



Matching Systems for Refugees

Will Jones
University of Oxford

Alexander Teytelboym
University of Oxford

Executive Summary¹

Design of matching systems between refugees and states or local areas is emerging as one of the most promising solutions to problems in refugee resettlement. We describe the basics of two-sided matching theory used in a number of allocation problems, such as school choice, where both sides need to agree to the match. We then explain how these insights can be applied to international refugee matching in the context of the European Union and examine how refugee matching might work within the United Kingdom, Canada, and the United States.

1. Introduction

Refugee resettlement is difficult. It is a complex administrative task, often conducted with inadequate funding and staffing, and, by necessity, often hurriedly. It frequently requires coordination across multiple agencies — public and private — operating on different sides of the globe.

Public attention and activism often focuses on the question of “how many?” In the context of incredibly limited supply of resettlement spaces, this is not unreasonable. The proportion of refugees in situations of “protracted displacement” (where more than 25,000 refugees have been in exile for more than five years) was estimated at two-thirds in 2009 (Loescher, Gil, and Milner 2009), prior to the current Syrian crisis. For refugees in protracted situations, the average length of stay is 17 years. There are now some 21.3 million refugees in the world, out of some 65.3 million forcibly displaced worldwide (UNHCR 2016a). In this context, global resettlement capacity is awesomely, ludicrously inadequate: in 2015, states admitted 107,100 refugees for resettlement, of which the United States accepted more than half (66,500) (UNHCR 2016b, 3). As one recent group of refugee experts uncompromisingly put it, “the current global system for refugee protection is broken” (Dauvergne and Hathaway 2016).

However, as with any scarce resource, it is also vital to consider how the limited resettlement capacity of the states can be used most effectively. In this paper, we focus on the question that arises after it has been decided that a given group of refugees will be resettled. Instead

¹ We are grateful to Alex Betts and especially to the participants of the Center of Migration Studies Rethinking the Global Refugee Protection System Conference (New York, July 2016) and the HCEO Market Design Perspectives on Inequality Conference (Chicago, August 2016) for their comments and suggestions.

of “how many?” we consider “who goes where, and how should it be decided?” Importantly, none of the methods discussed in this paper should *ever* be used to determine *which* refugees get resettled, only to determine where those who are being resettled could be best placed. Currently, national resettlement agencies or the United Nations High Commissioner for Refugees (UNHCR) conduct assessments in refugee camps and elsewhere to identify those most in need of resettlement, usually on the basis of vulnerability. Nothing in this paper suggests any change to that system. But *after* it has been determined that a refugee or refugee family is being resettled to a particular state, there is the need to decide the particular regional in which the refugee will be resettled.

This process is often committed, well-informed, and animated by a genuine desire to make the process work best for the refugees. However, it remains “bespoke”: refugee resettlement officers are manually attempting to collate and process vast amounts of information. Refugees are sometimes asked if they possess any particular preferences over where they might want to go, but their preferences are not comprehensively sought out and acted on in a systematic manner. In such processes, the preferences of refugees are in general *inferred*, rather than directly collected, which runs the risk that agencies make incorrect assumptions about what refugees want (for example, we often assume people wish to be near their extended family networks, but in some instances this could not be further from the case). Further, resettlement agencies constantly face trade-offs — for example, a given refugee has skills which suggests one area, but ensuring proximity to her co-linguists suggests another — but lack information as to how to prioritize the various factors. When such trade-offs arise between conflicting priorities, there is currently no transparent attempt to resolve the trade-offs, except by assuming in advance that a given dimension of priority (health or employability) is the most important one. In no systems that we have encountered in our research are refugees *themselves* asked directly how they would wish that trade-off to be made.

There is ample evidence, however, that placing refugees in the best areas for them has profound consequences for their long-term flourishing. Backing out of resettlement or simply outright refusal to participate in the process is more common than is generally thought: refugees may be understandably reluctant to embark on another long journey to a destination about which they know little, particularly if leaving the immediate region in which they have sought asylum involves losing contact with communities and economic networks which they have been relying on. Most extremely, when resettlement places refugees in areas where they do not want to go, they make the appalling choice to return, as when thousands of Iraqi refugees who arrived in Finland last year decided to cancel their asylum applications and to return home “voluntarily” (Forsell 2016). This is still comparatively rare: it is more common that refugees are placed and remain living in areas that lack the resources they need to embark on flourishing lives. For example, in the early 1990s Sweden deliberately placed refugees randomly around the country, hoping they would integrate better. The research done by Olof Åslund at Uppsala University based on this petri-dish human experiment has shown conclusively that the initial placement of refugees into less prosperous communities lowered their job prospects, health, income levels, and education (Åslund 2005; Åslund and Rooth 2007; Åslund and Fredriksson 2009; Åslund, Östh, and Zenou 2010; Edin, Fredriksson, and Åslund 2004). More recent evidence from the current refugee crisis in Germany suggests that refugees do not naturally gravitate

towards areas with more jobs and housing, but instead prioritize living close to their co-ethnics (Economist 2016). Getting resettlement right, therefore, requires paying attention to who goes where. Although this proposal is not designed to increase participation in resettlement directly, getting it right may make states, voluntary agencies, and communities more willing to participate in resettlement in the long run, thereby raising capacity.

