

PREJUDICE IN SWISS NATURALIZATION DECISIONS: THEORY AND EVIDENCE

Dragan Ilić*

February 15, 2013

Abstract

Recent empirical studies have revealed prejudice based on country of origin in the Swiss naturalization system before courts banned closed ballot voting in 2003. Although the switch to elected councils has ameliorated the situation for the discriminated applicant groups, little has been known about the issue since due to lack of microlevel data. This paper presents an alternative approach. I propose a simple model of a council deciding whether to grant applicants citizenships. The model implies an outcome test for relative prejudice using rejection rates grouped by country of origin. Importantly, this specific outcome test does not suffer from the well-known infra-marginality problem. The test is implemented with unique data from 2003 to 2012 raised from large Swiss municipalities to test the hypothesis of no relative prejudice.

*University of Basel, Faculty of Business and Economics, Peter Merian-Weg 6, 4002 Basel, Switzerland and University of Zurich, Center for Corporate Responsibility and Sustainability, Zähringerstrasse 24, 8001 Zürich, Switzerland (e-mail: dragan.ilic@unibas.ch). I am grateful to Georg Nöldeke and Brigitte Guggisberg for their help on this project. I would like to thank all participating councils for their time and effort in providing the data. Funding for this research was generously provided by the WWZ Forum.

1 INTRODUCTION

Roughly 40,000 immigrants - two percent of the foreign population - are currently naturalized in Switzerland every year. Since Swiss citizenship is not granted based on the place of birth, immigrants have to actively apply for citizenship. It is characteristic of Switzerland that municipalities themselves govern and decide on naturalization applications according to distinctive standards. This decentralized system keeps giving rise to intense debates and has sparked allegations that certain immigrant groups are treated unfairly. In particular, it is a vexing fact for critics of the current system that naturalization outcomes vary widely, both across applicants of different country of origin and, across municipalities, for applicants of a given country of origin.

Are these differences brought about by prejudiced decision-making or can they be traced to objective causes such as varying applicant characteristics or municipal-specific requirements? The answer is of high importance to the political and judicial discourse. If prejudice is the reason for the disparities in naturalization outcomes, it raises questions about the ability of the current federal structure to address this issue. Indeed, time and again there have been endeavors to centralize the Swiss naturalization system. Alleged prejudice was one motivation for this political venture, contending that the federal system allows too much discretionary leeway. On the other hand, if the differences can be traced to objective causes, current practice had better be cleared of allegations of prejudice. Instead, any remedial policies to address the disparities are better off concentrating on fostering the chances from the applicants' point of view.

Recent studies suggest that Switzerland used to discriminate heavily based on country of origin. Before 2003, many municipalities employed closed ballot votings to decide on naturalization applications. In line with the Swiss practice of municipal citizenship, all local Swiss citizens were eligible to cast their vote on any naturalization application in their municipality. Only applicants with a popular majority of "yes" votes received Swiss citizenship. In a comprehensive study, Hainmueller and Hangartner (2013) collected

data on 2,400 recorded naturalization referenda held between 1970 and 2003 in 44 Swiss municipalities using closed ballots. The analysis shows that, accounting for a myriad of applicant characteristics deemed relevant for naturalization, the proportion of "no" votes were about 40% higher on average for applicants from former Yugoslavia and Turkey when compared to applicants from countries like Germany or the UK. Hainmueller and Hangartner attribute the discriminatory treatment to lack of accountability associated with the nature of closed ballots. This conjecture is supported in a follow up study, where the same authors compare naturalization rates before and after 2003 (Hainmueller and Hangartner, 2012). That year the Swiss Federal Court declared closed ballot voting illegal for naturalization applications, mandating it to be replaced with elected councils. The ban was based on two grounds. First, the right to appeal rejected applications was not ensured. Closed ballots, by definition, lack the basis for any justification. Second, the court disapproved of the severe lack of privacy in the evaluation of the applicants, who had to reveal detailed background information to all eligible voters, usually in form of leaflets. Hainmueller and Hangartner demonstrate that the switch to elected councils increased naturalization rates of formerly discriminated applicants on average by 50%, more so in municipalities with strong anti-immigrant preferences. This is interpreted as evidence that, once accountable legislators instead of anonymous voters are in charge of naturalization requests, indulgence of discriminatory preferences seems to become more costly and thus less likely.

