Policy Analysis in Heterogeneous Agent Economies

Dissertation

Approved as fulfillment of the requirements for the degree of
Doctor rerum politicarum (Dr. rer. pol.)
by the
Department of Law and Economics
TECHNISCHE UNIVERSITÄT DARMSTADT

submitted by

Metin Akyol (Dipl. Volkswirt)
born on August 21st 1983 in Salzgitter

Date of submission: July 5th, 2016
Date of defense: October 20th, 2016

First Referee: Prof. Dr. Michael Neugart
Second Referee: Prof. Dr. Jens Krüger

Darmstadt, 2016

D 17

Contents

[1] This Chapter is based on Akyol (2016a)
[2] This Chapter is based on Akyol et al. (2015)

[3]This Chapter is based on Akyol (2016b)

List of Figures

List of Tables

Chapter 1

Introduction

In this dissertation I evaluate policy proposals for three different areas of the labor market, namely education, gender inequality, and social networks. The important role of education in the transition process of young adults into the labor market cannot be overestimated (see Goldberg and Smith (2008) for an overview occupational effects of educational systems). Achievement gaps in education between low income or minority students and their respective counterparts as they are observed in the United States (see OECD-Report (2012b)) can thus be the reason for a path dependency towards undesirable labor market outcomes of the individuals. Similarly, gender inequality constitutes another highly relevant aspect of the labor market, as gender gaps in wages and employment are still prevalent in many countries (for an overview see Altonji and Blank (1999a)) and are still a source of major policy debates. Attempts to alleviate the negative outcomes often include affirmative action policies such as mandatory board room gender quotas as they have been recently introduced in Norway and Germany. Social networks finally play a highly important role in affecting individual behaviour and the labor market outcomes of individuals fr instance through job referrals or educational choices (see (Jackson and Lopez-Pintado, 2013)). This embeddedness of the individual in the surrounding environment can be the reason for poverty traps from which individuals cannot escape easily. Policy measures that are directed at providing remedies for the challenges in these three areas often times have to be evaluated ex ante as empirical evidence is typically not available thus rendering a theoretical approach a necessary alternative. While traditional economic models are one viable approach, in this dissertation I analyze all policy measures using Agent Based models instead (see Neugart (2008) for an review of the use of Agent-Based models for the purpose of labor market policy evaluations). Unlike the standard economic models these provide multiple useful techniques that proof to be particularly useful in this context, especially when it comes to taking some of the complexities of the policies and their effect on the economy

into account. In order to provide some initial insight to the overarching methodological approach of all three analyses, Chapter 2 gives an overview of Agent Based models, and explains the motivation for using them and typical problems that lend themselves to be analyzed with Agent Based models.

Afterwards, in Chapter 3 I look into the first policy, namely the effect educational vouchers on educational inequality and inefficiency. The educational sector has displayed a large degree of heterogeneity with regard to student outcomes (see Book (2012)) in the past decades and public schools are often named as one of the culprits. The proposal of educational vouchers has been subject of debates in politics as well as academia. Proponents argue, that the introduction would serve two purposes at once: 1) Allow students from low income families to attend better (or more expensive schools) while at the same time 2) Exerting competitive pressure on public schools in order to increase their performance. Opponents argue, that this will cause a "*cream skimming*" effect in public schools, in the sense that students with higher ability will be the ones predominantly leaving public schools, thereby exposing the public schools to a deterioration of the average skills. It is therefore still incumbent on the research in this area to reconcile these seemingly contradicting predictions. As the educational vouchers have only been introduced in a few single school districts over the United States, the empirical evidence is limited. For this reason, I analyze in a computational model what the possible effects of such a policy could be. I find that for the traditional voucher as it was originally proposed (which would provide the entire student body with a subsidy) the effects are ambiguous. While the increased competitive pressure on public schools indeed incentivizes them to increase their performance (measured in expenditure on students) the deterioration in the form of a decreased mean ability in public schools counterbalances these welfare gains for students. Then, I analyze an alternative voucher proposal, which would target students with low ability and provide a higher voucher for students with low ability. I show that with this system it would be possible to reap the benefits of increased competition while avoiding the negative effects of "cream skimming".

Chapter 4 then analyzes an alternative policy instrument to the fixed gender quota for the boards of publicly listed companies as it was recently introduced in Germany and over other countries in Europe. Given the heterogeneous distribution of women over the various industry sectors, a fixed quota confronts companies with a different degree of difficulty to fulfil the quota. The case of Norway where the quota was introduced in 2003 provides empirical evidence of the negative possible effects of such a quota. The alternative proposal consists of a tradable quota similar to the tradable quotas for emmission certificates that would allow to alleviate the costs placed by the policy on the companies while simultaneously achieving the the goal to obtain a certain fraction of women on the boards industry wide. All of the targeted companies would be provided

with a certain number of certificates which would allow them to hire men. Companies operating in industries that have traditionally a higher share of men in the labor force, would be allowed to hire additional men by buying additional certificates from companies willing to sell their certificates. Companies having relatively few men on their boards, would be typically the ones selling certificates. Using an agent based simulation it is then shown that the alternative proposal is superior in terms of overall welfare to the fixed quota and achieves the goal at a lower adjustment cost by companies.

In Chapter 5, I evaluate the effects of a labor market policy that provides a subsidy for the purpose of investing into the education of individuals (e.g. job training, secondary education). In order to maximize any spillover effects on non-treated individuals, I examine how the effects of the policy can be increased by taking the position of an individual in the network into account. Similar policy proposals have been made by Ballester et al. (2006) who look into crime networks and try to determine how the overall activity among criminal individuals can be reduced the most by eliminating ties between the most central (and thereby influential) individuals in a network. Using a computational model, I show how the overal effect of the policy can be increased by targeting the most central individuals in the network. However, this comes at the cost of increasing inequality in the overall network as the agents in the periphery of the network (who tend to be the least active) fall further behind if the policies focus on the center of the network. Chapter 6 then concludes.

Chapter 2

Agent-Based Simulation Models

The exponential growth in computing power that characterized the past decades has introduced a vast array of new methods for analysing economic problems and rendered the available toolkit much more diverse. Theoretical models were suddenly not as much dependant on economic tractability but rather allowed for more complex assumptions regarding the behavior of agents. While the major fraction of theoretical models relied on so called standard economic models, that featured an all market encompassing general equilibrium and agents that are characterized by rational expectations, another strand of literature emerged that dispensed with the assumptions of equilibria and perfect rationality, namely agent-based computational economics (ACE). ACE is a computational approach to economies which are modelled as evolving systems comprised of autonomous interacting agents (see Tesfatsion and Judd (2006)).

What are some of the advantages of ACE models relative to more standard modeling approaches?

While the standard economic models focus on the equilibrium behaviour of agents, ACE models lay emphasis on the dynamics that describe the behaviour of agents apart from an equilibrium analysis, which may or may not lead to a stable outcome in the long run. Instead, ACE models comprise a dynamic system of interacting agents and attempts to gain more insight about about the aggregate outcomes resulting from the individual behaviours on the micro level. That is, ACEs allow to model the two-way feedback between the micro structure and macro structure which takes place in the economy (see Schelling (1978) and Olson (1965)) that was not possible before. A typical set of elements often featured in ACE models are inductive learning mechanisms, imperfect competition, endogenous trade network formation, the open-ended co-evolution of individual behaviors and economic institutions (see Tesfatsion (2002)). The approach of ACE models is often

guided by an attempt to detect an explanation for the emergence of certain global market phenomena that have not been obstructed via top-down planning and control. Put differently, what are the type of behaviours of individuals that may have lead to these outcomes from the bottom up is the typical question that is addressed. Another very popular approach (that is also employed in the following chapters) is the attempt to analyze the implications of certain economic entities or policies (see Dawid and Neugart (2011)) on the performance of the economy as a whole. Apart from the differences in equilibrium analysis, a second key departure of ACE models from the standard economic models is the facility of agents engaging in social communication, thereby allowing for interactions such as trading and exchanging of informations as it is observed in reality. This type of interaction allows to model in particular how individuals are affected by their peers in various ways (e.g. educational investment, job referrals etc.) and is part of the analysis in Chapter 5. Finally, it is often claimed that the core advantage that sets ACE models apart from the more traditional approach is the autonomy of the agents (see Jennings (2000)). While it could be argued that the typical representative consumer in neoclassical models is also characterized by a certain degree of autonomy, the agent based model extends this autonomy in various ways. It allows for instance to model the cognitive processes that guide the actions of the individual, e.g. a learning process through which individuals improve the outcomes of their interactions. This also includes the alterations of preferences as an outcome of such a learning process. This is of particular interest as it allows for individuals to have a limited ability to absorb all available information that is available to them, thereby rendering their behavior more similar to empirically observed behavioural patterns (see for instance Duffy (2006)). One important method of calibrating these behavioural patterns to be more realistic is the approach of using experimental data to parametrize the behavioural functions. A typical example is the modelling of learning processes through experimental data as done in the so called Experience Weighted Attraction (EWA) model (see Camerer and Ho (2008) and Brenner (2006) for a extensive review of the modeling of learning processes), a method that is also employed in the Chapter 3.

What are some of the disadvantages of ACE models relative to more standard modeling approaches?

The ability to model the interaction between agents in the model in a detailed fashion also represents one of the drawbacks of ACE models. It is very well conceivable that small variations in the specifications of the model can induce strong enough feedbacks within the model in oder to lead to radically different outcomes. Thus, in order to obtain sufficiently robust results from a model specification, a wide array of specifications must be included in the experiments (for an extensive discussion of this argument see Judd (2006) and Richiardi (2012)).

The deviation from the principle of an equilibrium to which the model converges introduces another complication for ACE models. The assumption of behavioural patterns when setting the micro behaviour (which are designed in order to match empirical patterns) causes outcomes that emerge often multi-peaked in the sense that several outcomes appear most likely to occur, rather than a central tendency. This makes the empirical validation of these outcomes difficult when it comes to a single observed outcome (e.g. stock market crash) rather than a distribution of multiple possible outcomes (e.g. a certain distribution of returns on the stock market). Richiardi et al. (2006) suggests that validation encompasses a large number of necessary verifications of the model structure, for instance: 1) Whether the specification still make sense even when inputs take on extreme values 2) Does the model generate the empirically observed characteristics of the real system.

Furthermore, Richiardi et al. (2003) point out the "interpretation problem" of ACE models, i.e the challenge to extract information from the observed phenomena in the simulated data regarding the underlying reasons why they occured in that way. One common approach is the use of econometric techniques to estimate reduced forms on the artificial data, where the researcher generally knows which parameters affect the outcomes of interest (since this is part of the specification of the micro rules). However, this approximation can not be used for further estimation on empirical data, since this is only a description of how the simulation model works for a given set of structural parameters. In more intricate models, the causal links between inputs and outputs might then be indirect and remain difficult to detect.

Some of the areas where ACE models provide a very useful methodological approach is the modelling of economic activity in exogenously determined networks. The particular design of a model consisting of individual autonomous agents lends itself to a network analysis where one can track the flow of information to and from any particular agent and look at the implications that results from the network structure (see for a review Wilhite (2006)). Chapter 5 employs such a network structure to analyze the diffusion processes that result from providing labor market policies to individuals embedded in a social network. A highly related field is modelling of endogenized interactions where the agents not only interact with other agent but also determine through their own preferences who they actually interact with (for a review see Vriend (2006)). A prominent example of such endogenized interactions is Schelling (1971)'s spatial proximity model of neighborhood segregation which features two types of agents. He demonstrates how even slightest preferences of the agents for neighbors of their own kind among the agents can lead to a highly segregated distribution of agents in the steady state even if the distribution is strongly integrated in the initial state. The set of possible applications also extends into models in finance, that allow to model the behaviour of agents in the stock market

while including agents that make trading decisions based on different heuristics or rule of thumb strategies which may not be perfect but are able to replicated stylized facts reasonable well (for an overview see Hommes (2006). ACE models are furthermore applied to topics in labor economics (for a survey of the agent based labor market models see Neugart et al. (2012)), typically to either explain stylized facts of labor market data (e.g. the Beveridge curve, the wage curve, the wage distribution etc.) or to analyze the implications of labor market policies (e.g. training policies, unemployment benefits). The latter approach is used in this dissertation, i.e. the evaluations conducted have in common the attempt to predict the outcomes of policies for which empirical evidence is either not available or very limited. Neugart et al. (2012) point out that one of the main advantages of ACE models in labor market models is the possibility of a flexible model design that allow to weaken many of the standard assumptions simultaneously.

What are the typical steps involved in an ACE model? An initiation, the ACE models is populated with a group of agents. These can represent anything from economic agents such as financial institutions or traders, or government institutions. These agents feature initial characteristic attributes (e.g. financial endowment, behavioural patterns). In a further step the behavioural patterns of all agents in the model are set using functions that model their behavior. As pointed out before, this could be based on empirical or experimental data where individuals react to the events in the model based on a function. This could also consist of allowing the agents to adjust their behavior based on information gained through a learning process (e.g. the agents could evaluate the outcomes of their previous actions in order learn through trial and error). In addition to the actual functional form of the behavioural patterns, the parametrization of these functions imposes a further challenge to the modeller. The proper calibration might require using empirical or experimental data that provides information on how the detailed reaction functions of agents might be best defined. The particular procedure in a model is often documented with pseudocode or chart flow diagrams. Then, the modeller either allows for the model to evolve for a while and then introduces a policy measure or does not intervene further . That is, the subsequent events arising, must arise from the interactions occurring within the model (i.e. without any exogenous impositions of market clearing conditions or optimality conditions). Once results are obtained, the interpretation and validation of these findings can include various steps. One possible approach is to use standard statistical or econometric methods, such as comparing means and distributions of those variables of interest. Another popular approach is the estimation of reduced form models as pointed out above to obtain an understanding of the underlying mechanics.

Chapter 3

Do Educational Vouchers Reduce Inequality and Inefficiency in Education?[1]

3.1 Introduction

The topic of poor public school performance is still subject of major policy debates and is ever more interesting in light of the great achievement gaps between low-income or minority students and their respective counterparts as suggested for instance by the PISA assessment for the United States. The survey finds that "[s]ocio-economic disadvantage has a notable impact on student performance in the United States" (OECD-Report (2012b)). In particular school systems that are highly stratified tend to perform worse in terms of student test scores (OECD-Report (2012a)). This is also reflected by the extensive reports by the National Center for Educational Statistics (Vanneman et al. (2009) and Hemphill and Vanneman (2011)). One proposed strategy for improving public school performance that has received increasing academic and public attention is to introduce competition through the provision of private school vouchers to students, which would grant them financial aid if they chose to attend a private school. The vouchers would provide a large student body with the opportunity to attend any school of their choice rather than having to attend the public neighborhood school, and thereby create competitive pressure for all public schools in the respective school district. The

[1]This Chapter is based on Akyol (2016a)

idea is based on the paradigm that, analogous to other markets, schools having to compete for students, will enhance efficiency and improve the quality of their "product", i.e. educational experience for the students attending.[2]

Several aspects particularly relevant to the introduction of a voucher system warrant special attention. Will the program actually benefit students from low income families or minorities who are attending low performing schools with lower quality peer groups, i.e. will the prevalent stratification and inequality in education be reduced? Will the public schools actually get better and increase their efficiency or rather be gradually be exposed to "*cream skimming*", i.e. a deteriorating peer group? This is of high relevance as it affects the peer group of both school systems. There is a rich empirical literature that presents strong evidence for the role of peer effects in primary and secondary education (for an extensive review see Epple and Romano (2011) or Sacerdote (2011)). Coleman et al. (1966) for instance use survey data and confirm the crucial role of peer effects and find that they are even more important for disadvantaged students. Summers and Wolfe (1977) also find that the peer group effects play a significant role in educational outcomes (see Calvo-Armengol et al. (2009a), McEwan (2003), Gaviria and Raphael (2001) and Zimmerman (2003) for results in a similar vein). It would thus be of interest to provide a comprehensive approach that would allow to engage into the complexities of the topic along multiple lines, i.e. efficiency concerns and distribution of students. The empirical literature on educational vouchers typically only focuses on certain aspects of the voucher program. Thus, as Nechyba (2000) puts it "...[empirical] work may not anticipate all the impacts from a large-scale policy...At the same time, theoretical models or school finance are also limited in that they often either focus on only one particular aspect of the general equilibrium school finance problem, or are they too rich and complex to yield crisp predictions. It is for this reason that there is great potential for simulation approaches...".[3]

In this paper, I implement an agent-based simulation model of a representative US school district which draws on the seminal contributions of Manski (1992) and Epple and Romano (1998, 2008). I use the model to analyze the effect of a voucher program on inequality and inefficiency in the educational system.[4] The contribution of this paper is a twofold approach that allows for an endogenized reaction of public schools to increased competition on the one hand and to track the distribution of students to schools in order to evaluate peer group effects on the other hand. In particular, I extend the above

[2]The idea dates far back, Thomas Paine for instance proposed a voucher plan in 1792, in *The Rights of Man*, for a discussion see West (1967). The more recent awakening of interest is typically credited to Friedman (1955, 1962, 1997)

[3]Two recent papers that allow for richer complexities to enter and that take into account the general equilibrium effects are Epple et al. (2013)and Fu (2014).

[4]See Spiro Maroulis and Wilensky (2014) for an approach to model the transition period for public choice in an agent-based model

mentioned models to account for the heterogeneity of the student body, as the agent based approach allows a much more detailed modeling of the individual behavior of the students.[5] Specifically, students have varying characteristics in multiple dimensions that are orthogonal to each other. Thus, I obtain distributions of students to schools that are only based on the inherent characteristics of students and their resulting school choices. The model then simulates how various student groups are affected not only by the introduction of the voucher program but also by the choices of their peers (and possible feedback effects). Using this approach I then I compare the effects of *universal vouchers* and *target vouchers* (which are a function of student characteristics), and compare outcomes for different distributions of income and ability.

The main findings indicate that the outcomes of a voucher program hinge on the design of the programs. First, I find that a higher number of students being able to afford to switch from public to private school (and thereby are able to exert competitive pressure on public schools as their children do not default into public school) causes an increase in public school performance. This effect can be created either by a change in distribution of family incomes (i.e. increasing the number of families in the middle class relative to the low income class) or through the distribution of vouchers in order to increase the available funds to families. This outcome is confirmed by empirical studies listed below, which find that public schools do react to higher competition.

Second, when looking at the effect of changes in the *universal voucher* level, the findings are similar to the results observed in (Manski, 1992) and Epple and Romano (1998, 2008). While students actually exercising choice profit from a voucher system, the effect on students remaining in their original schools is less clear. Given that the majority of students exercising choice and leaving into better schools have mostly higher ability levels, the students staying behind observe a decline in the peer group quality ("cream skimming"). Contrary to these papers, the endogenized reaction of public schools allows for them to react to a sudden decline in student enrollment by adjusting educational quality accordingly. Thus, the low performing schools actually increase their educational service to retain students, that is, the *"cream skimming"* effect, is partly alleviated through this effect.

Finally, when the model is extended to allow for so called *target vouchers*, which are a function of the ability level or the family income of the respective students, the observed deterioration effect is avoided in the case of *ability vouchers*. Public schools observe less of a decline in their mean ability while maintaining the increase in public school expenditure. Thus, *ability vouchers* allow to obtain the benefits of higher competition

[5] Lavecchia et al. (2014) provide a survey of recent applications of behavioral economics to the economics of education and provide a framework for modeling the decision making processes of parents and students while taking into account their limited rationality.

while circumventing the detrimental peer group effects. The concern to what extent this policy alternative would actually be feasible is discussed in that section.

The rest of the paper is structured as follows. Section 3.2 contains a literature review. In Section 5.2, I lay out the details of the model that I employ. Section 5.3 and section 5.3.3 present results from the case for *universal vouchers* and the case of *target vouchers*, respectively. Section 3.6 presents a number of robustness tests, section 6 then concludes.

3.2 Literature Review

There is no shortage of diverging arguments about possible effects of a voucher program on efficiency of public schools but also the distributional aspects, i.e. which students will actually benefit. Hoxby (2003b) for instance makes the case, that economic logic suggests that choice will enhance productivity, and render distributional concerns less relevant, as an overall improved school system will help to improve all students' educational experience ("...a rising tide that lifts all boats"). MacLeod and Urquiola (2012) argue, that the argument for more incentives through competition does not take the idiosyncrasies of the educational system into account, that would make this prediction fail. In fact, they show that if one applies the theory of incomplete markets and incomplete contracts to the educational system, the predictions must be much more nuanced, and that there is no a priori reason to conclude that the introduction will in fact benefit all students.

