



LUND UNIVERSITY

Dynamic Refugee Matching

Andersson, Tommy; Ehlers, Lars; Martinello, Alessandro

2018

Document Version:
Other version

[Link to publication](#)

Citation for published version (APA):

Andersson, T., Ehlers, L., & Martinello, A. (2018). *Dynamic Refugee Matching*. (Working Papers; No. 2018:7).

Total number of authors:

3

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Working Paper 2018:7

Department of Economics
School of Economics and Management

Dynamic Refugee Matching

Tommy Andersson
Lars Ehlers
Alessandro Martinello

March 2018
Revised: October 2018



LUND
UNIVERSITY

Dynamic Refugee Matching*

Tommy Andersson[†], Lars Ehlers[‡] and Alessandro Martinello[§]

First version: March 27, 2018. This version: October 29, 2018.

Abstract

Asylum seekers are often assigned to localities upon arrival using uninformed matching systems, which lead to inefficient and unfair allocations. This paper proposes an informed dynamic mechanism as an intuitive and easy-to-implement alternative. Our mechanism can be adopted in any dynamic refugee matching problem given locality-specific quotas and that asylum seekers map into specific categories. Any matching selected by the proposed mechanism is Pareto efficient, and envy between localities is bounded by a single asylum seeker. Our simulations show that the proposed mechanism outperforms uninformed mechanisms even in presence of severe misclassification error in the estimation of asylum seeker categories.

Keywords: forced migration, market design, refugee matching, dynamics, envy, efficiency.

JEL Classification: C71, C78, D71, D78, F22.

1 Introduction

According to UNHCR, 17.2 million persons obtained refugee status in 2016.¹ In an attempt to reduce pressure on countries sheltering large numbers of refugees, several international resettlement programs—the largest of which is operated by UNHCR—transfer a selected portion

*The authors are grateful to Atila Abdulkadiroğlu, Daniel Halvarsson, Fuhito Kojima, Johan Grip, John Hatfield, Jean-Jacques Herings, Thayer Morill, Umut Dur, Utku Ünver and Yeon-Koo Che as well as seminar audiences at North Carolina State University, the Ratio Institute (Stockholm), the Swedish Finance Ministry, Vanderbilt University, and the 2018 SAET conference (Taipei) for detailed comments and remarks. Andersson gratefully acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation (P2016-126:1) and the Ragnar Söderbergs Foundation (E8/13). Ehlers is grateful to SSHRC (Canada) and FRQSC (Québec) for financial support.

[†]Department of Economics, Lund University, P.O. Box 7082, SE-222 07 Lund, Sweden. E-mail: tommy.andersson@nek.lu.se.

[‡]Département de Sciences Économiques and CIREQ, Université de Montréal, C.P. 6128, Succursale Centre-Ville, Montréal, Québec H3C 3J7, Canada. E-mail: lars.ehlers@umontreal.ca.

[§]Department of Economics, Lund University, P.O. Box 7082, SE-222 07 Lund, Sweden. E-mail: alessandro.martinello@nek.lu.se.

¹Unless stated differently, all figures in this section appear at www.unhcr.org/figures-at-a-glance.html.

of refugees to countries that have agreed to admit them, and to ultimately grant them permanent settlement.² Although in practice these programs rarely take the characteristics and preferences of refugees and receiving countries into account, [Jones and Teytelboym \(2017a,b\)](#) note that these types of resettlement schemes share many fundamental properties with classical two-sided matching markets with capacity constraints such as the school choice problem ([Abdulkadiroğlu and Sönmez, 2003](#)). Even if the computational and modeling complexity of asylum seeker allocation is higher than that of a typical two-sided matching problem,³ results from the rapidly growing market design literature can lead to vast improvements in efficiency, equity and feasibility of refugee resettlement programs ([Andersson, 2017](#); [Bansak et al., 2018](#); [Delacretaz et al., 2016](#); [Jones and Teytelboym, 2017a,b](#)).

A crucial similarity between classical two-sided matching problems and refugee resettlement programs is that all available information on both sides of the market— asylum seekers and localities—is known when making the matching. Only once the full set of asylum seekers, their characteristics, and all preferences are known, the matching can take place. However, gathering such information necessarily requires holding asylum seekers in asylum centres for an undefined period of time. Moreover, less than 1 percent of all asylum seekers are part of resettlement programs.⁴ The vast majority of asylum seekers arrive directly to their host country, where the assignment to localities needs to take place in a short period of time, and without any information on future arrivals. The difficulty of predicting the characteristics of asylum seekers arriving in the near future necessarily introduces dynamics to the problem, i.e., asylum seekers must be assigned to localities directly upon arrival before the identity and characteristics of future arrivals are revealed.

This paper solves this dynamic refugee matching problem by proposing an intuitive, easy-to-implement and computationally efficient dynamic refugee matching mechanism that, by exploiting refugee-specific characteristics affecting the probability of successful integration at each locality, guarantees an efficient and unenvious matching at each arrival among localities that have not filled their quotas yet. The efficiency and the envy concepts are based on dynamic versions of Pareto efficiency and envy-by-one ([Budish, 2011](#)), respectively.⁵

²In recent years, the United States has been the world’s top resettlement country followed, by Canada, Australia and the Nordic countries. Since 2015, the European Union has also organized a series of resettlement programs in an attempt to reduce pressure on the countries located at the external border of the European Union (Greece, Hungary and Italy in particular).

³For example, refugees vary in several dimensions (education, spoken languages, household composition, need for healthcare, etc.), which introduce multidimensional constraints and potential complementarities to the matching problem. Complementarities have been studied in matching markets by, e.g., [Hatfield and Kominers \(2014\)](#), [Kojima et al. \(2013\)](#), [Pycia \(2012\)](#) and [Roth and Peranson \(1999\)](#).

⁴The UNHCR program, for example, contained files of only 163,200 asylum seekers for consideration by resettlement countries, and of these, only 126,200 actually departed to resettlement countries. Similarly, while more than 2.5 millions people applied for asylum in 2015 and 2016 alone in the European Union, the total number of relocations within the European Union program stands at 24,676 as of July 24, 2017 (European Commission, Press release, 26 July 2017, Brussels).

⁵The importance of these specific properties has also been stressed in several meetings that the authors of this

Despite its simplicity, the proposed mechanism strongly outperforms uninformed mechanisms such as those currently in use in, e.g., Sweden, Switzerland and the United States. Using aggregate data from Sweden and the United States to calibrate random flows of asylum seekers, this paper shows that even if the probability of successful integration for each asylum seeker and locality is unknown and needs to be empirically estimated (as in [Bansak et al., 2018](#)), the proposed matching mechanism increases efficiency by up to 50 percent, and guarantees a reduction in envy of between 7 and 45 percent.

Although the results presented in this paper are not specific to a particular country, Sweden represents a good example for describing the dynamic problem of matching asylum seekers to localities within a country. An asylum seeker who enters Sweden is temporarily placed at a Migration Board accommodation facility in anticipation of either a deportation order or a permanent residence permit. This facility is located in some locality and the problem for the Migration Board can, therefore, equivalently be described as a problem of matching asylum seekers to localities.⁶ In the process of matching asylum seekers to localities, asylum seekers can be mapped into specific groups, or categories (e.g., based on age, education, language knowledge, etc.). Asylum seekers belonging to different categories may not only have different probabilities to integrate overall, but also exhibit particular synergies with specific localities ([Bansak et al., 2018](#)).⁷

The current Swedish system to dynamically match asylum seekers to municipalities is approximately based on a random assignment coupled with municipality specific quotas. As described by [Bansak et al. \(2018, p.325\)](#), a similar mechanism is employed by the resettlement offices in the United States and by the federal government in Switzerland. Random mechanisms are uninformed and, consequently, cannot exploit the heterogeneity among localities and asylum seekers to achieve efficiency and fairness. In fact, if such heterogeneity is not taken into consideration, the resulting matching may be both inefficient and unfair (this point is illustrated in [Section 4](#)). In contrast to the random mechanism, the dynamic refugee matching mechanism proposed in this paper takes both locality-specific characteristics and the background of the asylum

paper have had with the Swedish Migration Board. In fact, in a recent report (Swedish Government, [SOU:2018:22, 2018](#)), the Swedish Migration Board is recommended to further investigate the feasibility of some of the ideas presented in this paper.

⁶Note that an asylum seeker may be transferred to a different locality once granted a residence permit. However, as the authors of this paper has discussed with the Swedish Migration Board, asylum seekers will ideally stay in the same locality as this creates incentives for the locality to start integrating the asylum seekers immediately after arrival to Sweden. This is, in particular, true for asylum seekers that are classified to belong to “Track 1” (Swedish “Spår 1”) according to the Migration Board classification systems. These are the asylum seekers that will be granted asylum with a probability close to 100 percent, e.g., because they come from a specific conflict region.

⁷That nationality matters for assimilation has previously been established by, e.g., [Ran et al. \(2014\)](#). [Edin et al. \(2003\)](#) and [Damm \(2009\)](#) show that matches between one’s nationality and that of other immigrants in a locality affects the probability of successful integration. If one proxies integration with employment, locality-specific heterogeneities in integration are also evident from aggregate data. 38.1 percent of all asylum seekers with more than 9 years schooling arriving in Sweden in 2014 had an employment at the end of 2016. However, employment probability in this group varies by gender (45.3 percent for males and 25.8 percent for females), and these numbers vary between Swedish municipalities. See the web page of Statistics Sweden (www.scb.se).

seekers into consideration and, consequently, exploits the heterogeneity among localities and asylum seekers to achieve efficient and fair assignments. The employed efficiency and the fairness concepts are dynamic versions of Pareto efficiency and an extended version of the concept of envy-by-one (Budish, 2011). The latter property means that whenever some locality envies the matching of some other locality, envy can be “eliminated” by removing a single asylum seeker either from the matching of envious locality or from the matching of the envied locality.

Envy between localities occurs whenever different categories of asylum seekers are “demanded” at different rates among localities. Conflicts between localities arise if an asylum seeker is either in high demand or not demanded at all. In such cases, there must be some rule that determines which locality the asylum seeker should be matched to. This paper resolves such conflicts by means of priority and rejection orders, where the former determines the matching of highly demanded asylum seekers and the latter determines the matching of non-demanded asylum seekers. These orders constitute the backbone in the considered dynamic mechanism that matches asylum seekers to localities. The main theoretical result demonstrates that if asylum seekers are matched to localities based on the proposed mechanism, any matching of asylum seekers at any point in time is Pareto efficient, and envy is bounded by a single asylum seeker (Theorem 1).

The theoretical results presented in this paper depend on the assumption that localities can perfectly classify asylum seekers as either acceptable or unacceptable. However, even if refugee characteristics are known, such classification will have to be empirically estimated and will, therefore, suffer from misclassification error. To investigate the consequences of such error, this paper simulates and evaluates the performance of the proposed matching mechanism in specific settings calibrated to resemble the United States and the Swedish situations. These simulations show that the proposed matching mechanism outperforms uninformed mechanisms even in presence of severe misclassification error.

At least two properties make our framework different from traditional matching models. First, the sequential arrival of asylum seekers introduces dynamics. Even if dynamic problems have been considered in the literature before, e.g., in a kidney exchange (Ünver, 2009), kindergarten allocation (Kennes et al., 2014), daycare assignment (Kennes et al., 2014), and house allocation (Bloch and Cantala, 2013; Kurino, 2009), a vast majority of the papers in the matching literature focus on static matching.⁸ In static settings, the set of agents, preferences, etc. are known when solving the matching problem and are not revealed during the process as in the dynamic problem considered in this paper.

Second, in many of the traditional market design applications, the agents can provide a strict ranking over the relevant objects. In, for example, the school choice problem (Abdulkadiroğlu and Sönmez, 2003), parents have access to information about the schools in their locality and can thus form preferences over schools. In the dynamic refugee matching problem, however, it may be difficult for local authorities to form preferences over asylum seekers and vice versa, even if

⁸There is also a small literature on dynamic notions of stability, e.g., Damiano and Lam (2005), Gudmundsson (2015), Kadam and Kotowski (2018b), and Kurino (2009). See also Section 2.

asylum seekers and localities can be characterized. As a consequence, the preferences of both the localities and the asylum seekers have to be estimated based on characteristics and historical data.⁹ Such approach has recently been considered in a refugee matching context by [Andersson and Ehlers \(2016\)](#) and [Bansak et al. \(2018\)](#).¹⁰

The remainder of the paper is outlined as follows. Section 2 highlights the paper’s contributions to the recent literature on dynamic matching. Section 3 introduces the theoretical framework and some basic definitions. Section 4 describes the dynamic notions of envy-freeness and Pareto efficiency. Section 5 presents and analyzes the dynamic refugee matching mechanism. Section 6 evaluates the performance of the proposed mechanism in settings resembling the US and the Swedish situations and in the presence of misclassification error. Section 7 concludes the paper. All proofs and additional simulation results appear in the Appendices.