Little has been known about the existence of malevolent discrimination in the naturalization process since. Disparate treatment does seem to have decreased by delegating legislation to elected councils, but has it vanished altogether? Ironically enough, the very same reason why the court opposed closed ballots allowed for the identification of discriminatory treatment before 2003 in the first place. Hainmueller and Hangartner (2013) overcome the causal identification problem in their empirical strategy because they have access to the same data as the voters themselves had. These data were meticulously

retrieved from the official voting leaflets that contain detailed applicant background information. This poses a problem for the identification of discriminatory treatment after 2003. In elected councils, applicant data is not readily accessible to the statistician. Even if, in contrast to closed ballots, most councils conduct personal interviews during the course of the application. It is unlikely that the researcher can gain access to all the relevant information gathered during these interviews, let alone quantify it. Omitted variable bias thus reappears in the empirical analysis, impeding the unambiguous attribution of disparate naturalization outcomes to differences in observed applicant characteristics.

This paper proposes an outcome test for prejudice in naturalization decisions that can be empirically applied in the presence of these inferential issues. The test is based on a simple model of a council's naturalization decision, adapting a theoretical approach by Anwar and Fang (2006). Applicants are either qualified or unqualified in terms of requirements for naturalization. The council weighs the benefit of naturalizing qualified applicants against the cost of mistakenly naturalizing unqualified applicants. In addition to country of origin, the council observes a noisy but informative signal about the merit of each applicant. Qualified applicants are more likely than unqualified applicants to send a favorable signal. The council's optimal decision is characterized by setting a minimum qualification standard. Only those applicants whose signal exceeds a specific level of required qualification for naturalization will be granted citizenship. The model implies that if councils are not prejudiced against applicants based on country of origin, the rankings of the rejection rates over applicants of given countries of origin should be the same in any council.

The proposed test has the advantage that it does not require comprehensive microlevel data. Instead, it is implemented with average rejection rates, data which are more easily obtainable in comparison. Notably, the notorious infra-marginality problem does not affect this specification of outcome tests. Infra-marginality cautions that in general, outcome tests are based on average outcome data. However, it is the treatment at the

margin that identifies malevolent discrimination. In the context of naturalizations, the question is whether the marginal applicant who was deemed just qualified enough to be granted citizenship has the same expected probability of qualification across all countries of origin. Only then the council has chosen the same benchmark for all groups. A mere comparison of average rejection rates, then, would be misleading as they do not imply equality of marginal rejection rates. My model offers a way around this issue by exploiting ordinality conditions of average rejections rates across councils.

The peculiarity of the Swiss naturalization system allows for an empirical implementation of this test. In Section 6, I apply the proposed test to data raised from large Swiss municipalities. Data collection is still in progress but is expected to be completed in the upcoming weeks. All Swiss municipalities with more than 20'000 residents as of 2010 have been contacted and asked for data on filed and eventually granted citizenships from 2003 to 2011 grouped by country of origin, if applicable on a yearly basis. Of the 22 German-speaking municipalities, six have provided detailed information, four more municipalities have agreed to do so. In addition to Hainmueller and Hangartner, data from the French-speaking region will also be analyzed. The 11 French-speaking municipalities have just recently been contacted in a second survey wave and have yet to respond. Results for existing data will be provided by the time of the conference.

The rest of the paper is structured as follows. In Section 2 I present the literature which is both contentually and formally related. Section 3 describes the Swiss naturalization system. Section 4 presents a simple model of how a council decides on naturalization applications. The model implies an outcome test for prejudice in these decisions. In addition, I review the infra-marginality problem and show why my outcome test is immune to this issue. Section 5 presents the data raised from large Swiss municipalities which are used for the empirical implementation of the test. Results of the empirical analysis are presented in Section 6. In that section, I also address two caveats associated with the test. First, I dwell on the power of the test by showing how the test fares with the

Hainmueller and Hangartner data. Second, I show why an assumption of the model - councils facing applicant pools of similar quality - is plausibly satisfied in my data by analyzing appropriate empirical proxies. Finally, Section 7 concludes.

2 RELATED LITERATURE

The related literature can be grouped into two branches. The first body of literature is related in content and addresses the general question whether immigrants are exposed to discrimination by natives. Hainmueller and Hangartner (2013) provide an excellent overview of recent studies attempting to answer this question. One of the main drawbacks of the existing literature is the dearth of consideration for revealed preferences as the empirical basis. Instead, many studies are based on survey data on reported attitude towards immigrants. What is more, most of these reports do not distinguish between immigrants of different background or educational skill. It stands to reason that studies based on these data provide insufficient evidence needed for a stringent argument of disparate treatment.

