

Struth, Lennart; Thon, Max

Working Paper

Discrimination, quotas, and stereotypes

ECONtribute Discussion Paper, No. 188

Provided in Cooperation with:

Reinhard Selten Institute (RSI), University of Bonn and University of Cologne

Suggested Citation: Struth, Lennart; Thon, Max (2022) : Discrimination, quotas, and stereotypes, ECONtribute Discussion Paper, No. 188, University of Bonn and University of Cologne, Reinhard Selten Institute (RSI), Bonn and Cologne

This Version is available at:

<https://hdl.handle.net/10419/262350>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

ECONtribute
Discussion Paper No. 188

Discrimination, Quotas, and Stereotypes

Lennart Struth

Max Thon

August 2022

www.econtribute.de



Discrimination, quotas, and stereotypes*

Lennart Struth and Max Thon**

August 9, 2022

Abstract

We analyze the effect of employment-quota policies on the development of uncertainty and stereotypes in a model of inaccurate statistical discrimination with ambiguity about worker abilities. We show that, even if group characteristics are identical, higher uncertainty about one group can result in discriminatory employment decisions. The success of a quota in correcting the beliefs then depends crucially on the firm's learning process. In particular, we find that the more confident the firm is in its initial priors, the longer a quota needs to be implemented until beliefs are sufficiently corrected such that discriminatory behavior vanishes.

JEL Codes: J71, K31, M51, D81

Keywords: Statistical discrimination, stereotypes, quotas, ambiguity

*We thank Anna Bindler, Simon Brandkamp, Zoë Cullen, Oliver Gürtler, Marius Gramb, Matthias Heinz, Bobak Pakzad-Hurson, Christoph Schottmüller, Heiner Schumacher, Dirk Sliwka, Tomasz Strzalecki, and participants of the WiP-Seminar at the University of Cologne for helpful comments. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 447657066 and funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1 – 390838866.

**Lennart Struth: University of Cologne, struth@wiso.uni-koeln.de; Max Thon: University of Cologne, m.thon@wiso.uni-koeln.de.

1. Introduction

Underrepresentation of certain groups in specific types of jobs, especially at high levels in firm hierarchies, is a prevalent topic of public debates and economic research.¹ To address the problem, quota policies have been proposed and implemented in many countries.² While the effects of quotas on labor-market outcomes have been studied extensively in the economic literature, their introduction still provokes controversial debates among policy-makers and researchers. In particular, once labor-market quotas are introduced, it is often not clear for how long they should be implemented. With this article, we further enhance the understanding of why workers from different groups are underrepresented in certain types of jobs, as well as providing hints for the determinants of the success of a quota, i.e., for how long a quota needs to be implemented until discriminatory employment practices vanish such that the quota becomes obsolete.

We propose a simple dynamic model of inaccurate statistical discrimination with ambiguity about worker abilities.³ There is a single firm that, for infinitely many periods, has to fill open positions in each period with workers stemming from two groups. Decision-makers in the firm are experienced with the first group, such that they know the distribution of abilities across the group and, since the ability of individual workers (and thus their value for the firm) is not observable, they face a decision under risk when employing a worker. The second group, however, has been underrepresented in the past, such that the decision-maker is more uncertain with such workers and potential stereotypes could not be corrected by past experiences.⁴ Hence, decision-makers hold inaccurate beliefs, which are modeled as

¹See [Holzer and Neumark \(2000\)](#), [Lang and Lehmann \(2012\)](#), [Blau and Kahn \(2017\)](#), and [Neumark \(2018\)](#) for an overview.

²Several countries have implemented quotas governing the gender distribution of corporate boards or election lists for their national parliament. See, for example, Law 120/2011 Golfo-Mosca in Italy, or Law 3/2007 on the Effective Equality between Women and Men in Spain.

³The term “inaccurate statistical discrimination” is borrowed from [Bohren et al. \(2019\)](#) and refers to statistical discrimination with incorrect beliefs about the underlying probability distributions. This will be explained in more detail when the model is introduced in Section 2.

⁴By using the term *stereotypes*, we refer to decision-makers having incorrect initial beliefs about the ability distribution within a group.

ambiguity about the distribution of worker abilities.

Our analysis provides three main results. First, we show that, without a quota, inaccurate beliefs and discriminatory employment decisions can persist in the long run. Similar to models of endogenous statistical discrimination (as, e.g., in [Coate and Loury 1993](#)), workers belonging to an underrepresented group find themselves in a reputation trap they cannot escape. They keep being excluded from the job as a result of the firm’s inexperience with their group, which again leads to the firm remaining inexperienced and workers being underrepresented.

Second, we show that a quota, in the long run, is successful in correcting ambiguous beliefs. Since the quota forces the firm to observe individual characteristics of workers from a group it initially is less experienced with, over time the firm learns about the distribution across the group such that beliefs converge to the true distribution in the long run.

Third, we characterize determinants of the duration for which a quota needs to be implemented in order to be successful. In contrast to a standard Bayesian model of uncertainty, in our model of ambiguous beliefs the success of a quota depends crucially on the parameters of the firm’s learning process. In particular, we find that the more confident the firm’s decision-makers are about the initial beliefs (the “stickier” stereotypes are), the longer a quota needs to be implemented.

Furthermore, to our knowledge, we are the first to introduce ambiguity into a model of statistical discrimination. In particular given the recent findings in the literature on inaccurate statistical discrimination ([Bohren et al., 2019](#)), we believe that it is important to bring ambiguity to the traditional models of statistical discrimination both from a methodological viewpoint, as well as verifying the robustness of the model’s implications.

Our analysis yields important results in light of many recent discussions about affirmative action. For instance the “Legge Golfo - Mosca”, a law governing gender quotas for companies listed on the Italian stock exchange, was recently revised.⁵ Like other types of affirmative

⁵See [Ferrari et al. \(2016\)](#) for a summary of the law.

action, the Italian law aims at changing beliefs about the ability of women to work in high-level positions. By being imposed temporarily, the idea is that employers will correct their wrong beliefs after being exposed to females in high-level positions after a sufficiently long period of time.⁶ However, lack of success led legislators to increase the quota and to extend its mandatory compliance period.

There is ample empirical evidence that affirmative-action policies in form of quotas have an impact on the development of beliefs about the abilities of individuals belonging to specific groups to perform in certain positions.⁷ The most prominent example is the study by [Beaman et al. \(2009\)](#). The paper provides evidence that repeated exposure to female chief councilors in Indian villages improves voters' perceptions about the ability of female politicians and caused a weakening of existing stereotypes. Another example is the work of [De Paola et al. \(2010\)](#), who find similar results for a reform reserving a percentage of the places in the party lists in certain Municipal Council elections in Italy for female candidates. They find that, even years later, female representation in politics is higher in these municipalities. Moreover, they are able to trace the effect back to a reduction in negative stereotypes.

However, the empirical evidence evaluating the effect of job-market quotas on the beliefs of employers or other stakeholders is far from vast. The reason is twofold. First, most of these laws were introduced rather recently. As pointed out by [Bertrand et al. \(2019\)](#), quota reforms are instruments which fully unfold their impact in the long run. Therefore, it is likely too early to draw conclusions about the effect.⁸ Moreover, there are no large-scale

⁶Another timely example is the recent announcement of Friedrich Merz (Chairman of the Christian Democratic Union in Germany) to introduce a temporary gender quota for public positions and mandates of the party. The reform comprises a staggered increase of the strength of the quota for a total time period of four years. According to party officials, the board expects that the accompanying cultural change will resolve discussions about a permanent quota automatically (for further information, see a recent newspaper article by [Spreter 2022](#)).