2 Related Matching Literature

Apart from the matching papers referred to in the above, there is also a small and quite recent literature on dynamic matching. This literature and the problems considered therein together with their associated assumptions will be described next.

One of the first papers on dynamic matching is due to [Ünver \(2009\)](#). In his investigated dynamic kidney exchange framework, a certain number of patient-donor pairs arrive at each point in time according to a Poisson process—meaning that each patient-donor type arrives with the same probability across all periods—and a specific matching is implemented among the existing patient-donor pairs in each period. The main contribution is the identification of a dynamically efficient (Markovian) matching mechanisms that minimizes the discounted cost of waiting time. The results are valid for a one-to-one dynamic matching market—as all patient-donor pairs only demand exactly one kidney and donate at most one kidney—and the functionality of the dynamic matching mechanism depends on detailed assumptions related to the arrival rates of the Poisson process.

[Kurino \(2009\)](#) considers a dynamic house allocation problem with overlapping generations, existing tenants and newcomers. In this framework, any agent lives for two periods and is a newcomer in the first period and an existing tenant in the second period. It is shown that, in general, no rule is dynamically Pareto efficient and strategy-proof, but that a seniority-based Top Trading Cycles Mechanism is dynamically Pareto efficient and strategy-proof under time invariant preferences.

⁹A different approach is taken by [Moraga and Rapoport \(2014\)](#) where preferences of asylum seekers and host countries are known a priori and taken into consideration when designing a system with “tradable immigration quotas”. Their proposed system exploits the comparative advantages of the hosting countries to efficiently match asylum seekers.

¹⁰[Haeringer and Iehlé \(2017\)](#) consider a similar approach to [Andersson and Ehlers \(2016\)](#) for deducing preferences. Their objective is, however, to obtain information on stable matchings from partial observation of preferences.

Kennes et al. (2014) model the daycare assignment problem as a dynamic version of the school choice problem in which agents enter and exit the economy over time (agents are assumed to live for two time periods). Daycare priorities are history-dependent in the sense that unmatched agents from the previous period have higher priority than newcomers. The main results show that no stable and strategy-proof mechanism exists but that a serial dictatorship mechanism is both Pareto efficient and strategy-proof.¹¹

Kadam and Kotowski (2018a) consider a dynamic matching problem over two periods with the same set of agents. They first show that dynamically stable matchings generally do not exist and then identify preference restrictions under which existence is guaranteed. Given their preference restrictions, a dynamic version of the Deferred Acceptance Mechanism is demonstrated to be strategy-proof (here, the second-period preferences are defined conditional on the first-period matching). Kadam and Kotowski (2018b) consider a generalization of this model to a finite number of time periods with the same set of agents being potentially matched in all periods, and show the existence of dynamically stable matchings under certain preference restrictions.¹²

Doval (2017) considers a general dynamic matching problem over two periods with the same set of agents where matches are irreversible and agent preferences over dynamic matchings are given by discounted utility.¹³ The adopted notions of dynamic stability and dynamic core differ from the one in Kadam and Kotowski (2018a,b). Nevertheless, one of the main results show that the dynamic core may be empty. Given this finding, different conditions that guarantee the non-emptiness of the dynamic core is presented.

A number of factors make this paper different from these matching papers. A first difference is that all theoretical results presented in this paper are independent of distributional assumptions on arrival rates and times. Second, most of the existing papers in the dynamic matching literature start by considering general models and, therefore, obtain mostly negative results for their corresponding dynamic frameworks. As a consequence, large efforts are devoted to finding additional assumptions and restrictions that lead to positive results. Third, this paper does not consider dynamically stable matchings in contrast to many of the above mentioned papers. Instead, the main focus is on dynamic notions of envy-freeness and Pareto efficiency.¹⁴ Finally, in the model considered in this paper, asylum seekers only require to be matched to one locality and localities are only active until their quotas are filled. Therefore, because the considered matching mechanism takes the characteristics of the asylum seekers into account, the size of the economy at a specific point in time is history-dependent, i.e., the size of the economy depends on the characteristics of

¹¹Kennes et al. (2015) show that the proportion of agents manipulating the dynamic Deferred Acceptance Mechanism is “small” in a large economy.

¹²See also Kotowski (2015) for a note on the dynamic stability notion used in this context. Kurino (2009) studies credible group stable matchings in a dynamic context with cardinal utilities.

¹³This model encompasses dynamic one-sided and two-sided matching markets. See Doval (2016) for an analysis of many-to-many matching markets when arrivals are stochastic.

¹⁴It can, however, be argued that Pareto efficiency, as defined in this paper, implies some sort of dynamic stability. For a recent paper on stability in a refugee matching context, see Aziz et al. (2018).

the asylum seekers that has arrived before that specific point in time.

3 The Model and Basic Definitions

This section introduces the basic ingredients of the dynamic refugee matching model and a number of important concepts and definitions.

3.1 Asylum Seekers and Localities

Asylum seekers arrive in a sequence $s = (s(1), \dots, s(n))$ to a country within a predetermined period. One can think of the period as a calendar year but the model can be applied to periods of arbitrary length.¹⁵ It is assumed that the n asylum seekers in the sequence s arrive one-by-one, meaning that the period can be divided into n (shorter) time periods where one asylum seeker arrives at each point in time $k \in \{1, \dots, n\}$. The k th asylum seeker to arrive is denoted by $s(k)$ and is sometimes referred to as arrival k .

Asylum seekers are of different categories (or “types”) describing their characteristics (e.g., in terms of age, education, spoken languages, etc.). The set of all possible categories is denoted by $T = \{t_1, \dots, t_{|T|}\}$. The category of asylum seeker $s(k)$ is denoted by $t(s(k)) \in T$. Throughout the paper, it is assumed that while the number of asylum seekers that arrive within the period (i.e., the cardinality of the sequence s) is known, how many asylum seekers of each category arriving within the period is unknown (i.e., the category distribution of the asylum seekers in the sequence s).

Localities are gathered in the set $M = \{1, \dots, |M|\}$. One can think of a locality as a state, a municipality or some other administrative subdivision of a country. Each locality $m \in M$ has a quota q_m specifying the number of asylum seekers to be assigned to the locality within the period. The quotas are gathered in the vector $q = (q_1, \dots, q_{|M|})$ and are chosen such that $\sum_{m=1}^{|M|} q_m = |s|$, i.e., such that all asylum seekers are assigned to some locality. In reality, one may not be able to predict the exact number of asylum seekers arriving during the period. If the quotas are not exhausted at the end of the period, then one may start the next period with new quotas. If the quotas are exhausted before the end of the period, then one may start a new period or an overhead (federal) facility may take any asylum seekers arriving after the first n arrivals.

Localities classify each category $t \in T$ as either acceptable or unacceptable and, consequently, also classify any given asylum seeker as either acceptable or unacceptable. This classification may be based on, e.g., historical observations and/or current labour market conditions. Formally, this means that each locality m partitions T into two disjoint sets, T_m^+ and T_m^- , where the former set contains all acceptable categories in T and the latter set contains all unacceptable categories in T . This partition is, for locality $m \in M$, denoted by $T_m = (T_m^+, T_m^-)$ where

¹⁵See Section 6.2 and the discussion following Definition 6 for additional remarks related to the choice of period.

$T_m^+ \cap T_m^- = \emptyset$ and $T_m^+ \cup T_m^- = T$. The partitions of all localities are gathered in the vector $\mathbb{T} = (T_1, \dots, T_{|M|})$.

3.2 The Dynamic Economy

Not all localities and not all asylum seekers are relevant during the entire period as asylum seekers arrive in a sequence s , and localities cannot be assigned additional asylum seekers once their quotas are filled. To formalize this, let, for any given $k \in \{1, \dots, n\}$, the set $M(k)$ contain all localities which have not yet filled their quotas when asylum seeker $s(k)$ arrives. Let the set $A(k)$ contain all asylum seekers in the set $\{s(1), \dots, s(k)\}$ who have been assigned to some locality in $M(k)$. An economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$ at arrival k contains the localities in $M(k)$ and their partitions of categories $\mathbb{T}(k) = (T_m)_{m \in M(k)}$ together with the asylum seekers in $A(k)$.

3.3 Preferences and Matchings

A bundle $x_m(k)$ contains all asylum seekers in $A(k)$ who have been matched to locality $m \in M(k)$ in economy $\mathcal{E}(k)$. A (dynamic) matching, at a given point in time k , is a vector $x(k) = (x_m(k))_{m \in M(k)}$ containing the bundles of all localities in $M(k)$. A matching is feasible, for the localities in $M(k)$, if no locality has been matched to more asylum seekers than its quota, i.e., if $|x_m(k)| \leq q_m$ for all $m \in M(k)$.¹⁶

Because localities classify each asylum seeker as either acceptable or unacceptable, any locality $m \in M(k)$ can partition all asylum seekers in a given bundle $x_i(k)$ into two disjoint sets, $A_m^+(x_i(k))$ and $A_m^-(x_i(k))$, where the former set contains all acceptable asylum seekers in bundle $x_i(k)$ and the latter set contains all unacceptable asylum seekers in bundle $x_i(k)$. Given this type of partitioning, it is assumed that a locality in $M(k)$ weakly ranks any two given bundles in economy $\mathcal{E}(k)$ based on the number of acceptable and unacceptable asylum seekers. More precisely, throughout the paper, it will be assumed that for any two bundles, $x_i(k)$ and $x_j(k)$, containing weakly fewer asylum seekers than the quota of locality m , locality m strictly prefers bundle $x_i(k)$ to bundle $x_j(k)$ if and only if the difference between the number of acceptable and the number of unacceptable asylum seekers is larger in the former bundle than in the latter, i.e., if and only if:

$$|A_m^+(x_i(k))| - |A_m^-(x_i(k))| > |A_m^+(x_j(k))| - |A_m^-(x_j(k))|. \quad (1)$$

¹⁶Note that the set $M(k)$ contains all localities that have not filled their quota when asylum seeker $s(k)$ arrives. This means that once $s(k)$ has been matched to a locality, the quota is binding for at most one locality at matching $x(k)$.

Similarly, locality m is indifferent between any two bundles if and only if:

$$|A_m^+(x_i(k))| - |A_m^-(x_i(k))| = |A_m^+(x_j(k))| - |A_m^-(x_j(k))|. \quad (2)$$

If locality $m \in M(k)$ weakly prefers bundle $x_i(k)$ to bundle $x_j(k)$, the notational convention $x_i(k)R_m(k)x_j(k)$ will be adopted to describe the relationship. The strict and indifference parts of $R_m(k)$ are denoted by $P_m(k)$ and $I_m(k)$, respectively.

A (deterministic) dynamic matching mechanism is a rule φ that for each $k \in \{1, \dots, n\}$ matches asylum seeker $s(k)$ to a locality in $M(k)$ as soon as the asylum seeker arrives.¹⁷

3.4 Notions of Demand

When the dynamic refugee matching mechanism is introduced in Section 5, it will be necessary to classify asylum seekers also in terms of aggregated demand. For this purpose, asylum seeker $s(k)$ is defined to be:

- non-demanded if $s(k)$ is unacceptable for all localities in $M(k)$,
- demanded if $s(k)$ is acceptable for exactly one locality in $M(k)$,
- overdemanding if $s(k)$ is acceptable for two or more localities in $M(k)$.

Note that the “demand status” of an asylum seeker $s(k)$ of category $t(s(k))$ may change throughout the dynamic process when localities fill their quotas. If, for example, a specific category is overdemanding in the first time period, this specific category will be demanded as soon as all but one of localities that find the category acceptable have filled their quotas. When all these localities have filled their quotas, the category will be non-demanded. However, it cannot be the case that a non-demanded (demanded) category becomes demanded or overdemanding (overdemanded) since the category partitioning \mathbb{T} is assumed to be constant throughout the entire period.