In this paper, we suggest a friendly amendment to the status quo based on a rich literature in economics on “market design.” What we propose is a matching system — *the refugee match* — that can be used to match refugees to countries, to local areas, to agencies, or to different forms of protection status. This system uses no money or trading, better respects the freedom and choices of refugees, and serves the core interests of the states, agencies, and communities participating in resettlement.

First, we consider the basics of “matching theory”: how it works, and what it can do (Section 2). Then, we consider two possible uses of matching systems for refugee resettlement, internationally (Section 3) and locally (Section 4). An international refugee match would involve a hypothetical burden-sharing system between some set of states (e.g., the European Union) which wished to adopt a joint resettlement scheme (Jones and Teytelboym 2017a; Moraga and Rapoport 2014; 2015a; 2015b) A local refugee match would consider how refugees are allocated across localities within one jurisdiction (Jones and Teytelboym 2017b; Delacrétaz, Kominers, and Teytelboym 2016). Local refugee matching systems look rather different depending on the institutional context of a given country. Therefore, we consider what matching might look like, and what benefits it could provide, in three different contexts: the United Kingdom, the United States of America, and Canada.

2. Matching in Theory and in Practice

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. A running analogy throughout this paper is going to be with the match between children and schools. In many school districts (in many large American cities and across several European countries, including the United Kingdom), parents are allowed to express preferences over schools they wish to send their children to. They submit a ranking reflecting that they prefer the Rydell to Bronson Alcott to Sunnysdale High and so on. Schools, on the other hand, have fixed (non-manipulable) priorities regarding children: for example, they would prioritize children who have siblings in the same school and live in the neighbourhood over children who only have a sibling or only live in the neighbourhood. Importantly, schools know in advance that they are going to get some fixed number of students, so they have only incentives to be as attractive as possible. Parents’ and children’s preferences clearly matter — the government would have no way of reliably figuring out which of the tens of thousands of children would fare best at what kind of school. This is the sort of information that only parents and their children have and one of the roles of a designed matching system is to aggregate this information similarly to how a competitive market does (Hayek 1945). A centralized application process is also important: it would be unfair if children missed out on a place in an appropriate school only because their parents did not spend the night before application day outside the school doors.

What makes a good school choice matching system? Economists have identified three key potentially desirable characteristics of a successful and fair allocation system, all of which cannot exist in a single system and thus must be prioritized:

- *Stability* — no student should miss out on a place in a school she prefers because a student with a lower priority has taken it. This is known as the “elimination of justifiable envy” in economics, because it means schools and students do not end up dissatisfied with their match and seeking to “re-match” under the table, thus leading the whole system to unravel.
- *Efficiency* — no student should be able to get a place in a more preferred school without another student’s ending up in a school he prefers less. Put simply, there is no potential “swap” of any two students that would make anyone happier without making someone else unhappier.
- *Strategy-proofness* (for students) — no student should have an incentive to lie about their preferences over schools in the hope of getting a place in a more preferred school: it is logically impossible to “game” the system.

In the past decade and a half, economists have closely studied how school choice systems work (Abdulkadiroğlu and Sönmez 2003; Abdulkadiroğlu, Pathak, and Roth 2005). Abdulkadiroğlu and Sönmez showed the allocation procedure used in the past by the Boston public school authorities was not strategy-proof and therefore unlikely to have been efficient. This meant that after the allocation happened, parents sought to rematch their children to different schools. Moreover, parents were actively gaming the system. The economists offered two alternative strategy-proof matching systems: one was stable and the other was efficient.² Unfortunately, the school district had to make a choice as it has been mathematically proven that no matching system can be stable, efficient, and strategy-proof: there is no perfect system (Roth 1982). As such, school districts must prioritize: in the end New York and Boston picked the former and New Orleans (for a while) and San Francisco opted for the latter.

School choice is just one of many settings that matching theorists have studied. Other two-sided matching systems include the allocation of residents to hospitals (Roth 1984) and cadets to army branches (Sönmez and Switzer 2013). We argue that insights from matching theory can play a crucial part in the global refugee resettlement process.