Recently, randomized experiments have been applied as a compelling methodology to address this question. Bertrand and Mullainathan (2004) present a seminal audit study that demonstrates how in the United States, job applicants with names predominantly found in the black community suffer from worse callback-rates than job applicants with "white" sounding names. Kaas and Christian (2012) use the same approach for skilled job openings in Germany and come to the conclusion that while Turkish names are a detrimental signal in the application process in comparison to German names, the handicap disappears once a letter of recommendation is included. Finally, Behagel, Crépon and Le Barbanchon (2011) conduct a similar study in France, where surprisingly minorities fare worse once their name was rendered anonymous. Randomized experiments have the advantage of a conclusive pathway describing the cause for disparate outcomes. Un-

fortunately, the methodology is not applicable in many areas of public interest in which discriminatory treatment is of concern, such as naturalization.

Hainmueller and Hangartner (2013) put forth a clever identification strategy to assess discrimination in naturalization decisions. They exploit the fact that in Switzerland, municipalities decide independently on naturalization applications. Until 2003, many did so in referendums using voting leaflets that, in detail, describe characteristics of the applicant. These referendums were held anonymously using secret ballots. All Swiss citizens of the deciding council were eligible to cast their vote. These data were raised to analyze how applicant and council characteristics affect the decision for naturalization. Among the applicant characteristics, country of origin turns out to be by far the most important factor for naturalization. Being of Turkish or (ex-)Yugoslavian descent roughly doubles the average probability of being denied citizenship. Hainmueller and Hangartner highlight three advantages of their empirical strategy. First, they have access to actual behavior towards immigrants. Moreover, voters did not have to fear any backlash of their decision because the voting process protected their identity. Second, the number of immigrant characteristics and the observed degrees described in the leaflets paint a sophisticated picture of the particular attributes that affect voting behavior across councils and time. And third, perhaps the most important advantage, the approach minimizes omitted variable bias. Because of the leaflets, the researchers have the same data at hand as the voters did when making their decisions.

Concern of omitted variable bias is also what sparked the second strand of related literature. In light of limited information that researchers usually face when analyzing potentially disparate treatment, many have turned to outcome tests. These tests have the unique advantage that they are not subject to the omitted variable bias critique. They rely on the assumption that decision-makers have already considered all possible information that affects the desired outcome. For this reason, any indication of disparate treatment is reflected in the outcome. On this note, Becker (1993) argued that the

fact that some banks deny black applicants mortgages at a higher rate does not prove prejudice. Instead, one should look at the outcomes of the granted mortgages. Do the default rates differ by applicant group? If yes, an unbiased bank would do better by adjusting the required level of creditworthiness. Biased banks, on the other hand, would seem to close the profitable gap in default rates with utility drawn from racial animus. This line of thinking has consequently been applied to other areas in which a principal expects certain outcomes from an agent, such as bail bond settings (Ayres and Waldfogel (1994)), paper acceptance rates (Smart, Shoven and Waldfogel (1996)), and even organ transplantations (Ayres (2005)). However, the so-called infra-marginality critique severely limits the inferences one can draw from outcome tests. This critique is explained in detail in Section 4.4. In a nutshell, it stresses that group-specific distributions of factors on which the decision in an outcome test is based on might not be equal, so average outcomes have little to say about marginal decision-making. Yet only disparate treatment at the margin implies prejudice.

The second strand of related literature embraces this critique and comprises economic models that imply outcome tests not susceptible to infra-marginality. These models describe circumstances under which average outcome data can be used to draw valid inferences about disparate treatment. A seminal model was proposed by Knowles, Persico and Todd (2001) and lets police officers and motorists interact rationally in motor vehicle searches aiming to uncover engagement in criminal activity. In the equilibrium of this game, the observable behavior of the average motorist is equivalent to the behavior of the marginal motorist, so that average outcome data is applicable in a test of disparate police treatment against white and black motorists. Related theoretical frameworks to tackle the infra-marginality problem in outcome tests have recently been applied to healthcare (Anwar and Fang (2013)), capital sentencing (Alesina and La Ferrara (2011)), and parole releases for prisoners (Mechoulan and Sahuguet (2011); Anwar and Fang (2012)).

In direct relation to the model presented in this paper, Anwar and Fang (2006) (hence-

forth AF) suggest a more general model to test for racial prejudice in motor vehicle searches. They caution that if a racially divided police force exhibits specific tastes for discrimination, existing tests may lead to false conclusions on the absence or presence of racial prejudice. Their model shows that if police groups have distinct costs of searching (and are thus nonmonolithic), the observed ranking of the search and search success rates against a given motorist race across officer race gives rise to an alternative test for racial prejudice. AF apply their test to data from the Florida Highway Patrol and cannot reject the hypothesis that officers of different races do not exhibit relative racial prejudice.

AF's approach lends itself particularly well to assess prejudice in Swiss naturalization decisions. By modeling the police force as potentially nonmonolithic, AF acknowledge autonomous decision-making based on distinctive characteristics, which is a good description of council autonomy in Switzerland. The next section provides a brief overview over the Swiss naturalization system to motivate the plausible adaption of AF's framework in Section 4.