To evaluate the outcome of the simulation study in Section 6, an additional partitioning of the demanded and the overdemanding asylum seekers is needed. More precisely, and as will be explained in Section 6, much of the efficiency gains from adopting an informed matching mechanism is due to the asylum seekers that are acceptable by some but not all localities. For this reason, an asylum seeker $s(k)$ will, in Section 6, be defined to belong to the set:

- \overline{D} if asylum seeker $s(k)$ is acceptable by all localities in M ,
- D if asylum seeker $s(k)$ is acceptable by some but not all localities in M ,
- \underline{D} if asylum seeker $s(k)$ is unacceptable for all localities in M .

¹⁷In the remaining part of the paper, it is understood that only feasible dynamic mechanisms are considered. These are mechanisms which always select a feasible matching.

4 Dynamic Notions of Envy-freeness and Pareto Efficiency

This section introduces notions of fairness and efficiency in order to evaluate the dynamic refugee matching model from the previous section.

Because the distribution of categories in any given sequence of asylum seekers s is *ex-ante* unknown, it is impossible to construct a mechanism that for any sequence s always selects a matching satisfying a set of predetermined properties *ex-post* after the arrival of the last asylum seeker in the sequence. For this reason, the fairness and efficiency properties, considered in this paper, will be defined for a given economy $\mathcal{E}(k)$ at a given arrival $k \in \{1, \dots, n\}$. This also means that predictions of future arrivals are not taken into consideration in the matching process. There is a variety of plausible notions that can be used to evaluate a given matching but the analysis in this paper is restricted to two classical properties, namely envy-freeness and Pareto efficiency.

A matching is envy-free (Foley, 1967) if no locality envies any other locality at a given matching. Modified for the considered dynamic setting, this property states that a matching $x(k)$ is envy-free in a given economy $\mathcal{E}(k)$ if $x_m(k)R_m(k)x_{m'}(k)$ for any $m, m' \in M(k)$. It is well-known that envy-free matchings generally do not exist when objects are indivisible (as the asylum seekers in this problem) and in the absence of monetary transfers.¹⁸ For this reason, the notion of envy-freeness will be slightly modified following Budish (2011). More specifically, in the remaining part of the paper, matchings that satisfy envy bounded by a single asylum seeker will be considered. This property means that whenever some locality m envies some locality m' , the envy can be “eliminated” by removing a single asylum seeker either from the bundle of locality m or from the bundle of locality m' .¹⁹

Definition 1. For a given economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, a matching $x(k)$ satisfies envy bounded by a single asylum seeker if, for any $m, m' \in M(k)$, at least one of the following conditions hold:

- (i) $x_m(k)R_m(k)x_{m'}(k)$,
- (ii) there exists some asylum seeker $a \in x_m(k)$ such that $x_m(k) \setminus \{a\}R_m(k)x_{m'}(k)$,
- (iii) there exists some asylum seeker $a' \in x_{m'}(k)$ such that $x_m(k)R_m(k)x_{m'}(k) \setminus \{a'\}$.

¹⁸To see this, suppose that there are two localities and that the first asylum seeker that arrives is acceptable for both localities. In this case, the locality that not is assigned the first asylum seeker will always envy the other locality at matching $x(1)$ according to condition (1).

¹⁹Note that this definition is different from the corresponding definition in Budish (2011) since some objects (asylum seekers in this paper) may be unacceptable and, in this case, envy bounded by a single object may be achieved by removing unacceptable objects from the agent’s (localities in this paper) *own* bundle. In Budish (2011), no agent is assigned unacceptable objects and it, consequently, suffices to remove acceptable objects from the bundles of *other* agents.

The notion of envy bounded by a single asylum seeker is important in a refugee matching context, because it guarantees that asylum seekers in high demand and asylum seeker that are non-demanded will be more fairly distributed across localities. This notion is also quite weak in the sense that it does not reveal anything about the number of acceptable and unacceptable asylum seekers who are matched to a specific locality. To make this point clear, suppose that locality m experiences that it has been matched to two acceptable and three unacceptable asylum seekers and that locality m experiences that some other locality m' has been matched to five acceptable and five unacceptable asylum seekers. In this case, locality m experiences that envy is bounded by a single asylum seeker. However, by isolating the asylum seekers and by only considering acceptable asylum seekers or by only considering unacceptable asylum seekers (again from the viewpoint of locality m), it is clear that locality m envies locality m' in terms of acceptable asylum seekers (and if localities m and m' share their views on unacceptable asylum seekers, then locality m' envies locality m in terms of unacceptable asylum seekers).

To evaluate envy in matchings from the perspective of only acceptable asylum seekers, the notion of envy bounded by a single acceptable asylum seeker will be adopted. This property is satisfied for locality m in relation to locality m' , if locality m experiences that it is matched to at most one fewer acceptable asylum seeker than locality m' at a given matching $x(k)$. The notion of envy bounded by a single unacceptable asylum seeker is defined in a corresponding fashion.

Definition 2. For a given economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, a matching $x(k)$ and two localities $m, m' \in M(k)$, locality m is unenvious of locality m' in terms of acceptable asylum seekers if $|A_m^+(x_m(k))| \geq |A_m^+(x_{m'}(k))|$; and m 's envy towards locality m' is bounded by a single acceptable asylum seeker if $|A_m^+(x_m(k))| \geq |A_m^+(x_{m'}(k))| - 1$.

Definition 3. For a given economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, a matching $x(k)$ and two localities $m, m' \in M(k)$, locality m is unenvious of locality m' in terms of unacceptable asylum seekers if $|A_m^-(x_m(k))| \leq |A_m^-(x_{m'}(k))|$; and m 's envy towards m' is bounded by a single unacceptable asylum seeker if $|A_m^-(x_m(k))| - 1 \leq |A_m^-(x_{m'}(k))|$.

To evaluate efficiency properties of matchings, the notion of Pareto efficiency is adopted. In the considered dynamic setting, this property states that a matching $x(k)$ is Pareto efficient in a given economy $\mathcal{E}(k)$ if it is impossible to reallocate the asylum seekers, which have been matched to the localities in $M(k)$, among the localities in $M(k)$ in such a way that all quotas are respected, all localities are weakly better off, and at least one locality is strictly better off.

Definition 4. For a given economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, a matching $x(k)$ is Pareto efficient if it is impossible to reallocate the asylum seekers in $A(k)$ among the localities in $M(k)$ to obtain a new matching $x'(k)$ where the quotas are respected for all localities in $M(k)$, $x'_m(k)R_m(k)x_m(k)$ for all $m \in M(k)$, and $x'_m(k)P_m(k)x_m(k)$ for some $m \in M(k)$.

The notion of Pareto efficiency is important in the refugee matching context as it guarantees that

an asylum seeker that is acceptable for some locality in $M(k)$ never can be matched to a locality in $M(k)$ that considers the asylum seeker as unacceptable.

5 Dynamic Order Mechanisms

This section introduces a dynamic matching mechanism and demonstrates that any matching selected by the proposed mechanism is Pareto efficient and satisfies envy bounded by a single asylum seeker. The basic observation is that asylum seekers must be matched to some locality directly upon arrival, and therefore conflicts between localities may arise because some asylum seekers are non-demanded and some asylum seekers are overdemanded. In either case, there must be some rule determining the locality that the asylum seeker should be matched to. These conflicts are resolved by means of priority and rejection orders where the former is used to resolve conflicts when an asylum seeker is overdemanded and the latter is adopted to match non-demanded asylum seekers to localities. Formally, an order τ is, for a given $M(k)$, a list $(\tau_1, \dots, \tau_{|M(k)|})$ such that $\{\tau_1, \dots, \tau_{|M(k)|}\} = M(k)$. This means that each locality, which has not yet filled its quota at arrival k , i.e., each locality in the set $M(k)$, is assigned a unique position in the list τ .

A priority structure π specifies for each arrival k an order in which overdemanded asylum seekers are assigned, i.e., it is a list of orders $\pi = (\pi(k))_{k=1}^n$ where $\pi(k)$ is an order of $M(k)$ for any given point in time $k \in \{1, \dots, n\}$. This means that if asylum seeker $s(k)$ is overdemanded, then the locality $\pi_1(k)$ has the highest priority over asylum seeker $s(k)$ among all localities in $M(k)$, the locality with $\pi_2(k)$ has the second highest priority among all localities in $M(k)$, and so on.

A rejection structure σ specifies for each arrival k an order in which non-demanded asylum seekers are assigned, i.e., it is a list of orders $\sigma = (\sigma(k))_{k=1}^n$ where $\sigma(k)$ is an order of $M(k)$ for any given point in time $k \in \{1, \dots, n\}$. This means that if asylum seeker $s(k)$ is non-demanded, then all localities in $M(k)$ find asylum seeker $s(k)$ unacceptable and $s(k)$ is assigned to locality $\sigma_1(k)$.

An order structure is a pair (π, σ) where π is a priority structure and σ is a rejection structure. Throughout it is assumed that the structure $(\pi(1), \sigma(1))$ is exogenously given, i.e., that the structure is in place when the first asylum seeker $s(1)$ arrives. The structure $(\pi(1), \sigma(1))$ may, for example, be chosen randomly or based on inheritance from the previous period (see also Section 6.2). A dynamic order mechanism is used to resolve conflicts of the above mentioned type.

Definition 5. A (dynamic) order mechanism φ is given by an order structure (π, σ) such that for $k \in \{1, \dots, n\}$ it selects a matching $x(k)$ where asylum seeker $s(k)$ is assigned to:

- (i) the locality in $M(k)$ with the highest position in $\sigma(k)$ if $s(k)$ is non-demanded,
- (ii) the only locality in $M(k)$ which finds $s(k)$ acceptable if $s(k)$ is demanded,

- (iii) the locality in $M(k)$ with the highest position in $\pi(k)$ which finds $s(k)$ acceptable if $s(k)$ is overdemanded.

Until this point, orders have been generally defined and it has not explicitly been stated if and how orders should be “updated” between any two arrivals k and $k + 1$. Given the interest in order mechanisms and matchings where envy is bounded by a single asylum seeker, a first observation is that the priority order *and* the rejection order must be updated between any two arrivals.

Example 1. Let $M(1) = \{m_1, m_2\}$ and suppose that $(\pi(1), \sigma(1))$ are such that locality m_1 has the highest position in $\pi(1)$ and locality m_2 has the highest position in $\sigma(1)$. Assume further that matchings are selected by an order mechanism where *only* the priority order is updated when an overdemanded asylum seeker is matched to a locality and *only* the rejection order is updated when a non-demanded asylum seeker is matched to a locality. Then if asylum seeker $s(1)$ is acceptable for both localities, $s(1)$ must be matched to locality m_1 , and if asylum seeker $s(2)$ is non-demanded, $s(2)$ must be matched to locality m_2 . But then envy is not bounded by a single asylum seeker at matching $x(2)$. The same conclusion holds if asylum seeker $s(1)$ is non-demanded and asylum seeker $s(2)$ is acceptable for both localities. \square

The above example demonstrates that if localities are not sufficiently “rewarded” when matched to a non-demanded asylum seeker and are not sufficiently “penalized” when matched to an overdemanded asylum seeker, envy need not generally be bounded by a single asylum seeker. To find the right compromise between rewards and penalties, the notion of rotation will be introduced.

To formalize the idea of rotation, suppose that asylum seeker $s(k)$ is matched to locality $m \in M(k)$. Let also $\sigma^O(k)$ be the position of the locality with the highest position in $\sigma(k)$ envying locality m at matching $x(k)$, and let $\sigma^O(k) = |M(k)|$ if no such locality exists. Let further $\pi^N(k)$ be the position of the locality with the highest position in $\pi(k)$ which is envied by locality m at matching $x(k)$, and let $\pi^N(k) = |M(k)|$ if no such locality exists.