Matching Theory in the Context of Refugee Resettlement

We now describe two types of two-sided matching systems for refugee resettlement. On the one side of the matching system are the refugees. We assume that families that do not wish to be split apart should not be. Therefore, the relevant unit of analysis is a refugee family. Moreover, we take as given the determination of refugee status and the decision to offer resettlement to a given family. On the other side are *countries* (Section 3) or *localities* (Section 4) that have stated their capacities to accept refugee placements in advance. The

² For those interested in reading further, these are algorithms known as the Gale-Shapley and Top Trading Cycles.

matching system only comes into effect after we know that a given population of refugees is to be resettled to a given set of placements. In this model, all refugees in the system would know that they were going *somewhere*, and all receiving countries and localities would know that they were receiving *some fixed number of refugees*.

The three potentially desirable properties of a matching system in the context of refugee resettlement would look like this:

- *Stability* — no refugee family should miss out on a resettlement place in a country/locality it prefers because a refugee family with a lower priority has taken it.
- *Efficiency* — no refugee family should be able to get a resettlement place in a more preferred country/locality without another refugee family's ending up in a country/locality it prefers less.
- *Strategy-proofness* (for families) — no country/locality should have an incentive to lie about its preferences over countries/localities in the hope of getting a resettlement place in a more preferred country/locality.

In the rest of this paper, we use “stability” and “efficiency” in these narrower technical senses rather than their vernacular meanings. Of course, there might be other properties that are equally important, such as maximizing the total number of refugees or refugee families that are resettled (Andersson and Ehlers 2016; Delacrétaz, Kominers, and Teytelboym 2016). As we pointed out above, there are trade-offs in satisfying these properties and therefore implementing one algorithm over another will depend on the particular context.

3. International Refugee Match: The Case of the European Union

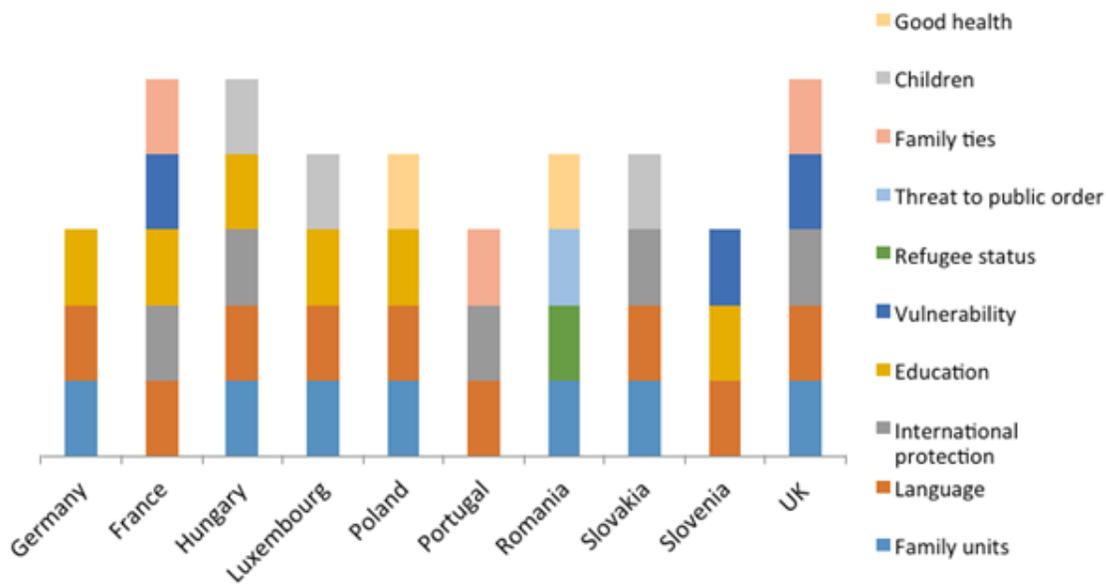
Resettlement still constitutes only a tiny fraction of the solution to the global refugee crisis: only one percent of all refugees get resettled annually. As such, it is an *exceptionally* scarce resource that is worth using as efficiently as is humanly possible. It is also worth recognizing that international refugee resettlement is already a two-sided matching system (Jones and Teytelboym 2017a). In order for a refugee family to be resettled, it must avail itself of the protection of the particular country that offers the family asylum. In fact, UNHCR already often acts as a matchmaker similar to the public school board by processing applications and suggesting refugees for resettlement in countries that have agreed to take the refugees. The current matching process, however, does not happen in a manner that has been systematically designed to achieve properties like stability, efficiency, and strategy-proofness. When UNHCR acts as a conduit for the vast majority of resettlement applications, it does not take into account the preferences of refugees and the priorities of countries in any systematic way. UNHCR typically suggests a family for resettlement in a particular country and the country processes the application it receives.