3 NATURALIZATIONS IN SWITZERLAND

The Swiss political system is characterized by direct democracy and a strong concern for federalism. Helbling (2008) and Hainmueller and Hangartner (2013) provide comprehensive overviews on citizenship policy in Switzerland. In this section, I emphasize the features which resonate with the model proposed in the next section.

Citizenship in Switzerland is regarded as an act of state that is to be primarily delegated to the deepest of the three political levels, the municipality. Interested immigrants have to apply at the municipal level. The requests are forwarded to the cantonal and federal level, which conduct formal checks whether the application fulfills the legal requirements.¹ Upon positive feedback, the applicant has yet to convince the municipality

¹Articles 14 and 15 of the Swiss Naturalization law leave some room for interpretation but list the following basic requirements: Overall 12 years of residency, integration into and familiarity with Swiss

of his or her merit. The municipality's decision is based on personal aspects of the applicant's characteristics and is by definition more subjective and open to interpretation. Integration to Swiss life assessed by familiarity with local laws, traditions, and customs is a key shared by all councils. These factors are proxied by measurable indicators like language skills, job status, or hobbies. It falls to each council, however, to which extent these requirements need to be met in order to be granted citizenship.

In economic terms, Swiss citizenship is a valuable signal of qualification for immigrants. Because even children of immigrants born in Switzerland do not receive citizenship by default, the needed qualifications associated with successful naturalization render citizenship indicative of underlying qualities. In this sense, a Swiss passport allows to classify immigrants of given origin. Particularly residents from disadvantaged backgrounds would plausibly profit from such a signal, for example in the labor or housing market (Bertrand and Mullainathan (2004); Kaas and Christian (2012); Yinger (1996*a*)).

Since the Swiss Federal court banned direct democracy in form of closed ballot votes as means for naturalization in 2003, most municipalities have switched to elected councils to decide on naturalization requests. Hainmueller and Hangartner (2012) show that by 2010, 60% of the municipalities in their sample have switched from direct democracy to elected councils. These councils are made up of local legislators and are elected to typically serve terms of four years (Helbling (2008)), upon which they are eligible for re-election. After thorough document revision, the council conducts face-to-face interviews with the applicant to complete their picture of naturalization merit. Because the council is elected by local citizens, it reflects their local attitude towards immigration and naturalization. It is therefore expected that municipalities still disclose distinct naturalization patterns despite the switch to elected councils. In contrast to a closed ballot vote, however, the council has to justify and document their decision-making process and must grant the applicant the right to appeal a rejection. Consistent with this heightened accountability, customs and tradition, adherence to Swiss law, and lack of threat for the internal and external security of Switzerland.

Hainmueller and Hangartner (2012) show that naturalization rates have soared after the ban of closed ballots.

It can be argued that while council members may still have a taste for discrimination, they act less on this taste because the costs of doing so have risen. It is unclear, however, whether disparate treatment has ceased entirely. Unfortunately, the empirical strategy employed by Hainmueller and Hangartner (2013) is not applicable when assessing naturalization decisions in elected councils. In contrast to the voting leaflets in closed ballots, microlevel data on applicant characteristics are not released to the public in order to protect the applicant's right to privacy. In fact, this was one of the reasons for the ban of closed ballots in the first place. The emphasis on personal interviews accompanied by the switch to elected councils exacerbates the application of Hainmueller and Hangartner's approach. Even with microlevel data at hand, it is arguably difficult to quantify the impression an applicant makes during the interview. This is particularly crucial because these interviews are likely decisive to the council's decision to naturalize.

In light of these limitations, in the next section I propose a simple model of a council's decision problem that implies an outcome test for prejudice in naturalization decisions. In contrast to a comprehensive empirical approach, the test only requires minimal data. The model relies on autonomous councils that draw distinctive benefits from naturalization, a feature that is reflected in the Swiss citizenship policy.

4 A MODEL OF NATURALIZATION DECISIONS

4.1 THE MODEL

Consider councils that separately evaluate their local immigrants which apply for naturalization. Let $c \in \{A, B\}$ denote the councils. In each council, there are continuums of applicants grouped into country of origin $e \in \{R, F\}$, where R and F stand for re-

lated and foreign, respectively.² Suppose that among applicants origin e , a fraction α_e is objectively unqualified for naturalization.

Councils evaluate the merit for naturalization based on the applicant's origin and a myriad of other applicant characteristics. A council may consider information such as gender, age, number of children, language skills, familiarity with local habits and law, duration of residence, employment status, or level of education and income. Importantly, a council also processes characteristics that are difficult to observe for a researcher, such as demeanor or congeniality during personal interviews.