Definition 6. Consider a given economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$. An order structure (π, σ) satisfies rotation if for any $k \in \{1, \dots, n - 1\}$, asylum seeker $s(k)$ is assigned to locality $m \in M(k)$, then:

- (i) all localities in $M(k) \setminus \{m\}$ have the same positions among themselves in $\pi(k + 1)$ and in $\pi(k)$,
- (ii) all localities in $M(k) \setminus \{m\}$ have the same positions among themselves in $\sigma(k + 1)$ and in $\sigma(k)$,
- (iii) if locality m has not filled its quota after being assigned asylum seeker $s(k)$, then:

- (a) if asylum seeker $s(k)$ is non-demanded, locality m gets the lowest position in the rejection order $\sigma(k+1)$ and locality m is placed immediately before $\pi^N(k)$ in the priority order $\pi(k+1)$ if $\pi^N(k)$ is smaller than the position of m in $\pi(k)$ and otherwise $\pi(k+1) = \pi(k)$ (i.e. m keeps the same position in $\pi(k+1)$ as in $\pi(k)$),
- (b) if asylum seeker $s(k)$ is demanded, locality m has the same position in priority orders $\pi(k+1)$ and $\pi(k)$ as well as in the rejection orders $\sigma(k+1)$ and $\sigma(k)$,
- (c) if asylum seeker $s(k)$ is overdemand, locality m gets the lowest position in the priority order $\pi(k+1)$ and locality m is placed immediately before $\sigma^O(k)$ in the rejection order $\sigma(k+1)$ if $\sigma^O(k)$ is smaller than the position of m in $\sigma(k)$ and otherwise $\sigma(k+1) = \sigma(k)$ (i.e. m keeps the same position in $\sigma(k+1)$ as in $\sigma(k)$).

Before illustrating an order mechanism that satisfies rotation (Examples 2 and 3), two remarks are in order. Note first that the priority order and, consequently, the order mechanism and the concept of rotation does not make any separation between asylum seeker categories, i.e., the same priority order is considered for each category in $\mathbb{T}(k)$. Both the order mechanism and the concept of rotation can be generalized by specifying a priority order for each category in $\mathbb{T}(k)$, i.e., a (generalized) priority order $(\pi_t(k))_{t \in \mathbb{T}(k)}$ for each $k \in \{1, \dots, n\}$. In such case, it is important to update each priority order in $(\pi_t(k))_{t \in \mathbb{T}(k)}$ between any two arrivals (otherwise, it is easy to see envy by a single asylum seeker may be violated). This can be achieved by identifying a vector of priority order specific constants $\pi^N(k) = (\pi_t^N(k))_{t \in \mathbb{T}(k)}$ for each arrival $k \in \{1, \dots, n\}$ exactly as in the above, and by adjusting the concept of rotation accordingly. For example, Part (a) in Definition 6(iii) needs to be written as:

- (a') if asylum seeker $s(k)$ is non-demanded, locality m gets the lowest position in the rejection order $\sigma(k+1)$ and for each $t \in \mathbb{T}(k)$, locality m is placed immediately before $\pi_t^N(k)$ in the priority order $\pi_t(k+1)$ if $\pi_t^N(k)$ is smaller than the position of m in $\pi_t(k)$ and otherwise $\pi_t(k+1) = \pi_t(k)$ (i.e. m keeps the same position in $\pi_t(k+1)$ as in $\pi_t(k)$).

This paper considers the less general approach without loss of generality since all results and conclusions presented in the paper continue to hold also for the more general setting. The reason for considering a less general framework is simply to avoid introducing additional notation. Moreover, because *all* priority orders and the rejection order must be updated between any two arrivals, based on envy, it follows that all priority orders will be identical after “sufficiently many” arrivals. Hence, the more general approach only makes a (small) difference in the very beginning of the dynamic process.

Note also that localities with smaller quotas will typically fill their quotas before localities with larger quotas so it may well be the case that certain localities are only active during a small part of the sequence. If the latter issue is seen as a problem, then one can resolve it by splitting the period into multiple shorter periods (e.g., instead of considering yearly quotas, one may split them in trimester quotas; see Section 6.2).

Example 2. This example illustrates how the constants $\pi^N(k)$ and $\sigma^O(k)$ should be interpreted in Definition 6. Suppose that the matching is selected by an order mechanism satisfying rotation, and that the priority order $\pi(k)$ is given by $m_4 < m_2 < m_1 < m_3$ and rejection order $\sigma(k)$ by $m_1 < m_2 < m_3 < m_4$. If asylum seeker $s(k)$ is non-demanded, then $s(k)$ is matched to locality m_1 since m_1 has the highest position in the rejection order $\sigma(k)$. If locality m_1 only envies locality m_3 after being matched to asylum seeker $s(k)$, it follows that $\pi^N(k) = 4$. Because the position of m_1 is smaller than 4 in $\pi(k)$, $\pi(k+1) = \pi(k)$ and $\sigma(k+1)$ is given by $m_2 < m_3 < m_4 < m_1$. If instead asylum seeker $s(k)$ is acceptable by localities m_1 and m_4 , asylum seeker $s(k)$ is overdemanding and, consequently, matched to locality m_4 since locality m_4 has the highest position in the priority-ordering $\pi(k)$ among all localities which find $s(k)$ acceptable. Thus, $\pi(k+1)$ is given by $m_2 < m_1 < m_3 < m_4$. If locality m_1 then envies locality m_4 , it follows that $\sigma^O(k) = 1$ and $\sigma(k+1)$ is given by $m_4 < m_1 < m_2 < m_3$. \square

Example 3. To illustrate an order mechanism satisfying rotation, suppose that $M(1) = \{m_1, m_2, m_3\}$ and $s = (1, \dots, 10, \dots)$. Assume further that the quotas for all localities in $M(1)$ are greater than 10. This assumption means that we do not need to take into account quotas during the first 10 arrivals. Let the initial priority order $\pi(1)$ and rejection order $\sigma(1)$ both be given by $m_1 < m_2 < m_3$. Suppose further that the “demand matrix” of the localities for the 10 first arrivals is given by Table 1 (the numbers 0 and 1 indicate that an asylum seeker is unacceptable and acceptable, respectively).

Table 1: Acceptable and unacceptable asylum seekers for the localities in Example 1.

s	1	2	3	4	5	6	7	8	9	10
m_1	0	0	1	1	0	0	1	0	0	1
m_2	0	0	1	1	0	0	1	1	0	1
m_3	0	0	1	1	0	0	0	1	0	1

Because asylum seeker $s(1)$ is non-demanded and locality m_1 has the highest position in the rejection order $\sigma(1)$, asylum seeker $s(1)$ is matched to locality m_1 . Consequently, locality m_1 envies both locality m_2 and m_3 at matching $x(1)$ and because locality m_2 has a higher position than locality m_3 in the priority order $\pi(1)$ (i.e., position 2), it follows that $\pi^N(1) = 2$. Hence, $\pi(2)$ is given by $m_1 < m_2 < m_3$. Furthermore, locality m_1 is placed at the bottom of the rejection order $\sigma(2)$, i.e., $\sigma(2)$ is given by $m_2 < m_3 < m_1$. The entire process for the first 10 arrivals is described in Table 2. From this table, it follows that, at matching $x(10)$, asylum seekers $s(1)$, $s(3)$, $s(6)$ and $s(7)$ are matched to locality m_1 , asylum seekers $s(2)$, $s(4)$, $s(9)$ and $s(10)$ are matched to locality m_2 , and asylum seekers $s(5)$ and $s(8)$ are matched to locality m_3 . \square

The first result establishes that any order mechanism satisfying rotation always selects a Pareto efficient matching where envy is bounded by a single asylum seeker.

Theorem 1. Let $\varphi = (\pi, \sigma)$ be an order mechanism where (π, σ) satisfies rotation. For each $k \in \{1, \dots, n\}$, if matching $x(k)$ is selected by φ for economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, then:

- (i) $x(k)$ satisfies envy bounded by a single asylum seeker, and
- (ii) $x(k)$ is Pareto efficient.

From Theorem 1, it follows almost immediately that each locality in $M(k)$ must be unenvious of any other locality in $M(k)$ either in terms of acceptable or unacceptable asylum seekers (or both) because if this is not the case, envy cannot be bounded by a single asylum seeker. If, for example, locality m envies locality m' by exactly *one* acceptable asylum seeker (the “minimal amount” of envy), then envy towards municipality m' will not be bounded by a single asylum seeker for any positive amount of envy in terms of unacceptable asylum seekers. In this situation, it thus follows that locality m must be unenvious of locality m' in terms of unacceptable asylum seekers. This finding is formally stated in the following corollary to Theorem 1.

Table 2: Description of an order mechanism that satisfies rotation for Example 3.

k	$s(k)$ matched to m_i	$\sigma^O(k)$	$\pi^N(k)$	$\sigma(k)$	$\pi(k)$
1	m_1	–	2	$m_1 < m_2 < m_3$	$m_1 < m_2 < m_3$
2	m_2	–	3	$m_2 < m_3 < m_1$	$m_1 < m_2 < m_3$
3	m_1	3	–	$m_3 < m_1 < m_2$	$m_1 < m_2 < m_3$
4	m_2	3	–	$m_3 < m_1 < m_2$	$m_2 < m_3 < m_1$
5	m_3	–	2	$m_3 < m_1 < m_2$	$m_3 < m_1 < m_2$
6	m_1	–	3	$m_1 < m_2 < m_3$	$m_3 < m_1 < m_2$
7	m_1	3	–	$m_2 < m_3 < m_1$	$m_3 < m_1 < m_2$
8	m_3	3	–	$m_2 < m_3 < m_1$	$m_3 < m_2 < m_1$
9	m_2	–	2	$m_2 < m_3 < m_1$	$m_2 < m_1 < m_3$
10	m_2	3	–	$m_3 < m_1 < m_2$	$m_2 < m_1 < m_3$
11				$m_3 < m_1 < m_2$	$m_1 < m_3 < m_2$

Corollary 1. Let $\varphi = (\pi, \sigma)$ be an order mechanism where (π, σ) satisfies rotation. For each $k \in \{1, \dots, n\}$, if matching $x(k)$ is selected by φ for economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, then for any two localities $m, m' \in M(k)$, it holds that locality m is unenvious of locality m' in terms of acceptable asylum seekers and/or unenvious of locality m' in terms of unacceptable asylum seekers.

Since an order that satisfies rotation $(\pi(k), \sigma(k))$ is updated based on the selected matching $x(k)$, it is possible to identify some very specific “envy patterns” in any structure $(\pi(k+1), \sigma(k+1))$ as reported in the following theorem.

Theorem 2. Let $\varphi = (\pi, \sigma)$ be an order mechanism where (π, σ) satisfies rotation. For each $k \in \{1, \dots, n\}$, if matching $x(k)$ is selected by φ for economy $\mathcal{E}(k) = (M(k), \mathbb{T}(k), A(k))$, then locality $m \in M(k)$ does not envy any locality with a higher position in $\pi(k+1)$ or any locality with a lower position in $\sigma(k+1)$.

Finally, as noted in Section 4, the notion of envy bounded by a single asylum seeker is quite weak since it does not say anything about the number of acceptable and the number of unacceptable asylum seekers contained in the bundle of locality m in the matching $x(k)$ and, furthermore, how these numbers relate to the corresponding numbers of the other bundles in the matching $x(k)$ (from the perspective of locality m). The above two results are also silent about such relation. Ideally, we would like to prove that, at any matching $x(k)$ selected by an order mechanism that satisfies rotation, envy is always bounded by a single acceptable asylum seeker *and* a single unacceptable asylum seeker. This means, for example, that if locality m at matching $x(k)$ experiences that it has been assigned two acceptable (three unacceptable, respectively) asylum seekers, then it cannot be the case that locality m experiences that some other locality m' at matching $x(k)$ has been assigned more than three acceptable (less than two unacceptable, respectively) asylum seekers.²⁰ We have, however, been unable to theoretically prove such result. Based on the simulations presented in Section 6 and Appendix B, we conjecture that this result holds since the maximum envy measure of acceptable and unacceptable asylum seekers never exceeds one in any of the 100 million observed envy measures.

6 Implementation: Misclassification and Quotas

The proposed order mechanism is simple to implement given locality-specific quotas and locality specific partitions of categories $A_m^+(x_i(k))$ and $A_m^-(x_i(k))$. However, the choice of these inputs has two practical implications. First, locality preferences need to be estimated, and thus locality-specific partitions will suffer from misclassification error. Second, a centralized clearinghouse can aggregate or split quotas over the period to affect whether localities with low quotas remain in the economy for multiple shorter spells, or for fewer longer ones.