In most general terms, a *centralized* international 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 centralized body, simultaneously specifying where, if successful, they

would wish to be relocated. A different system would be one where states have identified in advance a population to be resettled (either through a group determination of status, or some other special dispensation). The system 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 priorities they are seeking in refugees. The countries and refugees are then matched.

A *European* refugee match 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, whether it is compatible with minimum or maximum refugee quotas, or neither). This is necessary, because although the European Union has been working towards a Common European Asylum System since 1999, including a stated commitment to the harmonization of resettlement, there is still substantial heterogeneity in the priorities of EU states over the refugees they are seeking to resettle. While, for example, France and the United Kingdom prioritize family ties, Romania puts emphasis on concerns about threats to public security (see Figure 1). Hence, a part of any solution must take into consideration these remaining different priorities among states.

Figure 1. Source: EASO Fact Finding Report on Intra-EU Activities in Malta



Source: EASO (2012).

Refugees themselves must have very different preferences, but it is virtually impossible to know what they are from the available data. Much of the European public discourse assumes that refugees would be grateful to be resettled anywhere in Europe. But in February 2016, for example, hundreds of Iraqi refugees who were resettled in Finland asked to be sent home (Forsell 2016). To some extent, survey statistics mislead about how many refugees arriving in Europe genuinely prefer to go to Germany rather than, say, the

United Kingdom (IOM 2016). Of course, thousands of Syrian refugees set out for Germany precisely because Germany, unlike most other European countries, had virtually guaranteed them asylum there. It is not at all clear whether refugees would have so overwhelmingly preferred Germany *other things being equal*. And yet even imperfect anecdotal data suggest that there is a substantial degree of variation in the preferred destination countries between, say, Syrians and Eritreans.³ Currently, refugees overwhelmingly prioritize safety and ease of route and the likelihood of being granted status *somewhere*. In consequence, we know almost nothing about the full preferences of refugees once those considerations are taken out of the equation. Getting the match right would therefore considerably improve the welfare of refugees and ease the burden on the states.

A refugee matching system, therefore, has the following benefits. First and most obviously, rather than being directed by the internal procedures of bureaucracies, it can give refugees more choice than they currently possess about where they are resettled. Second, a matching system can make the process less arbitrary and contingent on ad-hoc bilateral deals by states, and dangerous journeys and lottery-like decisions to make journeys by refugees. Third, because the computer algorithm finds the matching outcome as soon as the preferences of both sides are submitted, the matching system can speed up resettlement. Fourth, because matching enables states to effectively share burdens whilst also empowering them to receive refugees that match their preferences, this mechanism may encourage more states to participate in resettlement.

There are numerous details that would need to be decided in order to implement such a system. We comment in detail on them in another paper (Jones and Teytelboym 2017a). There are also difficulties with implementing such a scheme. First, many EU countries are likely to require financial or other incentives to participate (Moraga and Rapoport 2014; 2015a; 2015b). It is worth noting, in passing, that if this can be overcome, matching produces a further advantage: if states know they are accepting some fixed number of refugees, and know that the choices of refugees will determine which refugees they receive, they would then have every incentive to be as *attractive* to refugees as possible. Second, even though this improved refugee match is likely to reduce secondary migration, it is unlikely to eliminate it entirely. Therefore, it is unclear how the stability desideratum can be easily maintained. Third, EU countries would need to surrender some administrative tasks of refugee status determination to a central authority. While that does not mean that refugee status determination or priorities themselves need to be entirely harmonized, EU countries might need to entrust the application of their rules and procedures to a central EU body.

Ultimately, however, a well-run international matching system delivers a sense of control for the states, value for money for the public purse, and has huge welfare benefits for refugees. This is precisely what might encourage EU states to eventually increase their quotas.

4. Local Refugee Match

Once a refugee family is given asylum in a particular country, there is still a question of where exactly they end up living within that country. This process tends to be highly

³ Author interviews, 2016.

centralized in many countries, but it is also one that almost entirely ignores both the preferences of refugees and the priorities of local areas. Indeed, Sweden experimented with *random* allocation of refugees in the 1990s. While the policy was eventually scrapped, it allowed researchers to study the causal effects of allocation to particular localities. The key finding was that refugees in *general* who are initially matched to less prosperous localities fared considerably worse over their lifetime (in terms of employment and education) than refugees who ended up in more affluent areas. It is very unlikely that this result would not extend to other countries (though to our knowledge no other country except Denmark has experimented with an explicitly random policy). However, in contexts where resettlement capacity is likely to continue to involve a mix of more and less affluent areas, it is important to make sure that refugees are placed as effectively as possible: not all areas will have the same bundle of public services, the same employment gaps, the same religious or community-based goods, and so there are potential “wins for free” in making sure that refugees are placed in the areas which best serve their *particular* needs.