⁷There are other studies showing that a forced exposure (not necessarily due to a quota), to a certain group can change attitudes and stereotypes against them in a positive way. Examples are [Carrell et al. \(2019\)](#) and [Schindler and Westcott \(2020\)](#).

⁸[Bertrand et al. \(2019\)](#) empirically evaluate the impact on economic outcomes of women of the introduced gender quota in Norway. They find that the quota has lowered the gender pay gap among board members. Although they do not find that the quota has been beneficial for other women employed by the companies subject to the quota, they suggest that the reform could have broader effects unfolding in the long run that could be found by similar studies in the future.

survey data available that measure the change in beliefs upon the introduction of a quota reform. Second, quotas are usually imposed on all businesses with certain characteristics at the same time, which makes their effects hard to identify. It is expected that the economic literature will address this question in the future when data availability increases. In any case, there is no reason to expect that mechanisms present in the political context are not similarly prevalent in labor markets.

With our model, we provide important insights for policy-makers and economic research alike, as it discusses mechanisms, optimal choice, as well as the interplay of different parameters of quota regulations. Therefore, our analysis provides guidance for the optimal design of future affirmative-action policies in the form of quotas. Moreover, by introducing a new modeling approach, our paper is a first step for future theoretical research in this area and provides a framework for future empirical studies.

Our work is related to several strands of literature. First, we contribute to the literature on (exogenous) statistical discrimination, rooted in [Phelps \(1972\)](#) and [Aigner and Cain \(1977\)](#). More specifically, as we consider a dynamic model, our paper is related to the work of [Bjerk \(2008\)](#) and [Kim and Loury \(2018\)](#). In particular, similar to how a lack of experience leads to discriminatory employment practices in our model, [Bjerk \(2008\)](#) shows that differential opportunities of workers to signal their individual ability lead to different likelihoods of being promoted. Furthermore, since we depart from the assumption of decision-makers having correct beliefs about group characteristics, our paper is most closely related to [Bohren et al. \(2019\)](#), who introduced the notion of inaccurate statistical discrimination.

Second, we contribute to the theoretical literature on the effects of affirmative-action policies, more specifically on the effects of employment quotas in firms. Thus, our work builds on [Welch \(1976\)](#), [Milgrom and Oster \(1987\)](#), [Coate and Loury \(1993\)](#), [Moro and Norman \(2003\)](#), [Gürtler and Gürtler \(2019\)](#) and [Bijkerk et al. \(2021\)](#).

Third, our model presents an application of existing concepts in the literature on ambiguity and Knightian Uncertainty that originated in [Knight \(1921\)](#). In particular, we build

on the multiple-prior framework of [Gilboa and Schmeidler \(1989\)](#) to model ambiguous beliefs and apply the learning process of [Epstein and Schneider \(2007\)](#) to our setting.⁹

The paper is organized as follows. Section 2 contains the description of the model, as well as more elaborate explanations of how we bring the notion of ambiguity to a model of statistical discrimination. In Section 3, we show that in such a model negative stereotypes and discriminatory employment practices may persist in the long run. The subsequent Section 4 analyzes the effects of a quota on the development of the firm’s stereotypes and the determinants of the success of a quota. Section 5 concludes. Unless stated otherwise, proofs are relegated to the Appendix.

2. Model

We consider a simple dynamic model of statistical discrimination in discrete time with an infinite number of periods, indexed by $t \in \{0, 1, \dots\}$. There are two groups of workers, A and B , denoted as $g \in \{A, B\}$. The distribution of group- g workers’ individual abilities a is given by the cumulative distribution function F_g and unchanged over time. In order to make the analysis as intuitive and simple as possible, we restrict attention to binary probability distributions, that is, a worker is either *talented* ($a = a_h$) or *untalented* ($a = a_l$). Denoting the fraction of talented workers in group g by $\bar{\theta}_g \in (0, 1)$, the probability mass functions f_g of the ability distributions are given by $f_g(a_h) = \bar{\theta}_g = 1 - f_g(a_l)$.

There is one firm. In each period t , there are $n \in \mathbb{N}$ open positions for which the firm needs to assign new workers. A worker’s value for the firm in that position is a strictly increasing function $y \in C^0(\mathbb{R})$ of ability.¹⁰ The timing of the job-assignment process in each period is as follows: The firm holds beliefs about the distribution of worker abilities. It then decides on the number of workers it employs from each of the two groups. There are

⁹For the results in which we restrict attention to binary ability distributions, our model can be seen as an extended application of the model of learning from ambiguous urns by [Marinacci \(2002\)](#).

¹⁰We also refer to a worker’s *value for the firm* as the worker’s *productivity* and use these terms interchangeably.

sufficiently many workers in both groups, such that the firm could fill all open positions with workers stemming from only one of the groups. At the end of the period, the firm observes the individual productivity of each employed worker and uses the obtained information to update its beliefs about the group characteristics. The updated belief is the firm's belief in the following period.¹¹

Furthermore, we assume that, in the first period, $t = 0$, the firm knows the distribution of group- A workers' abilities perfectly, while its belief about the second group B is different from the true distribution F_B , since it is more uncertain about their abilities, in a sense defined precisely below.¹² The firm is aware of the fact that its belief about group- A abilities is correct and that it is more uncertain about the second group. Therefore, it makes use of all available information to update its belief about group B in each period.

Uncertainty. Standard models of statistical discrimination introduce uncertainty about workers' individual characteristics by means of probability distributions. Decision-makers evaluate the expected value of a worker for the firm conditional on which group the worker belongs to. An employment decision then is a decision under risk in the sense of standard expected utility theory.

As we have argued in the introduction, we believe that uncertainty about group characteristics stemming from a lack of experience plays an important role in decision-makers' stereotypes and employment decisions. Therefore, we differentiate between the employment decision being – as a decision that uses information about the distribution of worker characteristics across the whole group as a proxy for the individual ability – a decision under risk, and the additional uncertainty about the distribution of characteristics in the group itself. We bring this to the model by distinguishing between measurable risk and unmeasurable

¹¹The updating process will be explained in more detail later in this section.

¹²Throughout the paper, we reason the assumption of incorrect beliefs by higher uncertainty due to a lack of experience with group- B workers. The model applies more generally, though in any context that justifies assuming that a firm is more uncertain about a specific group. Furthermore, our model includes equally uncertain, but inaccurate, beliefs about the second group B as a special case. Such a situation can be interpreted as the firm having (negative or positive) stereotypes about the second group.

ambiguity (or uncertainty in the Knightian sense).¹³

To clarify this distinction between risk and uncertainty in the Knightian sense, consider first group A . We assume that the firm has employed many group- A workers in the past. Therefore, it is very experienced with their abilities and views the distribution of abilities within the group to be F_A . When employing a group- A worker, the firm now faces a decision under risk, as it uses the distribution across the whole group as a proxy for the unobserved individual ability, and calculates the expected productivity EY_A of a worker with respect to the group's ability distribution F_A , that is,

$$EY_A = \int y(a)dF_A(a). \quad (1)$$

The firm compares the expected productivity to a threshold $\bar{y} \in \mathbb{R}$ and decides to employ group- A workers whenever $EY_A \geq \bar{y}$, that is, whenever the expected productivity exceeds the threshold.¹⁴ This is a decision under risk, there is no ambiguity regarding the individual ability of group- A workers.