This section begins by simulating and evaluating the performance of the proposed matching mechanism in several specific settings where these implications are likely to matter. Our simulations show that the proposed mechanism outperforms the uninformed naïve sequential matching mechanism (defined below) even in presence of severe misclassification error. Furthermore, with random flows of asylum seekers, the choice of splitting quotas across shorter time periods does

²⁰Note that this result does not follow from Theorem 1 and Corollary 1. The statements in these results are not violated if, e.g., locality m is matched to two acceptable asylum seekers and three unacceptable asylum seekers at matching $x(k)$ and locality m experiences that locality m' is matched to five acceptable and five unacceptable asylum seekers at matching $x(k)$. In this case, envy is bounded by a single asylum seeker (Theorem 1) and locality m is also unenvious of locality m' in terms of unacceptable asylum seekers (Corollary 1). However, envy is not bounded by a single acceptable asylum seeker.

not affect the performance of the mechanism if every time period is treated independently. However, splitting quotas can improve fairness at minor efficiency costs if measures of envy across municipalities are carried forward from period to period.

The section concludes by providing simulation results supporting our conjecture, stated in Section 5, that envy is always bounded *both* by a single acceptable asylum seeker *and* by a single unacceptable asylum seeker whenever the matching is selected by an order mechanism satisfying rotation.

6.1 Misclassification Error in Locality Partitions

Estimating locality-specific partitions of categories $A_m^+(x_i(k))$ and $A_m^-(x_i(k))$ is equivalent to predicting a demand matrix such as the one in Table 1, using asylum seeker characteristics as independent variables and desirable outcomes proxying integration (e.g., predicted probability of future employment) as dependent variables. For example, Bansak et al. (2018) use a machine learning algorithm to estimate such a matrix, using refugee data from the United States and Switzerland and allowing for locality-specific synergies with heterogeneous refugees. Nonetheless, out-of-sample prediction implies at least some misclassification error between the real and the estimated locality-specific partitions. Specifically, Bansak et al. (2018) report 24 percent and 14 percent misclassification error for their preferred models, estimated on US and Swiss data, respectively.²¹

This section assesses the sensitivity of the proposed mechanism to misclassification error by simulating random flows of asylum seekers calibrated to match aggregate data for the United States and Sweden. For each of these simulated flows, an increasingly high amount of misclassification error in the locality-specific partitions is introduced. The evolution of the measures of envy-freeness and efficiency are then compared through the dynamic matching process. The performance of the proposed mechanism is benchmarked against a naïve (and uninformed) sequential matching mechanism, which assigns asylum seekers according to their arrival order to a fixed list of rotating localities. By ensuring that, in the absence of quotas, each locality receives a similar amount of asylum seekers, a sequential matching mechanism tends to outperform truly random matching mechanisms such as the ones adopted in, e.g., the United States, Switzerland and Sweden. The naïve sequential matching mechanism, consequently, provides a more severe benchmark for the proposed mechanism than a truly random mechanism (simulation results for the truly random mechanism appear in Appendix B).

As explained in Section 3, asylum seekers are partitioned into three different sets in the analysis, i.e., the asylum seekers that are considered acceptable by all localities (\bar{D}), the asylum seekers that are acceptable by some but not all localities (D), and the non-demanded asylum seekers (\underline{D}). The asylum seekers in the set D thrive best in some environment and not in others,

²¹Note that misclassification error refers to observations wrongly (and not randomly) classified, and thus, for a binary variable, 50 percent misclassification error represents a zero-predictive-power scenario.

representing the empirical synergies documented by, e.g., [Bansak et al. \(2018\)](#), [Damm \(2009\)](#) and [Edin et al. \(2003\)](#). The simulations allow for autocorrelation of preferences across localities, so each asylum seeker in group D is acceptable by either a majority (on average 75 percent) or a minority (on average 25 percent) of the localities.

The proposed mechanism exploits the synergies of asylum seekers with heterogeneous localities to perform efficient and unenvious matchings. Therefore, relative to the naïve sequential matching mechanism, it produces larger efficiency gains and less envious matchings as the proportion of asylum seekers in the set D increases in the sample.²² Acceptability is approximated with employment, and the simulated flows of asylum seekers are calibrated to match the aggregate numbers from [Bansak et al. \(2018\)](#) for the United States. For Sweden, these numbers are based on the 2016 employment outcomes of the 2013 wave of asylum seekers. More precisely, 28 percent (45 percent) of the asylum seekers are assumed to belong to the set \bar{D} (the set \underline{D}). This number represents the lowest (highest) employment rate in any of the 21 Swedish regional entities of the 2013 wave.

Figure 1 compares the performance of the proposed mechanism (AEM) with increasing degrees of misclassification error in the locality-specific partitions with that of the naïve sequential matching mechanism, in terms of both envy and efficiency loss. The proportion of localities in the sample envying some locality by at least one and at least five asylum seekers are used as measures of envy. If the locality-specific partitions are correct, Theorem 1 shows that the proposed mechanism guarantees that envy is bounded by a single asylum seeker. Remarkably, Panel 1a shows that the share of localities envying another locality by at least one asylum seeker decreases with the number of asylum seekers per locality, and converges towards zero. Even with misclassification error, Panels 1a and 1b in Figure 1 show that the proposed matching mechanism always outperforms the naïve sequential matching mechanism in terms of envy. With the 24 percent misclassification error of [Bansak et al. \(2018\)](#), in the Swedish case the proposed mechanism guarantees already after 1,000 arrivals a decrease of 23 percent in the share of envying municipalities, and of 45 percent in the share of envying municipalities by more than five asylum seekers (here it should be noted that there were around 25,000 asylum seekers in Sweden in 2016 and 2017).²³


In Panel 1c in Figure 1 the efficiency loss is measured as the share of asylum seekers that, while acceptable for some localities, end up in a locality that considers them unacceptable. This measure reflects the preferences of asylum seekers of achieving the best possible outcomes for their future, and is in a Pareto efficient matching always equal to zero. Even with misclassifi-

²²In the extreme case in which all asylum seekers are either always acceptable or non-demanded, the naïve sequential matching mechanism produces matchings that are always Pareto efficient (however, envy does not need to be bounded by a single asylum seeker).


²³The corresponding numbers for the US case are 7 and 30 percent respectively. The difference is primarily due to the higher number of localities in the US, causing envy measures with misclassification error to decrease more slowly. A replication of these measures for different combinations of \bar{D} and \underline{D} proportions for 21 localities appears in Appendix B.

Figure 1: The proposed mechanism's sensitivity to misclassification error

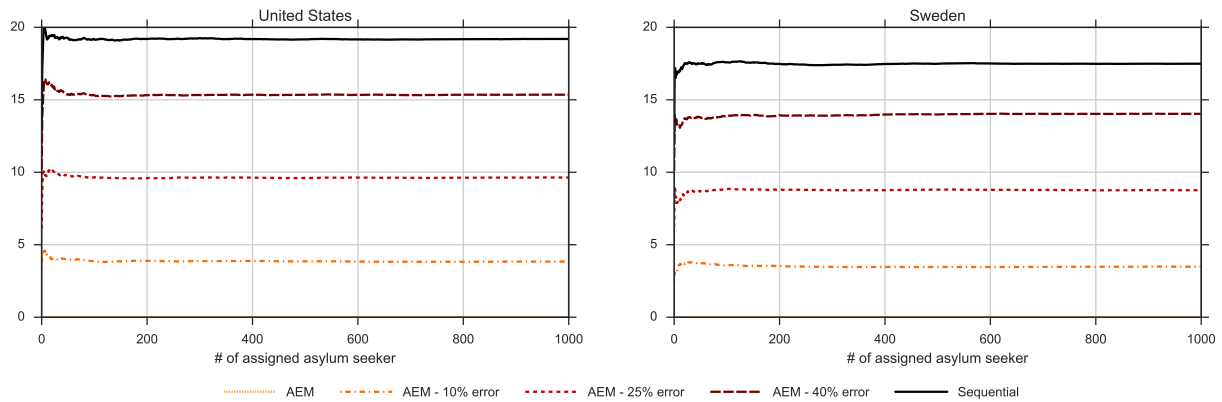
(a) Percentage of localities envying some locality by at least one asylum seeker



(b) Percentage of localities envying some locality by at least five asylum seekers



(c) Percentage of mismatched acceptable asylum seekers



NOTES: A thousand waves of asylum seekers of 1,000 asylum seekers each is simulated, distributed across 51 locations in the United States and the 21 Swedish regional entities (län). The measures of efficiency loss and envy are computed at each matching, comparing the performance of the proposed mechanism (with increasing degrees of misclassification error) with that of the naïve sequential matching mechanism. The flows of asylum seekers are calibrated to match the numbers reported by Bansak et al. (2018) for the United States (\bar{D} : 34 percent; \underline{D} : 39 percent) and average 2016 employment outcomes of the 2013 asylum seeker wave in Sweden (\bar{D} : 28 percent; \underline{D} : 45 percent).

cation error, the proposed mechanism produces strong efficiency gains with respect to the naïve sequential matching mechanism. With the 24 percent misclassification error of [Bansak et al. \(2018\)](#), the proposed mechanism decreases the share of mismatched acceptable asylum seekers by about 50 percent.

6.2 Choice of Quotas

The presence of heterogeneous quotas across localities introduces a further seasonality-related concern. For example, in Sweden each locality has by law the obligation of reserving housing units for asylum seekers (Swedish Law, [SFS:2016:38, 2016](#)). Yearly quotas coupled with the proposed mechanism will lead smaller localities to exhaust all their housing obligation in a few months, and thus will be obliged to provide housing units only at a given time of the year. Instead, larger localities will have to provide housing throughout the year. A solution to this disparity would be splitting yearly quotas in, e.g., trimester quotas, and dynamically match asylum seekers within each trimester. This section shows that such strategy is feasible and, as with random flows of asylum seekers, the performance of the proposed mechanism does not suffer from splitting quotas across shorter time periods.

A thousand random flows of asylum seekers is simulated and calibrated as in [Figure 1](#) for Sweden using the 2017 quotas.²⁴ For each flow of asylum seekers, quotas are split up in yearly, semestral, trimestral and monthly intervals.²⁵ For each interval, the matching mechanism restarts either blindly (i.e., resetting the aggregated envy to at zero at the start of the considered sub-period) or by initiating the mechanism in every sub-period after the first with the envy matrix calculated at the end of the previous sub-period for all localities (including those which filled their quotas). Therefore, even with no misclassification error in the partition, it is possible that some localities, which filled their quotas early in the sub-period, envy (or are envied by) others by more than one asylum seeker at the beginning of subsequent sub-periods. The process of transferring envy measures across periods is referred to as envy legacy. The measures of envy and efficiency loss are then compared at the end of the year, when the flow of asylum seekers has been completely assigned. All results are benchmarked against those obtained by the naïve sequential matching mechanism.

[Table 3](#) reports measures of efficiency and fairness loss computed at the end of the assignment year for each of the simulated asylum seeker flows across all localities, including those dropping out of the mechanism once their quota was filled. This exercise reveals three relevant properties of the proposed order mechanism when constrained by heterogeneous quotas. First, while the proposed mechanism can produce efficiency and fairness losses due to unequally sized quotas, it still strongly outperforms the standard naïve sequential matching mechanism. Second, with

²⁴These quotas are available at the website of the Swedish Migration Board (www.migrationsverket.se).

²⁵In 2017, the sum of all quotas was 23,600. For simplicity, each yearly regional quota is rounded to a number divisible by 12, and thus flows of 23,568 asylum seekers are considered.