The empirical research provides mixed results of voucher programs and supports the argument that expectations of educational vouchers improving the educational outcomes throughout all student groups should be tempered. Barrow and Rouse (2008) and MacLeod and Urquiola (2012), both provide a survey of empirical studies of the so called *direct effect*, i.e. the resulting changes in individual test scores when students switch from public to private schools. They find that a large number of programs only reveal a small or insignificant effect on student test scores when switching from public to private school (see for instance Neal (2002)). A sizable number of empirical studies also estimate the *indirect effect* or *general equilibrium effect*, i.e. the reaction of public schools to increased competition.[6] Levin and Belfield (2002) provide an extensive review of the empirical literature and find that a majority of studies report beneficial effects of competition, however the positive gains from competition are modest in scope with respect to realistic changes in levels of competition. Greene (2001), Greene and Winters (2003) and Figlio and Rouse (2006) for instance look at evidence from the voucher program introduced in Florida in 1999 and 2002 in which schools that did not have mean

[6] Albeit the number of studies is comparably small given the plethora of studies for the *direct effect*

test scores above a certain threshold were threatened to have their students become eligible for vouchers if they did not meet the criteria a second time. They find that the respective schools achieved test score gains that were twice as large as those achieved by other schools. West and Peterson (2006) also find positive effects on test performances of students at the threatened schools. Hoxby (2003b,a) looks at the program in Milwaukee and also finds a positive productivity response in public schools. Chakrabarti (2013) compares both programs, Milwaukee and Florida, and finds that the "threatening design" of the Florida program produced larger gains in productivity, and thus highlights that the incentive design of the voucher program is paramount in order to achieve productivity gains. Rouse et al. (2013) and Chiang (2009) argue that these positive results might not reflect genuine improvement but could result from "gaming" behaviors (focusing teaching efforts on students at the threshold of failing etc.). They find that threatened schools improved instructional practices which are explanatory for the observed higher test scores in the voucher program. Misra et al. (2012) use Geographical Information System (GIS) data to develop a school competition index for the state Mississippi and find that higher degrees of competition from private schools significantly increases public primary and high school efficiency in terms of test scores. Bayer and McMillan (2010) also find a positive effect of increased competition using data from the San Francisco Bay Area. From a review of the empirical literature it thus seems that there is potential for an improved public school system through pressure and also for small *direct effects* on students who switch to public school.[7]

Manski (1992) was among the first to develop a theoretical and computational analysis of vouchers and was highly influential for following research.[8] He finds a varying effect on the poor of a voucher program, as the ones staying behind in public school are worse off because their peer group deteriorates, while the few ones who switch to private schools are better off. At the same time, for the case of rent seeking public schools, he finds a decrease in their rents with the introduction of the voucher program. Epple and Romano (1998) analyze vouchers in a framework with peer effects that allows schools to put a price on the peer externality and charge students their reservation prices. They find a similar effect of a deteriorating peer group in public schools (or *"cream skimming"*), while a typical hierarchy of private schools results along with an "income and ability stratification" (in Epple et al. (2004) they provide empirical evidence for these findings). As the introduction of a voucher system hinges on the ability to gather enough political support, Epple and Romano (2014) examine the conditions under which such a system would be preferred. They show that the coefficient of variation in the wealth level of an

[7]Throughout this paper it is assumed that peer effects occur in a homogeneous fashion, i.e. individuals assigned to a certain peer group such as a classroom or school experience the same effects. Carrell et al. (2013) uses experimental data to show how individuals can be affected differently by their peer group due to a tendency to form homogeneous subgroups.

[8]For a recent survey of the theoretical modeling advances see Epple and Romano (2012)

economy is the relevant statistic determining whether a voucher is politically supported.[9] Epple and Romano (2008) extend the analysis and allow for a '*target voucher*' which is a function of ability and demonstrate that it can eliminate cream skimming and "reap the competitive efficiency gains ... with gains shared by all students", i.e. the voucher leads to an elimination of public schools and a homogeneous private schools which admit a cross section of all students. Thus, the voucher results in no stratification and the schools spending all the voucher money on expenditure. Nechyba (2000) also models target vouchers which are a function of income or are specifically targeted to low income neighborhoods and finds that overall variances on school quality variables could actually decline. Caucutt (2002) finds that target vouchers which are predominantly provided to low income students entail adverse effects on inequality and higher welfare losses compared to the general vouchers case.[10] This paper extends these models by allowing for an endogenous reaction of public schools to the educational vouchers and thereby evaluate the efficiency effects and the peer effects from the resulting redistribution of students to schools.

3.3 The Model

3.3.1 Model Specification

The model represents a typical school district in the United States that is comprised of a public sector and a private sector, which both compete with each other for students. Both sectors set control variables simultaneously at the beginning of each period in order to maximize their respective objective function, and are homogenous with respect to all characteristics. They have restricted knowledge of each others payoffs, i.e. there is no perfect information and either side can merely observe the other sides' policies after they have been implemented. Students then rank both sectors according to the possible utility they would gain from attending each sector and then choose the higher ranked sector. Schools are entirely transparent for households regarding the achievement of the students enrolled and the expenses per student. This appears to be a realistic assumption given the information that is made publicly available in many school districts in the United States. The schools then observe enrollment and use a learning mechanism in order to adjust their control variables in the next period.

[9] See Glomm and Ravikumar (1998) and Glomm et al. (2011) for a theoretical modeling of the voting process when public and private education alternatives exists

[10] See also Chen and West (2000) and Piolatto (2010), both papers show in a theoretical model how a system of selective vouchers would be superior to universal vouchers and would be preferred by the majority of voters

3.3.2 Household Characteristics

The model is populated with n heterogeneous households. Each household i has an income, denoted as y_i, and has one student who has an ability of b_i. Furthermore, each household has a private school preference p_i, which is a continuous variable measuring the idiosyncratic preference of a student for private school relative to public school, where p_i is distributed normal with variance one and mean μ, $p_i \sim \mathcal{N}$ (μ,1).[11] I assume that y, b, and p are distributed independently of one another. This simplification assures that the results are driven solely by our behavioral assumptions rather than any initial distributional assumptions. This assumptions serves the purpose of warranting that any of the distributional outcomes will solely emerge from the behavioral assumptions of the heterogeneous households and public and private schools, instead of being driven by any a priori assumptions about the initial distribution.[12]

In particular, b and y have uniform distributions, f(b) and f(y), which are assumed to be continuous and positive on their support, $S_b = [0, b_{max}]$ and $S_y = [0, y_{max}]$.

Households maximize their utility function, $U_{is} = U(\cdot)$, through their choice of school, s, for the household's student. $U(\cdot)$ is increasing in the consumption of a numeraire good, the mean ability of the student body, m_s, and the instructional expenditure of the school, e_s, and is continuous and twice differentiable in all arguments. Educational attainment for student i, $a_i = a(m_{s,}e_s)$, is a function of m_s and e_s and is a non negative, and increasing in both arguments. By taking m_s into account, households consider the peer-group effect in their choice of school.[13] Let t_s be the tuition charged at school s, and v_{is} the voucher that student i receives for attending schools s, then we have:

$$U_{is}(\cdot) = U(y_i - t_s + v_{is}, m_s, e_s, p_i) \tag{3.1}$$

with U'() positive for all arguments.[14]

[11] Similar to Manski (1992) I use p_i to adjust initial enrollment in public/private schools to reflect more realistically the empirically observed ratio of students in private vs. public schools, it can reflect various forces: preference for religious relative to secular schooling, the time required to commute to a private school relative to public school etc.

[12] This assumption is relaxed in the robustness section with the main results being unaffected.

[13] See Henderson et al. (1978), Summers and Wolfe (1977), and Hoxby (2000) who provide empirical evidence for a peer group effect in educational achievement

[14] As Epple and Romano (1998, 2008) I assume that $U_i(\cdot)$ satisfies everywhere the "single crossing" condition in income y_i and in ability b_i, that is:

$$\partial(\frac{\partial U/\partial e}{\partial U/\partial t_s})/\partial y_i > 0 \tag{3.2}$$

As mentioned above, the model abstracts away from informational issues, students are thus perfectly informed about school characteristics as they rank schools. In particular, households maximize utility, $U_{is} = U_{is}(\cdot)$, through their choice of school s. In order to preclude kickbacks by schools, that is schools that enter the market at the lower spectrum of educational quality and attract particularly low income students by kicking back monies of the voucher, the condition $t_s \geq v_{is}$ is introduced. This does not preclude topping up by participating schools, thus schools are allowed to charge a higher tuition than the voucher amount if they deem it appropriate in order to maximize their objective function. The utility function of household i in Equation 3.1 for attending school s is then given by:

$$U_{is} = \alpha_i log(e_s) + \beta_i m_s + \gamma log(y_i - max\{t_s + v_{is}, 0\}) + p_i \tag{3.4}$$

with $\alpha_i = \alpha(b_i)$, $\beta_i = \beta(b_i)$, and $\gamma > 0$, where $\alpha'()$ and $\beta'()$ are all positive, i.e. students with higher intrinsic ability, ceteris paribus, value the educational quality and peer quality of each respective school higher than students with lower intrinsic ability.[15]

3.3.3 School Characteristics

Schools are divided into public and private, with schools being homogeneous within each sector. Every student can choose between the public school sector (s = 0) and the private school sector (s = 1), whereas the former is available for free and preferred to no schooling, that is, tuition t_0 for public schools is set to zero, while tuition for private schools is endogenous. School finance policy warrants public and private schools a funding of v_{is} that is proportional to enrollment and a function of student characteristics (in case of target vouchers). Each school chooses their expenditure per student, e_s, and has a publicly known mean student ability, m_s, which is endogenous. Both, e_s and m_s constitute educational quality for the school. The amount of funding for the schools depends on the type of policy that is examined. For the sake of simplicity of notation, I will refer to all forms of funds from the government to schools (be it public or private) as a voucher, v_s. This voucher will be stable in size for public schools throughout all policy experiments, while it will be varied for private schools.

and:

$$\partial(\frac{\partial U/\partial e}{\partial U/\partial t_s})/\partial b_i > 0 \tag{3.3}$$

where e represents educational quality, either as expenditure per student or the mean ability of students enrolled in that particular school. Equation 3.2 implies that in the (e,t_s)-plane, the indifference curves of students with the same motivation are becoming steeper as income increases. This is equivalent to saying that income elasticity of demand for educational quality is positive, i.e. educational quality is a normal good. Analogously, 3.3 implies a positive motivation elasticity of demand for educational quality.

[15] See Rothstein (2006) for evidence of how parents are evaluating schools by peers and school quality.

3.3.4 Public Sector Schools

Public schools do not face within sector competition and offer free admission to all students.[16] They are local monopolists, in the sense that each each sector has only one public school, and given this lack of competition they are not required to expend all of their revenue in a way that benefits their students (this could be considered a profit that is not used for the benefit of the student body or a bureaucratic inefficiency).

Thus, the public sector can set expenditure e_0 such that their maximization problem yields a surplus:

$$\max_{e_s} \sum_{i=0}^{n} (1 - c_i)(v_s - e_s) \tag{3.5}$$

where c_i is an indicator variable that takes the value of 1 when student i is enrolled in private school, with n being the number of students.

While this allows each public school to accrue a surplus at the end of each period, it does not imply that public schools actually do so. It is very well possible for public schools to expend all of their surplus to become more attractive to students in their district. A number of empirical studies (see (Hoxby, 2003b)) find that increased competition in a school district actual leads to an increased performance of public schools (measured through students' achievement). The goal of this modeling approach is to examine if (under the assumption that public schools do have a surplus and thus ability to increase their expenditure) the introduction of increased competition actually does lead an increased expenditure in the sense that public school surplus is reduced. It should be also noted that it is indeed conceivable that the improved public school performance works through a "non-monetary" channel, i.e. public schools improve while maintaining expenditure. While this is not explicitly modeled here, the implications would be similar to the degree that public schools in both scenarios do have some slack in their ability to increase performance, be it through "monetary" or "non-monetary" channels.

3.3.5 Private Sector Schools

Private schools act competitively as the typical schools district consists of numerous private schools and students who are able to afford private school tuition are able to

[16] The fact that about 90 percent of the student body attends public schools, and 70 percent attend their local neighborhood school, supports this assertion (see Statistical Abstract of the US Book (2012)).

choose between these various options.[17] They set their tuition, t_s, and their expenditure, e_s, to maximize their preliminary objective function:

$$\max_{t_s, e_s} \sum_{i=0}^{n} c_i(t_s + v_s - e_s) \tag{3.6}$$

Since the entry into and exit from the private school market is not very costly, I assume that private school sector is forced to act competitively. This is corroborated by the relatively large number of students who default into their local public school as previously mentioned, while the remaining students choose a private school and thereby exert competitive pressure on these schools which precludes private schools from accumulating any profits within a period, i.e. they set their expenditure, e_1, such that equation 3.6 becomes zero[18] at all times, we thus have:

$$\sum_{i=0}^{n} c_i(t_s + v_s - e_s) = 0 \tag{3.7}$$

Private schools then choose t_s and thereby set e_s, i.e. the maximization problem reduces to the determination of t_s. They do not face any restrictions with the introduction of the voucher, i.e. they may charge tuition above the voucher level. I assume that private schools then set tuition in such a way, as to maximize private school sector enrollment, i.e. the maximization problem of the private sector becomes:

$$\max_{t_1} \sum_{i=0}^{n} c_i \tag{3.8}$$

Thus, private schools attempt to maximize their enrollment through their choice of tuition, t_1. One possible motivation for this modeling strategy is that the market structure of the private school sector can be thought of a as monopolistic competition, where the schools differentiate their services through advertising and other non-price strategies and have some pricing power. In the short run, it will then be possible for schools to accrue profits. However, because there are relatively low entry and exit costs, new schools will enter the market and in the long run, drive prices and revenues down toward an equilibrium similar to perfect competition, where profits are close to zero.

[17] An alternative modeling strategy could be to take the reputation of schools into consideration which could hinder new entries, see MacLeod and Urquiola (2009).

[18] this is analogue to the approach in (Manski, 1992). Epple and Romano (2008) use a similar condition to characterize profits of private schools in equilibrium close to zero.

3.3.6 School Behavior

The complexity of the presented model precludes a closed form solution for the endogenous variables, i.e. public school expenditure and private school tuition.[19] Instead, I employ the Experience-Weighted Attraction (EWA) model as suggested by Camerer and Ho (1999). The authors developed their model as an approach to determine how an equilibrium arises in non-cooperative games where it is conceivable that agents are not able to reason their way but rather adapt or evolve toward it. They diverged in their approach from the largely theoretical literature on evolutionary and adaptive process as they used experimental data to specify a model that would describe actual human learning and adaption best. The basic mechanism consists of two fundamental types of learning processes, *reinforcement learning* and *belief learning* and treats both as border cases for the model's parameters. *Reinforcement learning* is closely related to the concept from behavioral psychology and describes an individuals behavior as the result of the positive feedback received in the past. *Belief learning* on the other hand models how players engage in dynamic games in which they try to optimize their behavior using a prediction rule, which provides a forecast of other player's behavior as a function of their past history. As such, the EWA model constitutes a general learning model that features two traditional learning models as extreme cases. The model has experimental evidence supporting the algorithm and the choice of parameters (see Camerer and Ho (2008) and Brenner (2006)). In the context of this paper, the algorithm is used to model the path towards a stable outcome of the model in which the schools get as close as possible to maximizing their objective functions (given their control variables, i.e. public school expenditure and private school tuition).[20] The obtained outcome satisfies household utility maximization, and an approximate school's profit maximization given the respective other sector's choice. That is, both sectors do not have an incentive to unilaterally deviate from the given outcome.[21]

More specifically, the EWA Model assigns *attraction levels* to all possible strategies for the control variables based on previous payoffs or preconceived beliefs and maps those into probabilities for actual choices. In this context the attraction levels are a 1 to 1 mapping of the payoffs of the schools, that is, the payoffs for public schools are their *profits* while the payoff for private schools is their *enrollment*. These payoffs are a result

[19] See Scotchmer and Wooders (1987) and Scotchmer (1997), who show that this problem is pervasive to club economies, whereas the case of private schools can be seen as an example of clubs with "non-anonymous crowding". Manski (1992) employs a numerical simulation through all possible combinations of endogenous variables and then solves for a unique equilibrium using backward induction from the public school's perspective, given that they are dominant and the private sector consists only of small firms that take the actions of the public school sector as given. Fu (2014) employs a similar approach with a sequential game and then determines the sub-game perfect Nash equilibria that emerge.

[20] This is somewhat similar to the *"epsilon-equilibrium"* of Epple and Romano (1998)

[21] See Fowler (2011) who uses an Agent Based Model to replicate the results of a Core-Periphery Model and obtains a comparable equilibrium condition.

of the chosen strategies for the control variables, i.e. e_0 for public schools and t_1 for private schools. Each period after one strategy is drawn from the distribution, the attraction or payoff for that particular strategy is updated based on the current period's payoff and the discounted previous periods payoff. For instance, the public schools choose a certain e_0 from a probability distribution and observe the *profits* that result from this choice, while the private schools simultaneously choose t_1 and observe their *enrollment*. Both schools repeat this process in the subsequent periods, while gradually observing increased payoffs for certain strategies. Strategies with higher payoffs are chosen more frequently then. The attraction or payoff for each respective strategy j at time t for school s is given by:

$$A^j_{s,(t)} = \frac{\{\phi N_{t-1} A^j_{s,t-1} + [\delta + (1-\delta) * I(s^j_s = s_{s,t})]\pi_s(s^j_s, s_{-s,t})\}}{N_t} \tag{3.9}$$

with ϕ as a depreciation factor of previous attraction levels, and where δ represents the weight put on the fact whether a strategy has been employed before or not (where $\delta = 0$ and $\delta = 1$ constitute the border cases for reinforcement learning and belief learning). Then $s_{s,t}$ denotes the strategy employed by school sector s at period t, $s_{-s,t}$ is the strategy employed by the respective other school sector at period t, $I(s^j_s, s_{s,t})$ is an indicator function which equals one if $s^j_s = s_s t$ and zero otherwise. Finally, $\pi_s(s^j_s, s_{-s,t})$ is a payoff that sector s obtains if strategy s^j_s is chosen and the other sector chooses s_{-s}. N_t is the law of motion for the experience weight which depreciates payoffs from previous periods:

$$N_t = \rho N_{t-1} + 1 \tag{3.10}$$

where ρ is a discount factor.

The model is initialized with $N(t)$ and A^j_t according to beliefs about what experience the actors might be able to draw from, given similar other situations. The attraction levels of each strategy are then mapped into a probability for each strategy j to be chosen by sector s. Thus public schools have attraction levels for each possible level of expenditure and private schools have a attraction levels for each possible level of tuition. Each sector then draws from a Logit formulation to make a choice for the next period level for the respective control variable, which is then given by:

$$P^j_{i,t+1} = \frac{e^{\lambda A^j_t}}{\sum_{k=1}^{m_i} e^{\lambda A^j_t}} \tag{3.11}$$

where $\lambda > 0$ and measures the sensitivity of schools to the attraction of one particular strategy and m_i is the number of possible strategies that sector s can employ.

3.3.7 Policies

3.3.7.1 Universal Vouchers

I first simulate a *universal voucher* which is provided to all students independently of income or ability. The related literature has mainly focused on these type of "flat-rate" vouchers (see for instance Rangazas (1995), Epple and Romano (1996), Hoyt and Lee (1998), Caucutt (2002), Nechyba (1999, 2000, 2003), Ferreyra (2007)). In this setting the private school sector receives a voucher, v_1 for each student that is enrolled in the school, while not facing an upper boundary on the tuition that is set. As mentioned above, in order to obviate schools from offering particularly low educational expenditure and kicking back money to students,[22] private schools are required to charge tuition at least as high as the voucher amount, v_1. This scenario subsumes the case of no voucher, where only public schools receive government funding, v_0, for each student and private schools are only financed through tuition. Notice that the funding for public schools remains constant throughout all policy experiments.

3.3.7.2 Target Vouchers

In a second setting, I simulate the effects of *target vouchers*, which allow to make the tuition a function of student characteristics such as ability or income. While these do not have many real world counterparts, the idea behind *target vouchers* here is to devise a system that is more purposefully oriented towards the needs of the low income students than a *universal voucher* system. Epple and Romano (2008) study the effects of such an alternative voucher system that would allow to reap the benefits of increased competition between schools while curbing stratification effects. They show that voucher systems that condition awards on student ability can eliminate stratification and distribute benefits of increased competition. As they do not allow for public schools to react to the increased competition or for income vouchers, it is not obvious whether

[22] Epple and Romano (1998) use a similar restriction to exclude "*bottom feeder*" schools

their results for *ability vouchers* continue to hold under these more relaxed conditions or what the effects of income vouchers would be. It should also be noted that there are currently no real world counterparts of vouchers that are negatively correlated with student ability. One of the challenges of such a system would be the fact that the true ability of students is not observable to the policy maker, the feasibility thus hinges on the possibility of using an appropriate proxy for the unobserved ability. One possible approach could be to employ test scores as a proxy for ability and use them both, to assign the voucher and for admission decision. If ability is determined by examination in the process of admission to schools, there is no incentive to under-perform, i.e. this approach could be a viable attempt to render truthful ability revelation that is also incentive compatible.