Let us assume that a council condenses all this information into a one-dimensional index $\theta \in [0, 1]$ which reflects the likelihood that an applicant is unqualified for naturalization. This index is randomly drawn from a continuous density function $f_u^e(\cdot)$ if the applicant is actually unqualified. If the applicant is qualified, the index is drawn from $f_q^e(\cdot)$. For this index to be indicative of qualification, the two densities are assumed to satisfy the strict monotone likelihood ratio property, that is, for $e \in \{R, F\}$, $f_u^e(\theta)/f_q^e(\theta)$ is strictly increasing in θ . This implies $F_q^e(\theta) < F_u^e(\theta)$ for all θ . In words, higher values of the signal θ are more likely if the applicant is unqualified. Some applicants may produce signals which will unambiguously lead to rejections, such as a criminal record. I thus assume an unbounded likelihood ratio: $f_u^e(\theta)/f_q^e(\theta) \rightarrow +\infty$ as $\theta \rightarrow 1$.

After interviewing an applicant of country of origin e with signal θ , a council decides whether to grant citizenship or not. While councils only see imperfect signals during the course of the application, they do eventually realize whether their decision was correct. Councils derive the benefit $b(c, e)$ from naturalizing a qualified applicant. At the same time, councils bear the marginal cost t when mistakenly naturalizing an unqualified applicant. Note that the benefit can depend both on the deciding council and on the applicant's country of origin.³

Councils may have a taste to prefer applicants of certain origin to be naturalized (or

²The model readily extends to n councils and m immigrant groups.

³For the sake of simplicity, I assume t to be independent of council and country of origin.

conversely, dislike applicants of certain origin to become citizens). Based on this taste, a given council whose benefit depends on the applicant's country of origin is said to be prejudiced:

$$b(c, R) \neq b(c, F)$$

On the other hand, councils may derive different levels of benefit in general from naturalizing qualified applicants, levels that do not depend on country of origin. Such differences in benefits could stem, for instance, from varying identity preferences based on group distinction, high council standards associated with the merit of citizenship, or a general tendency towards xenophobia. Define councils to be heterogenous if

$$b(A, e) \neq b(B, e)$$

for some e . Councils that derive little benefit from naturalization are said to be strict. Likewise, councils that derive much benefit are said to be lenient. It is easy to see that heterogenous councils do not imply prejudice against a given country of origin. By the same token, homogenous councils do not imply the lack of prejudice as all councils might equally prefer applicants from certain backgrounds.

4.2 THEORETICAL IMPLICATIONS

Denote by U the undesirable event that a naturalized applicant turns out to be unqualified. The probability of this event depends both on the signal θ observed during the interview and on country of origin. Following Bayes' rule, this probability is given by

$$\Pr(U | e, \theta) = \frac{\alpha_e f_u^e(\theta)}{\alpha_e f_u^e(\theta) + (1 - \alpha_e) f_u^e(\theta)}. \quad (1)$$

The monotone likelihood ratio property implies that this probability strictly increases in θ . Since the signal is informative, a higher level correctly reflects in an increase in the

mistake probability.

Now consider the decision problem of a council faced with this information:

$$\max \{b(c, e) [1 - \Pr(U | e, \theta)] - t \Pr(U | e, \theta); 0\}$$

The first term describes the expected benefit from naturalizing a qualified applicant minus the cost of mistakenly naturalizing an unqualified applicant. Not naturalizing yields a benefit of zero. The costs associated with the naturalization process itself are considered fixed and are thus disregarded.

Thus, the council does best to naturalize if and only if

$$b(c, e) [1 - \Pr(U | e, \theta)] > t \Pr(U | e, \theta),$$

in words, whenever the expected benefit of naturalizing outweighs the expected cost.

This naturalization condition reduces to

$$\Pr(U | e, \theta) < \frac{b(c, e)}{b(c, e) + t} \tag{2}$$

Intuitively, a high benefit of naturalization makes for a riskier behavior in the sense that the council has more leeway for a higher probability of mistakes. On the other hand, high costs of making mistakes does not leave a lot of room for suspicious signals.

Because $\Pr(U | e, \theta)$ is strictly increasing in θ , the naturalization condition described by (2) implies that the council grants citizenship if and only if

$$\theta \leq \theta^*(c, e),$$

where the naturalization threshold $\theta^*(c, e)$ is pinned down by

$$\Pr(U \mid e, \theta^*(c, e)) = \frac{b(c, e)}{b(c, e) + t}. \quad (3)$$

The applicant sending the signal $\theta^*(c, e)$ is called the marginal applicant who is deemed just worthy enough to be granted citizenship. Any applicant with a higher signal will be rejected. Likewise, any applicant with a lower signal will be naturalized. It is straightforward to see that $\theta^*(c, e)$ is strictly increasing in $b(c, e)$ and strictly decreasing in t . The higher the benefit of naturalization, the worse the qualification signal is allowed to become. In turn, higher costs imply a stricter expected quality conveyed by the signal.