Table 3: Performance of the assignment algorithm if quotas are split over multiple sub-periods

	Percentage of mismatched acceptable asylum seekers				Number of localities envying some locality by at least one asylum seeker			
	Year	Semester	Trimester	Month	Year	Semester	Trimester	Month
Sequential	22.97 (0.30) [22.96]	22.94 (0.32) [22.95]	22.96 (0.29) [22.95]	22.96 (0.30) [22.97]	19.95 (0.63) [20.00]	19.93 (0.65) [20.00]	19.90 (0.65) [20.00]	19.91 (0.61) [20.00]
AEM	7.11 (0.22) [7.11]	7.10 (0.21) [7.10]	7.11 (0.21) [7.10]	7.12 (0.21) [7.12]	1.24 (0.43) [1.00]	1.19 (0.39) [1.00]	1.22 (0.41) [1.00]	1.19 (0.39) [1.00]
AEM with envy legacy	7.11 (0.22) [7.11]	7.20 (0.22) [7.21]	7.28 (0.22) [7.27]	7.33 (0.23) [7.33]	1.24 (0.43) [1.00]	1.00 (0.03) [1.00]	1.00 (0.00) [1.00]	1.00 (0.00) [1.00]

NOTES: A thousand waves of 23,568 asylum seekers each is simulated, distributed across 21 Swedish regional entities (län) according to 2017 Swedish regional quotas. The table reports the average, standard deviation (in parentheses) and median (in brackets) of fairness and efficiency loss measures calculated after the assignment of the 23,568th asylum seeker for (i) the proposed mechanism restarted blindly, (ii) the proposed mechanism allowing for dynamic legacies in envy across sub-periods, and (iii) the naïve sequential matching mechanism. Because the number of localities is constant, the table reports the number instead of the share of envying localities as fairness loss measure. Flows of asylum seekers are calibrated to match average 2016 employment outcomes of the 2013 asylum seeker wave in Sweden (\bar{D} : 28 percent; \underline{D} : 45 percent). Autocorrelation of locality preferences is allowed for each asylum seeker.

random refugee flows, partitioning the quotas across subperiods and restarting the mechanism each time with zero envy across localities does not affect its performance. Third, allowing for envy matrices to be transferred across subperiods can improve fairness at a small efficiency loss.

While this last property emerges in this particular manner for the specific Swedish situation, examining the sources of this trade-off delivers an intuitive understanding of the implicit trade-off between efficiency and fairness created by heterogeneous quotas. The largest regional entity in Sweden receives alone 30% of all asylum seekers in a given year, with the second and third largest localities receiving 16% and 12%, respectively. Such inequality in quotas imply that the largest localities will be the last to fill their quotas, and thus a mechanism satisfying rotation will always assign to those localities the last thousands asylum seekers, independently of their type. Because the calibrations of the asylum seeker flows imply that a random refugee is more likely to be unacceptable than acceptable for a given locality, it follows that the largest regional entities will accumulate a lot of envy towards smaller localities.

With yearly quotas, the largest entity envies some other smaller locality in every simulation, and the second or third largest 24% of the time. By partitioning yearly quotas into smaller intervals and allowing for envy legacy across subperiods, municipalities that have accumulated lots of envy in the previous subperiods are more likely to be at the first places in the priority order. Therefore, they receive relatively more acceptable asylum seekers in the following period, and fill their quotas slightly faster. As a consequence, the second and third largest locality never envy any other locality anymore. However, as they fill their quotas earlier, a longer tail flow of asylum seekers can only be assigned to the largest locality at the end of the period, thereby slightly reducing efficiency.

This process highlights that order mechanisms satisfying rotation with heterogenous locality-asylum seeker synergies perform better the longer a large enough amount of localities remains in the market. A higher number of available localities implies a higher likelihood that one of them finds an asylum seeker acceptable, thereby—by exploiting informed assignment—simultaneously improving efficiency and fairness.

6.3 The Envy Bounded by One Conjecture

As discussed in Section 5, we conjecture that envy is always bounded by a single acceptable asylum seeker *and* by a single unacceptable asylum seeker whenever the matching is selected by an order mechanism that satisfies rotation. To provide some empirical evidence for this conjecture, 5,000 waves of 1,000 seekers each are simulated, distributed across 20 localities, assuming that the characteristics of the refugee flow are random. At each assigned asylum seeker, the maximum envy a locality has towards the others in the market is calculated. This simulation gives a total of 100 million envy measures by locality. Envy measures are split across envy on acceptable asylum seekers (another locality has more asylum seekers considered acceptable) and envy on unacceptable asylum seekers (another locality has fewer asylum seekers considered unacceptable).

Table 4: Simulation results for the conjecture that envy is bounded by one in acceptable and unacceptable asylum seekers.

	p10	p50	p90	max	p10	p50	p90	max
Assigned asylum seekers	<i>Acceptable asylum seekers</i>				<i>Unacceptable asylum seekers</i>			
1-100	0	0	1	1	0	0	1	1
101-500	-2	0	1	1	-1	0	1	1
501-1000	-6	-1	1	1	-5	0	1	1

NOTES: Five thousand waves of one thousand asylum seekers each is simulated, distributed across 20 regional entities, delivering a total of 100 million locality-specific maximum envy measures (the highest envy towards any of the remaining 19 localities). The table reports the maximum, 90th, 50th, and 10th percentile of maximum envy at specific intervals of asylum seeker arrival. Envy on acceptable asylum seekers is separated from envy on unacceptable asylum seeker. The characteristics of the refugee flow (proportions of \bar{D} and \underline{D} asylum seeker and autocorrelation of preferences across municipalities) is randomized for each asylum seeker flow.

Table 4 shows the 10th, 50th and the 90th percentile of observed maximum envy across these 100 million observations, arranged by three intervals of number of assigned asylum seeker in the flow (1–100, 101–500, and 501–1,000). As can be seen from the table, the conjecture is not violated a single time. Moreover, the low percentiles of envy measures are decreasing with the number of arrivals, implying that the likelihood of a locality to be truly unenvious increases with the number of assigned asylum seekers.

7 Conclusions

This paper has investigated dynamic refugee matching as a market design application. A specific matching mechanism was proposed and it has been demonstrated that any matching selected by the proposed mechanism is Pareto efficient and satisfies envy bounded by a single asylum seeker. This theoretical finding hinges on the assumption that there is no misclassification error. For example, [Bansak et al. \(2018\)](#) report 24 percent misclassification error for their preferred refugee assignment models estimated on US data. However, even with a 24 percent misclassification error, the proposed mechanism (i) decreases the share of mismatched acceptable asylum seekers by around 50 percent and (ii) reduces the share of envying localities by between 7 and 45 percent.

This paper does not empirically estimate prediction functions, as we currently do not have access to the necessary data. By exploiting refugee data from the United States and Switzerland, [Bansak et al. \(2018\)](#) estimate locality-specific partitions by a machine learning algorithm, allowing for locality-specific synergies with heterogeneous refugees. Because their algorithm focuses on minimizing prediction error rather than uncovering structural parameters determining labor market success ([Mullainathan and Spiess, 2017](#)), machine learning is a particularly appropriate tool in this context. Nonetheless, such an unstructured approach highlights the low external validity of the estimated prediction function, which therefore needs to be re-estimated and recalibrated in different scenarios and as new data becomes available. The necessity to update the prediction function has the benefit of alleviating the problem of displacement and general equilibrium effects due to the field implementation of mechanisms as the one proposed in this paper ([Crépon et al., 2013](#)). If the labor demand for asylum seekers is very inelastic, then even an efficient assignment does not guarantee employment for asylum seekers predicted as acceptable. In practice, this mismatch would increase prediction error in our predictions. However, re-estimating the prediction function with updated data would capture these effects and return a better predicted “demand matrix” in a later stage.

A more theoretical remark is that the proposed mechanism assumes that asylum seekers arrive one-by-one. This means that the current form of the mechanism does not take families of asylum seekers into account (see, e.g., [Andersson and Ehlers, 2016](#), for a mechanism that keeps families intact in house allocation problems with asylum seekers). To also account for families, each asylum seeker can be described to be of a specific size representing the number of family members and some (limited) flexibility in the quotas must be introduced. In this case, envy cannot be guaranteed to be bounded by a single asylum seeker but is instead bounded by the maximal family size in the sequence.

Finally, this paper has abstracted from issues related to manipulability. It is well-established in the mechanism design literature that agents (localities in this paper) often cannot be prevented from manipulating mechanisms in their advantage by misrepresenting preferences. In the framework investigated in this paper, this means that localities may gain by classifying an acceptable asylum seeker as unacceptable or vice versa. To analyze this type of strategic behavior, some no-

tion of dynamic (non-)manipulability needs to be introduced. Such analysis is beyond the scope of this paper. Nevertheless, because the preferences exploited in this paper are derived from estimation and not first-party reports, manipulability is unlikely to constitute a major problem in the considered dynamic refugee matching problem.

References

- Abdulkadiroğlu, A. and T. Sönmez (2003). School choice – A mechanism design approach. *American Economic Review* 93, 729–747.
- Andersson, T. (2017). Refugee matching as a market design application. Lund University Department of Economics Working Paper 2017:16.
- Andersson, T. and L. Ehlers (2016). Assigning refugees to landlords in Sweden: Efficient stable maximum matchings. Lund University Department of Economics Working Paper 2016:18.
- Aziz, H., J. Chen, S. Gaspers, and Z. Sun (2018). Stability and Pareto optimality in refugee allocation matchings. International Conference on Autonomous Agents and Multiagent Systems (ICAAM) Paper.
- Bansak, K., J. Ferwerda, J. Hainmueller, A. Dillon, D. Hangartner, and D. Lawrence (2018). Improving refugee integration through data-driven algorithmic assignment. *Science* 359, 325–329.
- Bloch, F. and D. Cantala (2013). Markovian assignment rules. *Social Choice and Welfare* 40, 1–25.
- Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 1061–1103.
- Crépon, B., E. Duflo, M. Gurgand, R. Rathelot, and P. Zamora (2013). Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. *Quarterly Journal of Economics* 128, 531–580.
- Damiano, E. and R. Lam (2005). Stability in dynamic matching markets. *Games and Economic Behavior* 52, 34–53.
- Damm, A. P. (2009). Ethnic enclaves and immigrant labor market outcomes: Quasi experimental evidence. *Journal of Labor Economics* 27, 281–314.
- Delacretaz, D., S. D. Kominers, and A. Teytelboym (2016). Refugee resettlement. University of Oxford Department of Economics Working Paper.

- Doval, L. (2016). Dynamic and many-to-many matching markets. Northwestern University Working Paper.
- Doval, L. (2017). A theory of stability in dynamic matching markets. Caltech Working Paper.
- Edin, P.-A., P. Fredriksson, and O. A. slund (2003). Ethnic enclaves and the economic success of immigrants – Evidence from a natural experiment. *The Quarterly Journal of Economics* 118, 329–357.
- Foley, D. (1967). Resource allocation and the public sector. *Yale Economic Essays* 7, 45–98.
- Gudmundsson, J. (2015). Compromises and rewards: Stable and non-manipulable probabilistic matching. Lund University Department of Economics Working Paper 2015:32 (revised October 19, 2017).
- Haeringer, G. and V. Iehlé (2017). Two-sided matching with (almost) one-sided preferences. Working Paper.
- Hatfield, J. W. and S. D. Kominers (2014). Hidden substitutes. Harvard Business School Working Paper.
- Jones, W. and A. Teytelboym (2017a). The international refugee match: A system that respects refugees’ preferences and the priorities of states. *Refugee Survey Quarterly* 36, 84–109.
- Jones, W. and A. Teytelboym (2017b). The local refugee match: Aligning refugees’ preferences with the capacities and priorities of localities. *Journal of Refugee Studies*.
- Kadam, S. V. and M. H. Kotowski (2018a). Multi-period matching. *International Economic Review*, forthcoming.
- Kadam, S. V. and M. H. Kotowski (2018b). Time horizons, lattice structures, and welfare in multi-period matching markets. Harvard University Working Paper.
- Kennes, J., D. Monte, and N. Tumennasan (2014). The daycare assignment: A dynamic matching problem. *American Economic Journal: Microeconomics* 6, 362–406.
- Kennes, J., D. Monte, and N. Tumennasan (2015). Dynamic matching markets and the deferred acceptance mechanism. Aarhus University Working Paper.
- Kojima, F., P. A. Pathak, and A. E. Roth (2013). Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics* 128, 1585–1632.
- Kotowski, M. H. (2015). A note on stability in one-to-one, multi-period matching markets. Harvard University Working Paper.

- Kurino, M. (2009). House allocation with overlapping agents: A dynamic mechanism design approach. *Jena Economic Research Papers* 2009–075.
- Moraga, J. and H. Rapoport (2014). Tradable immigration quotas. *Journal of Public Economics* 115, 94–108.
- Mullainathan, S. and J. Spiess (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives* 31, 87–106.
- Pycia, M. (2012). Stability and preference alignment in matching and coalition formation. *Econometrica* 80, 323–362.
- Ran, A., L. P. Boustan, and K. Eriksson (2014). A nation of immigrants: Assimilation and economic outcomes in the age of mass migration. *Journal of Political Economy* 122, 467–506.
- Roth, A. E. and E. Peranson (1999). The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review* 89, 748–780.
- SFS:2016:38 (2016). Lag om mottagande av vissa nyanlända invandrare för bosättning. Svensk författningssamling.
- SOU:2018:22 (2018). Ett ordnat mottagande – gemensamt ansvar för etablering eller återvändande. Stantens offentliga utredningar.
- Ünver, U. M. (2009). Dynamic kidney exchange. *The Review of Economic Studies* 77, 372–414.