3.3.8 Model Calibration

The parameters of the household utility function, $U(\cdot)$ are $\alpha_i = \alpha(b_i)$, $\beta_i = \beta(b_i)$, which are assumed to be continuous and positive on their support, $S_\alpha = [2.5, 5]$ and $S_\beta = [1, 2]$, and $\gamma = 25$.[23] The private school preference parameter μ is set such that the initial enrollment ratio between public and private schools reflects the empirically observed ratio in the Unites States, i.e. about 92% of the student body is enrolled in public school with the remainder being enrolled in private school.[24] This type of initial distribution is obtained for $\mu = -2.5$, and simply shifts the distribution while leaving the qualitative and quantitative results unaffected. Ability, b, and the family income, y have a uniform distribution on their support, $S_b = [0, 100]$ and $S_y = [0, 100]$ (where y is denoted in multiples of one thousand).

The simulations are first performed under varying income distributions and then under varying distributions of student ability. For both cases the school district consists of families from five annual income groups, which are defined in Table 3.1:

The income is distributed uniformly within these intervals. The simulations are performed with the following distributions for the comparison of income distributions:

- all families are in the middle class (Gini coefficient: 0.05)
- all classes contain 20% of families (Gini coefficient: 0.33).

[23]The current parameter values imply that students at the upper end of the ability spectrum would like to spend 16% of their income on education in the optimum (for the students at the lower end this would be 9% of income), the derivation is provided in the Appendix. Epple and Romano (1998) point out that the current percentage of aggregate disposable personal income for the United States which is spent on education is close to 5.6 %

[24]See Digest of Education Statistics 2014

Table 3.1: Income Distribution

class	**lower boundary**	**upper boundary**
low	\$0	\$20,000
lower middle	\$20,001	\$40,000
middle	\$40,001	\$60,000
upper middle	\$60,001	\$80,000
high	\$80,001	\$100,000

Note that throughout both distributions the average income remains unchanged at \$ 50,000, thus all the observed effects result only from the change in the distribution. The number of students is set such that the results obtained are stable, particularly to warrant that the learning mechanism of the schools receive enough feedback (for N very small, the inflow of new students each period does not guarantee a correct feedback for the payoff function). For this purpose $N \geq 100$ is appropriate.

In order to compare the role of inequality in ability I examine the following two alternative settings:

- all students have an ability between 40 and 60.
- all students have an ability between 0 and 100.

Notice that the for each respective comparison, the other variable is set to the unequal distribution. That is, while income distribution is varied, I use the ability distribution for students between 0 and 100 and while varying student ability, the income distribution is set to the unequal case.

The public schools receive funding proportional to enrollment, i.e. they receive a "public voucher" equal to \$6000 ($v_0 = 6$). Voucher levels for private schools vary from zero ($v_1 = 0$), to \$2,000 ($v_1 = 2$), to \$4,000 ($v_1 = 4$) per student in the case of a *universal voucher* and vary from zero ($v_1 = 0$), to \$4,000 ($v_1 = 4$), to \$8,000 ($v_1 = 8$) per student in the case of a *target voucher*. The parameters of the EWA model are $\rho = 0.01$, $\phi = 0.9$, $\lambda = 5$, $N_0 = 1$, and $\delta = 1$.[25]

For each scenario the simulation iterates through 100 runs, while each single run consists of 500 time steps. For the analysis below, I record the values of all reported variables from the final time step and then show the respective boxplots or show the averages of the those final values from all 100 runs. Thus, for every policy experiment I have 100 observations.

[25] Camerer and Ho (1999) provide experimental evidence for a similar calibration. In the section 3.6 the simulations are run under numerous variations of the choice of parameters for EWA.

3.4 Computational Results

3.4.1 Universal Vouchers- Varying Income Distribution

I begin describing the results for the case of universal vouchers, which are provided to all the students independently of their characteristics. As described above, I examine three voucher levels: a baseline case with no vouchers, to a voucher level of $2,000 and a voucher level of $4,000. Additionally, I run all simulations with the two previously mentioned distributions, i.e. one case with all families having a family income within the middle class (uniformly within $40,0001 to $60,001) and one case where all classes contain 20% of the families. I refer to the previous one, as the *equal distribution*, whereas the latter will be referred to as the *unequal distribution*. I report all simulation results using box plots which are created using 100 observations for each treatment.

The first finding I describe is the effect of higher competition on the behavior of private and public schools. A change in competition can enter the school district in two ways, either in the form of voucher or through a change in the income distribution. Both increase the number of students who are close to the threshold of being able to choose between public and private school and thereby decrease the number of students who default into public schools. The control variables of public schools and private schools are expenditure per student and tuition, respectively. Both are depicted in Figure 3.1. The left panel shows the public school reaction to the introduction of vouchers under the two different distributions. The first column displays the expenditure level for a voucher equal to zero (baseline setting). Initially, public schools' mean expenditure is at $2,900 for the equal distribution and $2,200 for the unequal distribution.[26] Once the vouchers are introduced (columns two and three) we can observe the public schools increasing their expenditure along with the increase in voucher level for both distributions. This reaction can be explained by the relative attractiveness of private schools which increases as the voucher increases, as they make private schools more affordable to a larger group of students. Public schools observe that they are exposed to a higher degree of competition as more of their students are suddenly able to afford private schools as an alternative. This leads to a sudden decline in their enrollment and thus causes them to increase their "attractiveness" in response by adjusting their educational expenditure upwards. This is the effect of increased competition, the so called *indirect effect* or *general equilibrium effect*. Comparing the two distributions we can see that the spending of public schools is higher in the equal distribution at all times. This can be explained by the fact, that with more students in the middle income class, the number of students that default into public school (and thus do not have choice) is lower. Thus, public

[26] The mean values of all graphs in this section and the following are available upon request.

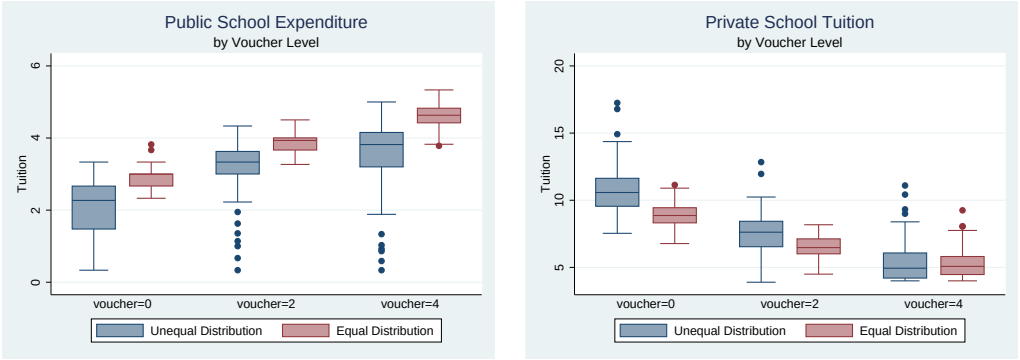

(a) (b)

FIGURE 3.1: Public School Expenditure(a) and Private School Tuition(b)

schools face a higher competitive pressure in the equal distribution and respond with higher educational quality. This effect is in line with the above cited empirical findings for schools facing higher competition.

The right panel in Figure 3.1 shows the values for the optimal private school tuition in all three voucher scenarios for both distributions. The first column displays the tuition values again for a voucher equal to zero, where tuition for the unequal setting is slightly above \$10,000 and slightly below \$9,000 for the equal distribution. Private schools are maximizing the enrollment in their sector by adjusting the tuition in order to cater to the needs of potential students. Private schools face a trade off with tuition as it enters student utility twice. It affects utility negatively as it lowers the family income (since the family has to pay the tuition) but it also increases utility as private schools use their entire tuition as expenditure which enters student utility positively. Thus private schools attempt to find a value for tuition (and thereby expenditure) in each scenario that maximizes enrollment. In order to find this optimal value, they employ the above mentioned learning algorithm. Now once the vouchers are introduced we can see that the tuition is lowered by a value which is very close to the voucher level as can be seen in the two columns to the right. The private schools now receive a voucher from each student of \$2,000 or \$4,000 and lower their tuition by this amount. Given that private school expenditure is the sum of all payments the schools receive (i.e. tuition plus voucher) this implies that the school expenditure remains fairly constant. The private schools have identified a value of school expenditure close to \$10,000 to be maximizing student utility and try to keep this value constant throughout all voucher levels. Notice also that initially in the case of no voucher, the tuition level is higher in the unequal distribution. The intuition behind this is that in the case of an unequal income distribution, the private schools have a higher portion of high income families whose preferences they are catering to. Since high income families have a higher willingness to pay for a better education this leads to an increased tuition.

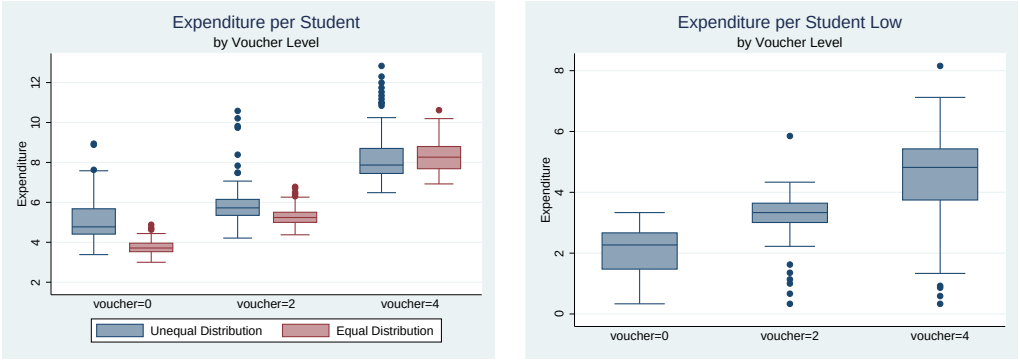

(a) (b)

FIGURE 3.2: Expenditure per student(a) and Expenditure per low income student(b)

Both results, the increased expenditure of public schools and the constant expenditure of private schools imply that the aggregate educational expenditure increases through the introduction of the vouchers. This overall increase in the expenses valued by students can be seen in both panels of Figure 3.2. In both distributions the expenditure per student increases with the voucher introduction, while the expenditure per student in the unequal distribution increases slower than the equal distribution. This can be explained by the fact that in the equal distribution there are more students close to the threshold of switching from public to private, i.e. as the voucher is introduced more students switch from public to private in the equal distribution thus rendering a higher number of students in private schools (which have higher educational expenditure).

This increase is also present with low income families, as depicted in the right panel of Figure 3.2, which is partly due to the increased expenditure e_0 by public schools (low income students are mostly enrolled in public schools) and partly due to the increased enrollment of students in private schools.[27] These findings support the argument that the introduction of the voucher system will increase the overall expenditure per student in the district due to more students being able to afford private schools and the public schools reacting to the increased pressure.

The next aspect of the introduction of a voucher system which deserves special attention is the distribution of students between public and private schools. Throughout all distributions we can observe that the fraction of students enrolling in public school decreases with the introduction of the voucher, see Figure 3.3, left panel, and consequently, increases in private schools. Observe also, that in the initial case without a

[27]Both effects can be separated by switching the objective function of the public schools to one which maximizes enrollment instead of surplus, then schools always set expenditure to the maximum possible amount of v_0 and thus one can measure the change in educational expenses solely attributed to an increase in private school enrollment.

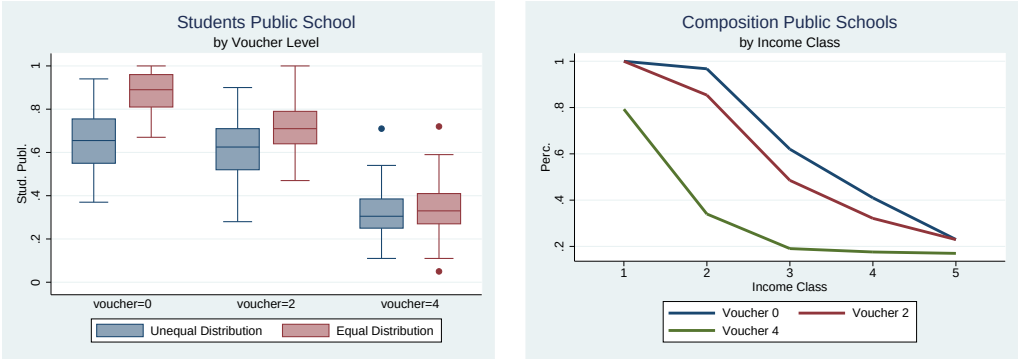

(a) (b)

FIGURE 3.3: Enrollment in Public Schools(a) and Composition of public schools(b)

voucher, there are more students enrolled in public school in the equal distribution, due to the higher quality of public schools in this setting. Since public schools have a higher expenditure per student in the equal distribution, more students find them attractive. As the voucher increases, the number of students enrolled in public school get closer for both distributions. However, low income students are mostly not among the students who are actually able to afford to switch to private schools and thus are usually staying in the public schools. In the right panel of Figure 3.3 we can see the enrollment in the public school sector as a function of income class for all three voucher levels. You can see that with rising income, the fraction of students enrolled in public school decreases, which does not change with the introduction of the voucher. You can also see that the largest difference between the voucher levels is for the income classes two and three. These are the classes closest to the income threshold where switching to private school is most likely, holding all other household characteristics constant. These two classes are the ones with the highest proportion of school changers and thus profit from the voucher the most in this regard. For the lowest income class the number of students exercising choice remains close to 20 %, i.e. only a small fraction of the low income students are actually exercising choice at the highest voucher level of $4000. We can thus conclude, that not all income classes do profit equally from the voucher, i.e. especially low income students are usually staying behind in the public schools and their benefits from a voucher program are closely correlated to the effect on public school expenditure. As shown above, public schools increase their expenditure which leads to an increase in the educational quality experienced by the students. The other aspect affecting the educational experience of students is the quality of peers, i.e. mean ability in the respective school.

Figure 3.4 depicts the mean ability in the public sector (left panel) vs. private sector (right panel) in the equal and the unequal distribution, respectively. In both school systems we can observe a decline in mean ability. This somewhat counterintiuitive

result can be explained if one looks at the type of students most likely to exercise choice in public schools, namely the high ability students. As they leave the public schools, mean ability falls, however they are on average at the lower end of the ability spectrum in private schools and thus diminish mean ability as they enter private schools. The decline for the public schools implies that the students staying behind in the public school are observing a decline in their peer group.

Now comparing the mean ability levels for both distributions we can see that the mean ability in the public schools is higher in the equal distribution before the voucher is introduced. Analogously to the higher educational quality in the equal distribution as explained above, this is the result of the public school sector being of higher quality in terms of educational expenditure. The higher expenditure per student attracts students which put higher weight on the educational expenditure, which are the ones with higher ability levels. However, due to the higher number of students switching schools in the equal distribution when the voucher is introduced, we can observe a stronger decline of ability in the public schools as it is mostly the high ability students who are exercising choice. This also explains how in the private school sector we can observe a stronger decline in mean ability in the equal distribution case. As more students switch from public to private, this lowers the mean ability since the high ability students from public schools are on the lower end of the ability spectrum in the private schools. We can thus conclude that there is an ambiguous effect on the educational attainment, $a_i = a(m_s, e_s)$, of students from the lower income groups. While the ones switching schools experience improved educational outcomes, the ones staying back are exposed to two effects. First, the educational expenses by public schools, e_0 increase and thus improve their educational experience. However, at the same time, as the peer group in public schools deteriorates, i.e. m_s declines, public school students have worse outcomes in this regard. This result is in line with findings from Manski (1992), Epple and Romano (1998), Epple and Romano (2008). As the universal voucher program discussed so far fails to increase both arguments of the educational attainment of students, it remains of interest whether an alternative program such as a target voucher system would be able to do so.

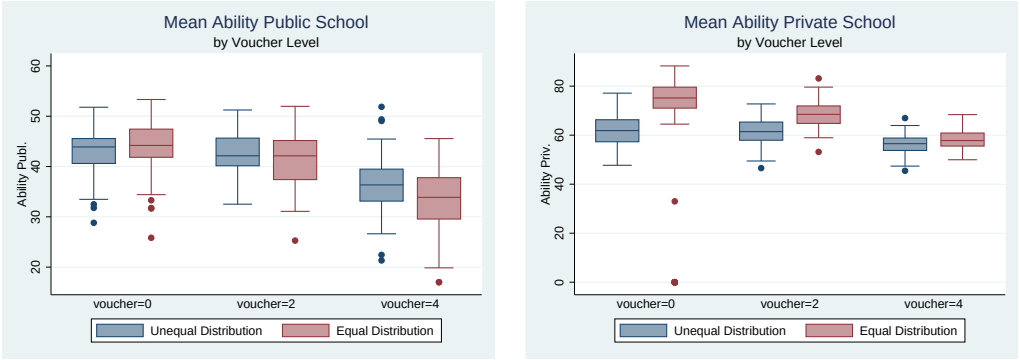

(a) (b)

FIGURE 3.4: Average Ability Public Schools(a) and Mean Ability Private Schools(b)

3.4.2 Universal Vouchers- Varying Income Distribution

In addition to the comparison between the equal and unequal income conducted in section 3.4.1 this section compares outcomes for a universal voucher for two different distribution of ability. Analogous to the two income distributions, in this section I use two distributions for ability, equal and unequal. In particular, for the equal distribution all students have an ability b which has a uniform distribution on its support, $S_b = [40, 60]$[28], whereas the unequal distribution ability has a uniform distribution on its support, $S_b = [0, 100]$. Family income y has a uniform distribution on its support $S_y = [0, 100]$ (where y is denoted in multiples of one thousand). The motivation behind this analysis is to see how far the problem of "*cream skimming*" is pervasive to this alternative setting where ability is not heterogeneous anymore. Walsh (2009) shows that "*cream skimming*" is very unlikely to be a concern when public schools exhibit very low levels of heterogeneity, i.e. the peer effects caused by vouchers would have to be unrealistically high to give any rise to concern of a deteriorating peer group.[29] Figure 3.3 in section 3.4.1 gives an indication of this lack of heterogeneity, as the student body within public schools consists to a large degree of low income families. Figure 3.5 shows how the results for the case of an equal distribution of ability are similar to the results in section 3.4.1 with regard toeh public school expenditure and private school tuition.

As in the case for the equal income distribution, equality of ability has a positive effect on the overall enrollment in public school as can be seen in Figure 3.7, due to the higher quality of public schools in this setting. Since public schools have a higher expenditure per student in the equal distribution, more students find them attractive. Also, similar to the previous comparison between equal and unequal income comparison, again the

[28]In order to check for robustness, I have also implemented a number of alternative 'equal distributions', including a case where all students have an ability $b = 50$, the results in all of these cases were similar to the one reported here

[29]I would like to thank an anonymous referee for encouraging this alternative specification

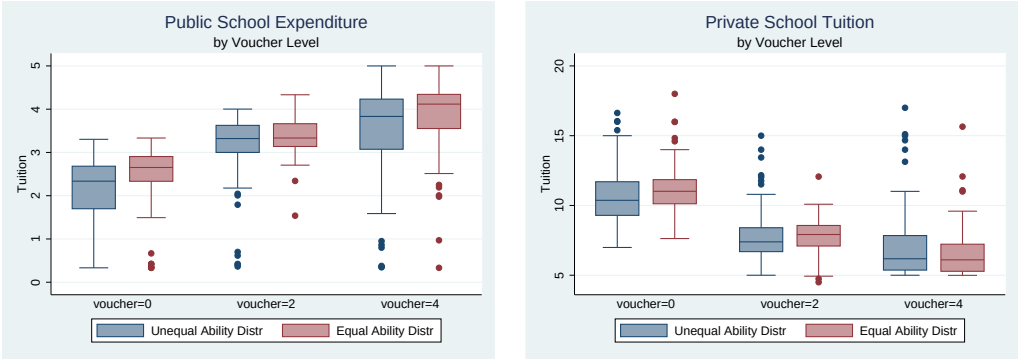

FIGURE 3.5: Public School Expenditure(a) and Private School Tuition(b)

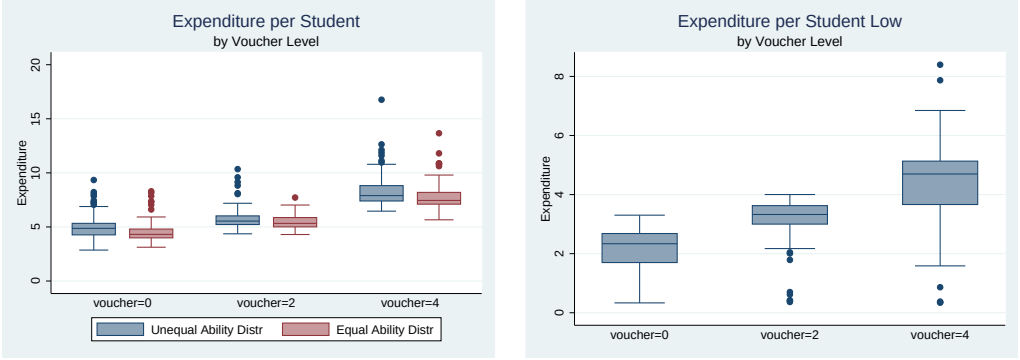

(a) (b)

FIGURE 3.6: Expenditure per student(a) and Expenditure per low income student(b)

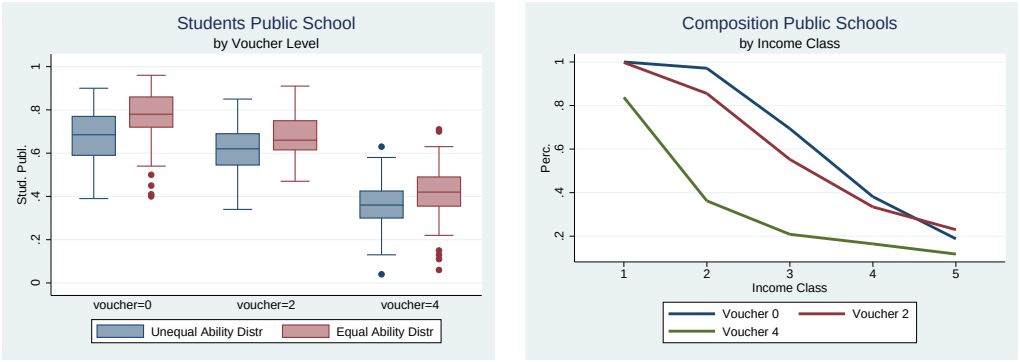

(a) (b)

FIGURE 3.7: Enrollment in Public Schools(a) and Composition of public schools(b)

low income students are mostly not among the students who are actually able to afford to switch to private schools and thus are usually staying in the public schools, as can be seen in the right panel of Figure 3.7.