Next, consider the underrepresented group B . We assume that the firm has not employed any (or at least very few) such workers in the past and, thus, it is not perfectly informed about the ability distribution of these workers. Hence, in addition to the risk in the decision of employing a worker whose individual ability is unobserved, the firm is uncertain about the distribution of abilities in the group and therefore faces a decision under risk and (Knightian) uncertainty. We bring ambiguity to the model by following [Gilboa and Schmeidler \(1989\)](#), and we assume that the decision-maker at the firm does not face a single probability distribution (the distribution of abilities across the group), but a set of probability distributions. Hence, the firm's belief in period t about an individual worker's ability is given by a set C_B^t of probability distributions, that is, the decision-maker sees all distributions within

¹³This notion of uncertainty has also been termed *severe uncertainty* or *deep uncertainty*. In this paper, we keep referring to it as ambiguity or Knightian uncertainty.

¹⁴Assuming that the decision-maker compares the expected productivity to an exogenous threshold $\bar{y} \in \mathbb{R}$ is mainly for convenience. If instead we assumed that the firm compares the groups with each other, our main results on the development of beliefs and the comparative statics would remain unchanged.

the set C_B^t as possible distributions of abilities across the second group B . Since the firm cannot assign probabilities across the elements of C_B^t , it cannot form an expectation about a group- B worker’s ability as it does for the first group. Therefore, we need to define how the firm evaluates its beliefs and on which value it bases the information whether employing a group- B worker is “expected” to be profitable. We follow [Gilboa and Schmeidler \(1989\)](#) who axiomatize preferences of an ambiguity-averse decision-maker, which leads to the notion of *maxmin* expected utility. This means, in our setting, that the firm evaluates the expected productivity of a worker for every prior F_B^t in C_B^t separately, and then it compares the worst value with the threshold. Hence, the “expected” productivity EY_B^t of a group- B worker in period t is given by

$$EY_B^t = \inf_{F_B^t \in C_B^t} \int y(a) dF_B^t(a). \quad (2)$$

Although EY_B^t is not an expectation in the usual sense, we slightly abuse the terminology and refer to EY_B^t as group- B workers’ *expected productivity*. Again, the firm decides to employ a worker whenever $EY_B^t \geq \bar{y}$, that is, whenever the expected productivity exceeds the threshold. Clearly, the firm’s decision now is one under ambiguity. Furthermore, it should be noted that the employment decision is based only on the payoff in the current period and that the decision-maker acts upon the currently available information. In particular, the decision-maker does not try to gather additional information at the expense of a loss in current payoff in order potentially to overcompensate the loss with future profits.¹⁵

Development of stereotypes. To model the development of the firm’s belief over time, we apply the model of learning under ambiguity of [Epstein and Schneider \(2007\)](#) to our setting.¹⁶

¹⁵This assumption to consider *one-step-ahead acts*, as they have been termed in the literature (see, e.g., [Marinacci 2002](#)), is only required for Proposition 1, and relaxing the assumption would not change any other results qualitatively. Furthermore, assuming that firm’s decision-makers do not voluntarily employ workers from underrepresented groups to develop the own perception of that group’s ability distribution seems reasonable in light of the observation that, as discussed in the introduction, quotas are seen to be necessary instruments to address discriminatory employment practices by policy-makers in many countries.

¹⁶Since ability distributions are binary and the firm is only interested in the fraction of talented workers, our setting is similar to the model of learning from ambiguous urns in [Marinacci \(2002\)](#).

Ambiguity in the first period $t = 0$ is then modeled as a compact set $C_B^0 \subset \Delta(\Theta)$ of prior probability distributions over a finite set $\Theta \subset (0, 1)$. It is further assumed that $\bar{\theta}_B \in \Theta$ and that, for all $\mu^0 \in C_B^0$, it holds that $\mu^0(\bar{\theta}_B) > 0$. This assumption ensures that for all priors the true fraction of talented workers has a positive probability.¹⁷ The expected productivity of a group- B worker can then be written as

$$EY_B^0 = \min_{\mu^0 \in C_B^0} \int \int y(a) dF(a|\theta) d\mu^0(\theta), \quad (3)$$

with the probability mass functions $f(\cdot|\theta) : \{a_l, a_h\} \rightarrow [0, 1]$ given by $f(a_h|\theta) = \theta$ for all $\theta \in \Theta$.¹⁸

If the firm decides to employ $n^t \in \mathbb{N}$ workers from group B , it learns about the group's ability distribution by independently drawing n^t times from the true distribution F_B .¹⁹ We denote the vector of draws in period t by $x^t = (x_1^t, \dots, x_{n^t}^t)$. Then, the vector $s^{t-1} = (x^0, \dots, x^{t-1})$ contains all information the firm has obtained when deciding which workers to assign to the open positions in period t .

Furthermore, the firm does not keep all priors in each period, but decides on which distributions to keep based on a likelihood-ratio test, that is, the firm tests all priors μ^0 against the one with the largest likelihood and only keeps those that fulfill the following condition:²⁰

$$\alpha \leq \frac{\int f(s^{t-1}|\theta) d\mu^0(\theta)}{\max_{\hat{\mu}^0 \in C_B^0} \int f(s^{t-1}|\theta) d\hat{\mu}^0(\theta)}. \quad (4)$$

The set of posteriors $\mu^t(\cdot) = \mu^0(\cdot|s^{t-1})$ is given by updating each of these priors by Bayes'

¹⁷A prior that assigns zero probability mass to the true parameter would not converge to the true distribution. Hence, as will be discussed in more detail later in this section, in this case the full Bayesian updating rule would lead to no convergence of the firm's beliefs to the true distribution. This is equivalent to the assumption of the true parameter being in the "strong support" of the set of priors in [Marinacci \(2002\)](#).

¹⁸The functions $F(\cdot|\theta)$ and $f(\cdot|\theta)$ denote the distribution function and probability mass function, respectively, of a binary ability distribution with the fraction of talented workers being θ .

¹⁹Note that, since y is a strictly increasing function of a worker's ability a , the firm, by observing a worker's output $y(a)$, also observes the individual ability.

²⁰Here we abuse notation slightly and denote by $f(\cdot|\theta)$ the probability mass function of the product measure, that is the probability of a multidimensional vector of draws.

rule. Hence, the firm's belief in period t is given by

$$C_B^t(s^{t-1}, \alpha) = \{\mu^t(\cdot) = \mu^0(\cdot|s^{t-1}) \in \Delta(\Theta) \mid \mu^0 \in C_B^0, \int f(s^{t-1}|\theta)d\mu^0(\theta) \geq \alpha \max_{\hat{\mu}^0 \in C_B^0} \int f(s^{t-1}|\theta)d\hat{\mu}^0(\theta)\}. \quad (5)$$

The parameter $\alpha \in [0, 1]$ can be interpreted as a measure for the confidence the firm has in its initial beliefs C_B^0 . The larger α is, the more priors are discarded, that is, the less confident the firm was about the initial beliefs before any information was available. In particular, it holds that if $\alpha_1 > \alpha_2$, then $C_B^t(s^{t-1}, \alpha_1) \subset C_B^t(s^{t-1}, \alpha_2)$. Another interpretation is that the smaller α is, the more hesitant the firm is about dissolving the initial uncertainty and the more “sticky” the stereotypes are about group B . Two special cases are included in this updating rule: If $\alpha = 0$, all priors pass the likelihood-ratio test. Therefore, the firm does not discard any priors and updates all of them using Bayes' rule. This rule has been termed *Full-Bayesian updating* in the literature. On the other hand, if $\alpha = 1$, the firm only keeps the priors that yield the maximum likelihood, hence this corresponds to the *Maximum-Likelihood updating*.