Localities would not rank refugees individually both for logistical and ethical reasons. To do so might enable unscrupulous localities, for example, 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 specialize 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 prioritized 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 from attempting to prioritize 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.

We now proceed to describe how best to match refugee to localities in three different countries. Each of these states allow refugees to migrate internally without restriction; offer considerable support to refugees in the early stages of resettlement; and have considerable control of the number of arrivals due to their geographical location, thereby allowing them to determine the optimal size and frequency of resettlement batches. Yet their particular institutional arrangements also offer interesting contrasts and comparisons in how the matching systems could be run. Moreover, as Delacrétaz, Kominers, and Teytelboym

(2016) point out, the local refugee match is a more complicated matching problem than the international refugee match because of the complex constraints on the public services that are required by refugee families.

United Kingdom

The simplest context for a local refugee match would be when refugees are relocated via a single scheme, operated by the central state. This is the case for the British government's Syrian Vulnerable Persons Resettlement (SVRP) program, extended by former Prime Minister David Cameron on September 7th 2015. By 2020, the SVRP program aims to resettle 20,000 Syrian refugees in families, alongside a further 3,000 unaccompanied minors from conflict situations⁴ in the Middle East and North Africa, and an unspecified number of unaccompanied refugee children currently in Europe. The program is managed by a joint unit between the Home Office and the departments for International Development and Communities and Local Government. The resettled refugees are given five years' humanitarian protection status, with permission to work and access public funds (paving a path to citizenship). As of March 2016, 1,602 people had been resettled in the United Kingdom under the scheme (Gower and Politowski 2016).

Under the SVRP program, refugees are matched to a *local authority* (which undertakes responsibility for that family for the duration of the visa), as well as to a specific house. Crucially, the participation of local authorities is entirely voluntary. Therefore, local authorities need incentives to participate in this program. This also means that the UK government is keen to prevent secondary movement of refugees from their initial placements, making it particularly important that refugees are matched successfully in the first instance.

In this program, most of the costs associated with refugees in the short-term are borne by the central state (e.g., benefits are paid directly by the Department of Work and Pensions), and the government has also undertaken to support local authorities with the additional costs they will bear. The costs have 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 supplements housing support and educational funding, which are separately funded. However, the subsequent costs of refugee resettlement are borne by local authorities⁵ or through a recently announced Full Community Sponsorship program (UK Home Office 2016).

There are non-trivial implementation issues for the British matching system. First, there are over 350 local authorities and the Home Office cannot possibly expect refugees to rank any or all of these localities without adequate information. One way to circumvent this is

4 The current British Prime Minister, Theresa May, currently appears to have rowed back on this commitment.

5 In another paper, we describe how heterogeneous financial incentives can be used to encourage more local authorities to participate (Jones and Teytelboym 2017b). Specifically, payments which scale according to the number of refugees hosted in an area will create very different patterns of resettlement based on whether the function determining how the total payment changes with the number of refugees hosted is concave or convex.

to elicit preferences over the *types* of local areas that refugees prefer (e.g., whether living in an urban area is more important than living near a place of worship). Second, refugees are being matched to housing procured on the private rental market, which can be volatile and unpredictable on a local level. Ensuring an adequate and reliable supply of housing is a crucial part of designing a well-functioning system.

Canada

Canada is one of the largest, regular destinations of refugee resettlement in the world. In 2014, for example, it resettled 7,233 refugees, which is around 10 percent of the UNHCR total. In November 2015, Prime Minister Justin Trudeau announced that Canada would resettle an additional 25,000 Syrian refugees and by late February 2016, this target was reached. Canada plans to resettle an additional 35,000 to 50,000 refugees by the end of 2016 alone.

There are three types of refugee resettlement schemes in Canada:

- Government assistance (GA) program — the government resettles refugees referred by the UNHCR.
- Private sponsorship program — private individuals or organizations can support resettlement of specific refugees.
- Blended Visa Office-Referral (BVOR) program — the government and private organizations share the costs of resettling refugees referred by the UNHCR.

The Canadian welfare system operates very similarly to the British system although the total share of government expenditures is closer to that of the United States at around 42 percent of GDP. The government (Ottawa or Quebec) covers the resettlement costs for up to one year in the GA program or six months in the BVOR program (the private sponsors cover the other six months). Healthcare is mostly free at the point of use and there is considerable government support for the unemployed, alongside other forms of welfare.