Appendix A: Proofs

This Appendix contains the proofs of all results in the paper.

Proof of Theorem 1. To prove part (i), suppose that envy not is bounded by a single asylum seeker at matching $x(k)$. This means that there are two localities in $M(k)$, say m and m' , and two integers, say k^0 and k^1 (where $0 < k^0 < k^1 \leq k$), such that:

- (I) locality m is indifferent between bundles $x_m(k^0 - 1)$ and $x_{m'}(k^0 - 1)$,
- (II) locality m envies locality m' by exactly one asylum seeker at matching $x(k^0)$,
- (III) locality m envies locality m' by exactly two asylum seekers at matching $x(k^1)$.

Let k^0 and k^1 be the largest and the smallest integers, respectively, satisfying conditions (I)–(III), and consider the subsequence $(s(k^0), \dots, s(k^1))$ of the sequence s . Because the integers k^0 and k^1 are chosen in the above way, locality m is not matched to any asylum seeker in the subsequence $(s(k^0 + 1), \dots, s(k^1 - 1))$, and locality m' is not matched to any non-demanded asylum seeker or any asylum seeker that is acceptable for locality m in the subsequence $(s(k^0 + 1), \dots, s(k^1 - 1))$. Hence, there four reasons to why conditions (I)–(III) hold:

- (a) Asylum seekers $s(k^0)$ and $s(k^1)$ are acceptable to locality m but both $s(k^0)$ and $s(k^1)$ are matched to locality m' ,
- (b) asylum seeker $s(k^0)$ is non-demanded but matched to locality m and asylum seeker $s(k^1)$ is acceptable to locality m but matched to locality m' ,
- (c) asylum seeker $s(k^0)$ is acceptable to locality m but matched to locality m' and asylum seeker $s(k^1)$ is non-demanded but matched to locality m ,
- (d) asylum seekers $s(k^0)$ and $s(k^1)$ are non-demanded but both are matched to locality m .

It is only proved that cases (a) and (b) cannot occur as the corresponding proofs of cases (c) and (d) are almost identical.

To prove that (a) cannot occur, note that asylum seeker $s(k^0)$ is acceptable for both m and m' . Hence, by rotation of (π, σ) , it follows that locality m' has a higher position than locality m in $\pi(k^0)$ and the lowest position in $\pi(k^0 + 1)$. Because asylum seeker $s(k^1)$ is matched to locality m' , it must also be the case that locality m' has a higher position than locality m in $\pi(k^1)$. The latter can, however, not occur. To see this, recall first that locality m is not matched to any asylum seeker in the subsequence $(s(k^0), \dots, s(k^1))$. Similarly, m' is not matched to any asylum seeker in the subsequence $(s(k^0 + 1), \dots, s(k^1 - 1))$. Hence, by rotation of (π, σ) , the relative positions among m and m' do not change for the priority orders $\pi(k^0 + 1), \dots, \pi(k^1)$, which is a contradiction to m' being matched to the demanded asylum seeker $s(k^1)$.

To prove that (b) cannot occur, recall that asylum seeker $s(k^0)$ is non-demanded and that locality m envies locality m' at matching $x(k^0)$. Hence, by rotation of (π, σ) , locality m has a higher position than locality m' in $\pi(k^0 + 1)$. Because the overdemanded asylum seeker $s(k^1)$ is matched to locality m' , it must be the case that locality m' has a higher position than locality m in $\pi(k^1)$. The latter can, however, not occur. To see this, recall first that locality m is not matched to any asylum seeker in the subsequence $(s(k^0), \dots, s(k^1))$. Similarly, m' is not matched to any asylum seeker in the subsequence $(s(k^0 + 1), \dots, s(k^1 - 1))$. Hence, by rotation of (π, σ) , the relative positions among m and m' do not change for the priority orders $\pi(k^0 + 1), \dots, \pi(k^1)$, which is a contradiction to m' being matched to the demanded asylum seeker $s(k^1)$.

To prove part (ii), suppose that $x(k)$ is not Pareto efficient. Consider the localities in the set $M(k)$ and the asylum seekers in the set $A(k)$. Because $x(k)$ is not Pareto efficient, by assumption, there exists a matching $x'(k)$ where the quotas are respected for all localities in $M(k)$, $x'_m(k)R_m(k)x_m(k)$ for all $m \in M(k)$, and $x'_m(k)P_m(k)x_m(k)$ for some $m \in M(k)$. The latter conditions imply that:

$$\sum_{j \in M(k)} (|A_j^+(x'_j(k))| - |A_j^-(x'_j(k))|) > \sum_{j \in M(k)} (|A_j^+(x_j(k))| - |A_j^-(x_j(k))|). \quad (3)$$

Note next that a locality in $M(k)$ can only be matched to an unacceptable asylum seeker if the asylum seeker also is unacceptable for all other localities in $M(k)$. This follows since φ is an order mechanism. But this also means that $\sum_{j \in M(k)} |A_j^-(x'_j(k))| = \sum_{j \in M(k)} |A_j^-(x_j(k))|$. Consequently, condition (3) reduces to:

$$\sum_{j \in M(k)} |A_j^+(x'_j(k))| > \sum_{j \in M(k)} |A_j^+(x_j(k))|, \quad (4)$$

i.e., that the total number of acceptable asylum seekers matched to the localities in $M(k)$ is greater at matching $x'(k)$ than at matching $x(k)$. But this is not possible since all asylum seekers in $A(k)$ are matched to some locality in $M(k)$ which finds them acceptable by construction of the order mechanism φ . Hence, $x(k)$ must be Pareto efficient. \square

Proof of Corollary 1. Because matching $x(k)$ is selected by φ , it follows from Theorem 1 that envy for locality m towards locality m' is bounded by a single asylum seeker, i.e., that:

$$|A_m^+(x_m(k))| - |A_m^-(x_m(k))| + 1 \geq |A_m^+(x_{m'}(k))| - |A_m^-(x_{m'}(k))|. \quad (5)$$

To obtain a contradiction, suppose now that the statement is not true. Then:

$$|A_m^+(x_{m'}(k))| > |A_m^+(x_m(k))|, \quad (6)$$

$$|A_m^-(x_m(k))| > |A_m^-(x_{m'}(k))|. \quad (7)$$

Inequality (6) implies $|A_m^+(x_{m'}(k))| - 1 \geq |A_m^+(x_m(k))|$. This condition together with inequality (5) implies $|A_m^-(x_{m'}(k))| \geq |A_m^-(x_m(k))|$. But the latter inequality contradicts inequality (7). \square

Proof of Theorem 2. It is only proved that locality m does not envy any locality with a higher position in $\pi(k+1)$ since the second part of the theorem is based on symmetrical arguments. The result is proved by induction. It is first demonstrated that the result is true for $\pi(2)$. Three cases must be considered:

- (1.a) *Asylum seeker $s(1)$ is non-demanded.* Because $s(1)$ is non-demanded, no locality in $M(1)$ will envy the locality which is matched to $s(1)$ and the locality which is matched to $s(1)$ will envy all localities in $M(1)$. Because (π, σ) satisfies rotation, the locality that is matched to $s(1)$ will have the highest position in $\pi(2)$. These arguments show that no locality in $M(1)$ envies any locality with a higher position in $\pi(2)$.
- (1.b) *Asylum seeker $s(1)$ is demanded.* In this case, $s(1)$ is matched to the only locality finding $s(1)$ acceptable, and no locality will envy this locality. Similarly, the locality which is matched to $s(1)$ will not envy any other locality since $s(1)$ is acceptable. Because (π, σ) satisfies rotation, it follows that $\pi(1) = \pi(2)$, and, consequently, no locality in $M(1)$ envies any locality with a higher priority in $\pi(2)$.
- (1.c) *Asylum seeker $s(1)$ is overdemanded.* In this case, asylum seeker $s(1)$ is matched to the locality in $M(1)$ with the highest position in $\pi(1)$ which finds asylum seeker $s(1)$ acceptable. This locality will be envied by all localities finding $s(1)$ acceptable, and the locality which is matched to $s(1)$ will not envy any locality in $M(1)$. Because (π, σ) satisfies rotation, the locality which is matched to $s(1)$ will have the lowest priority in $\pi(2)$. These arguments show that no locality in $M(1)$ envies any locality with a higher position in $\pi(2)$.

From the above arguments, it is concluded that the statement is true for $\pi(2)$. Consider now the induction hypothesis that the result holds for $\pi(k)$ for an arbitrary $k \in \{2, \dots, n-1\}$. Given this assumption, it remains to show that the result is true for $\pi(k)$. Again, three cases must be considered:

- (k.a) *Asylum seeker $s(k)$ is non-demanded.* Because asylum seeker $s(k)$ is not acceptable to any locality in $M(k)$ and because envy is bounded by a single asylum seeker at matching $x(k-1)$ by Theorem 1, no locality in $M(k)$ will envy the locality which is matched to $s(k)$ at matching $x(k)$. Because (π, σ) satisfies rotation, the locality which is matched to $s(k)$ will have a higher priority in $\pi(k+1)$ than all localities which the locality envies at matching $x(k)$. These arguments show that no locality in $M(k)$ envies any locality with a higher position in $\pi(k+1)$.
- (k.b) *Asylum seeker $s(k)$ is demanded.* In this case, asylum seeker $s(k)$ is matched to the only locality which finds $s(k)$ acceptable. Because (π, σ) satisfies rotation, it follows that

$\pi(k) = \pi(k + 1)$. Consequently, the result holds for $\pi(k + 1)$ since the result is true for $\pi(k)$ by the induction hypothesis.

(k.c) *Asylum seeker $s(k)$ is overdemandd.* In this case, asylum seeker $s(k)$ is matched to the locality in $M(k)$ with the highest position in $\pi(k)$ finding asylum seeker $s(k)$ acceptable. Because envy is bounded by a single asylum seeker at matching $x(k - 1)$ by Theorem 1, the locality which is matched to the acceptable asylum seeker $s(k)$ will not envy any locality at matching $x(k)$. Then by rotation of (π, σ) , the locality which is matched to $s(k)$ will have the lowest position in $\pi(k + 1)$. But then the result holds by the above arguments.

The above three cases, together with the result for $\pi(2)$ and the induction assumption, proves that locality m does not envy any locality with a higher priority in $\pi(k)$. \square

Appendix B: Additional Simulation Results

This Appendix collects additional simulation results and further checks on the robustness of the performance of the order mechanism as the characteristics of refugee flows vary. Tables 5 and 6 report the expected percentage improvement in measures of fairness and efficiency granted by the proposed mechanism over the naïve sequential matching mechanism for varying combinations of the proportion of non-demanded (\underline{D}) and always demanded (\overline{D}) asylum seekers in simulated arrival flows. For each combination of parameters, 1,000 refugee flows of 1,000 asylum seekers each are simulated and allocated to 21 localities. Each table row reports the average percentage difference across these simulated flows after the mechanism allocated the last refugee in the flow. All simulations allow for correlation of locality preferences for each asylum seeker.

The tables show that the relative improvement in fairness with respect to the naïve sequential matching mechanism is larger when misclassification error is lower and the proportion of asylum seekers in the set D is higher. The relative improvement in efficiency depends solely on the extent of misclassification error, and always ranges between 100 and 0 percent. The only exceptions consist in the situations in which all asylum seekers are either all in the set \underline{D} or all in the set \overline{D} , where both assignment strategies produce Pareto efficient and fair allocations. The naïve sequential matching mechanism always produces a Pareto efficient (but not necessarily fair) allocation if no asylum seekers of type D exist. In general, the proposed matching mechanism always outperforms the naïve sequential matching mechanism as long as some asylum seekers of type D exist and we some information about the locality-specific partitions exist (less than 50% misclassification error).

Figure 2 shows that the naïve sequential matching mechanism outperforms a truly random mechanism in terms of fairness. This result is due to the fact that the sequential assignment guarantees at least that the difference in the number of asylum seekers across localities is bounded by one.

Finally, Figure 3 provides evidence for the conjecture that envy is bounded by a single acceptable asylum seeker *and* a single unacceptable asylum seeker.