Finally, the expected productivity EY_B^t of a group- B worker in period t , given the updated beliefs based on observing the individual abilities s^{t-1} of previously employed workers and confidence α , can then be written as

$$EY_B^t(s^{t-1}, \alpha) = \min_{\mu^t \in C_B^t(s^{t-1}, \alpha)} \int \int y(a)dF(a|\theta)d\mu^t(\theta). \quad (6)$$

3. Persistence of uncertainty

In this section, we show that, given the larger uncertainty about group- B workers' abilities, the firm might not be willing to employ any of those workers such that they find themselves in a reputation trap and, even in the long run, keep being excluded from the job.

Suppose that the decision-maker's belief in the first period $t = 0$ is such that $EY_A > EY_B^0$.

This is a reasonable assumption in the sense that it is fulfilled whenever the firm’s belief about the second group B is not too optimistic. To illustrate this point, assume for a moment that $EY_A \leq EY_B^0$. Then, by definition, all priors the firm sees as possible distributions of abilities within group B lead to a larger expected productivity, that is, in other words, the firm is less experienced with group- B workers, but the optimism completely overcompensates the additional uncertainty for an ambiguity-averse decision-maker. In particular, all priors the firm sees as possible are at least as “good” as the known distribution F_A of the other group. Whenever the set of priors includes at least one “worse” prior, we obtain $EY_A > EY_B^0$. Now, additionally, suppose that the firm’s productivity threshold $\bar{y} \in \mathbb{R}$ is such that $EY_A > \bar{y} > EY_B^0$. Then, in period $t = 0$, the firm fills all n open positions with workers stemming from group A . By the assumption that the firm is aware that it already knows the ability distribution in group A perfectly, and since it does not get any additional information about group- B workers’ abilities, the belief remains unchanged and is identical in the next period, that is, $C_B^1 = C_B^0$. This argument applies iteratively in each period $t \in \mathbb{N}$. Therefore, the firm does not employ any group- B workers and stays uncertain about the ability distribution over time. Hence, the following Proposition 1 can be stated without further proof.

Proposition 1. *Suppose that in the initial period $t = 0$ the firm’s belief is such that $EY_A > \bar{y} > EY_B^0$. Then, it holds that $C_B^t = C_B^0$ for all $t \in \mathbb{N}$.*

Proposition 1 implies that, if in the initial situation stereotypes lead to the firm expecting group- B workers not to be profitable hires, the firm employs only group- A workers and stereotypes remain unchanged in all future periods. The intuition for the result is clear. If the firm’s uncertainty about group B , stemming from a lack of experience with such workers, is severe enough, it is unlikely to be willing to employ any group- B workers in the short run. This implies that the firm does not get any further information about the group charac-

teristics and stays uninformed in the long run.²¹ Therefore, workers keep being excluded from the job as a result of the firm’s inexperience with their group, which again leads to the firm remaining uncertain about the group. Workers are trapped in their disadvantageous group reputation.²² Note that this argument holds even if the true ability distributions are identical in both groups, since the firm’s decisions are based on its belief about the workers’ abilities rather than actual group characteristics.

Given this result, it is natural to think of a quota as a fitting policy, as one would expect it not only to address the underrepresentation of group- B workers in the specific job in the short run, but also to force the firm to learn about the underrepresented group, thus diminishing uncertainty and correcting stereotypes in the long run. We analyze the effects of a quota and the determinants of its success in the following Section 4.

4. Effects of a quota

While so far the firm could freely choose which workers to employ, we now enrich the model by introducing a quota that requires the firm for a number of periods to assign at least some of the jobs to workers of each of the groups.

Definition 1. A **quota** is a tuple $(q, T) \in (0, 1) \times \mathbb{N}$ that specifies the fraction $q \in (0, 1)$ of vacant jobs the firm has to fill with workers belonging to the underrepresented group in each period, as well as the number of periods $T \in \mathbb{N}$ the firm has to fulfill the quota.

Intuitively, one would expect that now, since the quota forces the firm to employ workers from the underrepresented group, the firm learns about group- B workers such that its belief

²¹While this result seems rather extreme, it qualitatively mirrors situations in which specific groups keep being underrepresented such that quotas are justified as a necessary affirmative action. Considering a model that, without a quota being implemented, allows for a small exogenously determined number of group- B workers being employed would leave our results on the belief development and the comparative statics qualitatively unchanged.

²²Note that this is different from the reputation trap in (dynamic) models of endogenous statistical discrimination such as Coate and Loury (1993) and Kim and Loury (2018), since in our model the workers have no actions to make (as, e.g., investing in qualification) and the group reputation is solely based on the firm’s beliefs and experiences.

converges to the true distribution of worker abilities in the group.²³ The following Proposition 2 confirms this intuition.

Proposition 2. *Let $\delta_{\bar{\theta}_B}$ denote the Dirac measure in $\bar{\theta}_B$. Then, for all $q \in (0, 1)$ and all $\alpha \in [0, 1]$, beliefs converge to the singleton $\{\delta_{\bar{\theta}_B}\}$, that is, it almost surely²⁴ holds that*

$$\sup_{\mu \in C_B^T(s^{T-1}, \alpha)} \left\{ \max_{X \subset \Theta} |\mu(X) - \delta_{\bar{\theta}_B}(X)| \right\} \xrightarrow{T \rightarrow \infty} 0. \quad (7)$$

Furthermore, the expected productivity of group- B workers with respect to the firm's belief in period T , given by

$$EY_B^T(s^{T-1}, \alpha) = \min_{\mu^T \in C_B^T(s^{T-1}, \alpha)} \int \int y(a) dF(a|\theta) d\mu^T(\theta), \quad (8)$$

for $T \rightarrow \infty$ converges almost surely to the true average productivity

$$EY_B = \int y(a) dF_B(a) = \int \int y(a) dF(a|\theta) d\delta_{\bar{\theta}_B}(\theta) \quad (9)$$

of group- B workers. Hence, ambiguity fades away and stereotypes are corrected in the long run.

While Proposition 2 states that a quota, if implemented for a sufficiently long time, eventually will be successful in correcting stereotypes, an essential question in policy is for how long exactly a quota needs to be implemented in order to have an impact on discriminatory employment practices after it has been removed again. Hence, the remainder of this section is devoted to analyzing the determinants of a quota's effect on resolving uncertainty and correcting stereotypes.

The answer to whether a quota is successful depends on the individual abilities of em-

²³Note that, since $q > 0$ and $n > 0$, the firm in each period at least has to employ at least $qn > 0$ workers. Hence, each period yields a strictly positive number of draws from which the firm learns about the group's ability distribution.

²⁴The convergence is almost sure with respect to the infinite product measure such that the coordinate draws are independent with the true distribution, that is, the true fraction $\bar{\theta}_B$ of talented group- B workers.

ployed workers, which are randomly drawn from the distribution of abilities across the set of workers. Hence, we have to define precisely what we mean by a quota's success from an ex-ante point of view.

