Unintended Impacts of Automation for Integration? Simulating Integration Outcomes of Algorithm-Based Refugee Allocation in Germany

Jakob Kappenberger[1], Clara Strasser Ceballos[2], Frederic Gerdon[1], Daria Szafran[1], Florian Rupp[3], Kai Eckert[3], Heiner Stuckenschmidt[1], Ruben Bach[1], Frauke Kreuter[2,4], Christoph Kern[2,4]

[1] University of Mannheim, Germany
[2] LMU Munich, Germany
[3] Technical University Mannheim, Germany
[4] Munich Center for Machine Learning (MCML), Germany
jakob.kappenberger@uni-mannheim.de, clara.strasserceballos@stat.uni-muenchen.de, frederic.gerdon@uni-mannheim.de

Abstract

The location to which refugees are assigned upon arrival in a host country is a critical factor for their integration prospects. Several research groups have developed algorithmic tools based on artificial intelligence (AI) to optimize refugee-location matching, with the overall aim of improving refugees' integration into the labor market. These tools are used in a highly sensitive context, and thus their design, social impacts, and potential long-term consequences need to be systematically assessed. To investigate such effects, we propose an agent-based simulation framework, grounded in sociological theory and real-world survey data. This framework allows for simulating different allocation mechanisms for refugees to locations and studying their impacts on integration outcomes. We illustrate the simulation framework in the German context by comparing the current approach of the *Königsteiner Schlüssel* (i.e., quasi-random allocation) with algorithm-based matching. We study each procedure's impact on both economic and social integration and assess structural effects on inequalities between subgroups of refugees. We find that (1) algorithmic assignment can improve both economic and social integration outcomes globally; (2) performance gains vary geographically and demographically, potentially reinforcing existing inequalities; (3) incorporating feedback loops—where each allocation round reshapes local community composition—is crucial for assessing the impacts of algorithmic allocation systems in dynamic social environments.

Code & Appendix —
https://github.com/JakobKappenberger/refugee-allocation-simulation

1 Introduction

The integration of refugees is heavily debated in the political arena and represents an important area of sociological research. For instance, one of the central topics of such political and scientific debates is the economic integration (see Harder et al. (2018)) of refugees into the labor market. Successful labor market integration becomes more likely if there is a match between refugees' characteristics and the characteristics of the resettlement location and/or population (Ahani et al. 2021; Aksoy, Poutvaara, and Schikora 2023; Auer 2018; Müller, Pannatier, and Viarengo 2023; Vogiazides and Mondani 2020), including the demographic composition of the local population and available ethnic networks (Kosyakova and Kogan 2022). Therefore, allocating refugees to locations upon arrival in a host country is a crucial component of integration policies and can critically shape refugees' integration outcomes.

Currently, governmental procedures that assign refugees to locations typically do not consider the match of refugees with host location characteristics. For instance, in Germany and Switzerland, the allocation process can be considered (quasi-)random. In Germany, the allocation follows the *Königsteiner Schlüssel* distribution key (BAMF 2024). Except for pre-existing family ties, no individual-level characteristics of the refugees are considered in the assignment process.

Recently, several research projects have developed algorithm-based matching tools that consider refugees'[1] characteristics for resettlement decisions, with the aim to improve their integration outcomes (e.g., Bansak et al. 2018; Ahani et al. 2021). One example of these algorithm-based tools is *GeoMatch* that combines supervised machine learning and optimal matching techniques to optimize the global average employment of newly arriving refugees (Immigration Policy Lab 2025). However, anticipating and assessing the social impacts of algorithmic matching tools requires a holistic perspective that takes the dynamic interplay between predictions, allocation outcomes, and their performative effects on local populations into account. Importantly, this includes understanding unintended "side-effects" of algorithmic matching on other key integration outcomes of interest. Labor market integration, for example, is just one among many dimensions of integration. One may question whether it is desirable to allocate refugees to locations mainly based on economic considerations and needs. Therefore, it is necessary to evaluate how optimizing solely for labor market integration may affect outcomes in other dimensions of integration. For instance, allocation decisions could correlate with and affect the demographic composition of local populations. This can in turn affect the chances of social inte-

[1] We use the term "refugee" in our study to refer to any person seeking asylum in a country other than their country of nationality, regardless of their legal asylum status (UNHCR 2024a,b).

gration, here understood as the level of social interaction between refugees and natives (Harder et al. 2018) as well as non-natives in the host society (Aksoy, Poutvaara, and Schikora 2023). Thus, feedback loops may result from how algorithm-based matching shapes the demographic composition of resettlement locations by affecting the integration chances of future refugees. Finally, there may be subgroups of refugees for whom the model works worse than for others, raising issues of fairness in allocation.

Contribution. Given these potential issues with algorithm-based allocations of refugees and following recent calls for more holistic, system-wide fairness assessments (Cruz Cortés and Ghosh 2020; Gerdon et al. 2022), we develop an empirically informed agent-based model, using data from the IAB-BAMF-SOEP Survey of Refugees in Germany (DIW 2024). Using this model, we simulate the allocation procedures based on refugees' and location characteristics. More concretely, we (1) examine whether optimizing for solely one integration outcome (e.g., labor market integration) leads to unintended consequences by impacting other integration outcomes (e.g., social integration), (2) show how the use of such matching procedures can reinforce social inequalities, i.e., have disparate impacts between subgroups of refugees, and (3) investigate feedback loops resulting from changes in local demographic compositions induced by matching, affecting future refugees.

With this research, we contribute to the growing literature on using algorithmic technologies in sensitive social contexts such as refugee allocation and, more broadly, on computational frameworks for evaluating the social impacts of algorithmic decision-making in dynamic environments, particularly relating to governmental decision-making. Furthermore, our study connects to research on dynamic fairness (Chouldechova and Roth 2020), which has shown the importance of dynamic interactions in the assessment of algorithmic tools. However, a common focus of this research is to study how prediction errors for protected groups evolve over time (D'Amour et al. 2020). Following related research that proposes the use of agent-based models to simulate the broader dynamics of socio-technical systems (e.g., Adelt et al. 2018; Cruz Cortés and Ghosh 2019; Scher et al. 2023), we aim to instead use agent-based simulations to explicitly model dynamic interactions, grounded in real-world survey data, to assess the potential impacts of algorithmic matching of refugees in a specific application context.

2 Background

Integration can be understood as "a dynamic two-way process that places demands on both the refugees and the receiving community" (UNHCR 2013, p. 8). Further, according to UNHCR (2013), successful integration of individuals goes beyond the simple provision of basic necessities and requires the creation of a supportive environment that helps refugees achieve economic stability, adapt socially, and develop a sense of belonging. Integration is multifaceted; for example, Harder et al. (2018) identify six dimensions of integration: economic, social, psychological, political, linguistic, and navigational. A large portion of existing studies focuses on the economic dimension when examining integration processes. To measure the degree of economic integration of refugees, researchers often rely on various labor market outcomes such as employment status, income, or job satisfaction (Harder et al. 2018).

There is a complex interplay between refugees' characteristics and features of the host country and resettlement location in shaping a refugee's labor market outcomes. Previous literature found an association between various individual and contextual characteristics of refugees and their employment probability in Germany. The characteristics include the refugees' age (Hahn et al. 2019; Kosyakova and Brenzel 2020), gender (Hahn et al. 2019; Kosyakova and Brenzel 2020; Kosyakova, Salikutluk, and Hartmann 2023; Aksoy, Poutvaara, and Schikora 2023), education level (Hahn et al. 2019; Aksoy, Poutvaara, and Schikora 2023; Hunkler, Edele, and Schipolowski 2021), country of origin (Kosyakova and Brenzel 2020), previous employment experiences (Kosyakova and Brenzel 2020), asylum process length and legal status (Kosyakova and Brenzel 2020), recognition of credentials (van Riemsdijk and Axelsson 2021), unemployment level (Aksoy, Poutvaara, and Schikora 2023), and negative attitudes of the host society towards immigrants (Aksoy, Poutvaara, and Schikora 2023; Schilling and Stillman 2024).