With rising income, the fraction of students enrolled in public school decreases again, which does not change with the introduction of the voucher. The largest difference between the voucher levels is for the income classes two and three as those are the classes

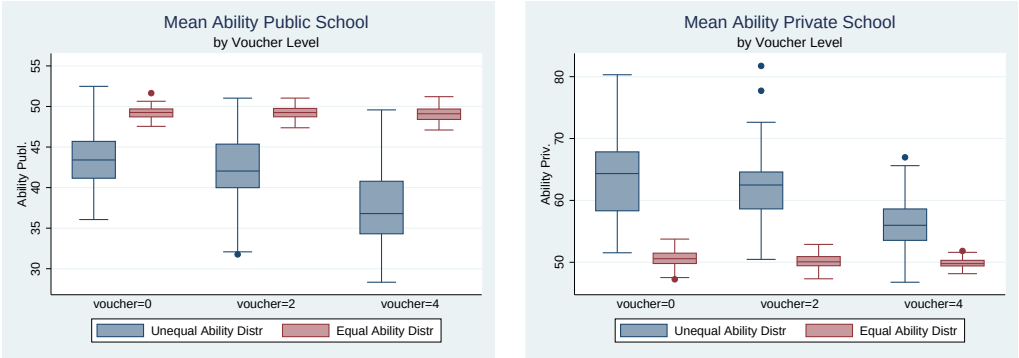

(a) (b)

FIGURE 3.8: Mean Ability Public Schools(a) and Mean Ability Private Schools(b)

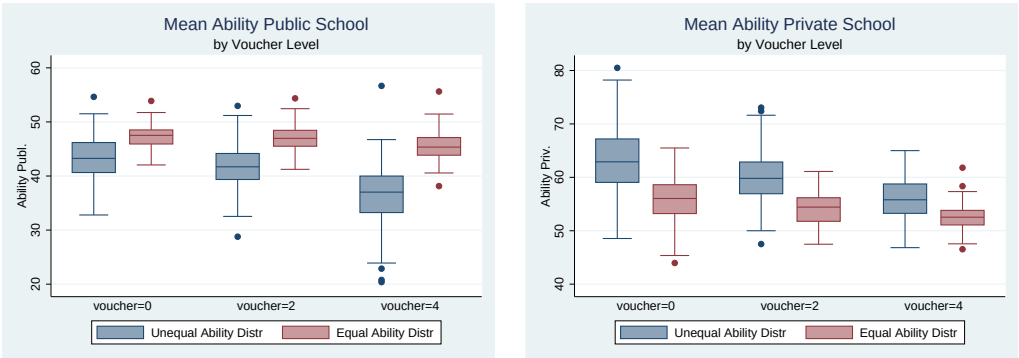

(a) (b)

FIGURE 3.9: Mean Ability Public(a) Schools and Mean Ability Private Schools(b)

closest to the income threshold where switching to private school is financially feasible, holding all other household characteristics constant. Thus, as before we can conclude, that not all income classes do profit equally from the voucher, i.e. especially low income students are still not able to afford the switch to private school. Meanwhile, due to the increased expenditure of public schools under the voucher system, the students who remain in the public schools experience only a beneficial effect in this setting. Unlike the comparison in section 3.4.1, here we do not have any decrease in Mean Ability in the public schools due to the homogeneous distribution of ability, i.e. mean ability remains almost constant through the voucher levels or only slightly decreases. Thus similar to Walsh (2009), the phenomenon of "*cream skimming*" is not observable for such a small degree of heterogeneity in ability.

In addition, Figure 3.9 plots the same results under a more heterogeneous distribution of ability, i.e. b is uniformly distribution on its support $S_b = [20, 80]$. Even here, the "*cream skimming*" effect appears to be more nuanced compared to section 3.4.1. Finally, the results for target vouchers under this homogeneous distribution of ability continue to hold and are very similar to section 3.5 and are not reported here.

3.5 Target Vouchers

In this section, I present the simulation results from a *target voucher* policy that allows the voucher to vary with student characteristics. The main target of such an alternative voucher program would be to avoid stratification or *cream skimming*, which was the result of the universal voucher system discussed above under the assumption of heterogeneous student ability.[30] For this purpose, I examine a *target voucher* system that allows the voucher to vary with household characteristics ability and income, i.e. $v = v(b, y)$. Target voucher systems have been discussed in simulation models by Nechyba (2000) and Caucutt (2002), who analyze the effect of targeting vouchers to income of the recipients. Epple and Romano (2008) alternatively examine the effects of a voucher that is a function of the ability of students in a general equilibrium model. Here, I examine both voucher alternatives, using a functional form that allows for the average voucher level to remain close to the universal voucher level. The functional form implies that for the student with either average family income or Mean Ability level the voucher level remains the same, while the students above or below the mean are affected.

The targeted vouchers I examine are of the following two functional forms for income and ability, respectively:

$$v_i(y_i) = v_0 + \frac{v_0}{100}(\bar{y} - y_i) \tag{3.12a}$$

$$v_i(b_i) = v_0 + \frac{v_0}{100}(\bar{b} - b_i) \tag{3.12b}$$

where v_0 can take three values, namely 0, \$2000, and \$4000. Notice that with the functional form of equations (3.12a) and (3.12b), the average voucher given to students remains the same and only the distribution and the maximum voucher level changes. Both voucher programs allow to subsidize students that are below the overall mean ability, $\bar{b}$ or below the overall mean income, $\bar{y}$, that is, these students receive a voucher higher than the average voucher v_0, with the students at the lower end of the spectrum receiving the highest possible voucher (e.g. for $y_i = 0$ or $b_i = 0$, we have $v_1 = 4$ or $v_1 = 8$). Meanwhile students above $\bar{b}$ or $\bar{y}$ receive a voucher gradually lower than v_0 with the students at the upper end of the spectrum receiving no voucher (e.g. for $y_i = 100$ or $b_i = 100$, $v_0 = 0$).[31] All of the policy experiments reported below are with an unequal distribution of income as used in the previous section, i.e. all classes contain 20% of families (Gini coefficient: 0.33).

[30] The finding presented in this section were also conducted under the assumption of heterogeneous distributions of ability and yielded similar results.

[31] This is contrary to the analysis in Fu (2014) who allows for vouchers to be positively correlated with ability.

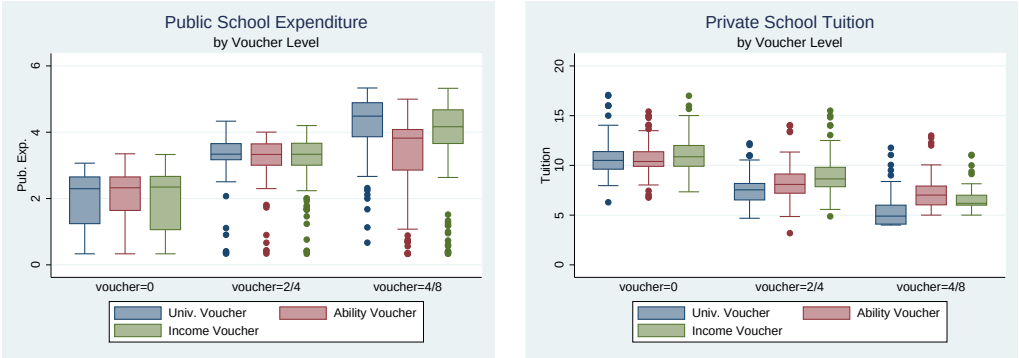

(a) (b)

FIGURE 3.10: Public School Expenditure(a) and Private School Tuition(b)

The first finding of this section is that the key results of the previous section for the variables private school tuition, public school expenditure, and overall expenditure reveal very similar patterns in the *target voucher* case compared to the *universal vouchers*. All graphs below contain the *universal voucher* from the previous section for ease of comparison. I briefly describe the outcomes for all three variables. Figure 3.10, left panel shows how public schools exhibit a similar behavior to the universal case, i.e. an increase in expenditure follows the introduction of the voucher. Private school tuition still falls in the two target voucher settings compared to the universal voucher by a value close to the voucher level, which in turn implies that again throughout all voucher levels, the expenditure in private schools remains constant, see right panel Figure 3.10.

As we look to the results in Figure 3.11, we can see that expenditure per student also develops similar to the *universal voucher* program, that is we can observe an increase in the expenditure per student by students. This also holds for the case of the expenditure per student for low income students. Notice that here, the effect is most pronounced in the case of $v_1 = 8$, as this specifically supports the low income students ability to switch schools. The enrollment patterns in both school systems also show a very similar development in the target and universal voucher case and are thus omitted here and left to the Appendix.

The second finding from this section is that one form of the target voucher, namely the *ability voucher* is actually able to reverse the negative peer effect, which we observed in the case of the universal voucher due to "*cream skimming*". Figure 3.12 plots the mean ability in public (left panel) and private schools (right panel). Mean ability falls in all voucher programs in the private school. For public schools we only observe a decrease in mean ability for the *universal vouchers* and the *income voucher*. The *ability voucher*, however, actually leads to an increase in the mean ability in public schools. This is caused by the incentive scheme that the *ability voucher* entails. The voucher amount

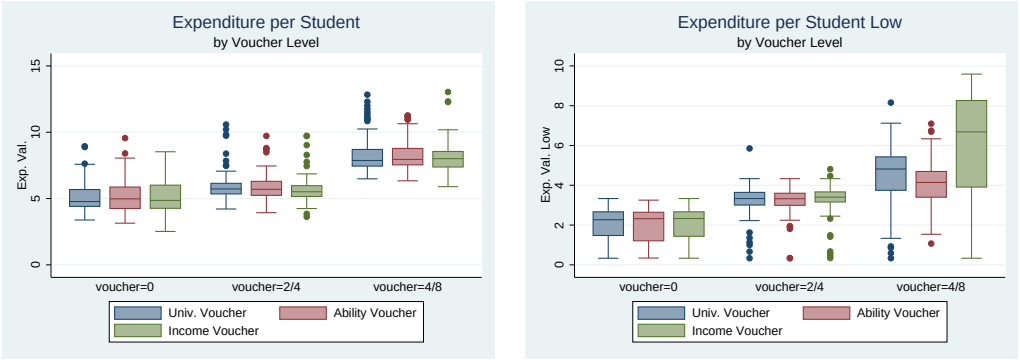

(a) (b)

FIGURE 3.11: Expenditure per student(a) and Expenditure per low income student(b)

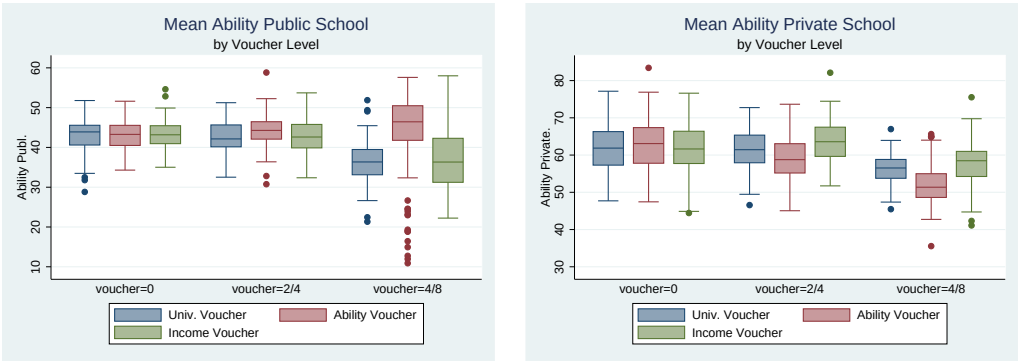

(a) (b)

FIGURE 3.12: Mean Ability Public Schools(a) and Mean Ability Private Schools(b)

is negatively correlated with the students' ability, i.e students with high ability obtain less voucher amounts than students with low ability. This reverses the incentives of high ability students who naturally put more weight on educational quality as is and thus are more likely to prefer private over public schooling. Unlike the universal voucher and the income voucher which do not consider the ability level, the *ability voucher* is able to circumvent the deterioration of the peer group (or to even slightly reverse it).

The target voucher thus allows to obtain the benefits of school competition without leading to detrimental outcomes for some students especially in the lower income groups. This finding suggests that a careful design of a voucher program would allow to reap the benefits of more competition (through an increase in productivity of public schools) while avoiding detrimental peer group effects.

3.6 Robustness Analysis

In order to ensure that the results obtained in the previous sections are robust, a series of robustness tests were run. Hereby, I focused on the main results, namely the increased public school expenditure and the peer group effect on public schools, i.e. whether the *ability voucher* system was indeed the only policy where the negative outcomes could be avoided. To this end, I first ran a number of combinations of parameter values for the EWA model. In particular, I systematically varied each of the parameters by small increments, ρ, ϕ, δ, λ to test whether the learning still took place in the same manner. None of the variations yielded substantial changes to the results. In a next step, I gradually changed the parameters of the utility functions of the agents, i.e. the boundaries of α, β, and γ were either increased or decreased by small amounts to adjust the optimal amount that a family would like to spend on tuition without a change in the key results. In a next step, instead of recording results after 1,000 iterations I increased the number of steps to 3000 in steps of 500. Then, both policy variables universal vouchers and target vouchers were varied. First, I increased the highest universal voucher level of \$4000 in increments of \$500 up to \$8000 without a change in key parameters. The same variations were run on the *target voucher*, where the value of v_0 was adjusted in the same vein. Another variation included the relaxation of distributional assumptions regarding income, by allowing for a Dagum Distribution (see (Dagum, 1975)) income which is positively skewed with mean larger then median. Finally, I relaxed the assumption that ability and income are independently distributed and instead allowed for various degree of correlation between those. In all of the performed variations the key results of a constant or increasing peer group in the public schools and an increased public school expenditure did not change.

3.7 Conclusion

I developed an agent-based model drawing on Manski (1992) and Epple and Romano (1998, 2008) and extended those models to study the distributional effects of a voucher program while incorporating the endogenous reaction of public and private schools to the voucher program. The model shows how the public school sector reacts with an increase in educational expenditure to the increased competition and thereby renders a positive educational outcome for students who remain in these schools in this regard. In particular, the paper shows that public schools increase their expenditure per student (public school rents fall) under a voucher system, adding to an overall increase in expenditure per student and thus leading to an increase in efficiency. However, students who remain in public school are exposed to a decline in peer group quality due to students with higher abilities or higher income are exercising choice with a higher probability. The model simulates this change in the student body, i.e. students exercising choice are either of high ability or from high income families. Thus, a *universal voucher* program aggravates the existing stratification within the school system.

In a further step, I have extended the model to allow for a targeted voucher program which allows to make the voucher a function of income or ability. Similar to Epple and Romano (2008) I examine the effects of these alternate voucher programs and find that only the ability targeted voucher is able to avoid the peer deterioration and even reverse it while still maintaining the increase in the educational expenditure of public schools. Thus, the *ability voucher* is able to introduce competition into educational system while avoiding the negative effects of of universal vouchers. That is, the ability targeted voucher is able to increase the educational expenditure of public schools while at the same time preventing a deterioration of peer quality (e.g. *cream skimming*). This comes, at a less stringent set of restrictions for public and private schools, i.e. allowing public schools to react to higher degrees of competition and while allowing private schools to set the tuition level above the voucher level. Thus, using such a voucher program would allow to avoid the negative effect of a voucher program on students from low income families and students with lower ability.

Another important caveat to these findings is that, in general when examining school performance and student outcomes, the expenditure of schools is only an imperfect proxy for the productivity of the school (or its ability to provide a valuable education to students). While the findings in this paper show that public school performance as measured through expenditure is improving, in the sense that the schools are exerting more "*monetary effort*", a number of empirical studies (see (Hoxby, 2003b)) find that increased competition leads to public school students' achievement rising significantly and rapidly while keeping spending in the schools unaffected, i.e. public schools increase

the "*non-monetary effort*". The model employed in this paper does not feature *non-monetary efforts*, i.e. the public schools in the model can only improve by increasing expenditure as this is the variable noticed by students, thus the effect of competition on public schools is observed via the *expenditure channel*. Yet, the findings in the model are still in line with these empirical observations to a certain degree as the public schools are using the funds already available to them in order to finance the increased funding, i.e. they are reducing their surplus and do not receive additional funds. In this sense the public schools are offering a better education to their students without requiring additional funding (or becoming more efficient by using the funds available to them). This result hinges on the assumption, that the public schools are indeed local monopolists and as such are able to accumulate surpluses while the private schools have zero economic profit due the competition they are exposed to. Given that a large fraction of students per school district default into public school and cannot choose their schools, this appears to be in line with what is observed empirically.

The effects of an introduction of a voucher based school system is still the subject of much debate but these results indicate that the educational outcomes of voucher programs are strongly influenced by the voucher design and a careful design could reduce the chances of further stratification.

3.8 Appendix

3.8.1 Calibration

Given that for private schools ($s = 1$) expenditure equals, $e_s = v_s + t_s$ the derivative of the student utility function with respect to tuition is:

$$\frac{\partial U_{is}}{\partial t_s} = \frac{\alpha_i}{t_s + v_{is}} - \frac{\gamma}{y_i - t_s + v_{is}} \tag{3.13}$$

For $v_{is} = 0$, this implies that in the optimum, households want to spend $\frac{\alpha}{\alpha + \gamma}$ of their income on education, or utility rises with tuition until tuition is optimal, namely (solving the equation above for t_1, after setting it equal to zero):

$$t_1 = \frac{\alpha}{\alpha + \gamma} y_i \tag{3.14}$$

For the given calibration this would mean that students at the upper end of the ability and income spectrum, i.e. $\alpha = 5$ and $\gamma = 25$, are willing to spend 1/6 of their family income in optimum.

3.8.2 Additional Simulation Results

Figure 3.13 plots a similarly decreasing enrollment in all voucher programs as the voucher level increases. The upper right panel and the lower panels display the switching pattern by income class as in the previous chapter. The upper right panel with the universal voucher is reproduced for comparison. The lower left panel shows the composition for *ability voucher* which displays a very similar development as the universal voucher, i.e. again the lowest class is merely exercising choice wile the classes two and three are the ones mostly switching schools. The lower right panel displays the composition for the income voucher, which unlike the two other programs display a very large number of low income students exercising choice.

Figure 3.14 plots the development of costs throughout the voucher programs. We can see that government costs per student increase from an average value of around \$ 4000 to \$ 4500. Overall cost of families remains stable while costs for low income families increase since they start to attend private schools in higher numbers now.