Matching systems could play a great role in government assistance and in blended program. In the private sponsorship scheme, the private organizations already have the ability to name particular refugees so there is no scope for matching with preferences.⁶ Since refugees could in principle express a preference over whether they enter the resettlement process via the GA or the BVOR programs, we can describe this matching problem as one of *matching with contracts* (Hatfield and Milgrom 2005; Sönmez and Switzer 2013). Hence, a refugee family might reasonably prefer to be resettled in Toronto by the Lutheran Church to being resettled in Vancouver by the government. The Canadian federal system lends itself to an expression of preferences among Canadian states in addition to the types of areas the refugees might prefer. Indeed, there are only 10 provinces (and three territories to which resettlement is unlikely) so it is possible to provide refugees with enough information that they are able to indicate a reasonable preference list. On the locality level, the overall service capacity would be jointly determined by the provision of the local governments

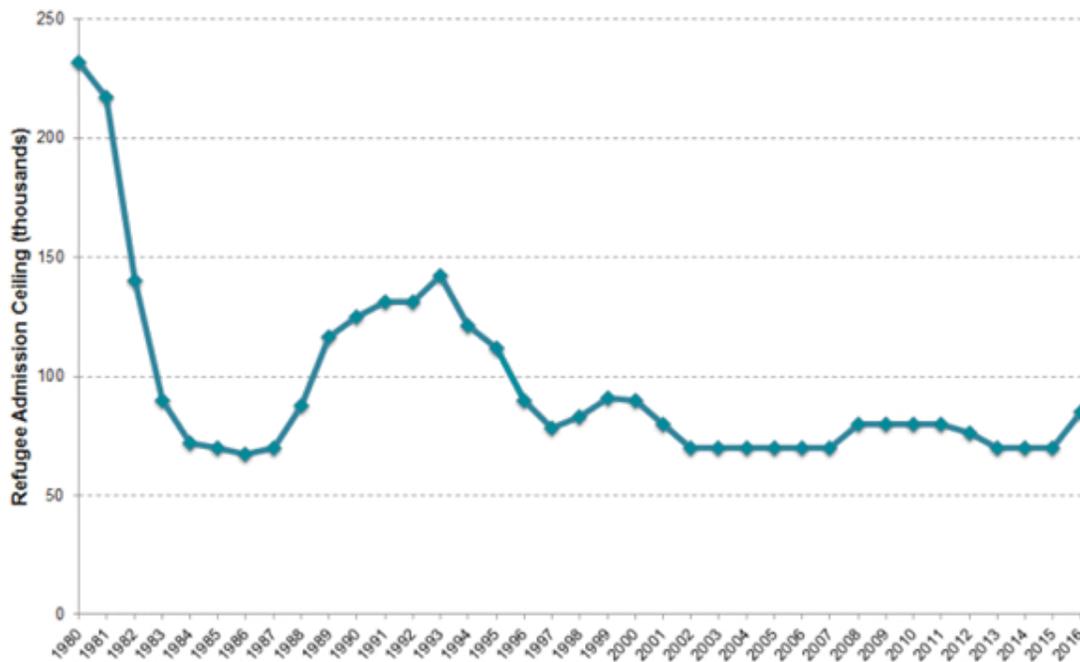
⁶ Private sponsorship schemes could, however, gain insights from optimizing the Pennsylvania Adoption Exchange (Slaugh et al. 2016).

and by the private organizations. Provinces would be able to coordinate the relevant information (including priorities) from individual localities and private organizations and encourage other localities to participate with extra funding. This makes it easier for the central government to negotiate province-level resettlement numbers that would act as total provincial quotas in the system (Kamada and Kojima 2015).

United States

The United States is, by a distance, the most generous participant in global resettlement. Since 1975, over three million refugees have been welcomed to the United States. The Department of State is the coordinating government body, and resettlement is funded by the Department’s Bureau of Population, Refugees, and Migration (the “Bureau”). The number of persons that may be admitted as refugees each year is established by the president in consultation with Congress. This “Refugee Admission Ceiling” is then divided into five global regions, alongside an unallocated reserve for unforeseen emergencies. Refugees are expected to apply for permanent residency (a green card) one year after being admitted to the United States, and may apply for US citizenship five years after admission as a refugee.⁷

Figure 2. US Annual Refugee Admission Ceilings, Fiscal Year 1980-2016



Source: US Department of State, Bureau of Population, Refugees, and Migration, “Proposed Refugee Admissions for Fiscal Year,” various years, with graph from Zong and Batalova (2015).

⁷ For a fuller account of the US refugee resettlement system and actual refugee admissions by year, see Kerwin (2015).