Algorithmic Matching Tools

Recently, researchers leveraged data and knowledge about such integration processes to predict and attempt to improve the labor market integration chances of refugees (Bansak et al. 2018; Ahani et al. 2021). For this purpose, they propose models that attempt to assign refugees to specific resettlement locations so that overall integration chances are maximized. Such algorithm-based allocation policies may thus be better at fostering integration than other commonly used assignment policies, such as (quasi-)random assignments of refugees to resettlement locations (for instance, according to the *Königsteiner Schlüssel* in Germany; see below).

To improve upon traditional allocation mechanisms, two major algorithms exist that combine supervised machine learning with optimal matching techniques. The first, *GeoMatch*, developed by the Immigration Policy Lab (Stanford and ETH Zürich), was introduced by Bansak et al. (2018). The algorithm is currently being piloted in the U.S., Switzerland, and the Netherlands (Global Compact on Refugees 2024). In the U.S., the resettlement agency Global Refuge uses the tool to assign refugees upon arrival to one of its local affiliates based on their individual-level characteristics. In Switzerland, the State Secretariat for Migration (SEM) employs the tool to resettle refugees across the 26 cantons. In both contexts, the tool is used to assign newly arriving refugees to locations such that employment is maximized, taking into account capacity constraints. The second algorithm, *Annie*$^{TM}Moore$, introduced by Ahani et al. (2021), works similarly to *GeoMatch*, but is used by a different resettlement agency, HIAS, in the U.S.

Challenges of Algorithmic Matching

While promising, these tools are accompanied by a number of challenges that deserve particular attention. We highlight three key issues: (1) the potential unintended consequences on other integration outcomes, in particular on social integration, resulting from the multi-dimensional nature of integration; (2) the non-use and/or unavailability of key variables that are important for predicting economic outcomes, as predictive models can only be based on data available at the time of arrival; and (3) potential disparities in allocation outcomes across different subgroups of refugees.

First, integration is a multi-dimensional process that extends beyond the economic dimension alone (Harder et al. 2018; Castles et al. 2003). It is therefore crucial to consider additional dimensions of integration, an aspect currently neglected by the tools. In particular, social integration has been shown to play a key role, as it is closely associated with refugees' overall well-being and the long-term success across multiple dimensions of integration (Landesmann and Leitner 2025; Kosyakova and Kogan 2024; Zwingel 2023; Hannafi and Marouani 2023; Kuschel et al. 2023; Walther et al. 2020; Kanas et al. 2012).

For social integration, as with economic integration, a range of individual and resettlement location-level characteristics play a significant role. These include, but are not limited to, socio-demographic characteristics like age (Hannafi and Marouani 2023) and country of origin (Landesmann and Leitner 2025), as well as contextual factors like the level of civic engagement (Barreto et al. 2022), local attitudes towards immigrants (Schilling and Stillman 2024; Aksoy, Poutvaara, and Schikora 2023), the number of co-ethics (Martén, Hainmueller, and Hangartner 2019), and housing regulations (Khalil and Tjaden 2025).

Importantly, employment can serve as a facilitator to develop a social network via the workplace (Landesmann and Leitner 2025). At the same time, social networks can aid in attaining employment (Martén, Hainmueller, and Hangartner 2019; Kanas et al. 2012). The latter effect was found to be also true for co-ethnic networks in the earlier phase after arriving in the resettlement location (Battisti, Peri, and Romiti 2022). However, individuals resettled to locations with larger co-ethnic networks were also found to be less likely to invest in their human capital, which in the long run reduced initial employment advantages brought by the co-ethnic network (Battisti, Peri, and Romiti 2022). Thus, there is an interplay of economic and social integration, and allocation regimes that only focus on the former may produce unintended consequences for the latter. Against this backdrop, we look at such potential consequences for social integration, i.e., the direction and strength of such effects in our simulation setup, which could point to potential issues when trying to optimize for labor market outcomes.

Second, the tools rely on a limited set of individual-level variables available at the time of arrival, approximately around ten, and currently omit any location-level information (Bansak et al. 2018; Ahani et al. 2021). Yet, research has shown that successful labor market integration depends on a wide range of factors (see above), some of which cannot be measured at the time of allocation. Evaluating matching systems using more comprehensive data can help identify which variables should be incorporated or excluded. To address this, we enrich our analysis by including location-level information, which is available at the time of resettlement and, as highlighted above, plays a critical role in shaping refugees' employment and social outcomes.

Third, focusing on aggregate employment gains risks overlooking subgroup differences in the model's allocation. That is, there may be subgroups of refugees for whom algorithmic location matching would actually lead to worse outcomes than other types of assignment, such as random assignment. Despite its importance, the allocation fairness of these tools across refugee subgroups remains largely under-researched, with only two notable exceptions (Freund et al. 2023; Bansak and Marten 2021). We recognize the need for systematic subgroup fairness and assess it in our research using a metric we define to measure inequity in integration outcomes achieved. We examine inequities of effects between several subgroups, for instance, relating to gender, age, world region of origin, education level, and religious affiliation.

3 Methods

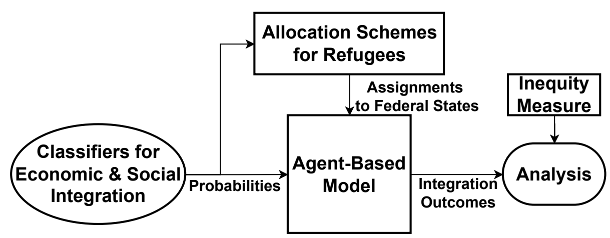

Figure 1: Overview of our method, from predicting economic and social integration for algorithmic assignment and the agent-based simulation of integration outcomes, to analyzing the results.

We propose a four-step method to tackle these challenges (shown in Figure 1). In particular, we (1) train classifiers on survey data to predict (a) labor-market and (b) social integration outcomes, (2) build an agent-based model that simulates these outcomes from 2016 to 2021 under various allocation schemes, (3) design two algorithmic allocation strategies that optimize each integration dimension relative to the *Königsteiner Schlüssel* baseline, and (4) define a metric to measure the inequity between outcomes achieved by different subgroups of agents. In the following, we detail each component of this method.

Data

As the main data source for our study, we draw on the IAB-BAMF-SOEP Survey of Refugees in Germany. This survey is integrated into the Socio-Economic Panel (SOEP) study and has been conducted annually since 2016. Random samples from the Central Register of Foreigners are drawn to gather information about refugees who immigrated to Germany since 2013 (DIW 2024). This survey is Europe's largest representative study of refugees (Schilling and Stillman 2024) and provides key data for analyzing their integration in Germany (Brücker et al. 2025). In addition to the survey data, we supplement the dataset with data on the unemployment rates (Destatis 2024a) as well as the community sizes of residents with foreign nationalities in the different federal states, respectively (Destatis 2024b). Crucially, due to the small sample sizes of a large number of nations, we utilize the United Nations' M49 geographical area coding, dividing the globe into 22 subregions instead of incorporating individual nationalities (United Nations 1999). Moreover, we draw on the European Social Survey (ESS) to derive the average (i.e., arithmetic mean of) attitudes towards immigrants in a given year and state via the item "Is your country made a worse or a better place to live by people coming to live here from other countries?" (ESS Round 11: European Social Survey European Research Infrastructure (ESS ERIC) 2024).

To obtain our final dataset, aside from dropping missing values for the features used (see below), we filter for respondents aged between 18 and 65 following existing research (e.g., Battisti, Peri, and Romiti (2022)) and discard all observations that lie more than 3 years after the immigration year of the respective individual. The latter step is necessary as the mandatory domicile obligation ("Wohnsitzauflage") of refugees in Germany lasts for three years, rendering the effect of the initial allocation on integration success potentially void afterward (Baba et al. 2024). We applied ordinal encoding to all categorical predictors so they could be used in our classifiers. The described preprocessing pipeline yields a dataset of 11,228 observations from 6,312 individuals. Tables 4 and 5 in the appendix show the distribution of the features used for our models.

As outlined before, we examine economic and social integration outcomes for refugees after their allocation in Germany. We operationalize economic integration as a binary indicator of current occupational status: respondents who are employed (full- or part-time) or enrolled in professional training are coded as 1, and 0 otherwise. For social integration, we follow Aksoy, Poutvaara, and Schikora (2023) by constructing an index from two items from the IAB-BAMF-SOEP survey—each scored 0 ("never") to 5 ("every day")—measuring the frequency of interactions with native and non-native individuals of other nationalities (DIW 2024). We then dichotomize this index so that respondents scoring 5 or above are coded as 1, and 0 otherwise.