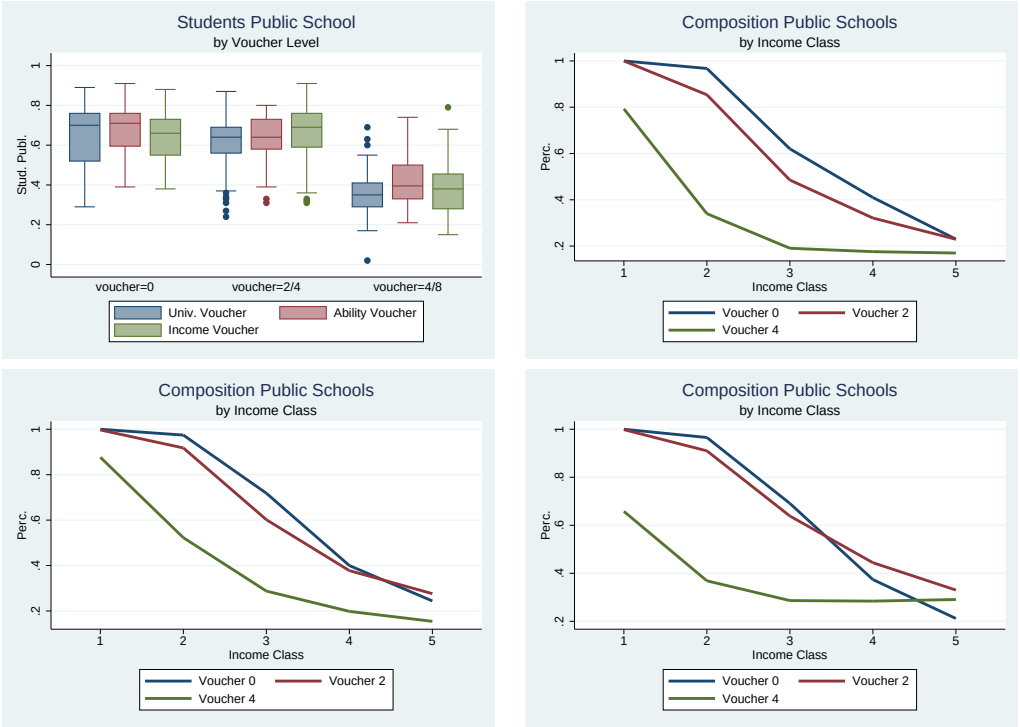

FIGURE 3.13: Public school enrollment and composition

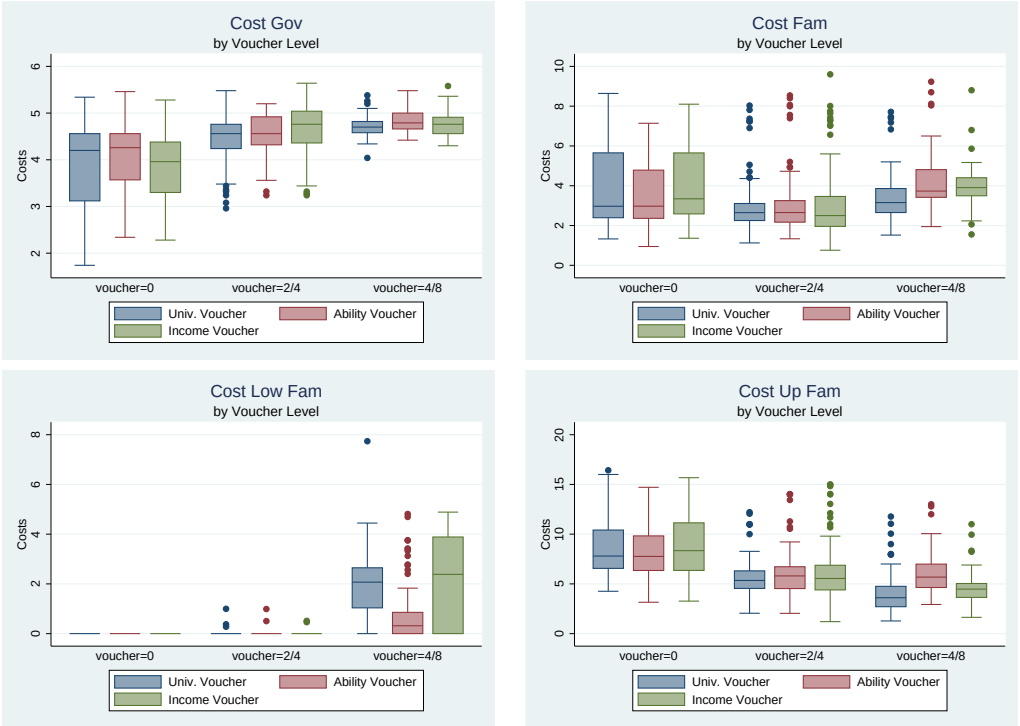

FIGURE 3.14: Costs by Voucher Level

Chapter 4

A Tradable Employment Quota[1]

4.1 Introduction

Affirmative action policies very often take the form of employment quotas. Norway introduced a quota in 2003 already. France, Iceland, and Spain have mandatory boardroom gender quotas forcing firms to have boards with 40% women by the years 2017, 2013, and 2015, respectively. German publicly listed firms subject to co-determination have to comply to a boardroom quota of 30% from 2016 onwards. Furthermore, the European Parliament passed a proposal by the European Commission to break the glass ceiling. According to this envisaged bill, European firms have to appoint female directors in order to make supervisory boards two-fifths female by 2020.[2]

As it is very likely that a uniform employment quota imposes larger adjustment costs on some firms than on others the question arises whether the goal of paving the way for more female employment can be achieved at lower costs. In particular, as the size of women's labor supply is heterogeneous across occupations, sectors, and regions, some of the firms forced to fulfill a fixed quota will find it more difficult to hire women who match the vacancies than others. A more flexible instrument is called for that does not compromise on the overall goal of achieving a certain share of female employment. Such an instrument should allow firms to fall short of the quota if costs of compliance would become unreasonable, while allowing others to gain from employing relatively more women.

In this article we propose and analyze the labor market effects of a tradable employment quota. Borrowing from the experience with environmental regulation policies to combat excessive emissions, we suggest to implement a mechanism that efficiently achieves a

[1] This Chapter is based on Akyol et al. (2015)

[2] See the European Commission Database on Women and Men in Decision Making.

fixed share of women working in relation to men. The idea is to issue permits to firms that give them the right to employ men and make these permits tradable. With such an affirmative action policy firms would only be allowed to employ men up to a number that matches the stock of permits that they hold. As a particular firm wants to employ an additional man it would only be able to do so by purchasing an additional employment right. Firms being in excess of permits because they find it more profitable to employ a woman than to hold a permit will want to sell this right. Trading of permits between those firms that want to buy and those firms willing to sell would yield a market price of a permit reflecting the profitability of employing an additional man. While the overall supply of permits of an issuing body would determine the share of female employment in the economy, single firms could adjust more flexibly and still comply.

We are aware of the fact that very often economists' ideas for resolving societal issues do not find widespread support outside of their own community. Sometimes even fierce opposition arises and we would not be surprised if such a reaction emerges as a response to our proposal. The public discussion of affirmative action policies is very much centered on equity considerations. It appears to us that efficiency or the loss of efficiency is of secondary importance, maybe because policymakers or those whom they represent are not willing to trade equity for efficiency. On this background we believe that it is important to stress that our proposal does not question the equity related aim of improving women's participation in the labor market. Personally, we also think that this should be an important goal for policymakers. What we suggest here is, however, a policy measure that has the potential to achieve equity at a lower cost for society. The way that this may be achieved is via a market-oriented instrument.

We expect that recurring to market mechanisms to resolve equity issues may become another reason why our proposal could be dismissed by the broader public upfront. In fact, it has been reported that initially there was heavy opposition to the introduction of tradable carbon dioxide emission rights coming from environmental groups. The hostility towards a marked-oriented instrument was mostly driven by moral or philosophical reasoning where it was argued that "It's Immoral to Buy the Right to Pollute" as the New York Times titled an article.[3] Similar objections may arise with respect to our proposal. Some critics may claim that it is immoral that firms can buy themselves out of the obligation to hire women by purchasing permits that allow them to employ men rather than women. We find it very difficult to resolve such kind of moral concerns. But we would like to point to a more recent discussion that has arisen as employment quotas got implemented. There is evidence that women feel stigmatized when their employer is subject to a quota (Heilman et al., 1992, 1997). The unease comes from the perception that fellow employees may think that a particular woman was only employed because of

[3]See Sandel (1997).

the quota policy rather than because of her qualification. Interestingly, in our proposal such a stigma is less likely to emerge because the mere fact that the employer chose to offer a contract to a woman signals appreciation. Alternatively, the employer could have employed a man by not selling a permit or buying an additional permit (which he did not do.)

In the context of environmental policies the use of tradable permits is based on Ronald Coase's idea that market participants may correct for negative externalities without the intervention of governments if legal rules of entitlement exist and transaction costs are negligible (Coase, 1960). Negative externalities which are arising from production through the emission of health and climate damaging gases are corrected for by giving firms the right to pollute (or those affected by the emission the right for unpolluted air). Consequently, emissions are only allowed if costly permits are held so that firms are confronted with the socially relevant marginal costs of production rather than their private marginal costs only. It follows that production is extended up to the socially desirable level only. Moreover, the abatement of (environmentally) unhealthy substances takes place at the lowest costs possible. Those firms with relatively low marginal abatement costs will choose to invest in clean technologies and sell their permits, whereas the firms with relatively high marginal abatement costs will want to expand production by purchasing permits. It is the cost reducing feature of permit trading which we borrow for developing our proposal of an affirmative action policy that increases female employment shares at relatively lower costs than one would have with a uniform quota. As in the case of environmental policies, we expect that those firms which will find it inherently difficult to hire women will rather purchase permits than leave vacancies unproductive or costly retrain their newly hired employees, whereas those firms able to hire women will sell their permits. All these cost saving decisions of firms leading to trades on the permit market, however, should not jeopardize the goal of achieving an overall female employment rate set by policymakers and implemented by issuing a corresponding number of tradable employment permits.[4]

We would like to give some empirical evidence on the current situation with respect to female participation in advisory and executive boards of mostly publicly listed firms. Restricting to these figures for illustrative purposes may be justified by the recent policy moves that started to regulate this particular part of the labor market. With 16% the

[4]We have been asked at various occasions whether our case of a discriminatory labor market entails some form of externality, which is then resolved through our proposed permit solution. In the sense that an externality is present whenever some economic agent's welfare is directly affected by the action of another agent (see, e.g., Hindriks and Myles, 2013) a discriminating firm does not constitute an externality. Choices of discriminating firms have effects on other agents' payoffs but they are mediated by prices, i.e. they are indirect or so-called pecuniary externalities. Thus, it is not an externatility that we correct for with tradable employment rights, but we rather target the heterogeneous adjustment costs of firms that can be handled more efficiently through our proposal.

U.S. and the 27 (by 2007) EU member states fare equally in terms of female representation on firm boards. Japan, as another major industrialized country, has only 1% women on firm boards. A closer look into single European countries reveals a large dispersion of female representation. In the three countries with the largest representation almost every third member is female, a share which, however, still falls short of Norway where a quota was introduced in 2003 already, forcing firms to comply by 2008. In the European countries that do worst less than one tenth of the positions are held by women. Interestingly, the countries doing relatively well in terms of female board membership hardly have women leading the board or being a CEO. Data for Germany allows for a closer look into the within country distribution of female representation. Again, we find a large variance between firms. Among the companies listed in the DAX (the major German stock market index), seven women served on the board of Henkel (a company producing personal care products) which was composed of 16 members at the time of data collection, while no woman was serving on the board of Fresenius (a medical care company).[5]

The introduction of a female board quota in Norway constituted a natural experiment that allowed for an analysis of firm reactions and their consequences more closely. At the time the law was introduced only 9% of women were on the boards of Norwegian firms. A legislated quota of 40% imposed a major change on the composition of Norwegian firm boards. Ahern and Dittmar (2012) use the pre-quota female representation across firms as an instrument for the changes of boards that followed the quota. For the days around the announcement of the law they find that stock returns fell by 3.52% for those firms with no female representation compared to firms that had at least one woman on the board. For the longer term, they estimate a decline in Tobin's Q of 12.4% as a response to a 10% forced increase in women representation on the boards. Overall they conclude that the imposed constraint had a large negative impact on firm value driven by the reorganizations of the boards. Drawing on the same policy change, Bohren and Staubo (2014) find that half of the firms that would have been affected by the gender quota chose to exit into another organizational form, thus avoiding exposition to the law. Also this piece of evidence suggests, at least indirectly, substantial costs of compliance that possibly could be diminished with a system of tradable permits to employ men.

In order to study whether a tradable employment quota is a feasible affirmative action policy and what labor market effects would possibly unfold we build a model that hosts a labor market and a market for permits. We study a labor market with a set of firms being allocated to multiple sectors. A fraction of the firms is characterized by taste discrimination against women as proposed by Becker (1957). The remaining fraction of

[5] A table summarizing these figures and some additional information can be found in the Appendix (see Table 4.7).

firms is indifferent between hiring men and women. Workers have sector specific skills meaning that their productivity does not fully unfold unless they work in a sector that requires precisely those skills. Female labor supply to sectors varies with the share of discriminating firms. In the words of Bergmann (1974) we assume "overcrowding" as firms systematically exclude women due to their distaste which ultimately leads to fewer women having invested in skills demanded by those firms. Firms post vacancies to which workers apply and are matched. In this framework taste discrimination leads to worse employment chances for women and lower wages. In order to balance this resulting gender gap we compare the labor market effects of two policies by studying the effects of a tradable as opposed to a non-tradable employment quota on welfare and various other labor market indicators.

Our main finding is that a tradable employment quota fares better in a heterogeneous labor market where firms are facing differentiated levels of female labor supply. The intuition why the tradable quota is superior is that given the heterogeneous labor supply of women, the degree to which firms are affected by a non-tradable quota varies, i.e. while firms residing in sectors with many women find it very easy to comply, the ones in sectors with less women find it more difficult. The tradable quota is a remedy for exactly this difference. It allows those firms facing high costs that would arise from a uniform quota to evade these costs by buying permits from those companies having an excess supply of permits because of the relatively higher labor supply of women to them. In a simulated version of our model we can disentangle the welfare effect of a tradable employment quota into changes of payoffs to firms and the wage sum and, furthermore, show how wages and payoffs change in discriminating and non-discriminating firms. This analysis shows that the advantage of the quota is mainly due to the fact that the tradable quota is a more flexible instrument and therefore better able to deal with an unequal female labour supply and discriminatory firms present in the market, without falling short of the equity goal of increasing female employment.

We analyze a simplified version of the model analytically and transpose a richer version of our framework into an agent-based simulation model to analyze whether a permit solution may actually work and what labor market effects might potentially emerge. Evaluating labor market policies using agent-based models has been suggested by Freeman (1998) already some time ago. Generally speaking an agent-based approach suits well for analyzing problems characterized by interacting heterogeneous agents. Moreover, agent-based modeling allows for a relatively detailed implementation of institutional arrangements. As we build a model with a sectoral structure hosting workers of different skill types to be employed by firms that may discriminate against women and, furthermore, will augment the labor market with a permit trading system, the agent-based approach seems to suit well for our purposes. Our approach may also be

subsumed under what Roth (2002, p.1341) called "design economics" where he argues that computational techniques should be seen as complementary to other tools applied to studying and designing markets, namely game-theory. One of the earliest attempts to analyze the effects of labor market institutions in an agent-based model can be found in Bergmann (1990). Others followed, with Tesfatsion (2001) working on wage setting or Neugart (2008) looking into training policies. Those and other contributions are surveyed in Neugart and Richiardi (2015).

Building our model we touch upon various strands of the literature. Women's wages, employment opportunities or occupations may be affected by labor market discrimination in various ways. Becker (1957) suggested that men and women being treated differently can be explained by taste discriminating firms. A second, alternative explanation, is statistical discrimination (Arrow, 1973; Phelps, 1972) which postulates that employers discriminate on the basis of expected differences in average productivity between men and women. Empirical work long evolved separately along these two explanations (Guryan and Charles, 2013). Only in recent years efforts have been made to test taste based explanations of discrimination against explanations pointing at statistical discrimination (See, e.g., Altonji and Pierret, 2001; Knowles et al., 2001). The jury still seems to be out and we do not make an attempt to resolve the issue here. Finally, a third explanation, typically referred to as "occupational crowding" argues that women are systematically excluded from "male jobs" (Bergmann, 1974). The origins of overcrowding may lie in a social climate where young women are told that some jobs are not suitable for "girls" or discrimination might already take place at the stage of human capital acquisition (Mechtenberg, 2009). But it may also be the results of a feedback from the discriminatory behavior of firms that offer worse pay and job conditions (Blau and Kahn, 2000, p. 82) or the expected lower likelihood to get a top position (Stark and Hyll, 2014) leading to a deliberate decision by women not to invest into skills demanded by those discriminating firms. Drafting our model to analyze policies we heavily draw on the idea of taste discriminating firms and the feedback arising with respect to labor supply decisions by women.

Welch (1976) was probably among the first dealing with affirmative action policies from a theoretical point of view. One concern was to elaborate the consequences of a quota arising from skill bumping, i.e. the upgrading of skills of workers so that firms would comply to the affirmative action policy. Equal employment opportunities as one proliferation of affirmative action policies have been studied in the framework of a search and matching model of the labor market by Kaas and Lu (2010). They find that if an imperfectly monitored equal employment opportunity legislation is combined with an

equal pay obligation inequality increases.[6] Contrary to them we do not look into a fixed quota but a tradable one.

Tradable permits solutions have gained widespread attention in the area of environmental policies (see, e.g., Stavins, 1998), but were also proposed as alternative policy instruments in the context of debt policies in the European Union (Casella, 1999), immigration (Moraga and Rapoport, 2014), birth control (De la Croix and Gosseries, 2009), or noise control at airports (Bréchet and Picard, 2010). With most of these proposals as with our application welfare consequences remain a theoretical discussion because the policy instruments were hardly implemented. The exceptions are, of course, the permit markets for carbon dioxide and sulfur dioxide. According to Schmalensee and Stavins (2013) these policy experiments mostly achieved the cost reducing goals. It is claimed that cost savings were at least 15 percent and perhaps even up to 90 percent compared with command-and-control approaches. While tradable permit solutions have been applied to other policy areas, there is, however, little to no work in the area of labor market policies.[7] This contribution may be seen as a first step into the analysis of such kind of affirmative action policies.

In the following Section 4.2 we lay out our model and introduce the reader to the two affirmative action policies we are going to compare. As a benchmark for the simulation exercise we derive an analytical solution of a simplified version of our model with and without policies in Section 4.3. In Section 4.4 we introduce the reader to the parametrization and the simulation set-up. In Section 4.5 we present our results on the effects of a tradable employment quota in comparison with a non-tradable quota on welfare and other labor market indicators, and also report on various robustness tests. In the last section we conclude, discuss our proposal in relation to other policies, and point toward possible extensions.

4.2 The Model

4.2.1 A General Description

Our model consists of a labor market with heterogeneous firms, male and female workers, and a permit market. The labor market has a sectoral structure where the sectors hosting

[6] Other contributions looking into the labor market consequences of equal pay legislation are Bowlus and Eckstein (2002); Coate and Loury (1993); Kaas (2009). Further search and matching models with taste discrimination but without policy analyses can be found in Black (1995); Lang et al. (2005); Rosen (2003).

[7] Winker (2000) sketches the idea of how collective wage agreements could be made more flexible using tradable permits.

the firms are allocated on a Salop circle (Salop, 1979), and the share of female workers differs in sectors.[8]

Firms sitting in a particular sector have distinct skill needs. Workers are equipped with different sector specific skills. A worker's productivity unfolds fully if she is employed in a sector that matches her skills. Workers employed by firms in other sectors lose part of their productivity. Labor demand for each firm is fixed. Vacancies are posted and workers apply.

A fixed share of firms discriminates against women. Due to taste discrimination these firms will only employ women instead of men if wages of the former are sufficiently lower (or productivity is sufficiently higher). As a consequence of the taste discriminating firms and in accordance with theories of discrimination related to "occupational crowding" outlined earlier on, labor supply of women is modeled as being negatively correlated with the share of discriminating firms in a sector.

Firms send non-binding wage offers to the applicants with highest profitability. Workers know that the offer is not binding as the firm has to comply with an affirmative action policy once all workers decided on which offer to accept (if any).

Specifically, we consider two affirmative action policies. Under a non-tradable employment quota every firm is allowed to only employ a share of men that does not exceed the quota. Alternatively, we implement a tradable quota issuing permits that give a firm the right to employ men.

Once a firm knows about the gender composition of applicants who accepted its offers, it is able to determine the actual number of men it can employ (in relation to its female workforce) to comply with a fixed quota, or in the case of a tradable quota, how many permits it needs to purchase. Consequently, in both scenarios the affirmative action policies magnify the problem of coordination failure giving rise to labor market frictions (See, e.g. Petrongolo and Pissarides, 2001, sec. 3.2). Due to the policies, men may have to be turned down and vacancies stay empty.

The permit market is modeled with a central market maker to whom each individual firm submits its individual supply and demand schedule of permits. The central market maker aggregates these bids and asks and determines the market clearing price at which the permits are reallocated.

[8]The Salop model has been used as a framework for studying the effects of discrimination and labor market policies with imperfect competition by others before. Examples are Bhaskar et al. (2002), Bhaskar and To (2003), Kaas (2009), or Berson (2014). Contrary to the existing work we add tradable permits to a model of imperfect competition with discriminating firms whereas many of our other modeling assumptions are in line with respect to this strand of the literature.