Matching Systems for Refugees

However, the Bureau's role in the operation of resettlement is, in contrast to the British system, extremely limited. Instead, the Department of State works with nine private voluntary agencies (or "Volags"). These agencies sign cooperative agreements with the Department of State to provide reception and placement services to all refugees arriving in the United States. The current list of Volags includes:

- Church World Service
- Episcopal Migration Ministries
- Ethiopian Community Development Council
- Hebrew Immigrant Aid Society
- International Rescue Committee
- Lutheran Immigration and Refugee Service
- US Committee for Refugees and Immigrants
- United States Conference of Catholic Bishops
- World Relief

These agencies meet weekly to review a batch of refugees based on biographical information sent to them by the Resettlement Support Centers (the overseas offices which identify refugees for resettlement). During this weekly meeting, decisions are made as to which agency will take responsibility for which refugees. Volags take turns every week in their priority to pick the cases that they want to take on. Then, the agencies themselves decide where to place refugees across some 190 communities throughout the United States. As such, in the US context, there is a two-stage match: allocating refugees to Volags, and then to local communities.

The American system is extremely well-institutionalized and can tap into a vast well of expertise and a huge network of supportive actors because it is embedded in (largely religious) community networks. However, the system also expects refugees to rapidly transition to self-reliance (within about six months) and operates in a context with a minimal welfare safety net. After that point, support and services are largely coordinated by the federal Health and Human Services's Office of Refugee Resettlement, which similarly will have a complex and uneven national map of provision capacities. For these reasons, the priority in the match is likely to be designing a system which maximizes the likelihood that refugees can integrate, join the labor market, and cease to need extensive support as soon as possible. Furthermore, the sheer scale of resettlement in the United States makes preference-free matching particularly likely to miss important information.

As in the British case, it is unlikely that refugees will have adequate information to process and rank localities directly. Therefore, it is also likely here that the best system would encourage refugees to express preference over *types* rather than specific named areas. As in the Canadian case, refugees could have preferences over not only where they are resettled but also which Volag they are resettled by, opening another application of matching with

contracts. Refugees could be asked to rank their most preferred states (as in the Canadian refugee match) as well as properties of areas within these states (as in the British refugee match) alongside preferences over Volags. That process would greatly simplify the weekly allocation task for the Volags themselves.

5. Conclusion

This paper has argued that matching must be an important part of any durable and comprehensive response to the global refugee crisis. We have shown that centralized matching in which the preferences of refugees and the priorities of states (international) or of localities (local) are elicited directly can be used in a variety of contexts to make the resettlement process more fair, efficient, and humane. The next step would be to understand empirically the effect of matching with preferences on the lifetime outcomes of refugees and development of community relations.

While we are confident that major host countries for refugee resettlement will eventually adopt comprehensive matching systems for resettlement, there is a greater need for matching systems in places where most refugees currently are — the developing world. Designing and implementing matching systems in countries with weak institutions pose a substantial challenge to researchers and policymakers. However, this challenge must be the most fruitful area for further research.

REFERENCES

- Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2005. “The New York City High School Match.” *American Economic Review* 95(2): 364-67. <https://doi.org/10.1257/000282805774670167>.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. “School Choice: A Mechanism Design Approach.” *The American Economic Review* 93(3): 729-47.
- Andersson, Tommy, and Lars Ehlers. 2016. “Assigning Refugees to Landlords in Sweden: Stable Maximum Matchings.” Working paper. Lund, Sweden: Lund University Department of Economics.
- Åslund, Olof. 2005. “Now and Forever? Initial and Subsequent Location Choices of Immigrants.” *Regional Science and Urban Economics* 35(2): 141-45. <https://doi.org/10.1016/j.regsciurbeco.2004.02.001>.
- Åslund, Olof, and Peter Fredriksson. 2009. “Peer Effects in Welfare Dependence Quasi-Experimental Evidence.” *Journal of Human Resources* 44(3): 798-825. <https://doi.org/10.3368/jhr.44.3.798>.
- Åslund, Olof, John Östh, and Yves Zenou. 2010. “How Important Is Access to Jobs? Old Question—improved Answer.” *Journal of Economic Geography* 10 (3): 389-422. <https://doi.org/10.1093/jeg/lbp040>.