Predictive Models

To simulate how different allocation schemes might influence refugees' integration success, it is crucial to predict the relevant integration dimensions as accurately as possible. We therefore build two machine learning models, one for labor market integration and one for social integration, and employ both to forecast outcomes in our simulation and to devise an allocation strategy inspired by *GeoMatch*.[2] While the labor market model estimates an individual's probability of finding work in a given federal state and year, the social integration model predicts whether an individual has established social contacts outside the community originating from the respective world region. As discussed above, these two constructs are not entirely independent. Having a job, for instance, may foster personal relationships and thereby support social integration. Indeed, in our data, they exhibit a correlation of 0.23, suggesting a weak to moderate positive relationship (Cohen 2009).

Both models are trained using the *GPBoost* framework, which combines gradient-boosted decision trees for fixed effects with a mixed-effects structure for random effects, thereby incorporating the panel nature of the survey data (Sigrist 2023). The two components play distinct roles. Fixed effects (modeled via gradient boosting) learn the relationships between observed predictors, such as age, gender, and other individual attributes, and the integration outcomes. This component benefits from the excellent performance characteristics of tree-based methods on tabular data (McElfresh et al. 2023). Random effects capture unobserved, group-level variability by introducing random intercepts for each individual (person ID), federal state, and survey year. These intercepts are assumed to follow a Gaussian distribution, allowing the model to account for within-subject and within-group correlations. Ignoring these random effects would typically underestimate standard errors and overstate confidence in the fixed-effect estimates (Laird and Ware 1982).

We performed a random train–validation–test split to evaluate model performance on unseen data, applying early stopping on the validation set and assessing the final model performance on the test set. Because the economic integration dimension is heavily imbalanced, we oversampled the minority (positive) class to increase its influence during training. Hyperparameters were tuned using Bayesian optimization for both models, with model selection guided by 5-fold cross-validation (Table 6 in the appendix shows all hyperparameters).

Table 1 shows the performance of both models on the test data across various metrics. Despite overall accuracy appearing promising for economic integration—primarily due to the imbalanced class distribution—the classifier struggles to correctly identify positive cases (i.e., respondents reporting an occupation), as indicated by its low precision and recall. In contrast, the social integration model performs better on all metrics except accuracy and ROC AUC, with the former being well-known for overestimating performance in imbalanced settings (Provost and Fawcett 2001). To put these results into perspective, it is essential to note that, despite us adding additional data on the location level, many potentially relevant factors (e.g., local job availability or unobserved personal circumstances) are still absent from the dataset, and all environment-specific information is limited to the federal-state level, potentially overlooking finer regional nuances that most likely significantly affect integration outcomes. Therefore, while the mediocre performance of the predictive models (particularly for economic integration) is a limitation of this study, both classifiers appear suf-

[2] Our implementation mainly differs from *GeoMatch* in terms of the model architecture chosen as we estimate one global model instead of one per federal state (see chapter 4 for more details).

Classifier	Accuracy	Precision	Recall	F1 Score	ROC AUC
Economic Integration	0.80	0.43	0.58	0.49	0.71
Social Integration	0.63	0.65	0.68	0.67	0.63

Table 1: Performance Metrics for Classifiers.

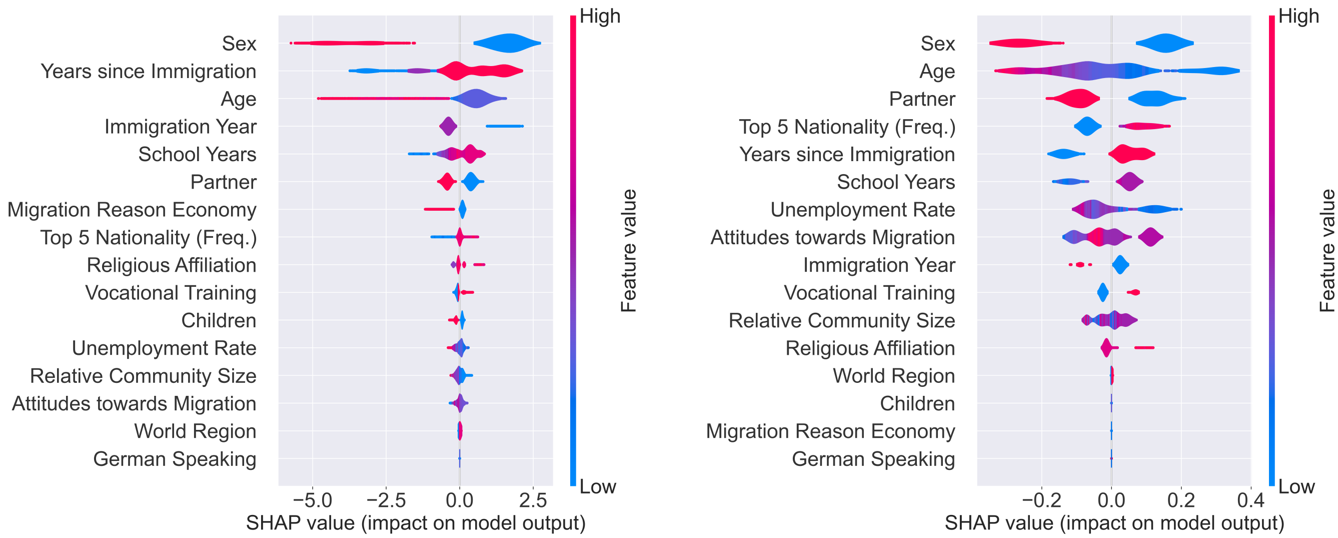

(a) SHAP values for features used to predict economic integration.

(b) SHAP values for features used to predict social integration.

Figure 2: SHAP values for predictive models for economic and social integration.

ficiently robust—given the difficulty of the task—to investigate how different allocation schemes might influence the predicted outcomes.

We rely on *SHAP* (SHapley Additive exPlanations) values to further examine the models' outputs to estimate how each feature influences the respective classifier's predictions. SHAP values extend Shapley values—a method from cooperative game theory—to machine learning, allocating a model's output to each input feature. They measure how the prediction changes when a feature is "absent" by comparing the full model's output to the output with that feature removed (Shapley 1953; Lundberg and Lee 2017). Figure 2 illustrates the SHAP values for the fixed effects in both the social and economic integration classifiers. The features are ordered by their cumulative SHAP values across all individuals, indicating their relative importance to each classifier's output. Note that the SHAP values for encoded categorical variables with many different values (e.g., religious affiliation) are difficult to interpret, as it is not easily determinable what higher or lower values correspond to.[3]

The most important feature for both classifiers is the respondent's sex, showing strong negative effects as women receive lower predictions on average (women are coded as 1, men as 0), which is in line with existing research on the labor market success of refugees (e.g., Kosyakova, Salikutluk, and Hartmann 2023; Aksoy, Poutvaara, and Schikora 2023). In a similar vein, a respondent's age, the unemployment rate in the given federal state, and the year of immigration have negative effects on the predicted integration success. Conversely, the years since arrival, the number of school years, and vocational training are positively associated with both integration dimensions. Interestingly, both having children and the capability to speak German before arrival have little impact on the classifiers' predictions. For the latter variable, this may be caused by its skewed distribution, as only very few respondents report levels above not being able to speak German at all before arriving (see Table 4 in the appendix). Finally, the relative size of one's world region group community positively predicts social integration but, surprisingly, correlates negatively with economic integration predictions. This finding departs from general evidence on co-ethnic networks and labor market integration of immigrants (e.g., Battisti, Peri, and Romiti (2022)) but aligns with research in the German context for refugees (Stips and Kis-Katos 2020). We conjecture that this divergence may stem from Germany-specific policies—such as the mandatory domicile obligation—that restrict refugees' ability to leverage co-ethnic networks, thereby limiting short-term economic gains and exposing them to potential longer-term drawbacks (e.g., reduced investment in human capital, as discussed above). Further research is needed to test this hypothesis and uncover the precise mechanisms at work, as it does not fully explain differences to other contexts (e.g., Switzerland)