4.2.2 Labor Market Environment

We consider a partial labor market where firms' fixed labor demand is derived from aggregate product demand. Firms reside in $s \in S$ sectors. Workers are equipped with skills for a specific sector k but may be employed at a firm in a sector other than k. A worker i employed at a firm in a sector requiring her specific skills unfolds full productivity A. Worker specific productivity $P_i(k, s)$ declines as she is employed in a more distant sector. Sectors are allocated on a Salop circle, i.e. a worker's productivity with specific skills in sector k, working in sector s follows

$$P_i(k, s) = \begin{cases} A - a \cdot |k - s| & \text{if } |k - s| \leq S/2 \\ A - a \cdot (S - |k - s|) & \text{else,} \end{cases} \tag{4.1}$$

with $0 < a < 1$, $k = 1, ..., K$, and $s = 1, ..., S$. As there are as many skill types as sectors we have $K = S$.

4.2.3 Workers

The fixed total labor supply of size I can be decomposed into a fraction σ of female workers and a fraction $(1 - \sigma)$ of male workers . The share of female labor supply may differ between sectors according to σ_k. A worker i has an individual reservation wage w_i^r that is drawn from a uniform distribution with support $[0, A]$ resulting in an upward sloping labor market supply curve.

Workers send out a fixed number m (with $m > 1$) of applications preferably to firms which value their sector specific skills. Unobservable characteristics orthogonal to workers' skill endowment make some sectors more attractive to a specific worker than others. Thus, she may send an application to a firm in a sector which does not value her specific skills most. More formally, we recur to a discrete choice specification postulating that worker i with specific skills for sector k sends out an application to a firm in sector $\tilde{s}$ with probability

$$Prob_i(k, \tilde{s}) = \frac{e^{\lambda P_i(k,\tilde{s})}}{\sum_s e^{\lambda P_i(k,s)}}, \tag{4.2}$$

where $\lambda \geq 0$ drives the intensity of choice and the denominator sums up the exponential of worker specific productivities in all sectors.

Workers accept job offers with attached wage offers above their reservation wage w_i^r. If a worker receives more than one job offer, she chooses the job offer with the highest wage.

4.2.4 Firms

There is a total of L firms which we denote with $l = 1, ...L$. A firm l has J_l vacancies to fill. Each sector is populated with N firms. Overall, there is a share of $\mu > 0$ of discriminating firms. The share of discriminating firms may differ across sectors. We denote the share of discriminating firms in a sector s with $\mu_s > 0$. A discriminating firm has a dis-utility from hiring a female worker which is modeled with a discrimination coefficient d as suggested by Becker (1957). The firm l residing in sector s has payoffs calculated as the sum of the productivity net of the wage and the discrimination cost over all workers $\tilde{J}_l$ it finally employs:

$$\pi_{l,s} = \sum_{\tilde{J}_l} (P_i(k,s) - d_{g,l} - w_{l,k,g,s}), \tag{4.3}$$

where $g = M, F$ is the gender of the worker. For discrimination costs we have $d_M = d_F = 0$ for a non-discriminating firm and $d_F > 0$ for a firm that discriminates against an employed woman.

Firms set male and female wages to maximize payoffs. Job offers conditional on the firm having to comply with the affirmative action policy are sent out including worker specific wage offers.

A firm finding a wage that maximizes payoffs faces the following trade-off: higher wage offers increase the likelihood that a vacancy can be filled and becomes productive. Higher wage offers, however, also increase the wage bill and depress payoffs. Firms learn given their past experience on payoffs and wage offers how to best place themselves on this trade-off. To this end, each firm runs regressions of payoffs per job offer on the wage offers of the past τ iterations. For a positively estimated slope coefficient $\hat{\beta}$, a firm l that resides in sector s will adjust the wage offer w^o for a worker coming from sector k and of gender g upwards by $\epsilon > 0$ in iteration t with respect to the previous iteration $t-1$. For a negatively estimated slope coefficient the wage is adjusted downwards. Formally, we have

$$w^o_{l,k,g,s,t} = \begin{cases} w^o_{l,k,g,s,t-1} + \epsilon & \text{if } \hat{\beta}_{l,k,g,s,t} > 0, \\ w^o_{l,k,g,s,t-1} - \epsilon & \text{if } \hat{\beta}_{l,k,g,s,t} < 0, \\ w^o_{l,k,g,s,t-1} & \text{else.} \end{cases} \tag{4.4}$$

Wage offers are adjusted if they are within bounds $[0, A]$.

4.2.5 Policies

We consider an affirmative action policy where every single firm has to employ at least $\bar{\sigma}\tilde{J}_l$ women where $\bar{\sigma}$ is the quota and $\tilde{J}_l$ is the number of workers employed at a firm l. In order to comply with the quota firms rank applicants by gender and payoffs. Again, the firm will send offers – which are declared conditional on the firm being able to comply with the quota – to the best workers taking into account the quota.

Alternatively, we simulate a market with permits for employing men. In this case, a firm is only allowed to employ as many men as it holds permits for employing male workers. There is a fixed number for permits C for the whole economy. Initially every firm gets an equal share of the total number of permits. These permits can be sold and bought at a central clearing agency. For the clearing agency, we consider a central market maker who collects ask and bid prices, determines the market clearing price in every period and reallocates the permits between the buying and selling firms.

If a firm owns permits, its offer curve for selling permits is constructed as follows. All unused permits are offered at reservation price zero. The offer of that single firm increases by one more permit at a price equal to the payoff of the least profitable male worker. The second least profitable worker determines the price of yet an additional permit. As we move to even more profitable workers the full schedule of the offer curve for that particular firm is derived.

Turning to the demand side the central market maker looks into a single firm that will ask for as many permits as there are more men employed than the firm holds permits currently. Permits are used to employ the most profitable male workers in the firm. The bid price for the first additional permit is the payoff of the most profitable male worker for whom the firm does not yet have a permit. The bid price for the second additional permit is the payoff of the second most profitable male worker for whom the firm does not yet have a permit, and so on.

4.3 Analytical Solution

We solve a simplified version of the model analytically in order to pin down some of the mechanisms which will underlie our simulation results that are going to be presented in the later sections. To this end, we set the number of firms hosted by a single sector to $N = 1$ and let workers only apply to the firm of the sector to which their skill endowment fully matches. This implies setting $m = 1$, $Prob(k, \tilde{s}) = 1$, and $P(k, s) = A$. Essentially, we eliminate the part of labor market frictions arising through coordination failure in

TABLE 4.1: A firm's profits without policies

w_l	$\phi(w_l)$	π_l	$\Delta\pi_l$
0	1	$2\phi(w_l)(A - w_l) = 2A$	
$\frac{A}{J/2}$	2	$4A - \frac{8A}{J}$	$2A - \frac{8A}{J} = 2A - \frac{1}{J}8A(\phi(w_l) - 1)$
$2\frac{A}{J/2}$	3	$6A - \frac{24A}{J}$	$2A - \frac{16A}{J} = 2A - \frac{1}{J}8A(\phi(w_l) - 1)$
...	...	...	...
$(J/2 - 3)\frac{A}{J/2}$	$J/2 - 2$	$6A - \frac{24A}{J}$	
$(J/2 - 2)\frac{A}{J/2}$	$J/2 - 1$	$4A - \frac{8A}{J}$	$-2A + \frac{16A}{J} = -2A + \frac{1}{J}8A(\phi(w_l) - 1)$
$(J/2 - 1)\frac{A}{J/2}$	$J/2$	$2A$	$-2A + \frac{8A}{J} = -2A + \frac{1}{J}8A(\phi(w_l) - 1)$

the process of applying and hiring of workers. For the moment we will also assume an even number of sectors S and jobs J at each firm, and an equal number of male and female workers $\sigma = 0.5$ in each sector summing up to J workers. Furthermore, for each gender reservation wages are distributed according to $[0, \frac{A}{J/2}, 2\frac{A}{J/2}, ..., (J/2-1)\frac{A}{J/2}]$. We denote with $\phi(w)$ male and female labor supply in a sector, respectively. Finally, we let there be no taste discriminating firms in the market ($\mu = 0$).

4.3.1 No Employment Quota

Labor supply for a firm l is upward sloping due to the distribution of reservation wages. In this monopsonistic setting a firm l chooses as many workers, and posts wages accordingly, to maximize profits. Let w_l be the wage paid by a firm l, then we may write for the firm's profits

$$\pi_l = 2\phi(w_l)(A - w_l), \tag{4.5}$$

i.e. a worker's productivity A less the wage paid times $\phi(w_l)$ male and $\phi(w_l)$ female workers.

Table 4.1 lists profits of a firm l (third column) as a function of the wage offer w_l and the corresponding labor supply $\phi(w_l)$. The firm will offer higher wages and expand production as long as profits are increasing in wages ($\Delta\pi_l > 0$). This gives optimal number of workers $\phi_l^* \equiv \phi(w_l^*)$ for firm l as a function of overall supply J, and optimal wages w_l^* picked from the distribution of reservation wages (such that the firm gets the optimal number of workers ϕ_l^*). Total output Y in this economy becomes

$$Y = S \cdot 2 \cdot \phi_l^* \cdot A. \tag{4.6}$$

4.3.2 Employment Quota

What decisions will firms make if they are confronted with an employment quota that forces them to have a workforce where every second worker is female? A firm l in sector s still wants to have $2\phi(w_l^*)$ vacancies filled and offers w_l^*. It faces a labor supply of $\phi(w_l^*)$ women and $\phi(w_l^*)$ men and complies to the quota. Overall output in this economy is

$$Y^{ntq,eqFemDis} = S \cdot (\phi_l^* + \phi_l^*) \cdot A = S \cdot 2 \cdot \phi_l^* \cdot A. \tag{4.7}$$

Now, let us assume that women are not equally distributed over sectors S. Rather we inquire a labor market in which all women are allocated to sectors $s = 1, ..., S/2$ and all men to sectors $s = S/2+1, .., S$. Although being subject to an employment quota, firms in "female" sectors will be able to fill all of their $2\phi_l^*$ profit maximizing jobs. Firms in "male" sectors will produce nothing due to the constraint imposed by the employment quota that half of the workforce has to be female. Output with an employment quota becomes

$$Y^{ntq,UneqFemDis} = S/2 \cdot 2\phi_l^* \cdot A + 0 = S \cdot \phi_l^* \cdot A < Y^{ntq,eqFemDis}. \tag{4.8}$$

4.3.3 A Tradable Employment Quota

Alternatively to the non-tradable employment quota permits are issued that allow the firms to employ men now. We keep the assumption of an unequal distribution of women across sectors. In order to make a labor market with permits comparable to the one discussed without policies we issue $C = S \cdot \phi_l^*$ permits which is the number of male workers that all firms employed when maximizing profits. Initially, permits are equally distributed among firms. Thus, each firm is endowed with ϕ_l^* permits.

Firms in sectors $s = 1, ..., S/2$ that face a female labor supply only can run all jobs $2\phi_l^*$ without using a single permit. For any permit price $p \geq 0$ they will be willing to sell their permits. Thus, overall supply of permits is $S/2 \cdot \phi_l^*$.

Firms in sectors $s = S/2 + 1, .., S$, facing a male labor supply only, have to purchase permits as they want to employ more men than they initially were allocated permits. The profit function of a firm l in a "male" sector becomes

$$\pi_l = \phi(w_l^p)(A - w_l^p) - (\phi(w_l^p) - \phi(w_l^*))p \quad if \quad \phi(w_l^p) - \phi(w_l^*) > 0 \tag{4.9}$$

where w_l^p is the wage paid by a firm l in a labor market with permits. The first part of the profit function constitutes the revenues of workers net of wage costs. The second

part gives the costs of the permits that have to be bought in addition to the initially allocated number of permits.

Table 4.2 lists profits as a function of the permit price having substituted in the wage that has to be paid by a firm to draw on a certain labor supply. The table is drawn starting off with the wage w_l^* as the firm will never want to employ more workers than in the case where it did not face a permit requirement even if permits would be available for free.

As two men have equal reservation wages, a firm offering a marginally higher wage w_l^p will face a labor supply of two additional men. It may want to employ both of them or only one of them. By comparing profits it can easily be shown that they will always choose to employ two additional men as long as the product A net of the wage for the marginal worker is larger than the permit price.

Comparing profits of the firm for a labor demand of $\phi(w_l^p) = 2\phi(w_l^*)$ rather than $\phi(w_l^p) = 2\phi(w_l^*) - 2$ yields an upper permit price for the firm to be willing to employ $2\phi(w_l^*)$ of

$$A - (w_l^* - \frac{A}{J/2}) > p + \frac{A}{J/2}\phi(w_l^*). \tag{4.10}$$

For the firm to employ $2\phi(w_l^*)$ men the additional net product, i.e. the worker's productivity minus the wage to be paid has to cover the permit price plus the increase in the wage bill arising from the fact that all currently employed have to also receive the higher wage. As, moreover, the marginal profits are positive and increasing with lower demands for labor, the firm will demand $2\phi(w_l^*)$ workers if permit prices fulfill

$$\overline{p} = A - (w_l^* - \frac{A}{J/2}) - \frac{A}{J/2}\phi(w_l^*) > p. \tag{4.11}$$

Thus, for prices $0 \leq p < \overline{p}$ an exchange of ϕ_l^* permits between firms in "female" and "male" sectors is mutually advantageous.[9] Firms in the "male" sectors will be able to purchase additionally needed permits to employ $2\phi_l^*$ men in total. With a permit market output becomes

$$Y^{tq,UneqFemDis} = S/2 \cdot 2\phi_l^* \cdot A + S/2 \cdot 2\phi_l^* \cdot A = S \cdot 2\phi_l^* \cdot A \tag{4.12}$$

[9] $\overline{p} > 0$ as it is equal to the marginal product which the firm faces as it chooses the optimal wage without a policy constraint.

TABLE 4.2: A firm's profits facing male labor supply in a market with permits

Labor demand of firm l	w_l^p	π_l
$\phi(w_l^p) = 2\phi(w_l^*)$	w_l^*	$2\phi(w_l^*)(A - w_l^*) - \phi(w_l^*)p$
$\phi(w_l^p) = 2\phi(w_l^*) - 1$	w_l^*	$(2\phi(w_l^*) - 1)(A - w_l^*) - (2\phi(w_l^*) - 1 - \phi(w_l^*))p$
$\phi(w_l^p) = 2\phi(w_l^*) - 2$	$w_l^* - \frac{A}{J/2}$	$(2\phi(w_l^*) - 2)(A - (w_l^* - \frac{A}{J/2})) - (2\phi(w_l^*) - 2 - \phi(w_l^*))p$
$\phi(w_l^p) = 2\phi(w_l^*) - 3$	$w_l^* - \frac{A}{J/2}$	$(2\phi(w_l^*) - 3)(A - (w_l^* - \frac{A}{J/2})) - (2\phi(w_l^*) - 3 - \phi(w_l^*))p$
$\phi(w_l^p) = 2\phi(w_l^*) - 4$	$w_l^* - 2\frac{A}{J/2}$	$(2\phi(w_l^*) - 4)(A - (w_l^* - 2\frac{A}{J/2})) - (2\phi(w_l^*) - 4 - \phi(w_l^*))p$
...	...	...

4.3.4 Results

Comparing the output effects of the two policies gives

$$Y = Y^{tq,UneqFemDis} > Y^{ntq,UneqFemDis}. \tag{4.13}$$

The tradable employment quota is the superior policy instrument to the non-tradable quota when female labor supply is heterogeneous across sectors. Jobs that in male sectors are not productive, due to the employment quota constraint imposed on the firms there, turn productive as those firms can purchase permits that allow them to employ men.

4.3.5 Why Move on with Simulations?

In a next step, one may ask what happens to output in comparison of the two policies as we also allow for lower wage offers to women by discriminating firms, i.e. we have $\mu > 0$. Again, a strongly simplifying assumption may help to gain some intuition. Assume that all discriminating firms would be located in the "male sectors" and the non-discriminating firms in the "female sectors". Then, we would have no wage effects at all arising from the taste discrimination and results would stay the same. Obviously, this particular assumption as the ones previously employed are very strong, perhaps even too strong to derive robust claims on the allocative effects of a tradable employment quota. A more realistic scenario is that a particular sector hosts discriminating and non-discriminating firms. Wage offers and hiring decisions of the non-discriminating firms will affect the discriminating firms and vice versa. Gender specific wages by the type of firm will emerge as well as gender specific employment rates. Moreover, firms unable to fill vacancies with workers fully matching their skill requirements will very likely start searching for workers in more distant sectors although those workers do not fully match in terms of their human capital endowment. As a consequence, the firms' and workers' decisions in a particular sector will have spill-overs on the labor market conditions of workers and firms in the adjacent sectors. In the following sections we will study those interrelated decisions and simulate the outcomes of the more general model introduced in Section 4.2. The intuition derived on the basis of the simplified model will help us to better understand the findings of the more general case.

4.4 Simulation Set-Up

4.4.1 Sequencing

The pseudocode outlined in Algorithm 2 gives the timing of the various actions for the simulated version of the model. A particular iteration starts with each firm l posting J vacancies. Workers apply to firms with a positive number of vacancies, with each worker sending out m applications. Firms evaluate how high their wage offer should be to optimize on the trade-off of actually attracting workers and not letting the wage bill increase by too much. Firms make non-binding job offers obeying the non-tradable employment quota including the wage they are willing to pay to the most profitable applicants. Workers choose the job with the best wage offer conditional on it being above their individual reservation wage. Firms that were not able to hire enough women to fulfill the quota withdraw their offers to male workers.

For the case where we are looking into an economy with a tradable employment quota firms observe how many workers they are able to attract and compare their stock of permits with the number of male workers willing to work for them. Each firm draws its individual supply and demand schedule for the permits. The market maker aggregates these up and determines the market clearing price at which the permits are reallocated between firms. Firms not able to purchase the required number of permits for all the men who wanted to work for them withdraw their offers to the least profitable men.

Finally, the firms produce and observe their payoffs. At the end of each iteration all workers are dismissed and the cycle restarts.

4.4.2 Parametrization

We simulate a labor market with six sectors ($S = 6$) each hosting $N = 20$ firms. Each firm has $J = 10$ positions to fill. Total supply of workers is $I = 1,200$ and reservation wages are equally distributed on $[0, A]$. There is an equal number of male and female workers. When applying for jobs workers send $m = 5$ applications. A firm that fills a vacancy with a worker that has sector specific skills produces $A = 2$. The wage adjustment parameter $\tau = 10$ implies that firms learn over the past 10 iterations. They adjust wages from one iteration to the other with $\epsilon = 0.05$ which equals 2.5% of a worker's maximum productivity.

In one version of the simulation model firms do not discriminate ($\mu = 0$) and women are equally distributed across sectors ($\sigma = 0.5 = \sigma_k$). In the other version firms taste discriminate ($\mu = 0.5$ and $d_F = 0.5$). In the latter case, sectors are split into two types.

```
create sectors;
create workers;
create firms;
if tradable quota then distribute permits;
for t = 0 to T do
    for all firms l = 1 to L do
        post vacancies
    end
    for all workers i = 1 to I do
        send m applications
    end
    for all firms l = 1 to L do
        adjust wage offer given past experience
        send non-binding job offers
    end
    for all workers i = 1 to I do
        if wage offer above reservation wage then
            accept best wage offer
        else
            decline
        end
    end
    if tradable quota then  for all firms l = 1 to L do
        draft supply and demand schedules for permits
    end
    central market makers aggregates supply and demand schedules central market
      makers determines clearing price central market makers reallocates permits ;
    for all firms l = 1 to L do
        if affirmative action policy is not fulfilled then firm withdraws offers to excess
          workers;
        produce
        dismiss workers
    end
    for all workers i = 1 to I do
        if wage offer above reservation wage then
            accept best wage offer
        else
            decline
        end
    end
end
```

Algorithm 1: Pseudocode of model implementation

There is one half of evenly numbered sectors where the share of taste discriminating firms is $\mu_s = 0.1$ and the share of female labor supply $\sigma_k = 0.7$. In the oddly numbered sectors the share of discriminating firms is $\mu_s = 0.9$ and the share of female labor supply $\sigma_k = 0.3$.

The parametrization for the policies is as follows. First we look into an affirmative action policy which prescribes every single firm to employ at least 50% women. Then, we introduce a tradable employment quota. 180 permits are equally distributed among the firms at the beginning of a run. Table 4.3 summarizes all these parameters.

Justifying the choice of parameters is inherently difficult as some of these parameters are hardly observable and empirical estimates are lacking. A firm's taste for men is hardly measurable so that one has to recur to its consequences arising in terms of the wage differentials we observe between men and women. But even those raw differentials which often amount to about 25% lower wages for women (see, e.g., Altonji and Blank, 1999b) have to be corrected for various other characteristics of the firms and workers. That skills and industry are among those control variables in empirical studies of wage differentials may be seen as justifying our assumption on the declining productivity of workers coming from adjacent sectors. The choice of parameters on workers' application behavior was based on the notion of a labor market with frictions and adjusted such that workers apply mostly in their skill-specific sector but also consider adjacent sectors. The larger share of non-discriminatory firms coming with higher shares of female labor supply shall reflect the occupational crowding of women in sectors where they expect not to be discriminated. The quota policy parameter is in the range of the policies we actually observe or have been proposed in the ongoing policy debates (see our discussion in the introduction). Given the limited evidence and the stylized nature of the simulation model, it becomes most important to evaluate the findings against parameter changes which we do in a section on robustness (see 4.5.3) once the main results have been presented and discussed.