The signal threshold $\theta^*(c, e)$ determines the average rejection rate of council c against applicants of origin e in equilibrium:

$$r(c, e) = \alpha_e [1 - F_u^e(\theta^*(c, e))] + (1 - \alpha_e) [1 - F_q^e(\theta^*(c, e))] \quad (4)$$

This rejection rate is monotonically decreasing in $\theta^*(c, e)$: The worse the signal is allowed to become, the less applicants are rejected.

Note that a council that is not prejudiced might nevertheless use different signal thresholds for the applicant groups. An unprejudiced council does not aim to set equal signal thresholds but equal probabilities of qualification for the marginal applicants. If one applicant group is known to have a larger unqualified fraction or if the signal θ is distributed differently between the two applicant groups, councils take this information into account by making use of statistical discrimination. This implication is an inherent part of models of statistical discrimination. Equation (1) illustrates that because the applicant's qualification is not perfectly observed, a council's optimal assessment about the qualification of a given applicant does not solely depend on that applicant's signal. The assessment also considers the fraction of qualified applicants in that group.

Also note that if councils are homogenous they all derive the same benefit from a

given applicant group. The threshold condition (3) then implies that all councils set the same signal threshold for this group. Consequently, the rejection rate defined by (4) against this particular applicant group is the same for all councils.

Based on this insight, consider the following simple transitivity example. Assume $b(A, R) > b(B, R)$ applies so that council A derives a larger benefit from naturalizing related immigrants than council B does. If the two councils are not prejudiced, it must be true that $b(A, R) = b(A, F)$ and $b(B, R) = b(B, F)$. It immediately follows that $b(A, F) > b(B, F)$ so that council A also derives a larger benefit from naturalizing foreign immigrants. In other words, if the councils are heterogenous but not prejudiced, the ranking of $b(A, e)$ and $b(B, e)$ does not depend on country of origin.

What can be said about the signal thresholds? We know that $\theta^*(c, e)$ is strictly increasing in $b(c, e)$, so above example implies $\theta^*(A, R) > \theta^*(B, R)$ and $\theta^*(A, F) > \theta^*(B, F)$. Moreover, because the rejection rate $r(c, e)$ is monotonically decreasing in $\theta^*(c, e)$, it also follows that $r(A, R) < r(B, R)$ and $r(A, F) < r(B, F)$. The rejection rates of the lenient council A will be lower for both applicant groups.

To sum up, if councils are heterogenous but not prejudiced, the ranking of the rejection rates across councils does not depend on country of origin. If this rank order is violated, we can deduce prejudice among the councils. However, because there is no objective rank order that defines impartiality, one cannot pinpoint the blamable council(s).

4.3 AN OUTCOME TEST FOR PREJUDICE IN NATURALIZATION

DECISIONS

The model provides a test for prejudice that is applicable with average outcome data. Importantly, we do not need specific data on the involved decision-making. The theoretical implications predict that under the null hypothesis of no prejudice among the councils, the ranking of the average rejection rates for a given country of origin e across councils $c \in \{A, B\}$ does not depend on country of origin $e \in \{R, F\}$. Heterogeneity

across councils is a prerequisite for empirical applicability as the test exploits the model predictions for the average rejection rates under this circumstance. This necessary condition is implicitly verified via the test for prejudice if equality of the average rejection rates across councils for a given country of origin is rejected.

4.4 DISCUSSING THE MODEL

Outcome tests have notorious problems with infra-marginality. Generally one cannot infer disparate treatment from (average) outcome data. Instead, it is the outcome of marginal decision-making that is informative of animus.⁴ It is useful to elaborate on this distinction. Recall that the councils only naturalize applicants with a signal that is below the naturalization threshold (3) for that group. In other words, a council only naturalizes applicants who are deemed qualified enough. If a council is not biased, at the margin it requires the same probability of qualification no matter the country of origin. But we know that depending on the group-specific distributions of the signals and the fraction of unqualified applicants in that group, the average rejection rates may vary despite the same marginal decision-making process. Since empirical data only provides information on average outcomes, the infra-marginality issue poses a key obstacle for inferences of disparate treatment.