[3]Although one-hot encoding can mitigate this issue, we found that ordinal encoding produced superior predictive performance for our models.

where similar restrictions apply.[4]

Simulation

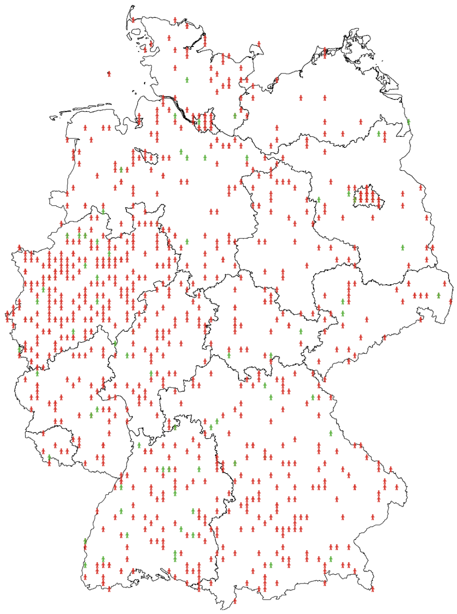

Figure 3: Visual interface of the simulation. Green icons denote agents who are predicted to achieve economic integration, while red icons denote those who are not. Agents are positioned randomly in their assigned federal state.

We present an agent-based model to compare the integration outcomes of refugees in Germany following their allocation via different assignment schemes and capture dynamic interactions caused by these placements. Figure 3 shows the visual interface of the simulation, and algorithm 1 in the appendix details the core principles of the simulation. The environment of the agent-based model consists of Germany's 16 federal states. Every time step in the simulation corresponds to one year, starting in 2016 and finishing in 2021 (the period for which we have data from the IAB-BAMF-SOEP survey).

The agents in the simulation represent refugees who, after their arrival and initial registration, are allocated to one of the federal states. All agents in the model own a series of attributes, the values of which are derived from the survey data (e.g., age, world region of origin, sex, education, German language skills).

Every iteration of the simulation first involves creating new agents. Throughout the simulated time frame, the volume of agents spawned is proportional to the number of refugees officially recorded in Germany for that specific year. The newly spawned agents are then distributed according to one of the three allocation schemes defined below, and the current iteration is concluded. After three time steps (i.e., years) in the simulation, agents are removed as the mandatory domicile obligation ("Wohnsitzauflage") of refugees in Germany ends (Baba et al. 2024).

After allocation, we compute the integration outcomes of agents in the federal state to which they were assigned. For the dimension of labor market integration, this means estimating the probability of a given individual finding work in the respective federal state in the current time step, which we derive from the *GPBoost* classifier described above. It is crucial to note that, since we performed oversampling during training, the probabilities output by the classifier are somewhat higher than justified by the empirical data (Figure 7 in the appendix plots the calibration curves for both predictive models). Likewise, for social integration, the respective classifier is deployed to predict whether or not an individual will reach the cumulative threshold defined above for the frequency of contacts with people not from the same world region. Consequently, at every iteration, predictions are made for all agents that have not yet achieved the respective dimension, and a random draw is conducted reflecting the estimated probability (i.e., a probability of 0.5 does not simply equate success, but rather a 50% chance of success). Agents cannot "lose" successful integration in any of the two dimensions.

Crucially, the community sizes in each federal state—which Figure 2 shows to be relevant for both classifiers—change as a result of the simulated allocations. These feedback loops lead to emergent changes in the demographic composition found in the simulation, potentially affecting the integration success of subsequent refugees. At the start of the simulation, each state is initialized with the empirical distribution of communities per world region in 2016. After each round of allocations, we adjust these shares toward the newly simulated distribution for each world region, moving at most ± 5 percentage points of the previous share to prevent unrealistic volatility. In practice, this cap means that even if a region's simulated share deviates by 100% from its empirical value, the community's size can change by no more than 5 percentage points in a single time step, thus reflecting that the simulation only covers a small part of all changes to the non-German subpopulations in German states. In experimental settings where we turn off these feedback loops (see chapter 4), the relative community sizes for a federal state are simply queried from the empirical data for any given year.

Finally, after the final iteration of the model, all outcomes achieved by the agents are saved, enabling us to compare the labor market and social integration outcomes between agents with different attributes, such as sex and world region of origin, and how inequalities in integration outcomes change depending on the allocation scheme.

Assignment Schemes

We use the simulation to evaluate two families of assignment schemes for refugees upon their arrival in Germany:

- The procedure based on the *Königsteiner Schlüssel* serves as our baseline. This scheme governs the distribution of refugees in Germany by quasi-randomly assign-

[4]It is also worth noting that the framing of the relevant community attribute differs slightly between our approach (world region) and other works (mostly either nationality or language).

ing them to federal states according to a fixed distribution key derived from each state's population size and tax revenue (BAMF 2024). Individual distribution shares range from 21% for the largest state, North Rhine-Westphalia, to just under 1% for the smallest city-state, Bremen.

- Our second approach draws on recently proposed methods, such as *GeoMatch* (Bansak et al. 2018) and *Annie*$^{TM}Moore$ (Ahani et al. 2021). We model refugee assignment as a constraint-based linear sum assignment problem (LSAP). First, our classification models estimate counterfactual probabilities for each integration dimension for every refugee across all federal states (i.e., for each state, what is the predicted probability of achieving the given integration dimension for that individual?). Since we are simulating the allocation of refugees upon their arrival, we set the variable *years since immigration* to 0. These probabilities form a matrix that is then used as input to an LSAP, and a mixed-integer linear programming solver is deployed to determine an optimal assignment scheme under the constraint of closely fulfilling the distribution determined by the *Königsteiner Schlüssel*. We apply this pipeline to search for optimal allocations for both integration dimensions, yielding two algorithmic schemes—*algorithmic economic* and *algorithmic social*—that respectively maximize the average expected labor market and social integration outcomes.

Measuring Fairness in Allocation Decisions

Because we are interested in examining whether different allocation schemes disproportionately affect individual subgroups of the model population, it is necessary to define an appropriate measure. To address this, we operationalize fairness using the concept of *equity*. Unlike equality—which advocates treating all citizens identically (e.g., providing the same level of services)—equity considers individual circumstances and aims to achieve similar outcomes for all (Deutsch 1975). In practice, this means that citizens with greater needs might receive a higher level of services so that their outcomes align more closely with those of others. Thus, while equality focuses on uniform treatment, equity targets comparable results.

Equity-based metrics are particularly appropriate for our approach as it centers on the outcomes achieved by agents from different subgroups under various allocation schemes. Specifically, we adopt the following inequity measure (Kappenberger, Stuckenschmidt, and Gerdon 2025):

Definition 3.1 (Inequity)**.** Let N be the set of all individuals and let $\sim$ be an equivalence relation on N, partitioning N into $m = |N/\sim|$ equivalence classes denoted $g_1, g_2, \dots, g_m$. For each class g_j, we define

$$f_{g_j} \;=\; \text{avg}\{\, O_i : i \in g_j \},$$

as the mean outcome O_i of all individuals i in class g_j. Let

$$F(\sim) \;=\; \big[f_{g_1},\, f_{g_2},\, \dots,\, f_{g_m}\big].$$

represent the average outcomes of the m equivalence classes.

Finally, the inter-group inequity with respect to the equivalence relation $\sim$ is defined as:

$$\text{inequity}(\sim) = JSD\left(\frac{F(\sim)}{\sum_{j=1}^{m} f_{g_j}} \,\middle\|\, \mathcal{U}(n)\right)$$

This measure quantifies the similarity between the observed outcome distribution across equivalence classes (or subgroups) and the discrete uniform distribution $\mathcal{U}(n)$, which represents equal outcomes for all groups. The distance is computed using the Jensen-Shannon Distance (JSD), a divergence measure widely used in statistics to compare probability distributions.[5] We selected the JSD because it is strictly bounded and allows for an intuitive interpretation (Endres and Schindelin 2003). Lower JSD values indicate greater similarity—that is, less inequity between the observed outcomes and the ideal equal distribution.