Each iteration t as described in the Pseudocode is replicated for $T = 1,010$ times. This we call a single run. Every treatment consists of 100 runs. For our analysis of the simulation outcome we record the average of the last 10 observations of every run. Thus, we have 100 observations for every treatment.

4.4.3 Difference-in-Difference Approach

A simple comparison of a tradable with a non-tradable quota would yield flawed results with respect to their labor market effects. Only if it was possible to issue the number of permits which exactly matches the restrictions that firms are facing from a uniformly

Table 4.3: Parameter choices

Parameters	
Number of sectors	$S = 6$
Number of firms per sector	$N = 20$
Vacancies per firm	$J = 10$
Full productivity	$A = 2$
Sectoral productivity decline	$a = 0.5$
Learning period	$\tau = 10$
Wage adjustment	$\epsilon = 0.05$
Number of workers	$I = 1,200$
Overall share of female workers	$\sigma = 0.5$
Number of applications per worker	$m = 5$
Intensity of choice	$\lambda = 1$
Share of discriminating firms	$\mu = \{0; 0.5\}$
Discrimination coefficient	$d_M = 0$, $d_F = 0.5$
Quota	$\bar{\sigma} = 0.5$
Number of permits	$C = 180$

applied quota the two policies would be comparable. However, as overall employment and the structure of employment by gender are endogenous this particular number of permits which makes the policies comparable cannot be determined in advance.

As a solution we apply a difference-in-difference approach to analyze the policy effects in comparison. Later on, in the robustness section we implement an alternative approach. For now, the procedure will be to compare the difference of the variables of interest arising from a non-tradable employment quota comparing a labor market with an equal distribution of women and firms without discrimination to a labor market with an unequal distribution of women and of discriminating firms, with the difference between the same variables under a tradable employment quota. Denote with v the labor market variable of interest, with ntq the policy of a non-tradable employment quota, with tq the tradable employment quota, and with eq equally distributed female workers and firms and with $uneqDis$ unequally distributed female workers and discriminating firms across sectors. Then, the policy effect T writes

$$T = (v_{(ntq,eq)} - v_{(ntq,uneqDis)}) - (v_{(tq,eq)} - v_{(tq,uneqDis)}). \tag{4.14}$$

4.5 Simulation Results

4.5.1 Baseline Scenarios Without Policy

We start off with the version of our simulation model to which no policies, neither the tradable nor the non-tradable quota are applied and compare two cases. In one case, female workers and firms are equally distributed across all sectors, and there is no taste discrimination. In the other case, taste discriminating firms and workers are unequally distributed across sectors as explained in Section 4.4.2. The simulation results are shown using box plots drawing on the 100 observations for each treatment.

Figure 4.1 (a) shows the employment rates for both versions of the model. As expected male and female employment match on the left hand side as there are no discriminating firms and allocation of female workers and firms is homogeneous across sectors. Turning to the discriminatory version of the simulation model ($uneqDis$) reveals that male employment increases while female employment decreases. The employment pattern of the discriminatory version of the model is reflected in the wages by gender. Panel (b) shows that female wages fall short of the male wages. Note, that within this baseline scenario a quota could be a way to improve on the labor market prospects of women in the unequal case. However, we will argue that tradable permits will reach the same goal with lower distortions of welfare.

The employment and the wage effects are triggered by the introduction of discriminatory firms. First, due to their taste discriminating behavior those firms offer lower wages to women. Women who were willing to work at going labor market conditions withdraw from the labor market as wage offers decrease. Consequently female employment declines. Discriminatory firms looking for male workers to fill their vacancies will have to increase their wage offers to draw on an additional male labor supply. Male workers who did not choose to work formerly accept job offers now and, thus, male employment increases. Secondly, however, choices of the discriminatory firms have spill over effects on the behavior of the non-discriminating firms in the same and adjacent sectors. As can be seen in Figure 4.1 (c) the non-discriminating firms have a larger share of their vacancies filled than the discriminating firms. The difference in overall employment by the two types of firms arises as the discriminating firms have a substantially lower body of female workers. To the contrary, the non-discriminating firms employ significantly more female workers and less male workers. The non-discriminatory firms adjust their gender specific employment pattern as a response to the discriminatory behavior of their competitors. The gender wage gap (panel d) at discriminating firms reflects the dis-taste of those firms against women and manifests itself in the employment pattern due to the

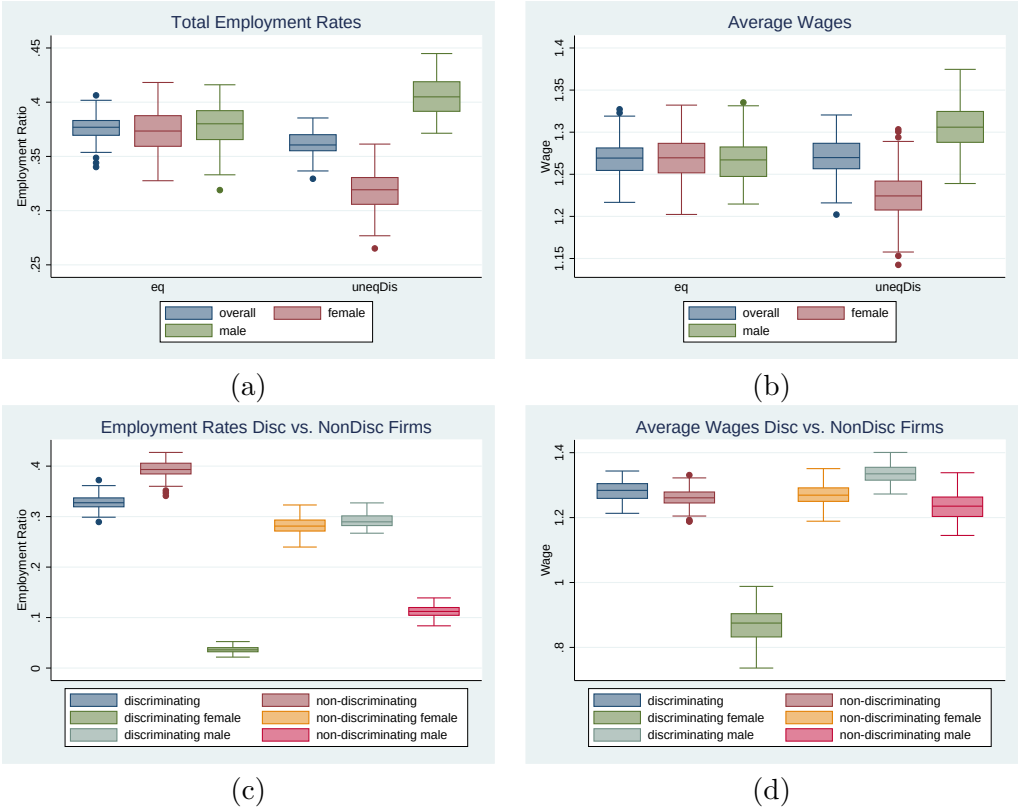

FIGURE 4.1: Employment and wage effects; (a) total employment rates by distribution of female workers and firms, (b) average wages by distribution of female workers and firms, (c) employment rates at disc. vs. non-disc. firms by gender, (d) average wages at disc. vs. non-disc. firms by gender.

distribution of the reservation wages of workers. A further consequence of the taste discriminating firms with respect to the wage pattern is that the non-discriminating firms pay lower male wages than the discriminatory firms because for them a female worker is a perfect substitute for a male worker.

Overall these simulations are in line with the findings of the existing literature on taste discrimination in search models of the labor market. As in Kaas (2009) and Kaas and Lu (2010) our simulation model generates employment segregation. The non-discriminating firms employ more women than men and the discriminating firms employ substantially more men. Comparable to Bowlus and Eckstein (2002) wages for women are lower than those for men. Moreover, we can relate to the search and matching model with heterogeneous reservation wages by Burdett and Mortensen (1998) who argue for the emergence of a wage dispersion in such frameworks. As we re-run our simulation model we do not get a single market wage, neither by gender nor by the type of firm for the 100 repetitions.

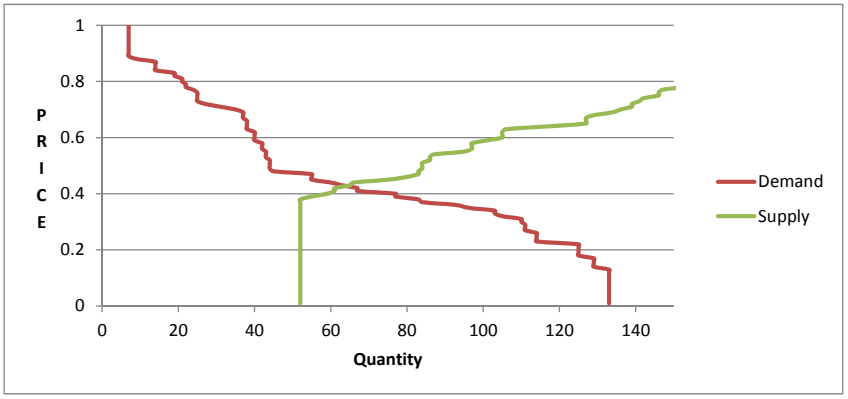

FIGURE 4.2: Bid-ask-diagram for permit trading market with an unequal female labor supply and discriminating firms across sectors at iteration 1,005.

4.5.2 Policy evaluation

4.5.2.1 Trading of Permits

Let us first look into the functioning of the permit market. Figure 4.2 shows a snapshot of the supply and demand schedules that the market maker is facing for a discriminatory labor market at a particular iteration. The downward sloping market demand stems from aggregating up the individual firms' demands for permits given prices. Similarly the upward sloping supply is the sum of permits that firms are willing to sell at given prices. As explained earlier, the market maker chooses the price where supply and demand schedules cross and reallocates the permits from those willing to sell at the market clearing price to those firms willing to buy.

The price for the permits and the traded volume as shown in Figure 4.2 refer to one iteration whereas the observations entering the box plots in Figures 4.3 (a) and (b) are the averages over ten iterations for the 100 runs. At the median about 64 permits are traded between firms. Prices are in the order of 0.44 which is slightly more than one fifth of the maximum productivity of a worker.

4.5.2.2 Welfare Effects

We define welfare as the sum of all wages paid in the economy and all payoffs accruing to firms. The welfare effects of the two policies using the difference-in-difference approach from equation (4.14) can be seen in Table 4.4. The first row shows a decline in welfare of 155 units as we move from a non-discriminatory labor market to a discriminatory labor market applying a non-tradable employment quota. Welfare also decreases for a

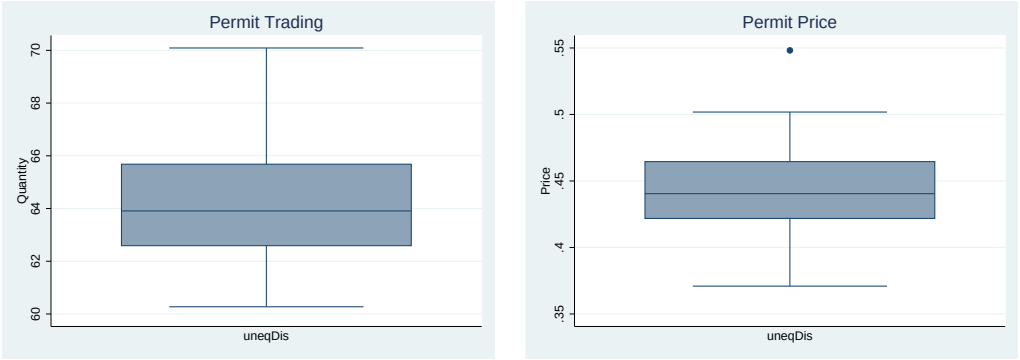

(a) (b)

FIGURE 4.3: (a) Permit trading, and (b) permit prices

tradable employment quota as we make the labor market discriminatory by 75 units. However, the decline is smaller so that the comparison of the welfare losses between the two policies (80) speaks for a tradable employment quota as the superior policy instrument.

These units may be interpreted in relation to total possible output. With 1,200 workers and a maximum per worker productivity of 2, the maximum of total units possible to produce in this economy is 2,400. Measured welfare falls short of the output potential due to the wage offers of profit maximizing firms below reservation wages which drives workers out of the market, and the allocation of workers to sectors where their productivity does not fully unfold. In Table 4.4 the output loss due to skill mismatch is calculated as the difference between actual output and the output that could have been achieved if all workers were employed in their own sector unfolding full productivity.

Overall the output loss due to skill mismatch is lower under the tradable quota. This is because firms can adjust to their labor market situation by engaging on the permit market, while under a non-tradable quota firms will try to fulfill the quota even if this means they have to hire somebody from a different sector. However, as we move from an equal to an unequal distribution of firms and workers over sectors the reduction in output loss due to skill mismatch is larger for the non-tradable quota. This is due to the fact that discriminatory firms are not able to hire women from other sectors as their distaste for women and the sectoral productivity decline induces them to offer too low wages to women from other sectors. Given the limited labor supply of women in their own sector, hiring women from neighboring sectors would be an option in order to fulfill the quota. However, the wages that these firms are able to offer are not attractive enough. Under a tradable quota such problems do not occur, as discriminatory firms in sectors with a smaller female labor supply will simply engage on the permit market. The payoff loss due to the distaste for women of discriminating firms is the wage equivalent of the

Table 4.4: Welfare analysis

	eq	uneqDis	diff			
non-tradable quota						
welfare	711.038	555.952	-155.086			
output loss due to skill mismatch	46.960	33.294	-13.667			
cost of distaste for women	0	23.425	23.425			
output loss due to unfilled vacancies	1,641.984	1,787.326	145.342			
output potential	2,400	2,400				
tradable quota						
welfare	772.798	697.779	-75.019			
output loss due to skill mismatch	40.479	31.803	-8.676			
cost of distaste for women	0	10.366	10.366			
output loss due to unfilled vacancies	1,586.686	1,660.038	73.353			
output potential	2,400	2,400				
				T(diff-in-diff)		**Std. err**
welfare				80.067	***	4.969
output loss due to skill mismatch				4.990	***	0.752
cost of distaste for women				-13.060	***	0.314
output loss due to unfilled vacancies				-71.989	***	5.116

Note: *** denote significance at the 1% level.

distaste (d_F) multiplied with the female employment at discriminating firms. As there is less female employment at discriminating firms those costs decrease with a tradable employment quota. The largest effect on welfare of a tradable employment quota comes from a reduction of unfilled vacancies as the permits give firms a higher flexibility in terms of the gender composition of their workforce. The sum of these differentiated effects gives the welfare effect of the two policies in comparison.

4.5.2.3 Effect on Employment, Wages, and Payoffs to Firms

An alternative way of looking at these welfare measures is to split welfare in firm payoffs and the wage sum (see Table 4.5). Here we find that the wage sum is approximately twice as large as the payoffs to firms. When looking at the differences-in-differences result in the lowest third of the table we find that more than 70% of the effect of a tradable employment quota on welfare accrue to changes in the wage sum. This effect we can disentangle into an employment and a wage effect. The remaining effect on welfare is due to changes in the payoffs to the firms .

The drop in employment is less pronounced for the tradable quota as one moves from a non-discriminatory labor market to a discriminatory labor market if compared to the non-tradable quota. Why is that? Essentially, it is driven by the economic mechanisms that was already detected in the simplified version solved analytically. Without permits a firm wanting to produce may be constrained by the employment quota. There is a vacancy to be filled, and while there is no female applicant to fill the vacancy, a male worker would be willing to accept the job offer. However, the firm cannot employ him because the employment quota has to be obeyed. A tradable employment quota gives firms facing such a situation more flexibility. They may purchase a permit allowing them to employ an additional man, and they will do so as long as the additional male worker's profitability covers the price of the permit. Thus, more vacancies can actually become productive with a tradable employment quota in a discriminatory labor market as reflected in the non-decline of the male employment rate for the tradable quota.

Moreover, for the tradable employment quota there is only a small wage effect while wages decrease for the non-tradable quota. Looking into the gender composition of the wage effect reveals that this occurs mainly because female wages decrease more for the non-tradable quota. As discriminating firms are facing a fixed quota they are required to hire women in order to balance their workforce. However, due to their distaste for women they are only willing to do this at a lower wage. This effect is non-existent for the more flexible tradable quota because discriminating firms can avoid to hire women by purchasing permits.

Table 4.5: Effect on employment, wages, and payoffs to firm

	eq	uneqDis	diff
non-tradable quota			
payoffs	228.937	183.868	-45.069
wage sum	482.101	372.084	-110.017
empl.	0.316	0.255	-0.061
empl. female	0.452	0.388	-0.064
empl. male	0.180	0.123	-0.057
ave. wage	1.272	1.215	-0.057
ave. wage female	1.287	1.214	-0.073
ave. wage male	1.233	1.215	-0.018
tradable quota			
payoffs	266.377	242.537	-23.840
wage sum	506.421	455.242	-51.179
empl.	0.339	0.308	-0.031
empl. female	0.378	0.317	-0.061
empl. male	0.300	0.300	0.000
ave. wage	1.245	1.230	-0.015
ave. wage female	1.269	1.229	-0.039
ave. wage male	1.215	1.231	0.016

	T(diff-in-diff)		**Std. err**
payoffs	21.228	***	1.959
wage sum	58.838	***	4.070
empl.	0.030	***	0.002
empl. female	0.003		0.004
empl. male	0.057	***	0.001
ave. wage	0.042	***	0.005
ave. wage female	0.034	***	0.006
ave. wage male	0.034	***	0.007

Note: *** denote significance at the 1% level.

Finally, the increase in the employment rate also explains the effect of a tradable employment quota on payoffs as an otherwise unproductive vacancy is filled now.

One may even go further and analyze the distributional consequences of the two policies in comparison. How are the policies affecting wages paid at discriminating and non-discriminating firms, and do we observe effects on wages paid to workers employed in their own sector or in a sector which does not fully match their skills? Furthermore, how do the firms payoffs adjust with respect to the two policies in comparison. This is analyzed in Table 4.6.

Wages paid to workers who found jobs in the sector for which their skills match best decrease less if we move to an unequal distribution of firms and workers for the tradable quota if compared to the non-tradable quota. The same but larger effect can be observed

Table 4.6: Distributional effects

	eq	uneqDis	diff
non-tradable quota			
ave. wage own sector	1.371	1.312	-0.059
ave. wage other sector	0.937	0.828	-0.109
ave. wage non-disc.	1.272	1.282	0.010
ave. wage non-disc. female	1.287	1.301	0.014
ave. wage non-disc. male	1.233	1.202	-0.031
ave. wage disc.	*1.272*	1.004	-0.268
ave wage disc. female	*1.287*	0.868	-0.420
ave. wage disc. male	*1.233*	1.232	-0.001
payoff disc.	*1.908*	0.680	-1.228
payoff non-disc.	1.908	2.384	0.476
tradable quota			
ave. wage own sector	1.331	1.308	-0.024
ave. wage other sector	0.861	0.818	-0.043
ave wage non-disc.	1.245	1.249	0.004
ave. wage non-disc. female	1.269	1.274	0.005
ave. wage non-disc. male	1.215	1.166	-0.049
ave. wage disc.	*1.245*	1.202	-0.043
ave. wage disc. female	*1.269*	0.862	-0.407
ave. wage disc. male	*1.215*	1.256	0.040
payoff disc.	*2.220*	1.646	-0.574
payoff non-disc.	2.220	2.396	0.176

	T(diff-in-diff)		**Std. err**
ave. wage own sector	0.035	***	0.000
ave. wage other sector	0.066	***	0.000
ave. wage non-disc.	-0.006		0.005
ave. wage non-disc. female	-0.008		0.006
ave. wage non-disc. male	-0.018	**	0.008
ave. wage disc.	0.224	***	0.005
ave. wage disc. female	0.013	*	0.007
ave. wage disc. male	0.041	***	0.007
payoff disc.	0.654	***	0.016
payoff non-disc.	-0.300	***	0.020

Note: ***, **, * denote significance at the 1%, 5%, 10% levels, respectively. Entries in *italics* are approximated values for the firms which are currently non-discriminating but become discriminating firms for an unequal distribution of firms and workers.

for average wages of workers not employed in their own sector. Thus, there is some heterogeneity in the overall wage effect already detected in Table 4.4. Disentangling the overall wage effect along discriminating and non-discriminating firms reveals the following picture: the average wage of discriminating firms falls quite dramatically when moving from an equal to an unequal distribution with a non-tradable quota. This is because discriminating firms are forced to hire females, but pay much lower wages. For the tradable quota the share of females hired by discriminatory firms is quite low and thus the lower wages for females have a lower weight when calculating average wages. Actually, decomposing those wages by gender reveals that it is the wages paid to male workers by the discriminating firms driving this effect. Those firms are not constrained by a fixed quota anymore. Rather they are allowed to hire more men if they own the corresponding number of permits. But hiring those additional men requires to offer higher wages.