- 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. <https://doi.org/10.1111/j.1468-0297.2007.02024.x>.
- Dauvergne, Catherine, and James Hathaway. 2016. "Report of an Expert Workshop on Understanding Refugee Law As Shared Responsibility ('The Vancouver Principles')." Vancouver: Allard School of Law, University of British Columbia.
- 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>.
- Delacrétaz, David, Scott Duke Kominers, and Alexander Teytelboym. 2016. "Refugee Resettlement." *Mimeo*, August 4.
- EASO (European Asylum Support Office). 2012. *EASO Fact Finding Report on Intra-EU Relocation Activities from Malta*. Valletta, Malta: EASO. www.refworld.org/pdfid/52aef8094.pdf.
- Economist. 2016. "Daily Chart: Refugees in Germany May Be Seeking Asylum in the Wrong Places." *The Economist*, April 25. <http://www.economist.com/blogs/graphicdetail/2016/04/daily-chart-8>.
- Edin, Per-Anders, Peter Fredriksson, and Olof Åslund. 2004. "Settlement Policies and the Economic Success of Immigrants." *Journal of Population Economics* 17(1): 133-55. doi:10.1007/s00148-003-0143-4. <https://doi.org/10.1007/s00148-003-0143-4>.
- Forsell, Tuomas. 2016. "Thousands of Iraqi Refugees Leave Finland Voluntarily." *Reuters*, February 12. <http://www.reuters.com/article/us-europe-migrants-finland-idUSKCN0VL0UE>.
- Gower, Melanie, and Ben Politowski. 2016. "Commons Library Briefing: Syrian Refugees and the UK Response." London: House of Commons Library. <http://researchbriefings.parliament.uk/ResearchBriefing/Summary/SN06805>.
- Hatfield, John William, and Paul R. Milgrom. 2005. "Matching with Contracts." *American Economic Review* 95(4): 913-35. <https://doi.org/10.1007/s00148-003-0143-4>.
- Hayek, Friedrich August. 1945. "The Use of Knowledge in Society." *American Economic Review* 35(4): 519-30.
- IOM (International Organization for Migration). 2016. *Analysis: Flow Monitoring Surveys. DTM Flow Monitoring Survey*. <http://migration.iom.int/docs/Analysis%20-%20Flow%20Monitoring%20and%20Human%20Trafficking%20Surveys%20in%20the%20Mediterranean%20and%20Beyond%20-%202011%20August%202016.pdf>.
- Jones, Will, and Alexander Teytelboym. 2017a. "The International Refugee Match: A System That Respects Refugees' Preferences and the Priorities of States." *Refugee Survey Quarterly* 36(2): 84-109. <https://doi.org/10.1093/rsq/hdx004>.
- . 2017b "The Local Refugee Match: Aligning Refugees' Preferences with the Capacities and Priorities of Localities." *Mimeo*.

- Kamada, Yuichiro, and Fuhito Kojima. 2015. "Efficient Matching under Distributional Constraints: Theory and Applications." *American Economic Review* 105(1): 67-99. <https://doi.org/10.1257/aer.20101552>.
- Kerwin, Donald. 2015. "The US Refugee Protection System on the 35th Anniversary of the Refugee Act of 1980." *Journal on Migration and Human Security* 3(2): 205-54. <https://doi.org/10.14240/jmhs.v3i2.49>.
- Loescher, Gil, and James Milner. 2009. "Understanding the Challenge." *Forced Migration Review* 33: 9-11.
- Moraga, Jesús Fernández-Huertas, and Hillel Rapoport. 2014. "Tradable Immigration Quotas." *Journal of Public Economics* 115(2014): 94-108. <https://doi.org/10.1016/j.jpubeco.2014.04.002>.
- . 2015a. "Tradable Refugee-Admission Quotas and EU Asylum Policy." *CESifo Economic Studies* 61(3-4): 638-72.
- . 2015b. "Tradable Refugee-Admission Quotas (TRAQs), the Syrian Crisis and the New European Agenda on Migration." *IZA Journal of European Labor Studies* 4(1).
- Roth, Alvin E. 1982. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research* 7(4): 617-28. <https://doi.org/10.1287/moor.7.4.617>.
- . 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *The Journal of Political Economy* 92: 991-1016.
- Slaugh, Vincent W., Mustafa Akan, Onur Kesten, and M. Utku Ünver. "The Pennsylvania Adoption Exchange Improves Its Matching Process." *Interfaces* 46(2): 133-53. <https://doi.org/10.1287/inte.2015.0828>.
- Sönmez, Tayfun, and Tobias B. Switzer. 2013. "Matching With (Branch-of-Choice) Contracts at the United States Military Academy." *Econometrica* 81(2): 451-88. <https://doi.org/10.3982/ECTA10570>.
- UK Home Office. 2016 "Community Sponsorship Scheme Launched for Refugees in the UK." London: UK Home Office. <https://www.gov.uk/government/news/community-sponsorship-scheme-launched-for-refugees-in-the-uk>.
- UNHCR. 2016a. "Figures at a Glance." Geneva: UNHCR. <http://www.unhcr.org/uk/figures-at-a-glance.html>.
- . 2016b. *Global Trends 2015*. Geneva: UNHCR. <http://www.unhcr.org/en-us/statistics/unhcrstats/576408cd7/unhcr-global-trends-2015.html>.
- Zong, Jie, and Jeanne Batalova. 2015. "Refugees and Asylees in the United States." Washington, DC: Migration Policy Institute. <http://www.migrationpolicy.org/article/refugees-and-asylees-united-states>.