4 Results

To obtain robust results, we conducted 100 independent simulations for each allocation strategy. Before presenting overall results, we examine the impact of feedback loops—the adjustments to community sizes after each allocation—on integration outcomes. Because agents from Western Asia constitute the largest subcommunity in our data and simulation, their respective subcommunities in each federal state—and consequently their outcomes—are likely to be most affected by these demographic shifts. Figure 4 plots the average relative difference in predicted integration success (%) between Western Asian agents and all others, across all assignment schemes and years (2016–2021). Enabling feedback loops systematically lowers the simulated economic integration outcomes for Western Asian agents—reflecting the negative effect of growing co-regional communities on labor market outcomes observed in our data—while it improves their social integration outlook, consistent with the positive influence of larger community networks on social integration overall. These results highlight that dynamic changes in local demographics can meaningfully alter both economic and social integration outcomes, underscoring the importance of modeling feedback mechanisms when evaluating allocation policies.

Table 2 summarizes the overall performance of our two algorithmic schemes relative to the *Königsteiner Schlüssel* baseline (Table 7 in the appendix lists the overall performance without feedback loops). Both the *algorithmic economic* and *algorithmic social* strategies yield significant improvements—around 5%—in their respective target dimensions. Moreover, *algorithmic economic* also delivers a modest but significant gain in social integration compared to the baseline, whereas the economic integration results under *algorithmic social* do not differ significantly from the baseline. Even though an average increase of 5% in outcomes achieved most likely constitutes a substantial improvement for any social policy, it is worth discussing why

[5]The JSD is the square root of the Jensen-Shannon divergence (see Menéndez et al. (1997) for a detailed definition and discussion of the Jensen-Shannon divergence).

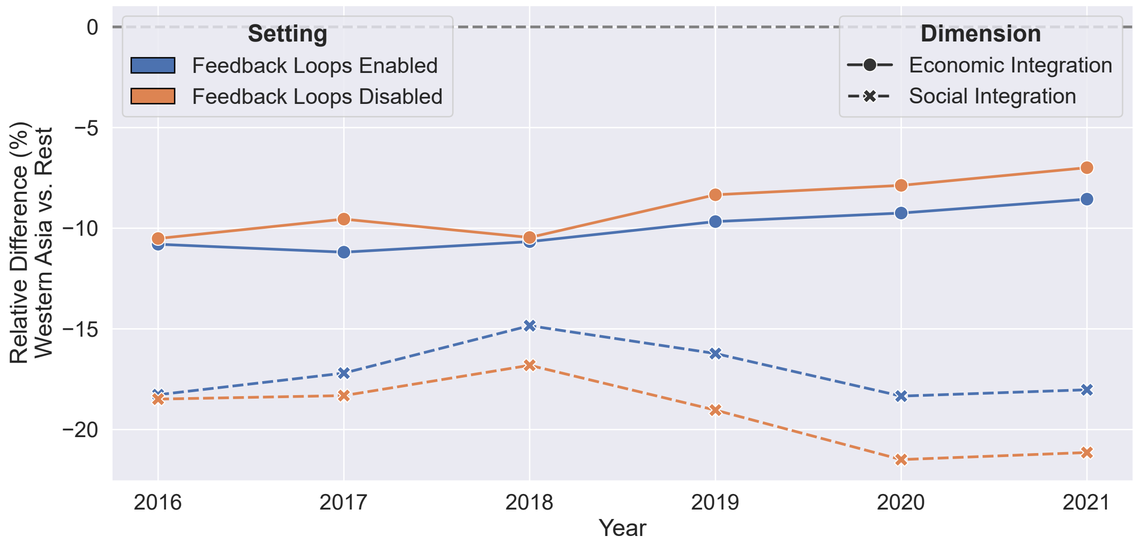

Figure 4: Average relative difference in predicted integration outcomes (%) between refugees from Western Asia and those from all other world regions, across all assignment schemes and over time (2016–2021) after 100 simulations per scheme. Results are shown separately for economic integration (solid lines) and social integration (dashed lines), and for simulations with feedback loops enabled (blue) versus feedback loops disabled (orange).

Assignment	**Avg. Economic Integration**	**% Diff. Baseline**	**Avg. Social Integration**	**% Diff. Baseline**
Algorithmic Economic	**0.31**	+4.62**	0.47	+1.16**
Algorithmic Social	0.30	-0.15	**0.49**	+4.81**

$^{**}p < 0.01$; $^{*}p < 0.05$

Table 2: Results of algorithmic assignment schemes versus Königsteiner baseline after 100 simulations per strategy.

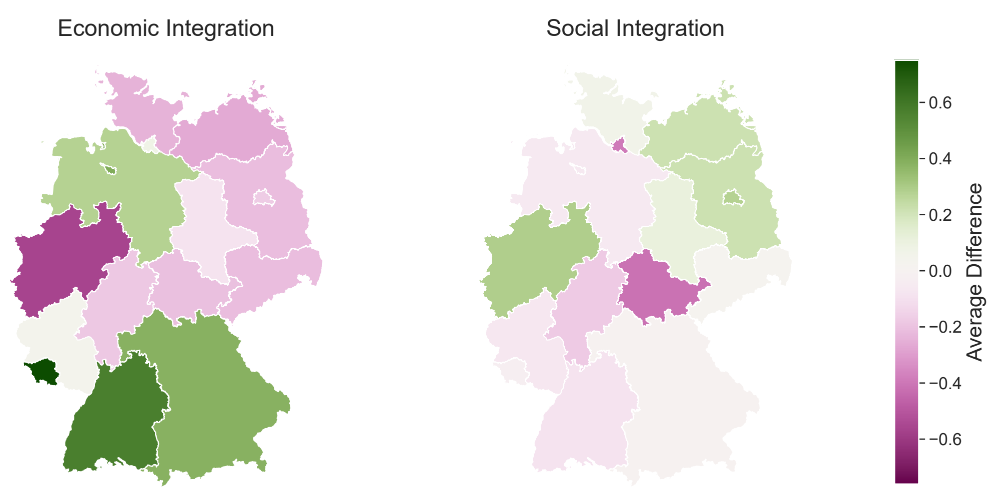

Figure 5: Average relative difference in predicted integration outcomes by federal state, comparing algorithmic versus *Königsteiner Schlüssel* allocations.
Left: Economic integration (*algorithmic economic* vs. baseline). **Right:** Social integration (*algorithmic social* vs. baseline). States shaded in green represent a relative gain over the baseline; states shaded in pink represent a relative loss.

these gains are relatively modest compared to those predicted by the authors behind the *GeoMatch* tool (41% for the U.S. and 73% for Switzerland for economic integration, respectively) (Bansak et al. 2018). Apart from the different