Definition 2. A quota is **successful with probability** $p \in (0, 1)$, if, by the end of its implementation, the firm with probability p voluntarily employs group- B workers, that is, if $\mathbb{P}[EY_B^T \geq \bar{y}] \geq p$.²⁵

We denote by $T_p^* \in \mathbb{N} \cup \{\infty\}$ the first period for which the quota is successful with probability $p \in (0, 1)$.²⁶ Clearly, T_p^* depends on the number $n \in \mathbb{N}$ of vacant jobs in each period and the quota parameter $q \in (0, 1)$. Furthermore, since T_p^* depends on the development of the firm's belief, it also depends on the measure α of the firm's confidence in its initial belief. Hence, T_p^* is a function of these parameters and given by $T_p^*(q, n, \alpha) = \inf \{T \in \mathbb{N} : \mathbb{P}[EY_B^T \geq \bar{y}] \geq p\}$. In the following, we show how T_p^* depends on these variables, or, in other words, how the minimum duration for which a quota needs to be implemented in order to be successful depends on the specifics of the learning process.

Proposition 3. *If $EY_B > \bar{y}$, then for all $p \in (0, 1)$ it holds that $T_p^*(q, n, \alpha)$ is finite for all $q \in (0, 1)$, $n \in \mathbb{N}$ and $\alpha > 0$, and is decreasing in q and n .*

Proposition 3 states that the more strict a quota is and the larger the number of open positions in each period is, the sooner a quota will be successful. This result is intuitive. If, given the initial belief in period $t = 0$, it is not expected to be profitable for the firm to employ group- B workers, the firm hires as many workers as to comply exactly with the quota, that is, it employs exactly qn group- B workers.²⁷ Hence, the number of observations

²⁵We slightly abuse notation and denote by \mathbb{P} the one-dimensional probability measure as well as the product measure. As it should be clear from the context which measure is denoted, this should cause no confusion.

²⁶Note that this does not necessarily mean that the firm will be voluntarily employing group- B workers in all future periods, if the quota is removed in T_p^* . Even if at that point in time beliefs are sufficiently good, the decision-maker keeps learning and updating beliefs such that a number of negative draws could potentially lead to worse beliefs again.

²⁷In general, the product qn is not a natural number. In that case, the firm is required to employ at least $\lceil qn \rceil$ workers. To simplify notation, we assume parameters to be such that $qn \in \mathbb{N}$.

the firm learns from in each period is increasing in the required fraction q and the number of vacancies n in the firm. Therefore, since beliefs converge to the true ability distribution, for which the group- B workers' expected value is sufficiently large, the threshold is sooner met the larger q and n are.

Proposition 4. *If $EY_B > \bar{y}$, for all $p \in (0, 1)$, it holds that $T_p^*(q, n, \alpha)$ is decreasing in α .*

Proposition 4 states that the more confident (the smaller α) the firm is in its initial beliefs, the longer a quota needs to be implemented in order to be successful.²⁸ This result, too, is intuitive. The more hesitant the firm is to discard initially biased priors (stereotypes), the slower it updates its belief about the underrepresented workers' abilities and the more observations, thereby the more periods, it takes until beliefs are sufficiently corrected and group- B workers are expected to be of sufficient value to the firm.

Finally, in Example 1 in the Appendix, we numerically and graphically illustrate the results. In addition to exemplifying our general results on the determinants of a quota's success, we also consider how the degree of initial uncertainty affects the success of a quota. As intuition suggests, we find that the more uncertain (ambiguous) the decision-maker initially is, the longer the quota needs to be implemented in order to be successful.

5. Conclusion

In this paper, we have analyzed the effects of a quota on the development of uncertainty and learning behavior of a firm in a dynamic model of inaccurate statistical discrimination. In particular, we have combined a model of statistical discrimination with the literature on ambiguity to model that a firm is less experienced with a group as a result of underrepresentation of workers in the past. The model yields a number of results. First, we have shown that, even if the groups' true ability distributions are identical, uncertainty and stereotypes,

²⁸Note that such a result could not be obtained in a standard Bayesian model, as the updating rule is unique and leaves no room for parameters such as the confidence in initial beliefs that characterize the decision-maker's uncertainty beyond a single probability distribution.

which are a result of a firm's lack of experience with one of the groups, can persist in the long run such that workers belonging to the underrepresented group, due to the negative group reputation, keep being excluded from the job in the long run. Second, we have found that a quota, if it is implemented for a sufficiently long time, resolves the decision-makers' uncertainty such that the firm voluntarily employs workers who have been discriminated against before. Third, the length of time a quota needs to be in place in order to be successful in correcting beliefs in the long run has been shown to depend on a number of determinants. In particular, the more confident a firm is in its initial priors, the longer the quota needs to correct the initial beliefs sufficiently.

In practice, quotas are often implemented for top-level jobs. It is typically argued that better representation of initially underrepresented groups in high levels of a firm's hierarchy also affects employment and promotion decisions in lower levels through a trickle-down effect. In our model, we have shown that a quota could have an additional effect, as it forces decision-makers to gain experience with workers' value for the firm, and thus it resolves uncertainty and eliminates stereotypes. These theoretical predictions are in line with findings of previous empirical studies, e.g., [Beaman et al. \(2009\)](#) and [De Paola et al. \(2010\)](#), which demonstrated that quotas influence beliefs and may lead to a reduction of existing stereotypes.

Throughout the paper, we imposed the assumption that the distribution of abilities within each group is exogenous and invariant over time. This ensures that ambiguity fades away in the long run and stereotypes are corrected eventually. While an analysis of exogenously time-varying distributions involves technical difficulties that are beyond the scope of this paper, it would also be interesting to consider a model in which workers can influence their own ability, and thus the reputation of their group.

Furthermore, it is important to complement the research of [Beaman et al. \(2009\)](#) and [De Paola et al. \(2010\)](#) by investigating the impact of job quotas on belief development and stereotypes empirically, as this will help to understand better the magnitude of the suggested theoretical effects, as well as being of great importance for policy-makers in improving the

design of quota policies in the future.

References

- Aigner, D. J. and G. G. Cain (1977). Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review* 30(2), 175–187.
- Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova (2009). Powerful women: Does exposure reduce bias? *Quarterly Journal of Economics* 124(4), 1497–1540.
- Bertrand, M., S. E. Black, S. Jensen, and A. Lleras-Muney (2019). Breaking the glass ceiling? The effect of board quotas on female labour market outcomes in Norway. *The Review of Economic Studies* 86(1), 191–239.
- Bijkerk, S. H., S. Dominguez-Martinez, J. Kamphorst, and O. H. Swank (2021). Labor market quotas when promotions are signals. *Journal of Labor Economics* 39(2), 437–460.
- Bjerk, D. (2008). Glass ceilings or sticky floors? Statistical discrimination in a dynamic model of hiring and promotion. *Economic Journal* 118(530), 961–982.
- Blau, F. D. and L. M. Kahn (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature* 55(3), 789–865.
- Bohren, J. A., K. Haggag, A. Imas, and D. G. Pope (2019). Inaccurate statistical discrimination: An identification problem. Working Paper 25935, National Bureau of Economic Research.
- Carrell, S. E., M. Hoekstra, and J. E. West (2019). The impact of college diversity on behavior toward minorities. *American Economic Journal: Economic Policy* 11(4), 159–82.
- Coate, S. and G. C. Loury (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review* 83(5), 1220–1240.
- De Paola, M., V. Scoppa, and R. Lombardo (2010). Can gender quotas break down negative stereotypes? Evidence from changes in electoral rules. *Journal of Public Economics* 94(5), 344–353.

