew and Alexander Teytelboym. 2017. 'Refugee Resettlement via Integer Programming'. In Preparation.
- UNHCR. 1999. 'Kosovo Crisis Update'. *UNHCR*. March 30. <http://www.unhcr.org/3ae6b80dc.html>.