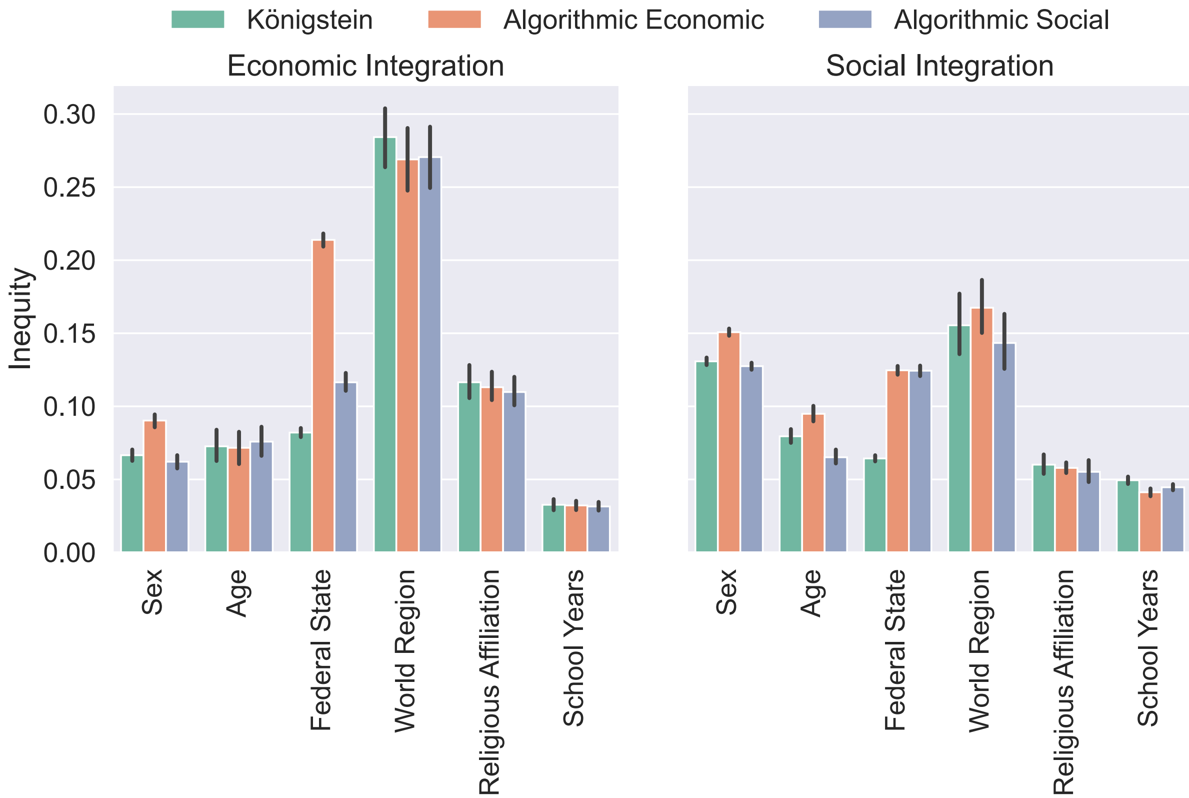

Figure 6: Average inequity for different subgroups achieved by algorithmic assignment versus *Königsteiner Schlüssel* regarding economic and social integration (error bars indicate 95% confidence intervals). Lower values indicate greater equity (more similar outcomes across members of different subgroups). We use binning for age (18-36, 36-57, 57-) and school years (1-7, 7-12, 12-).

context from which our data originates, and the additional data that we supply on the assignment locations, we hypothesize that the main cause of this disparity lies in the model architectures chosen. Unlike our approach, which relies on random effects to capture the variance of integration outcomes due to the assigned federal state within one predictive model, *GeoMatch* estimates separate models for each location, which—given varying sample sizes—likely introduces greater variance in the predicted outcomes and thus larger potential gains from optimized assignment. For our data, following this approach would have only yielded a dataset of 177 respondents assigned to the federal state with the smallest sample, Mecklenburg-Vorpommern, to create a training, validation, and test set, rendering the resulting model susceptible to statistical noise. Therefore, we opted to build a single unified predictive model for each integration dimension to ensure robust estimates across all states, given our sample-size constraints.

While the algorithmic schemes—especially *algorithmic economic*—yield promising overall gains, it is important to examine their geographic distribution. Figure 5 visualizes, for each federal state, the relative change in integration outcomes under the algorithmic allocation schemes compared to the *Königsteiner Schlüssel* baseline. The overall improvements are not uniform across states. Instead, *algorithmic economic* in particular concentrates high-potential agents in states where the model predicts the greatest success—boosting average economic integration by up to 60 % over the baseline in the top-performing states (i.e., those that had the highest share of successful agents in the dataset). Conversely, individuals with lower predicted probabilities are assigned to less favorable states (featuring, e.g., higher unemployment rates), further depressing average labor market outcomes in those locations. These disparities are less pronounced in the case of *algorithmic social*. While there are still sizable differences in the average outcomes recorded in the individual states, the strategy does not favor already previously successful locations to the same degree. For instance, North Rhine-Westphalia—which has the greatest loss of economic integration versus the baseline under *algorithmic economic*—sees the highest increase in social integration outcomes despite only having the seventh highest average under baseline allocation. Overall, however, the spatially targeted strategy of *algorithmic economic* underscores how algorithmic allocation can amplify regional disparities—improving outcomes where prospects are already strong while exacerbating challenges in lagging areas—highlighting the need to incorporate equity in the evaluation of assignment schemes.

For this purpose, Figure 6 compares the average intergroup inequity—as defined by equation 3.1—across key subgroups for both economic and social integration under the two algorithmic schemes versus the *Königsteiner Schlüssel* baseline. For the variables world region, religious affiliation, and school years, inequity levels remain statistically indistinguishable across all three allocation methods, indicating no increased unfairness for these attributes. In contrast, assignment through *algorithmic economic* significantly amplifies inequity by sex (outcomes further favor men), federal state (echoing the geographic disparities in Figure 5), and age (younger agents gain disproportionately higher social integration). The *algorithmic social* scheme produces more moderate inequity overall but still significantly increases interstate disparities for both integration dimensions. These results hint at potential unintended consequences of maximizing overall integration outcomes, suggesting that algorithmic allocation policies potentially lead to increased inequity across selected demographic groups for the integration dimensions examined here.

5 Discussion

Using algorithms in general and machine learning in particular to allocate refugees upon their arrival in a host country is an increasingly popular approach. In this paper, we presented an agent-based model calibrated using survey data to simulate the integration outcomes of refugees upon their allocation to federal states in Germany. We then used this simulation to conduct extensive experiments evaluating how algorithmic assignment schemes—optimized separately for economic and social integration—compare to Germany's standard *Königsteiner Schlüssel* regarding refugees' integration outcomes. Our key findings are:

1. Algorithmic assignment can improve both economic and social integration outcomes. In particular, optimizing for labor market success not only raises average employment probabilities by roughly 5% but also yields additional gains in social integration. Although these improvements are statistically significant, they remain considerably lower than those estimated for the U.S. or Swiss contexts (Bansak et al. 2018).
2. Performance gains are uneven across federal states and agent subgroups, reinforcing existing inequalities. The algorithmic schemes disproportionately benefit specific regions and, in the case of *algorithmic economic*, selected demographic groups (males and younger individuals), thereby amplifying subgroup inequities rather than eliminating them. This underscores that optimizing aggregate outcomes without fairness constraints risks entrenching disparities when implementing such systems in practice.
3. Feedback loops are essential to accurately modeling integration outlooks. Incorporating feedback loops—where each allocation round alters local community composition—profoundly affects subsequent integration predictions. Ignoring these dynamics risks overlooking significant shifts in refugees' prospects over time.

There are several limitations of our work to note. First, like all simulations, our model relies on simplifying assumptions that constrain its empirical validity. In particular, the geospatial resolution is limited to the federal-state level, which masks potentially significant variation at finer spatial scales (e.g., municipalities or districts) that is likely more impactful than differences on the state level (Tjaden and Spörlein 2025). Moreover, the mechanism implemented to adjust community sizes and model emergent changes caused by allocation decisions is relatively simple and could be improved based on more data and future research. Lastly, our study was conducted using survey data, which, in many relevant areas, most likely is less extensive and accurate than the official records that would be available if an algorithmic matching system were actually implemented in practice.

Second, although our predictive models replicate effects found in the literature, their ability to identify positive cases of labor market integration remains suboptimal, especially due to the class imbalance of economic integration in our data. This also impacts the probabilities fed to our simulation, thus further underscoring the relevance of the first limitation mentioned above.

Third, we evaluate model performance using a random train–validation–test split, which likely overestimates real-world accuracy compared to the time-based or rolling-window splits that a live matching system would require. Overall, our simulation and its results should be viewed as an indication of potential effects of different assignment schemes rather than a fully accurate representation of the complex and bidirectional social phenomenon that is the integration of refugees into a host society.

Beyond addressing these limitations, it remains for future research to systematically examine the potential gains and inequity challenges brought by algorithmic matching systems in different contexts and using different model architectures (i.e., one predictive model overall versus one per location). As our paper only included two (correlated) dimensions of integration, it would be interesting to examine other aspects of integration and how they affect each other following the assignment to a federal state. Nonetheless, it is important to note that the matching of refugees with locations to optimize any metric most likely will always come with disparate impacts for different subgroups (even though there have been solutions proposed that take refugees' (simulated) preferences into account (Jones and Teytelboym 2018)). Therefore, any allocation tool for refugees should be developed with a multitude of stakeholders (e.g., refugees themselves, practitioners, researchers, etc.) and tested rigorously before any rollout.