- Epstein, L. G. and M. Schneider (2007). Learning under ambiguity. *The Review of Economic Studies* 74(4), 1275–1303.
- Ferrari, G., V. Ferraro, P. Profeta, and C. Pronzato (2016). Gender quotas: Challenging the boards, performance, and the stock market. CESifo Working Paper Series 6084, CESifo.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Gürtler, M. and O. Gürtler (2019). Promotion signaling, discrimination, and positive discrimination policies. *The RAND Journal of Economics* 50(4), 1004–1027.
- Holzer, H. and D. Neumark (2000). Assessing affirmative action. *Journal of Economic Literature* 38(3), 483–568.
- Kim, Y.-C. and G. C. Loury (2018). Collective reputation and the dynamics of statistical discrimination. *International Economic Review* 59(1), 3–18.
- Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- Lang, K. and J.-Y. K. Lehmann (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature* 50(4), 959–1006.
- Marinacci, M. (2002). Learning from ambiguous urns. *Statistical Papers* 43(1), 143–151.
- Milgrom, P. and S. Oster (1987). Job discrimination, market forces, and the invisibility hypothesis. *Quarterly Journal of Economics* 102, 453–476.
- Moro, A. and P. Norman (2003). Affirmative action in a competitive economy. *Journal of Public Economics* 87(3), 567–594.
- Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature* 56(3), 799–866.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review* 62(4), 659–661.

Schindler, D. and M. Westcott (2020). Shocking racial attitudes: Black G.I.s in Europe. *The Review of Economic Studies* 88(1), 489–520.

Spreter, J. (2022). ‘Friedrich Merz will innerparteiliche Frauenquote einführen’. Zeit Online. [Online], Last accessed on 2022-06-30.

Welch, F. (1976). Employment quotas for minorities. *Journal of Political Economy* 84(4), 105–141.

6. Appendix

Example 1. Let the production function $y = \text{id}_{\mathbb{R}}$ be the identity. In each period, the firm has $n \in \mathbb{N}$ open positions. Workers are either talented ($a = a_h = 1$) or untalented ($a = a_l = 0$). In each group, two thirds of the workers are talented, that is, $\bar{\theta}_A = \bar{\theta}_B = \bar{\theta} = 2/3$. The firm's productivity threshold is given by $\bar{y} = 6/10$. Initially, the firm believes that the fraction of talented workers is either $1/3$ or $2/3$, that is, $\Theta = \{1/3, 2/3\}$. Priors are given by $C_B^0 = \{\mu^0 \in \Delta(\Theta) \mid \mu^0(\bar{\theta}) = \eta, \eta \in [\varepsilon, 1]\}$ with $\varepsilon > 0$.

We start to analyze the firm's employment decisions. Note that the group- A ability distribution is sufficiently good such that an individual worker's expected productivity, given that the firm uses the group's distribution as a proxy, meets the employment threshold, that is, $EY_A = 2/3 > 6/10 = \bar{y}$. The firm's initial belief C_B^0 about group- B abilities in the first period $t = 0$, on the other hand, is such that the productivity threshold for the employment decision is not met, since, for $\varepsilon > 0$ sufficiently small, it holds that

$$EY_B^0 = \min_{\mu^0 \in C_B^0} \int \int a dF(a|\theta) d\mu^0(\theta) = \min_{\eta \in [\varepsilon, 1]} \left\{ \frac{1-\eta}{3} + \frac{2\eta}{3} \right\} = \frac{1+\varepsilon}{3} < \frac{6}{10} = \bar{y}. \quad (10)$$

Hence, we have $EY_A > \bar{y} > EY_B^0$, and therefore Proposition 1 applies. The firm does not employ any group- B workers in the first period, it does not update its beliefs, and it keeps excluding them for all periods $t \in \mathbb{N}$.

Now, a quota is introduced. The firm is required to assign to the n open positions at least a fraction $q \in (0, 1)$ of group- B workers in each period $t \in \{0, 1, \dots, T\}$ and observes their abilities. Hence, in each period, the firm learns about the group- B ability distribution by independently drawing at least qn times from a Bernoulli distribution with parameter (unknown by the firm) $\bar{\theta} = 2/3$. Therefore, the information s^{t-1} is a vector of $N^t := qnt$ independent draws from $Ber(\bar{\theta})$. Accordingly, the firm's beliefs in period t are given by the

set

$$C_B^t(s^{t-1}, \alpha) = \{\mu^t(\cdot) = \mu^0(\cdot|s^{t-1}) \in \Delta(\Theta) \mid \mu^0 \in C_B^0, \int f(s^{t-1}|\theta)d\mu^0(\theta) \geq \alpha \max_{\mu^0 \in C_B^0} \int f(s^{t-1}|\theta)d\mu^0(\theta)\}, \quad (11)$$

where, by Bayes' rule, for each $\mu^0 \in C_B^0$, and hence for all $\eta \in [\varepsilon, 1]$, it holds

$$\mu^t\left(\frac{2}{3}|\eta\right) = \frac{\left(\frac{2}{3}\right)^{M^t} \left(\frac{1}{3}\right)^{N^t-M^t} \eta}{\left(\frac{2}{3}\right)^{M^t} \left(\frac{1}{3}\right)^{N^t-M^t} \eta + \left(\frac{1}{3}\right)^{M^t} \left(\frac{2}{3}\right)^{N^t-M^t} (1-\eta)}, \quad (12)$$

with $M^t = \sum_{i=0}^{t-1} \sum_{j=1}^{q^n} X_j^i$ denoting the sum of all observations. Since $a_h = 1 = 1 - a_l$, the variable $M^t \sim \text{Binom}(N^t, 2/3)$ equals the number of workers who are talented, out of the N^t employed workers.

The expected productivity of a group- B worker in period t as a function of the confidence in beliefs is then given by

$$EY_B^t(s^{t-1}, \alpha) = \min_{\mu^t \in C_B^t(s^{t-1}, \alpha)} \left\{ \frac{2}{3} \mu^t(2/3) + \frac{1}{3} \mu^t(1/3) \right\}. \quad (13)$$

Note that, by Bayes' rule, for each $\mu^0 \in C_B^0$, and hence for all $\eta \in [\varepsilon, 1]$, it holds

$$\begin{aligned} \mu^t\left(\frac{2}{3}|\eta\right) &= \mathbb{P}_\eta\left[\theta = \frac{2}{3} \mid s^{t-1}\right] \\ &= \frac{\mathbb{P}\left[s^{t-1} \mid \theta = \frac{2}{3}\right] \mathbb{P}_\eta\left[\theta = \frac{2}{3}\right]}{\mathbb{P}\left[s^{t-1} \mid \theta = \frac{2}{3}\right] \mathbb{P}_\eta\left[\theta = \frac{2}{3}\right] + \mathbb{P}\left[s^{t-1} \mid \theta = \frac{1}{3}\right] \mathbb{P}_\eta\left[\theta = \frac{1}{3}\right]} \\ &= \frac{\binom{N^t}{M^t} \left(\frac{2}{3}\right)^{M^t} \left(\frac{1}{3}\right)^{N^t-M^t} \eta}{\binom{N^t}{M^t} \left(\frac{2}{3}\right)^{M^t} \left(\frac{1}{3}\right)^{N^t-M^t} \eta + \binom{N^t}{M^t} \left(\frac{1}{3}\right)^{M^t} \left(\frac{2}{3}\right)^{N^t-M^t} (1-\eta)} \\ &= \frac{\left(\frac{2}{3}\right)^{M^t} \left(\frac{1}{3}\right)^{N^t-M^t} \eta}{\left(\frac{2}{3}\right)^{M^t} \left(\frac{1}{3}\right)^{N^t-M^t} \eta + \left(\frac{1}{3}\right)^{M^t} \left(\frac{2}{3}\right)^{N^t-M^t} (1-\eta)}, \end{aligned}$$

with

$$M^t = \sum_{i=0}^{t-1} \sum_{j=1}^{qn} X_j^i$$

denoting the sum of all observations.