The proposed outcome test in this paper circumvents this issue. The test does not compare the mere average rejection rates for a given country of origin across councils. Instead, it makes use of the rankings implied by the model, an indirect identification strategy so to speak. These rankings exploit the simple fact that under the model assumptions, the direction in which the average rejection rate moves is uniquely determined by the direction the council adjusts the required marginal probability of qualification. A higher (lower) marginal probability of qualification always implies a higher (lower) average rejection rate. Put simply, the average moves with the marginal because they are

⁴For an extended description of this issue see Becker (1993), Yinger (1996*b*), or Ayres (2002).

strictly monotonically related. So although we cannot infer directly if the required marginal probabilities of qualification are equal when looking at the average rejection rates, we do know which rank order of the average rejection rates would reject the marginal probabilities being equal.

This identification strategy bears a caveat, however. Like in AF's test for racial prejudice, there is some leeway in the average rejection rates due to the ordinal nature of the test. Imagine that one council is prejudiced against applicants from a certain country of origin, which raises the rejection rate. But the proposed test will fail to detect prejudice if this rate remains within the allowed range which is consistent with the null hypothesis of an independent rank order across councils. This is the case if the prejudice is not too strong, where strong is relative and depends on the magnitude of prejudice required for a violation of the rank orders and on the differences in benefit across councils. The larger these differences, the more leeway there is. In statistical terms, there is a high probability of a type-II error, not rejecting the null hypothesis of no prejudice. All the same, as AF point out, if the test does indicate prejudice, it does so with high confidence.

Another issue looms in the assumption that all councils face pools of applicants which, for a given country of origin, are sampled from the same population. I address the validity of this assumption in the empirical data in the next section using three different Swiss data sets to proxy for qualification. First, data from the Program for International Student Assessment (PISA) will help answering the question whether students of a given country of origin have equal educational skill sets across the observed councils. Second, I analyze data from the Swiss Labor Force Survey (SLFS) to compare labor market qualifications across councils. Finally, the data raised by Hainmueller and Hangartner (2013) will provide an even deeper insight into applicant qualification across municipalities at the micro level.

5 DATA

All Swiss municipalities with at least 20'000 citizens as of 2010 were invited to list their total number of both filed and eventually granted citizenship applications by country of origin, if applicable on a yearly basis. Because the statistical implementation of the outcome test described in the next section requires a sufficient number of observations, the test lends itself for application in municipalities which conduct and consequently reject a reasonable number of naturalization requests. It is important to emphasize that these kind of data are often not readily available and must be specifically raised. Depending on the available resources and existing data format, some municipalities cannot provide this information promptly, if at all. Data collection is still partly in progress, so the description in this section is at a very preliminary stage. It can be reasonably expected that the data collection and processing phase will be completed within the upcoming weeks of the project.

Of the 35 municipalities, 22 are located in the German-speaking region of Switzerland. Of those 22, six have already provided sufficient data. Four more have agreed to do so shortly. Four have declined due to resource constraints. Two municipalities have recently merged and were dropped from requests as data on earlier naturalizations were hard to obtain.

11 of the 35 municipalities are located in the French-speaking region of Switzerland and were just recently contacted in a second wave. As of now, there are no replies. Finally, one of the 35 municipalities is located in the Italian-speaking region and was dropped for efficiency.

6 EMPIRICAL ANALYSIS

The empirical analysis of the existing data will be provided by the time of the conference. The empirical test is implemented by comparing the rejection rates by country of origin

for their observable rank order across councils. To assess the ordinal rankings, I conduct pairwise tests for the null hypothesis of equal rejection rates against the observed one-sided alternative hypothesis by using

$$Z = \frac{\widehat{r(X, e)} - \widehat{r(Y, e)}}{\sqrt{\frac{\text{SVar}_A}{n_A} + \frac{\text{SVar}_B}{n_B}}}$$

for any two councils X and Y (descending in overall levels of rejection rates) against e , a given country of origin. $\widehat{r(X, e)}$ and $\widehat{r(Y, e)}$ denote the estimated rejection rates, SVar_X and SVar_Y the sample variances of the rejection rates, and n_A and n_B the number of applications (all by councils X and Y , respectively). The null is rejected in favor of the alternative hypothesis at significance level α if $Z \geq z_\alpha$ where Z follows a standard normal distribution, so $\Phi(z_\alpha) = 1 - \alpha$.

If the null is rejected and rank orders of the rejection rates across councils seemingly depend on country of origin, relative prejudice based on country of origin can be deduced.

7 CONCLUSION

The ban of closed ballot voting on naturalization decisions in Switzerland has decreased discrimination against applicants from Turkey and (former) Yugoslavia, but has it abolished disparate treatment altogether? The lack of microlevel data used from ballot decisions that is associated with the ban in 2003 renders existing empirical strategies not applicable. This paper proposes an alternative test to assess this question. I describe a simple model of a council deciding on naturalization requests. Councils benefit from naturalizing qualified applicants, but qualification assessment is limited by a noisy but informative signal.