Overall, our findings highlight the critical trade-offs between maximizing overall integration outcomes and safeguarding equity, pointing to the need for allocation frameworks that explicitly attempt to balance efficiency with fairness. Furthermore, more broadly, they demonstrate how simulations can aid the evaluation of AI-based (and more generally, algorithm-based) social policies before implementing them. This is especially true in high-stakes domains—such as the assignment of refugees to locations in host countries—where decisions can have far-reaching consequences.

Ethics Statement

During our research, we confronted three key ethical challenges:

1. **Modeling simplifications**: Our agent-based simulation of refugee integration necessarily relies on simplifying assumptions—such as irreversibility of "integration success"—that cannot capture the full complexity of individual experiences. Likewise, grouping respondents by broad world-region categories assumes regional homogeneity but fails to reflect the diverse cultural and social distinctions that exist within those regions.
2. **Proxy measures**: The empirical research in this paper required operationalizing measures of integration with a specific data set, which included discussing and deciding on critical trade-offs and eventually constructing proxy variables that represent only a subset of possible dimensions and cannot encompass the full, multifaceted nature of integration.
3. **Data privacy**: In handling individual-level SOEP data on refugees, we strictly adhered to the SOEP's access protocols, signed a nondisclosure agreement, and refrained from publishing any identifying information to safeguard respondent privacy.

Adverse Impact Statement

We see two crucial potential unintended impacts of our research:

1. **Overconfidence in validity.** Ignoring simplifying assumptions made for building our agent-based model (see above) and the moderate accuracy of our classifiers risks readers taking simulated outcomes at face value and overestimating their real-world applicability. To combat this, we detail the limitations of our approach and stress that our findings should be viewed as an indication of potential effects instead of a fully accurate representation of the processes modeled.
2. **Misinterpretation of benefits.** The substantial predicted gains in employment and social integration under algorithmic assignment could be read as an endorsement of implementing such tools in practice. Therefore, we acknowledge that our simulation cannot capture the full complexity of integration processes and show in the course of our equity analysis that algorithmic schemes may exacerbate existing disparities among subgroups. To be clear, with the present paper, we do not advocate for using such technologies, but point out potential benefits and challenges that may arise from their implementation.

Acknowledgments

This work was conducted as part of the grants "sMArt roots – Smart City Modellstadt Mannheim", funded by the German Federal Ministry for Housing, Urban Development and Building under Grant No. 13622847, and "Consequences of Artificial Intelligence for Urban Societies (CAIUS)", funded by Volkswagen Foundation. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant INST 35/1597-1 FUGG. Furthermore, we would like to thank Irena Kogan for her valuable input, and Lena Arnold and our colleague Lea Cohausz, as well as the reviewers, for their helpful feedback on this paper.

References

Adelt, F.; Weyer, J.; Hoffmann, S.; and Ihrig, A. 2018. Simulation of the governance of complex systems (SimCo): basic concepts and experiments on urban transportation. *Journal of Artificial Societies and Social Simulation*, 21(2).

Ahani, N.; Andersson, T.; Martinello, A.; Teytelboym, A.; and Trapp, A. C. 2021. Placement Optimization in Refugee Resettlement. *Operations Research*, 69(5): 1468–1486.

Aksoy, C. G.; Poutvaara, P.; and Schikora, F. 2023. First time around: Local conditions and multi-dimensional integration of refugees. *Journal of Urban Economics*, 137: 103588.

Auer, D. 2018. Language roulette – the effect of random placement on refugees' labour market integration. *Journal of Ethnic and Migration Studies*, 44(3): 341–362.

Baba, L.; Schmandt, M.; Tielkes, C.; and Weinhardt, F. 2024. Evaluation der Wohnsitzregelung nach § 12a AufenthG. https://www.bamf.de/SharedDocs/Anlagen/DE/Forschung/Beitragsreihe/beitrag-band-13-evaluation-wohnsitzregelung.html?nn=283560. Accessed: 2025-05-17.

BAMF. 2024. Initial Distribution of Asylum Seekers (EASY). https://www.bamf.de/EN/Themen/AsylFluechtlingsschutz/AblaufAsylverfahrens/Erstverteilung/erstverteilung-node.html. Accessed: 2025-05-17.

Bansak, K.; Ferwerda, J.; Hainmueller, J.; Dillon, A.; Hangartner, D.; Lawrence, D.; and Weinstein, J. 2018. Improving refugee integration through data-driven algorithmic assignment. *Science*, 359(6373): 325–329.

Bansak, K.; and Marten, L. 2021. Algorithmic Decision-Making, Fairness, and the Distribution of Impact: Application to Refugee Matching in Sweden. *SOFI Working Papers in Labour Economics* 6/2024, Stockholm University, Swedish Institute for Social Research.

Barreto, C.; Berbée, P.; Torres, K. G.; Lange, M.; and Sommerfeld, K. 2022. The Civic Engagement and Social Integration of Refugees in Germany. *Nonprofit Policy Forum*, 13(2): 161–174.

Battisti, M.; Peri, G.; and Romiti, A. 2022. Dynamic effects of co-ethnic networks on immigrants' economic success. *The Economic Journal*, 132(641): 58–88.

Biewald, L. 2020. Experiment Tracking with Weights and Biases. https://www.wandb.com/. Accessed 2025-05-14.

Brücker, H.; Kosyakova, Y.; Rother, N.; Zinn, S.; Liebau, E.; Gider, W.; Schwanhäuser, S.; and Siegert, M. 2025. Exploring Integration and Migration Dynamics: The Research Potentials of a Large-Scale Longitudinal Household Study of Refugees in Germany. *European Sociological Review*, jcaf032.

Castles, S.; Korac, M.; Vasta, E.; and Vertovec, S. 2003. *Integration: Mapping the Field.* Home Office. Research, Development and Statistics Directorate.

Chouldechova, A.; and Roth, A. 2020. A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5): 82–89.

Cohen, J. 2009. *Statistical power analysis for the behavioral sciences*. New York, NY: Psychology Press, 2nd ed., reprint edition.

Cruz Cortés, E.; and Ghosh, D. 2019. A Simulation Based Dynamic Evaluation Framework for System-wide Algorithmic Fairness. arXiv:1903.09209.

Cruz Cortés, E.; and Ghosh, D. 2020. An Invitation to System-wide Algorithmic Fairness. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, 235–241. New York, NY, USA: Association for Computing Machinery.

D'Amour, A.; Srinivasan, H.; Atwood, J.; Baljekar, P.; Sculley, D.; and Halpern, Y. 2020. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 525–534. New York, NY, USA: Association for Computing Machinery.

Destatis. 2024a. Arbeitslose, Arbeitslosenquoten, Gemeldete Arbeitsstellen: Bundesländer, Jahre. https://www-genesis.destatis.de/datenbank/online/statistic/13211/table/13211-0007. Accessed: 2025-01-29.

Destatis. 2024b. Ausländer: Bundesländer, Stichtag, Geschlecht/Altersjahre/ Familienstand, Ländergruppierungen/Staatsangehörigkeit. https://www-genesis.destatis.de/datenbank/online/statistic/12521/table/12521-0021/. Accessed: 2025-01-29.

Deutsch, M. 1975. Equity, Equality, and Need: What Determines Which Value Will Be Used as the Basis of Distributive Justice? *Journal of Social Issues*, 31(3): 137–149.

DIW. 2024. IAB-BAMF-SOEP-Befragung Geflüchteter 2022, Daten der Jahre 2016-2022. 10.5684/soep.iab-bamf-soep-mig.2022.

Endres, D.; and Schindelin, J. 2003. A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7): 1858–1860.

ESS Round 11: European Social Survey European Research Infrastructure (ESS ERIC) . 2024. ESS11 - integrated file, edition 2.0 [Data set]. Sikt - Norwegian Agency for Shared Services in Education and Research. https://doi.org/10.21338/ess11e02_0.

Freund, D.; Lykouris, T.; Paulson, E.; Sturt, B.; and Weng, W. 2023. Group fairness in dynamic refugee assignment. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, 701. New York, NY, USA: Association for Computing Machinery.

Gerdon, F.; Bach, R. L.; Kern, C.; and Kreuter, F. 2022. Social impacts of algorithmic decision-making: A research agenda for the social sciences. *Big Data & Society*, 9(1): 20539517221089305.