The expected productivity of a group- B worker in period t as a function of the confidence in beliefs is then given by

$$\begin{aligned} EY_B^t(s^{t-1}, \alpha) &= \min_{\mu^t \in C_B^t(s^{t-1}, \alpha)} \int \int a \, dL(a|\theta) \, d\mu^t(\theta) \\ &= \min_{\mu^t \in C_B^t(s^{t-1}, \alpha)} \int a_h \theta + a_l(1 - \theta) \, d\mu^t(\theta) \\ &= \min_{\mu^t \in C_B^t(s^{t-1}, \alpha)} \int \theta \, d\mu^t(\theta) \\ &= \min_{\mu^t \in C_B^t(s^{t-1}, \alpha)} \left\{ \frac{2}{3} \mu^t(2/3) + \frac{1}{3} \mu^t(1/3) \right\}. \end{aligned}$$

In the following, we graphically illustrate how, for a population with ability distribution F_θ , the minimum number T^* of periods varies with different parameters of the model.²⁹ First, we consider the implications of Proposition 3, which states that the larger the firm is (the larger n) and the more strict a quota is (the larger q), the sooner the quota will be successful. Figure 1 illustrates these findings by showing the graph of the minimum time T^* a quota needs to be implemented in order to be successful as a function of the firm size n in the left panel and as a function of the quota parameter q in the right panel.

²⁹Parameters are chosen such that all but the running variable are held constant at $n = 5$, $q = 0.2$, $\alpha = 0.3$, $\varepsilon = 10^{-2}$ and $p = 0.99$. The values of the function $T^*(\cdot)$ that are shown in the graphs are calculated as follows: For each combination of parameters we simulate the development of beliefs 1000 times by drawing random sequences from the ability distribution. The number $T^*(\cdot)$ is then calculated as the smallest time period such that the quota has been successful for a share of $p = 0.99$ of the 1000 simulations.

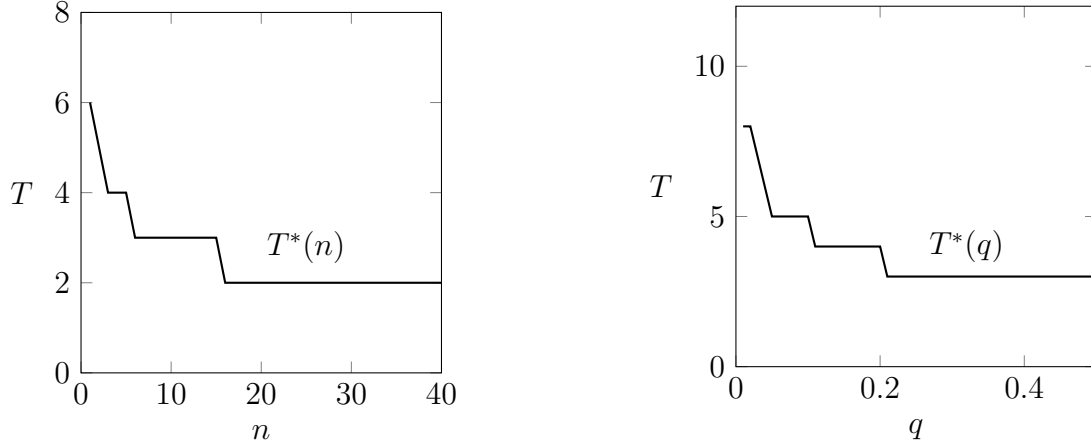


Figure 1: Graph of T^* as a function of the firm size n and the quota parameter q

Clearly, T^* is decreasing in both parameters, that is, the larger the firm size or the stricter the quota, the sooner uncertainty is resolved sufficiently, such that the quota becomes obsolete.

Second, we illustrate the results on the dependence of a quota's success on the decision-maker's confidence α in initial beliefs. The corresponding Proposition 4 states that T^* is decreasing in α , that is, the more confident the firm is in its initial beliefs, the longer a quota needs to be implemented in order to be successful. Figure 2 shows the simulation results for T^* as a function of the confidence parameter α . As the graph of minimum period T^* is decreasing in α , this simulation is in line with the results of Proposition 4.

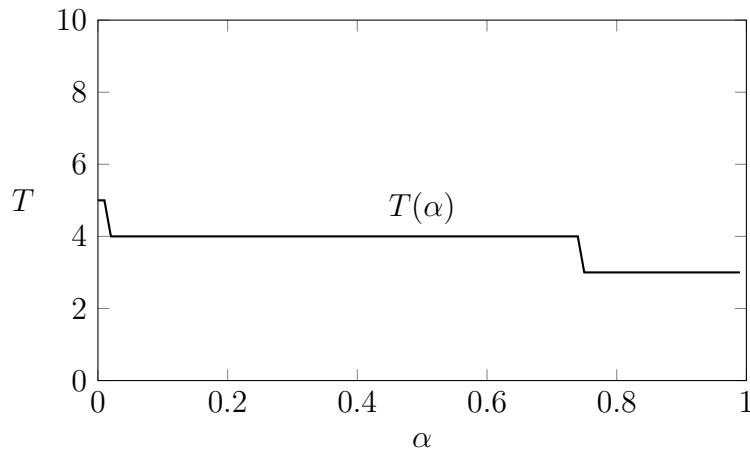


Figure 2: Development of required time T as function of the confidence parameter α

Third, although we have not provided any formal results on this effect, we illustrate how the amount of uncertainty affects the quota's success. Figure 3 shows the graph of T^* as a function of ε . In this example, the parameter ε is a measure of uncertainty, as for any $\varepsilon_1 > \varepsilon_2$ it holds that $C_B^0(\varepsilon_1) \subset C_B^0(\varepsilon_2)$. Hence, the smaller ε is, the larger is the set of initial beliefs the decision-maker considers to be possible.

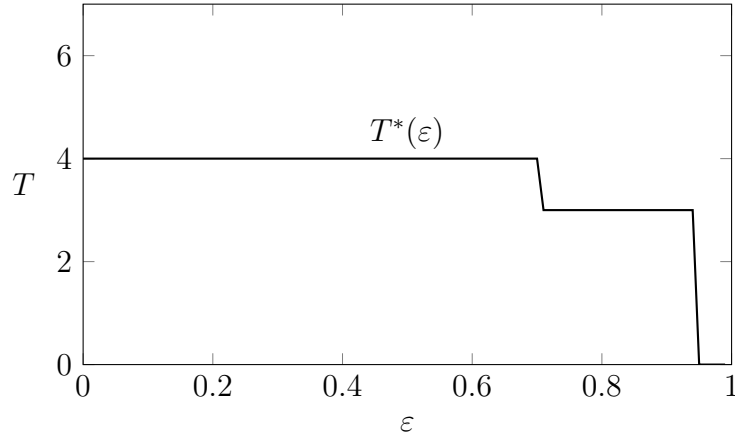


Figure 3: Development of required time T as function of the uncertainty parameter ε

We can see that, as one would intuitively expect, the less uncertain the decision-maker initially is (the larger ε), the sooner the quota is successful in resolving uncertainty and correcting stereotypes.