The models shows that a council's benefit and its required probability of qualification are strictly inversely related. This relationship gives rise to ordinal rank order conditions

of average rejection rates grouped by country of origin. If the councils are not prejudiced, the rank order of the grouped rejection rates across councils must be independent of country of origin. Importantly, this specified outcome test is not subject to the inframarginality problem as it attempts no inferences about equality of required probability of qualification at the margin. Instead, the test exploits variation in rejection proclivity on municipal level in order to make predictions about the implied ranking of average rejection rates.

Data from large Swiss municipalities are used for the empirical implementation of the test. Data collection and processing is still in progress, but there is already sufficient data raised for an empirical analysis. It is expected to present results in due time.

The results have important political implications for the Swiss naturalization system and inform current policy debates. If prejudice can be detected despite the recent transition to elected councils, it raises questions about the suitability of the current federal structure to tackle disparate treatment by granting too much discretionary leeway. On the other hand, if the differences can be traced to objective causes, current practice had better be cleared of allegations of prejudice. Instead, any remedial policies to address the disparities are better off concentrating on fostering the chances from the applicants' point of view.

REFERENCES

- Alesina, Alberto, and Eliana La Ferrara.** 2011. “A Test of Racial Bias in Capital Sentencing.” *Harvard Institute of Economic Research, Discussion Paper No. 2192*.
- Anwar, Shamena, and Hanming Fang.** 2006. “An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence.” *American Economic Review*, 96(96): 127–151.
- Anwar, Shamena, and Hanming Fang.** 2012. “Testing for Racial Prejudice in the Parole Board Release Process: Theory and Evidence.” *NBER Working Paper 18239*.
- Anwar, Shamena, and Hanming Fang.** 2013. “Testing for the Role of Prejudice in Emergency Departments Using Bounceback Rates.” *BE Journal of Economic Analysis & Policy*, forthcoming.
- Ayres, Ian.** 2002. “Outcome Tests of Racial Disparities in Police Practices.” *Justice Research and Policy*, 4: 131–142.
- Ayres, Ian.** 2005. “Three Tests for Measuring Unjustified Disparate Impacts in Organ Transplantation.” *Perspectives in Biology and Medicine*, 48(1, supplement): S68–S87.
- Ayres, Ian, and Joel Waldfogel.** 1994. “A Market Test for Race Discrimination in Bail Setting.” *Stanford Law Review*, 46(5): 987–1047.
- Becker, Gary S.** 1993. “The Evidence Against Banks Doesn’t Prove Bias.” *Business Week*, 18.
- Behagel, Luc, Bruno Crépon, and Thomas Le Barbanchon.** 2011. “Do Anonymous Resumes Make the Battlefield More Even? Evidence From a Randomized Field Experiment.” *mimeo.* mimeo.

- Bertrand, Marianne, and Sendhil Mullainathan.** 2004. “Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination.” *American Ec*, 94(4): 991–1013.
- Hainmueller, Jens, and Dominik Hangartner.** 2012. “Does Direct Democracy Hurt Immigrant Minorities? Evidence from Naturalization Decisions in Switzerland.” *Working Paper*.
- Hainmueller, Jens, and Dominik Hangartner.** 2013. “Who Gets a Swiss Passport? A Natural Experiment in Immigrant Discrimination.” *American Political Science Review*, forthcoming.
- Helbling, Marc.** 2008. *Practising Citizenship and Heterogenous Nationhood: Naturalizations in Swiss Municipalities*. Amsterdam University Press.
- Kaas, Leo, and Manger Christian.** 2012. “Ethnic Discrimination in Germany’s Labour Market: A Field Experiment.” *German Economic Review*, 13(1): 1–115.
- Knowles, John, Nicola Persico, and Petra Todd.** 2001. “Racial Bias in Motor Vehicle Searches: Theory and Evidence.” *Journal of Political Economy*, 109(1): 203–229.
- Mechoulan, Stéphane, and Nicolas Sahuguet.** 2011. “Assessing Racial Discrimination in Parole Release.” *mimeo, HEC Montréal*.
- Smart, S., J. Shoven, and J. Waldfogel.** 1996. “A Citation-based Test for Discrimination at Economic and Finance Journals.” *NBER Working Paper*, 5460.
- Yinger, John.** 1996a. “Discrimination in Mortgage Lending: A Literature Review.” In *Mortgage Lending, Racial Discrimination, and Federal Policy*, ed. John Goering and Ron Wienk. Washington, D.C.: Urban Institute Press.

Yinger, John. 1996*b*. “Why Default Rates Cannot Shed Light on Mortgage Discrimination.” *Cityscape: A Journal on Policy Development and Research*, 2(1): 25–31.