Global Compact on Refugees. 2024. GeoMatch: Connecting People to Places Using Artificial Intelligence. https://globalcompactrefugees.org/good-practices/geomatch-connecting-people-places-using-artificial-intelligence. Accessed: 2024-11-05.

Hahn, E.; Richter, D.; Schupp, J.; and Back, M. D. 2019. Predictors of refugee adjustment: The importance of cognitive skills and personality. *Collabra: Psychology*, 5(1): 23.

Hannafi, C.; and Marouani, M. A. 2023. Social integration of Syrian refugees and their intention to stay in Germany. *Journal of Population Economics*, 36(2): 581–607.

Harder, N.; Figueroa, L.; Gillum, R. M.; Hangartner, D.; Laitin, D. D.; and Hainmueller, J. 2018. Multidimensional measure of immigrant integration. *Proceedings of the National Academy of Sciences*, 115(45): 11483–11488.

Hunkler, C.; Edele, A.; and Schipolowski, S. 2021. The role of educational resources in the labor market integration of refugees. The case of Syrian asylum seekers in Germany. *Journal for educational research online*, 13(1): 157–181.

Immigration Policy Lab. 2025. GeoMatch. https://immigrationlab.org/geomatch/. Accessed: 2025-05-17.

Jaxa-Rozen, M.; and Kwakkel, J. H. 2018. PyNetLogo: Linking NetLogo with Python. *Journal of Artificial Societies and Social Simulation*, 21(2): 4.

Jones, W.; and Teytelboym, A. 2018. The Local Refugee Match: Aligning Refugees' Preferences with the Capacities and Priorities of Localities. *Journal of Refugee Studies*, 31(2): 152–178.

Kanas, A.; Chiswick, B. R.; Van Der Lippe, T.; and Van Tubergen, F. 2012. Social Contacts and the Economic Performance of Immigrants: A Panel Study of Immigrants in G ermany. *International Migration Review*, 46(3): 680–709.

Kappenberger, J.; Stuckenschmidt, H.; and Gerdon, F. 2025. Pricing parking for fairness — A simulation study based on an empirically calibrated model of parking behavior. *Transportation Research Part A: Policy and Practice*, 193: 104389.

Khalil, S.; and Tjaden, J. 2025. How reception centers affect the integration of asylum seekers and recognized refugees. *European Societies*, 1–32.

Kosyakova, Y.; and Brenzel, H. 2020. The role of length of asylum procedure and legal status in the labour market integration of refugees in Germany. *Soziale Welt*, 71(1/2): 123–159.

Kosyakova, Y.; and Kogan, I. 2022. Labor market situation of refugees in Europe: The role of individual and contextual factors. *Frontiers in Political Science*, 4: 977764.

Kosyakova, Y.; and Kogan, I. 2024. The role of social capital in immigrants' and refugees' labour market integration: evidence from Germany. In McDonald, S.; Côté, R.; and Shen, J., eds., *Handbook on Inequality and Social Capital*, 288–304. Edward Elgar Publishing.

Kosyakova, Y.; Salikutluk, Z.; and Hartmann, J. 2023. Gender employment gap at arrival and its dynamics: The case of refugees in Germany. *Research in Social Stratification and Mobility*, 87: 100842.

Kuschel, A.; Hansen, N.; Heyse, L.; and Wittek, R. P. 2023. Combining Language Training and Work Experience for Refugees with Low-Literacy Levels: A Mixed-Methods Case Study. *Journal of International Migration and Integration*, 24(4): 1635–1661.

Laird, N. M.; and Ware, J. H. 1982. Random-Effects Models for Longitudinal Data. *Biometrics*, 38(4): 963.

Landesmann, M.; and Leitner, S. M. 2025. Various Domains of Integration of Refugees and Their Interrelationships: A Study of Recent Middle Eastern Refugee Inflows in Austria. *European Journal of Population*, 41(1): 6.

Lundberg, S. M.; and Lee, S.-I. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, 4768–4777. Red Hook, NY, USA: Curran Associates Inc.

Martén, L.; Hainmueller, J.; and Hangartner, D. 2019. Ethnic networks can foster the economic integration of refugees. *Proceedings of the National Academy of Sciences*, 116(33): 16280–16285.

Marton, S.; Lüdtke, S.; Bartelt, C.; and Stuckenschmidt, H. 2024. GRANDE: Gradient-Based Decision Tree Ensembles. In *The Twelfth International Conference on Learning Representations*.

McElfresh, D.; Khandagale, S.; Valverde, J.; C., V. P.; Ramakrishnan, G.; Goldblum, M.; and White, C. 2023. When do neural nets outperform boosted trees on tabular data? In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, 76336–76369. Red Hook, NY, USA: Curran Associates Inc.

Menéndez, M.; Pardo, J.; Pardo, L.; and Pardo, M. 1997. The Jensen-Shannon divergence. *Journal of the Franklin Institute*, 334(2): 307–318.

Mitchell, S.; O'Sullivan, M.; and Dunning, I. 2011. PuLP: A Linear Programming Toolkit for Python. https://optimization-online.org/2011/09/3178/. Accessed: 2025-05-13.

Müller, T.; Pannatier, P.; and Viarengo, M. 2023. Labor market integration, local conditions and inequalities: Evidence from refugees in Switzerland. *World Development*, 170: 106288.

Provost, F.; and Fawcett, T. 2001. Robust Classification for Imprecise Environments. *Machine Learning*, 42(3): 203–231.

Scher, S.; Kopeinik, S.; Trügler, A.; and Kowald, D. 2023. Modelling the long-term fairness dynamics of data-driven targeted help on job seekers. *Scientific Reports*, 13(1): 1727.

Schilling, P.; and Stillman, S. 2024. The impact of natives' attitudes on refugee integration. *Labour Economics*, 87: 102465.

Shapley, L. S. 1953. A Value for n-Person Games. In Kuhn, H. W.; and Tucker, A. W., eds., *Contributions to the Theory of Games (AM-28), Volume II*, 307–318. Princeton University Press.

Sigrist, F. 2023. Latent Gaussian Model Boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 1894–1905.

Stips, F.; and Kis-Katos, K. 2020. The impact of co-national networks on asylum seekers' employment: Quasi-experimental evidence from Germany. *PLOS ONE*, 15(8): e0236996.

Tjaden, J.; and Spörlein, C. 2025. How Much Do "Local Policies" Matter for Refugee Integration? An Analytical Model and Evidence from a Highly Decentralized Country. *International Migration Review*, 59(3): 1194–1218.

UNHCR. 2013. The Integration of Resettled Refugees–Essentials for Establishing a Resettlement Programme and Fundamentals for Sustainable Resettlement Programmes. https://www.unhcr.org/media/integration-resettled-refugees-essentials-establishing-resettlement-programme-and. Accessed: 2025-04-15.

UNHCR. 2024a. Asylum-Seekers. https://www.unhcr.org/about-unhcr/who-we-protect/asylum-seekers. Accessed: 2024-11-11.

UNHCR. 2024b. Refugees. https://www.unhcr.org/about-unhcr/who-we-protect/refugees. Accessed: 2024-11-11.

United Nations. 1999. Standard Country or Area Codes for Statistical Use. https://unstats.un.org/unsd/publication/SeriesM/Series_M49_Rev4(1999)_en.pdf.

van Riemsdijk, M.; and Axelsson, L. 2021. Introduction "Labour market integration of highly skilled refugees in Sweden, Germany and the Netherlands". *International Migration*, 59(4): 3–12.

Vogiazides, L.; and Mondani, H. 2020. A geographical path to integration? Exploring the interplay between regional context and labour market integration among refugees in Sweden. 46(1): 23–45.

Walther, L.; Fuchs, L. M.; Schupp, J.; and von Scheve, C. 2020. Living Conditions and the Mental Health and Well-being of Refugees: Evidence from a Large-Scale German Survey. *Journal of Immigrant and Minority Health*, 22(5): 903–913.

Wilensky, U. 1999. NetLogo. http://ccl.northwestern.edu/netlogo/. Accessed: 2025-05-13.

Zwingel, A. M. 2023. Co-Ethnic Networks and the Labor Market Integration of Refugees: Empirical Evidence from a Natural Experiment in Germany. Available at SSRN. https://dx.doi.org/10.2139/ssrn.4566650l. Accessed: 2025-05-17.