Proof of Proposition 2. Before we proceed to prove the convergence results in the Proposition, observe that the convergence with respect to the number of periods is implied by a convergence with respect to the number of draws the decision-maker observes from the true ability distribution, since the quota requires the firm to employ, in each period $t \in \{1, \dots, T\}$, at least $\lceil qn \rceil \geq 1$ group- B workers. Hence, for all $T \in \mathbb{N}$, the decision-maker observed at least T realizations of draws from the true ability distribution, and therefore T can be seen as the minimal number of observations in period T .

We begin by proving the convergence in equation (7), that is, we want to show that

$$\sup_{\mu \in C_B^T} \left\{ \max_{X \subset \Theta} |\mu(X) - \delta_{\bar{\theta}_B}(X)| \right\} \xrightarrow{T \rightarrow \infty} 0, \quad (14)$$

where $\delta_{\bar{\theta}_B}$ denotes the Dirac measure in $\bar{\theta}_B$ and the convergence is \mathbb{P} -almost surely.

Suppose for now that $\alpha = 0$. By assumption, for every $\mu^0 \in C_B^0$ it holds that $\mu^0(\bar{\theta}_B) > 0$. In other words, the true parameter $\bar{\theta}_B$ is an element of the strong support of the set of priors C_B^0 . Hence, the convergence in equation (14) follows immediately by Lemma 1 of Marinacci (2002).

Now, consider the case $\alpha > 0$. While in case of $\alpha = 0$ posteriors are formed for all priors $\mu^0 \in C_B^0$, the updating process in the general case is defined in equation (5) and given by

$$C_B^t(s^{t-1}, \alpha) = \{\mu^t(\cdot) = \mu^0(\cdot | s^{t-1}) \in \Delta(\Theta) \mid \mu^0 \in C_B^0, \\ \int f(s^{t-1} | \theta) d\mu^0(\theta) \geq \alpha \max_{\mu^0 \in C_B^0} \int f(s^{t-1} | \theta) d\mu^0(\theta)\}.$$

Hence, for all $\alpha > 0$ and all $t \in \mathbb{N}$, it holds that $C_B^t(s^{t-1}, \alpha) \subset C_B^t(s^{t-1}, 0)$. Therefore, we obtain

$$0 \leq \sup_{\mu \in C_B^T(s^{T-1}, \alpha)} \left\{ \max_{X \subset \Theta} |\mu(X) - \delta_{\bar{\theta}_B}(X)| \right\} \leq \sup_{\mu \in C_B^T(s^{T-1}, 0)} \left\{ \max_{X \subset \Theta} |\mu(X) - \delta_{\bar{\theta}_B}(X)| \right\} \xrightarrow{T \rightarrow \infty} 0.$$

Hence, the convergence in equation (14) follows for all $\alpha > 0$ almost surely.

Next, we show the convergence of the expected productivity of group- B workers, given by equation (8) to the true productivity, that is, we show that, for $T \rightarrow \infty$, it holds that

$$EY_B^T(s^{T-1}, \alpha) = \min_{\mu^T \in C_B^T(s^{T-1}, \alpha)} \int \int y(a) dF(a | \theta) d\mu^T(\theta) \\ \longrightarrow \int \int y(a) dF(a | \theta) d\delta_{\bar{\theta}_B}(\theta) = EY_B \quad (15)$$

almost surely with respect to the probability measure \mathbb{P} .

Consider first the case $\alpha = 0$. Since C_B^0 is compact and $\bar{\theta}_B$ is, as argued before, an element of the strong support of C_B^0 , the convergence in equation (15) follows immediately from Theorem 1 in Marinacci (2002).

Suppose now that $\alpha > 0$. Then, by again noting that for all $\alpha > 0$ and all $T \in \mathbb{N}$, it

holds that $C_B^T(s^{T-1}, \alpha) \subset C_B^T(s^{T-1}, 0)$, the result follows from the former case. \square

Proof of Proposition 3. By Proposition 2, for all $\alpha \in [0, 1]$ it holds that $EY_B^T(s^{T-1}, \alpha) \rightarrow EY_B$ almost surely for $T \rightarrow \infty$. Hence, since the convergence is almost surely with respect to the probability measure \mathbb{P} , the variable $EY_B^T(s^{T-1}, \alpha)$ also converges to EY_B in probability with respect to that measure, that is, for all $\varepsilon > 0$, it holds that

$$\lim_{T \rightarrow \infty} \mathbb{P} \left[|EY_B^T(s^{T-1}, \alpha) - EY_B| \geq \varepsilon \right] = 0.$$

Since, by assumption, $EY_B > \bar{y}$, it thus follows that

$$\lim_{T \rightarrow \infty} \mathbb{P} \left[EY_B^T(s^{T-1}, \alpha) \geq \bar{y} \right] = 1,$$

that is, for every $p \in (0, 1)$, there exists a $T_p > 0$ such that for every $T > T_p$ it holds that $\mathbb{P} \left[EY_B^T(s^{T-1}, \alpha) \geq \bar{y} \right] \geq p$. Denoting, for each $p \in (0, 1)$, the smallest such $T_p \in \mathbb{N}$ by T_p^* , we obtain

$$T_p^* = T_p^*(q, n, \alpha) = \min \left\{ T \in \mathbb{N} : \mathbb{P}[EY_B^T(s^{T-1}, \alpha) \geq \bar{y}] \geq p \right\}.$$

Clearly, $T_p^*(q, n, \alpha)$ is finite for all $q \in (0, 1)$, $n \in \mathbb{N}$ and $\alpha > 0$.

By the definition of the quota, the minimum number of draws (that is, the length of the information vector s^{T-1}) the firm has to learn from in each period is given by $\lceil qn \rceil \in \mathbb{N}$. Since $q \in (0, 1)$ and $n \in \mathbb{N}$, the product $qn > 0$ is strictly increasing in both factors. Therefore, for each $T \in \mathbb{N}$, it holds that $\mathbb{P} \left[EY_B^T(s^{T-1}, \alpha) \geq \bar{y} \right]$ is increasing in q and n . Hence, $T_p^*(q, n, \alpha)$ is decreasing in both parameters. \square

Proof of Proposition 4. Since $EY_B > \bar{y}$, by analogous arguments as in the proof of Proposition 3, for every $p \in (0, 1)$ there exists a $T_p > 0$, such that for every $T > T_p$ it holds that $\mathbb{P} \left[EY_B^T(s^{T-1}, \alpha) \geq \bar{y} \right] \geq p$, and the smallest such T_p is given by

$$T_p^*(q, n, \alpha) = \min \left\{ T \in \mathbb{N} : \mathbb{P}[EY_B^T(s^{T-1}, \alpha) \geq \bar{y}] \geq p \right\}.$$

Fix parameters $n \in \mathbb{N}$ and $q \in (0, 1)$ and let $\alpha_1, \alpha_2 \in [0, 1]$ with $\alpha_1 > \alpha_2$. Then, for all $T \in \mathbb{N}$, it holds that $C_B^T(s^{T-1}, \alpha_1) \subset C_B^T(s^{T-1}, \alpha_2)$. Since

$$EY_B^T(s^{T-1}, \alpha) = \min_{\mu^T \in C_B^T(s^{T-1}, \alpha)} \int \int y(a) dF(a|\theta) d\mu^T(\theta),$$

for every $T \in \mathbb{N}$ it therefore must hold that $EY_B^T(s^{T-1}, \alpha_1) \geq EY_B^T(s^{T-1}, \alpha_2)$. Hence, $T_p^*(q, n, \alpha)$ is decreasing in α . □