Table 5: Simulated relative reduction in efficiency and envy measures with the proposed algorithm, by misclassification error in locality partitions, percentage of asylum seekers in set \underline{D} and percentage of asylum seekers in set \overline{D} in the simulated flows (I).

Miscl. error:	Percentage decrease in the share of localities envying some locality by at least one asylum seeker					Percentage decrease in the share of mismatched acceptable asylum seekers				
	0%	10%	25%	40%	50%	0%	10%	25%	40%	50%
% \underline{D} in flow						% \overline{D} in flow: 0.00				
0.00	100.00	100.00	93.80	29.99	-1.77	100.00	79.96	49.96	19.82	-0.02
10.00	100.00	99.97	89.16	25.02	-2.03	100.00	80.07	49.96	19.80	-0.07
20.00	100.00	99.93	82.88	21.44	-1.69	100.00	80.05	50.07	20.03	0.16
30.00	100.00	99.51	74.54	18.07	-2.17	100.00	79.98	49.90	19.95	0.05
40.00	100.00	98.23	64.96	15.30	-1.61	100.00	79.96	49.82	20.15	0.11
50.00	100.00	94.75	53.64	11.78	-1.75	100.00	79.92	49.81	20.01	-0.19
60.00	100.00	86.85	43.82	9.80	-1.98	100.00	79.97	50.17	19.85	0.02
70.00	100.00	72.09	32.74	7.03	-1.95	100.00	79.94	50.12	20.06	0.50
80.00	99.72	51.67	20.74	3.89	-2.96	100.00	80.13	49.92	19.68	-0.20
90.00	90.28	25.70	10.31	0.53	-3.56	100.00	80.25	50.10	20.55	0.92
% \underline{D} in flow						% \overline{D} in flow: 10.00				
0.00	100.00	100.00	90.27	26.49	-1.92	100.00	79.95	50.04	20.17	-0.21
10.00	100.00	99.92	83.43	21.65	-1.65	100.00	80.03	49.97	19.99	0.04
20.00	100.00	99.60	74.64	18.06	-1.32	100.00	79.94	50.17	20.20	0.10
30.00	100.00	98.36	62.94	13.86	-1.80	100.00	79.94	49.71	19.82	-0.10
40.00	100.00	94.06	50.76	10.42	-1.72	100.00	79.91	49.81	19.62	-0.07
50.00	100.00	84.62	37.44	7.56	-2.10	100.00	79.91	49.82	19.76	-0.25
60.00	99.97	67.50	26.39	5.09	-1.80	100.00	80.08	49.91	20.24	0.52
70.00	99.40	42.33	16.03	2.78	-2.34	100.00	80.26	50.22	20.05	-0.23
80.00	89.35	16.53	5.85	0.15	-2.25	100.00	80.09	50.14	20.37	0.33
90.00	89.58	-2.63	-2.79	-2.90	-2.96	nan	nan	nan	nan	nan
% \underline{D} in flow						% \overline{D} in flow: 20.00				
0.00	100.00	99.95	85.30	22.92	-1.57	100.00	80.14	50.01	20.05	0.16
10.00	100.00	99.67	75.77	18.02	-1.70	100.00	79.94	50.05	20.06	-0.15
20.00	100.00	98.42	63.35	13.86	-1.29	100.00	80.10	49.91	19.97	0.32
30.00	100.00	94.39	50.20	10.41	-1.64	100.00	79.99	50.03	19.90	0.03
40.00	100.00	84.21	36.35	7.00	-1.78	100.00	80.06	49.61	19.78	-0.09
50.00	99.97	65.45	24.54	4.38	-1.87	100.00	79.99	50.10	20.16	0.15
60.00	99.39	39.40	14.36	2.48	-1.67	100.00	80.17	50.16	19.90	0.03
70.00	86.44	14.85	4.96	-0.03	-1.69	100.00	79.77	49.78	19.61	-0.06
80.00	38.23	-1.26	-1.78	-2.00	-2.04	nan	nan	nan	nan	nan
% \underline{D} in flow						% \overline{D} in flow: 30.00				
0.00	100.00	99.77	79.59	19.00	-1.56	100.00	80.10	50.04	20.04	0.06
10.00	100.00	98.80	66.18	14.78	-1.96	100.00	80.02	49.98	20.03	0.04
20.00	100.00	95.74	51.91	10.29	-1.88	100.00	80.01	50.29	20.05	0.18
30.00	100.00	85.55	37.79	6.69	-1.84	100.00	80.00	50.15	20.10	-0.07
40.00	99.96	66.31	23.75	4.07	-2.05	100.00	79.94	49.91	20.00	-0.39
50.00	99.03	39.31	12.48	2.36	-1.76	100.00	79.99	50.24	19.86	-0.18
60.00	84.66	14.75	4.16	-0.20	-1.98	100.00	79.88	49.84	20.60	0.13
70.00	94.87	-0.98	-1.72	-1.84	-1.82	nan	nan	nan	nan	nan

NOTES: A thousand waves of asylum seekers of 1,000 asylum seekers each is simulated, distributed across 21 localities and for varying proportions of \overline{D} and \underline{D} . The measures of efficiency loss and envy are computed at the end of each, benchmarking the performance of the proposed mechanism (with increasing degrees of misclassification error) against that of the naïve sequential matching mechanism.

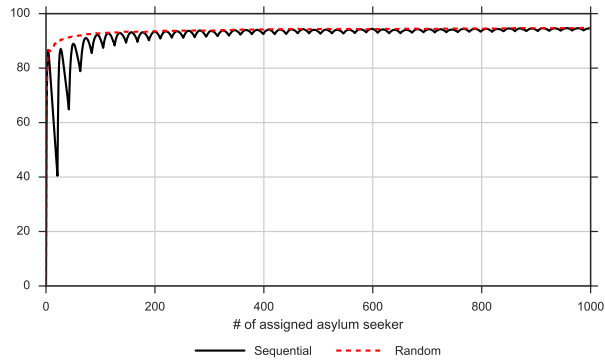
Table 6: Simulated relative reduction in efficiency and envy measures with the proposed algorithm, by misclassification error in locality partitions, percentage of asylum seekers in set \underline{D} and percentage of asylum seekers in set \overline{D} in the simulated flows (II).

Miscl. error:	Percentage decrease in the share of localities envying some locality by at least one asylum seeker					Percentage decrease in the share of mismatched acceptable asylum seekers				
	0%	10%	25%	40%	50%	0%	10%	25%	40%	50%
% \underline{D} in flow	% \overline{D} in flow: 40.00									
0.00	100.00	99.38	71.88	15.99	-1.77	100.00	80.06	50.23	20.24	0.26
10.00	100.00	96.43	56.11	11.75	-1.33	100.00	79.95	50.31	20.00	-0.07
20.00	99.99	88.05	39.95	7.67	-1.81	100.00	80.12	50.16	20.28	-0.06
30.00	99.96	69.05	24.79	4.36	-1.86	100.00	79.99	50.13	19.89	-0.13
40.00	99.30	41.47	13.17	2.23	-2.19	100.00	79.97	50.13	19.77	-0.00
50.00	87.44	15.13	3.78	0.02	-1.84	100.00	80.07	49.75	19.92	0.26
60.00	43.68	-0.67	-1.54	-1.64	-1.69	nan	nan	nan	nan	nan
% \underline{D} in flow	% \overline{D} in flow: 50.00									
0.00	100.00	98.12	62.27	13.03	-2.23	100.00	79.88	50.10	20.04	-0.15
10.00	100.00	91.49	44.67	8.52	-2.20	100.00	80.02	50.08	20.17	0.02
20.00	99.96	73.88	27.25	4.58	-2.12	100.00	79.85	49.70	19.89	-0.15
30.00	99.01	45.12	13.60	2.21	-1.88	100.00	80.18	50.19	20.04	0.25
40.00	83.06	17.05	4.91	-0.06	-2.01	100.00	79.77	50.21	20.16	-0.28
50.00	100.00	-0.57	-1.54	-1.80	-1.81	nan	nan	nan	nan	nan
% \underline{D} in flow	% \overline{D} in flow: 60.00									
0.00	100.00	94.72	53.24	9.85	-2.67	100.00	79.87	49.90	20.11	-0.12
10.00	99.94	79.99	31.48	5.27	-2.28	100.00	80.06	49.94	20.07	-0.03
20.00	99.27	50.25	15.87	2.54	-2.09	100.00	79.91	50.10	20.25	-0.04
30.00	88.58	18.48	4.44	-0.32	-2.03	100.00	79.85	50.09	19.74	-0.13
40.00	48.58	-0.35	-1.81	-2.07	-2.05	nan	nan	nan	nan	nan
% \underline{D} in flow	% \overline{D} in flow: 70.00									
0.00	99.95	87.61	40.74	6.48	-3.07	100.00	80.17	49.97	19.86	0.07
10.00	98.89	58.51	19.56	2.33	-2.84	100.00	80.08	49.97	19.71	-0.59
20.00	80.76	22.64	6.18	0.09	-2.16	100.00	80.01	50.10	20.28	0.06
30.00	-3.19	-0.01	-1.89	-2.23	-2.30	nan	nan	nan	nan	nan
% \underline{D} in flow	% \overline{D} in flow: 80.00									
0.00	99.08	72.15	28.20	3.09	-4.27	100.00	80.15	50.06	19.96	-0.45
10.00	89.40	30.67	7.57	-0.43	-2.91	100.00	80.07	49.47	19.69	-0.62
20.00	53.27	1.37	-1.98	-2.49	-2.69	nan	nan	nan	nan	nan
% \underline{D} in flow	% \overline{D} in flow: 90.00									
0.00	78.12	42.62	12.35	-3.69	-9.24	100.00	79.94	49.98	20.03	-0.09
10.00	-0.66	7.53	-2.77	-3.93	-4.21	nan	nan	nan	nan	nan


NOTES: A thousand waves of asylum seekers of 1,000 asylum seekers each is simulated, distributed across 21 localities and for varying proportions of \overline{D} and \underline{D} . The measures of efficiency loss and envy are computed at the end of each, benchmarking the performance of the proposed mechanism (with increasing degrees of misclassification error) against that of the naïve sequential matching mechanism.

Figure 2: Comparison of sequential and truly random assignments .

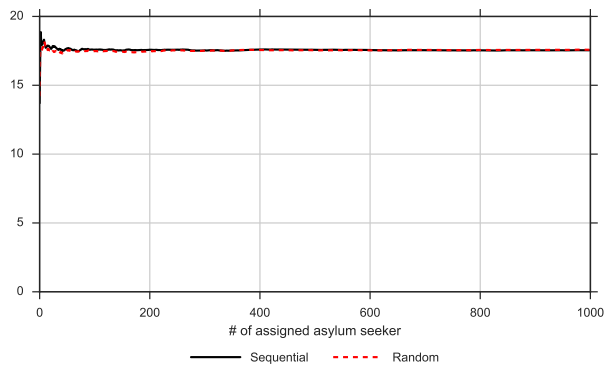
(a) Percentage of localities envying some locality by at least one asylum seeker



(b) Percentage of localities envying some locality by at least five asylum seeker




(c) Percentage of mismatched acceptable asylum seekers



NOTES: 1,000 waves of asylum seekers of 1,000 asylum seekers each is simulated, distributed across 21 Swedish regional entities (län). The measures of efficiency and envy is computed at each matching, comparing the performance of the naïve sequential matching mechanism with that of the truly random assignment mechanism. The flows of asylum seekers are calibrated to match the average 2016 employment outcomes of the 2013 asylum seeker wave in Sweden (\bar{D} : 28 percent; \underline{D} : 45 percent).

Figure 3: Simulation for the conjecture that envy is bounded by one in acceptable and unacceptable asylum seekers.



NOTES: 5,000 waves of 1,000 asylum seekers each is simulated, distributed across 20 localities. The characteristics of the refugee flow are random. This simulation gives us a total of 100 million maximum envy measures by locality (the maximum envy that each locality has towards any of the remaining 19 localities). Envy measures are split across envy on acceptable asylum seekers (another municipality has more asylum seekers considered acceptable) and envy on unacceptable asylum seekers (another municipality has less asylum seekers considered unacceptable). The figure plots the maximum, median and minimum observed maximum envy across these 100 million observations, arranged by number of assigned asylum seeker in the flow. The maximum measure never exceeds one for either acceptable and unacceptable asylum seekers, while the minimum of maximum envy decreases with the number of assigned asylum seekers.