On average the discriminating firms have lower payoffs than the non-discriminating firms in both policy scenarios. This is due to their distaste for women and the lower wages they pay. As a consequence they draw on a smaller female labor supply which allows them to only fill fewer vacancies. The wedge between the payoffs of the discriminating and non-discriminating firms decreases with a tradable quota as the discriminating firms may purchase permits and thereby avoid idle vacancies due to their distaste for women. Relatively speaking, the discriminating firms profit from a tradable employment quota. It may be noteworthy at this point that looking into the distribution of payoffs to firms over time reveals that even the least profitable firms have non-negative payoffs (and will stay in the market.)

Comparing the results from the simulation model with what we derived earlier on within our analytical framework shows that similar economic mechanisms are driving the main outcomes. This is so although the simulation model rests on adaptively behaving firms and analytical results were derived under the usual assumption of rationally behaving firms. Rationality does not seem to be the crucial assumption being responsible for the effects of the two policies in comparison. Rather it occurs that adaptively behaving firms trying to find a wage offer which maximizes their profits is a somehow equivalent assumption to rationality in this context. One may even claim that it is also more appropriate as in a new policy environment firms will also have to find out for themselves how to behave best.

4.5.3 Robustness

We ran a series of robustness tests involving (a) the change of particular modeling assumptions, (b) changes of key parameters of our model, (c) a different initial allocation of permits to firms, and (d) an alternative evaluation of the effects of a tradable permit system when compared to a fixed quota. In all four instances we were interested in whether our main result of the positive welfare effect of a tradable employment quota prevails.

In our baseline model we let firms post wage offers to prospective workers without taking into account the upcoming costs that may arise with the need to purchase permits to employ men. Alternatively, one may imagine that firms make wage offers already taking into account that they may have to go to the permit market. To this end we ran a series of robustness tests where firms anticipate that they will have to hold permits in order to employ men and thus adjust the wage taking into consideration an estimate of the permit price. In particular, this estimated permit price p_t^e in period t was calculated as the average price over the last ten periods: $p_t^e = \left(\sum_{\tau=t-10}^{t-1} p_\tau\right)/10$. It turned out that this led to lower wages and lower employment for men, while employment of women increased even further. Our main result, however, that a permit system is superior to a fixed quota was not affected.

In order to analyze whether our results are sensitive to changes in the parametrization of the model we combined the two key policy parameters, i.e. the number of permits issued and the quota with all of the remaining parameters, one-by-one, and re-ran the simulations. The results of this exercise are summarized in Figures 4.4, 4.5, and 4.6. In all left columns of those figures we combine changes in the number of permits with one of the remaining parameters. In the right columns we combine changes in the quota with one of the remaining parameters. The grids are over three values of each parameter. The welfare effects of those nine parameter constellations are combined to a plane. The dark blue planes show the mean effect on the diff-in-diff welfare measure of 100 repetitions. Upper and lower planes in gray constitute the confidence intervals. The black dots refer to the diff-in-diff value for the welfare effect at the standard parameter constellation chosen in Table 4.5.

Figure 4.4 collects all parameters related to the adjustment behavior of agents and the number of iterations of each run. In order to exemplify how to read those graphs let us turn to Figure 4.4 (a). As we increase the number of iterations there is no change in the welfare measure which suggests that we have chosen an appropriate length for our simulation runs, i.e. one where the adjustment processes have worked themselves out. The flat plane furthermore indicates that this observation holds for the higher and also

the lower value of permits. Turning to Figure 4.4 (b) replicates the analysis for a fixed quota showing that also in this case increasing the number of iterations does not alter the results. Changing the wage adjustment parameter or the learning period of the firms does not affect our main finding, either.

Figure 4.5 brings together parameters related to frictions in the labor market. Here, we were mainly interested in whether the number of applications sent out by the workers, their attachment to the skill-specific sector, and the closeness of sectors in terms of how suitable workers are for firms that come from adjacent sectors affect our simulation outcomes. Again, the results confirm our main findings.

Finally, Figure 4.6 assembles all robustness analyses related to parameters describing the discriminatory behavior of firms. This set of analyses may also be interpreted as a test on the robustness of the results given that firms' behavior may change over time. One could imagine that firms confronted with affirmative action policies or an evolving market reconsider their behavior and become non-discriminatory. Although we opted for a modeling of discriminatory behavior in the sense of taste discrimination, such behavior may, in particular, arise if one thinks of statistical discrimination as an alternative explanation of discriminatory behavior where firms may update their beliefs over time. Then the question arises whether a tradable quota system is still working. To this end, we alter the firms' costs of discrimination, the distribution of discriminatory firms over sectors, and the distribution of women over sectors. Once more, the permit system yields welfare gains.

Another look at the right columns of all three Figures 4.4, 4.5, and 4.6 reveals that the welfare effects of a tradable quota relative to a non-tradable quota measured along the diff-in-diff methodology increases as the non-tradable quota becomes smaller. This result is driven by a lower difference in welfare for a fixed quota comparing an equal with an unequal distribution of firms at decreasing quotas (as the number of permits is unaffected by this change). A lower non-tradable quota increases welfare for an equal distribution of firms and women across sectors as it is easier for firms to meet this requirement. However, the increase becomes smaller as the quota is less binding (the slope when moving from a 60% quota to a 50% quota is steeper as compared to moving from a 50% quota to a 40% quota). For an unequal distribution of discriminating firms and women across sectors a less binding quota does not have the same effect on additional production. This is due to the discriminating firms offering lower wages because of their distaste for women, thus depressing female labor supply so that those firms will not be able to hire women which would allow them to fill vacancies with additional men even under a relaxed quota.

In our baseline simulation we allocated permits on a per-capita basis to the firms initially. One may wonder how alternative allocations of permits affect the results. We experimented with various other rules including a random and a highly skewed distribution of permits. As time evolves we get similar results in terms of welfare differences between the two policy instruments when compared to the per-capita rule.

Finally, we opted for an alternative evaluation of the policy effects. Rather than employing a diff-in-diff approach we modeled a government that endogenously adjusts the number of permits issued in every period so that policies become comparable. More specifically, for this robustness test permits are valid for only one period and the government every period issues a number of permits that would achieve a ratio of male to female employment compliant with a quota. What we expected from such an exercise is that on average welfare is higher for the policy simulations with tradable permits than with a fixed quota. And indeed, that is what the simulations showed which makes us more confident that the results are not driven by our methodology chosen to evaluate the policy effects.

4.6 Discussion

Discrimination of women (as well as other labor market groups) calls for affirmative action policies. In an economy where women's labor supply is heterogeneous across sectors or regions, an employment quota applied uniformly to firms may cause avoidable costs to society. Firms located in sectors where the supply of female labor is relatively scarce may find it inherently difficult to comply with an employment quota whereas firms in other sectors where female labor supply is relatively strong will do better in terms of filling vacancies with female workers. Output losses may occur in sectors with relatively weak female labor supply as firms subject to the employment quota cannot fill up vacancies with men even if a woman cannot be found.

We propose and analyze a flexible quota solution. As we argue, a tradable employment quota gives firms additional flexibility to hire men if female labor supply is insufficient. By issuing permits to firms allowing them to hire men and making these permits tradable across firms, firms in shortage of women will not be forced to abandon output. Rather they will try to purchase a permit that allows them to hire a man up to the point where this additional man's profitability covers the costs of the permit. Equally, we will have firms in this market that will find it profitable to sell permits as they can easily fill their vacancies with women. An advantage of a tradable permit system is that it allows for a flexible adjustment at the firm level without having to compromise on the overall policy goal to achieve a certain share of female employment in the labor market. The scope

of female employment in the economy can be managed at the aggregate level by issuing or withdrawing permits. Our results suggest that a market for a tradable employment quota may emerge, and that a more flexible policy solution is actually improving welfare. Those results have to be interpreted as arising from a comparison between two different policy measures trying to achieve higher female employment shares. Running our model without any restrictions on employer behavior yields higher welfare levels than what one gets in the policy experiments. From our point of view, however, a meaningful discussion of affirmative action policies starts from policy goal that female employment should improve, and then tries to answer the question which policy measures are most suitable to achieve this goal. Our exercise should be seen as a contribution in this spirit.

Although we did make an effort to implement features of the labor and permit market in considerable detail, we have been silent about the occupational or regional scope of a permit system. Given our discussion in the introduction on the recent moves to make shares of female workers legally binding for boards of firms, one may at first think of a permit market for that segment. In principle, however, we believe that a tradable employment quota would also work for other groups of occupations, may it be introduced nationally or for a set of countries (as it may work for other under-represented groups in the labor market).

An outcome of our analysis is that a tradable employment quota achieves an overall female employment rate at lower costs for society but results in an unequal distribution of women across firms. One may question such an outcome on the backdrop of a more normative assessment which posts that women should be equally represented also on a firm basis. In this case, evaluation of policies would certainly be tilted towards a non-tradable quota. Moreover, in a comparison of the two policies one might be inclined to favor a fixed quota if one believes that forcing firms to employ women may help to reduce prejudices against women as a possible cause of discriminatory behavior. We abstained from modeling such a feedback process as there seems to be no strong evidence currently available that points towards such a channel (see, e.g., Bertrand et al., 2010). Moreover, it is not clear to us if a fixed employment quota would only help changing preferences of "male" firms towards becoming more "neutral". Actually, there is evidence reporting that women being protected by fixed employment quotas are stigmatized (Heilman et al., 1992, 1997). Thus, a fixed quota might actually lead to additional discrimination as women are disrespected by "ordinary" employees.

If Norwegian firms did not comply with the employment quota by January 2008 they were planned to be dissolved (Nygaard, 2011, p.23). Rather than forcing firms to go out of business, one may also think of an implementation where non-complying firms are fined. Whether such a fining system is functionally equivalent to a cap-and-trade system that

we propose can be discussed very much along the well known lines of practical problems arising with the implementation of a Pigouvian tax.[10] Very likely the marginal costs for employing women differ for firms which would require to install tailored penalties (or taxes). Collecting the necessary information and administrating such a system may be prohibitively difficult, and a permit trading system a viable alternative.

We also did not compare the two policies with respect to their medium or longer run consequences. Our focus was rather on how costs for firms that become subject to a uniform quota and have to change the gender composition of their workforce can be mitigated right after an affirmative action policy is introduced. While we have shown that within the short time horizon a tradable quota fares better it is also conceivable that a tradable quota is the better policy choice in the medium and longer run. It gives governments a flexible tool that allows them to react to changes in firms' and workers' behavior as time evolves by adjusting the supply of permits and thereby the restrictions that firms are facing. Such a policy response could be easier to implement as changes of the law which seem to be quite often accompanied with intense public debates.

Our analysis and the results we derive rest on various assumptions we had to make. Although we ran robustness checks changing the parametrization of our model as well as the model implementation, one may be concerned about the partial nature of the model which abstracts from various feedback processes. In particular labor demand of firms is derived from a fixed product demand. Introducing a fixed or a tradable quota may, however, change the production costs of firms which firms may pass through to product prices. Thus, with heterogeneous firms relative prices and firms' market shares may change as a consequence of the policies. Actually, there is evidence on cost-pass through but estimated magnitudes to which extent this happens vary widely, see Goldberg and Hellerstein (2013) and Fabra and Reguant (2014) for two extreme cases. Given this rather large variety in estimates with a larger share of estimates finding relatively small effects of cost-pass through, we believe that a partial model should not bias our main results substantially. But it may be a worthwhile exercise to embed a comparison of affirmative actions policies in a macroeconomic model of the economy in future work, even more so as a larger share of firms becomes subject to these policies.

[10] Arguments for gender-based taxation have already been brought forward by Rosen (1977), Boskin and Sheshinski (1983), or, more recently, Alesina et al. (2011). There, differentiated taxes may also boost female employment at lower losses in welfare because labor supply elasticities differ beween men and women.

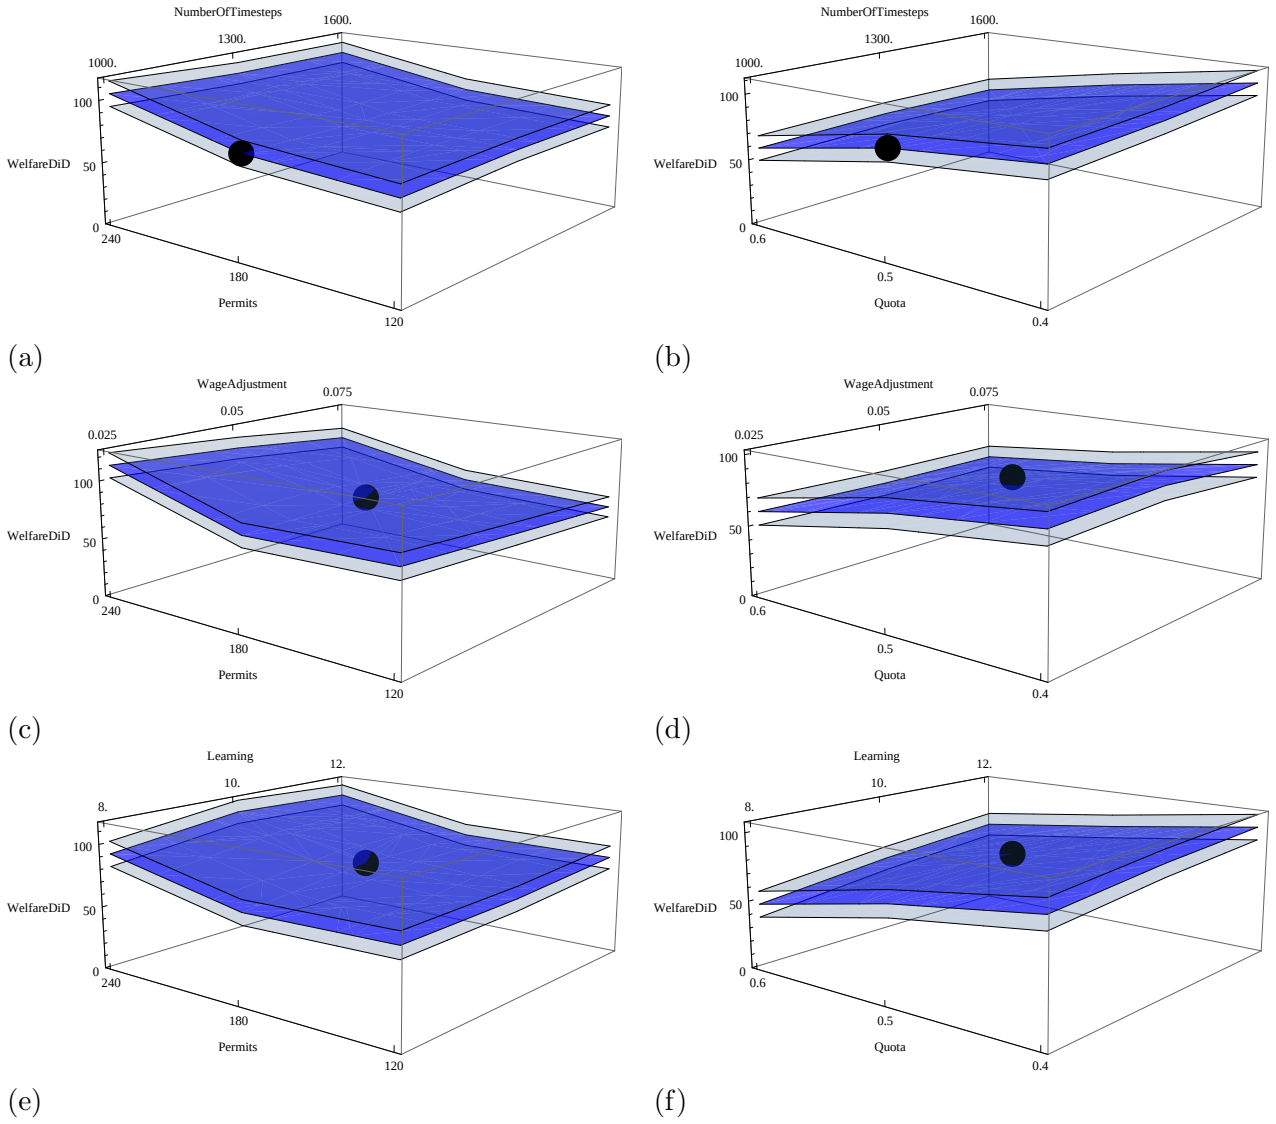

FIGURE 4.4: Robustness with respect to dynamic behavior of agents

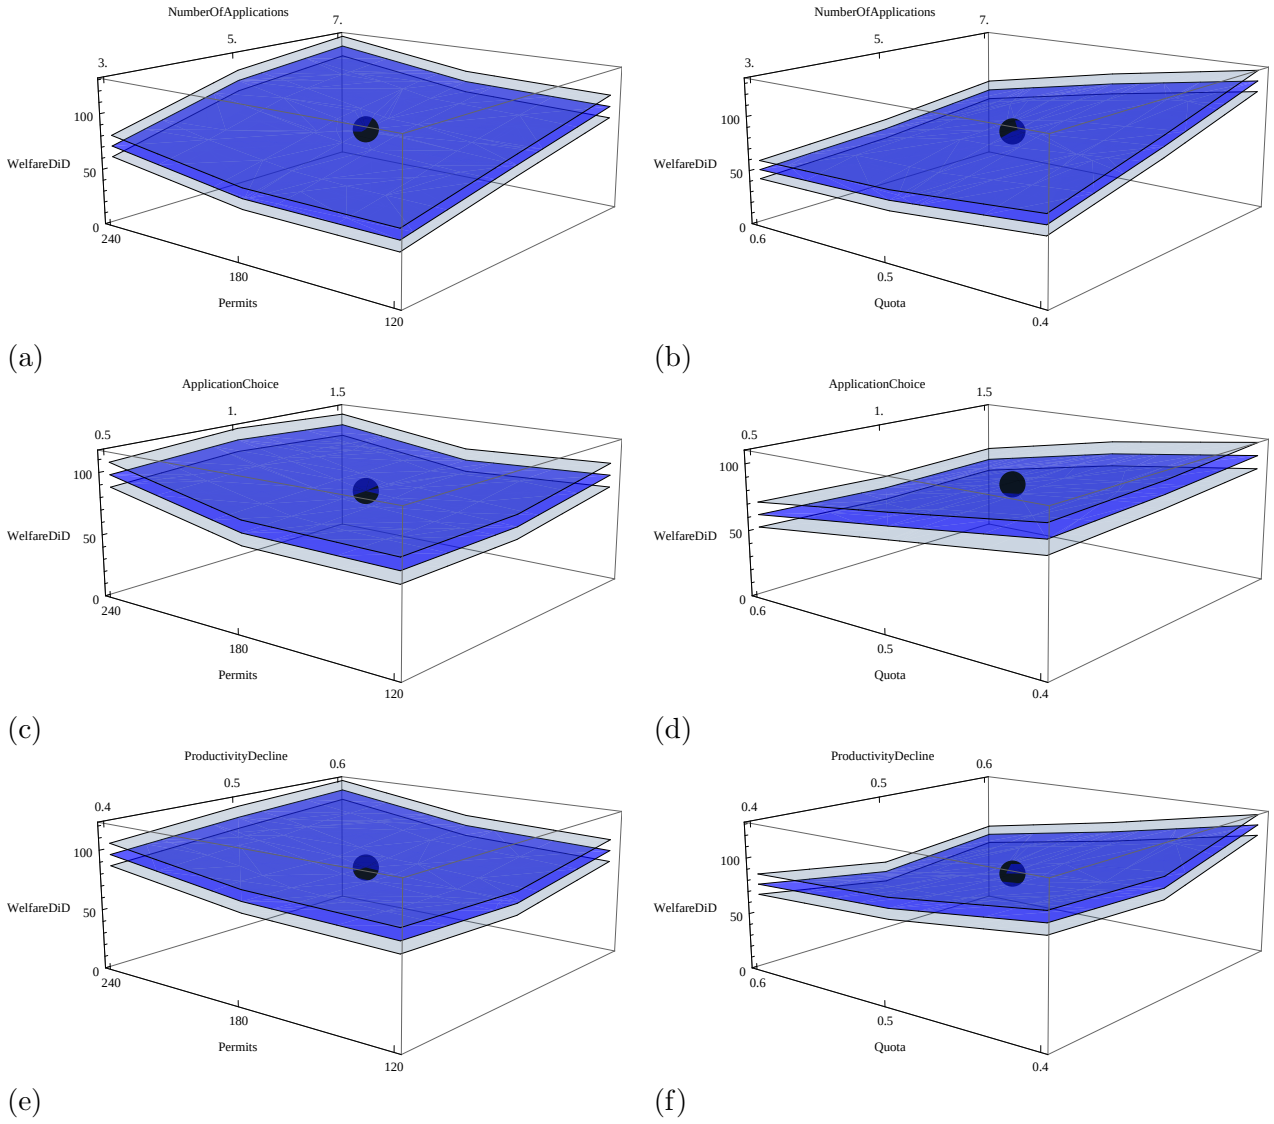

FIGURE 4.5: Robustness with respect to labor market frictions and heterogeneity

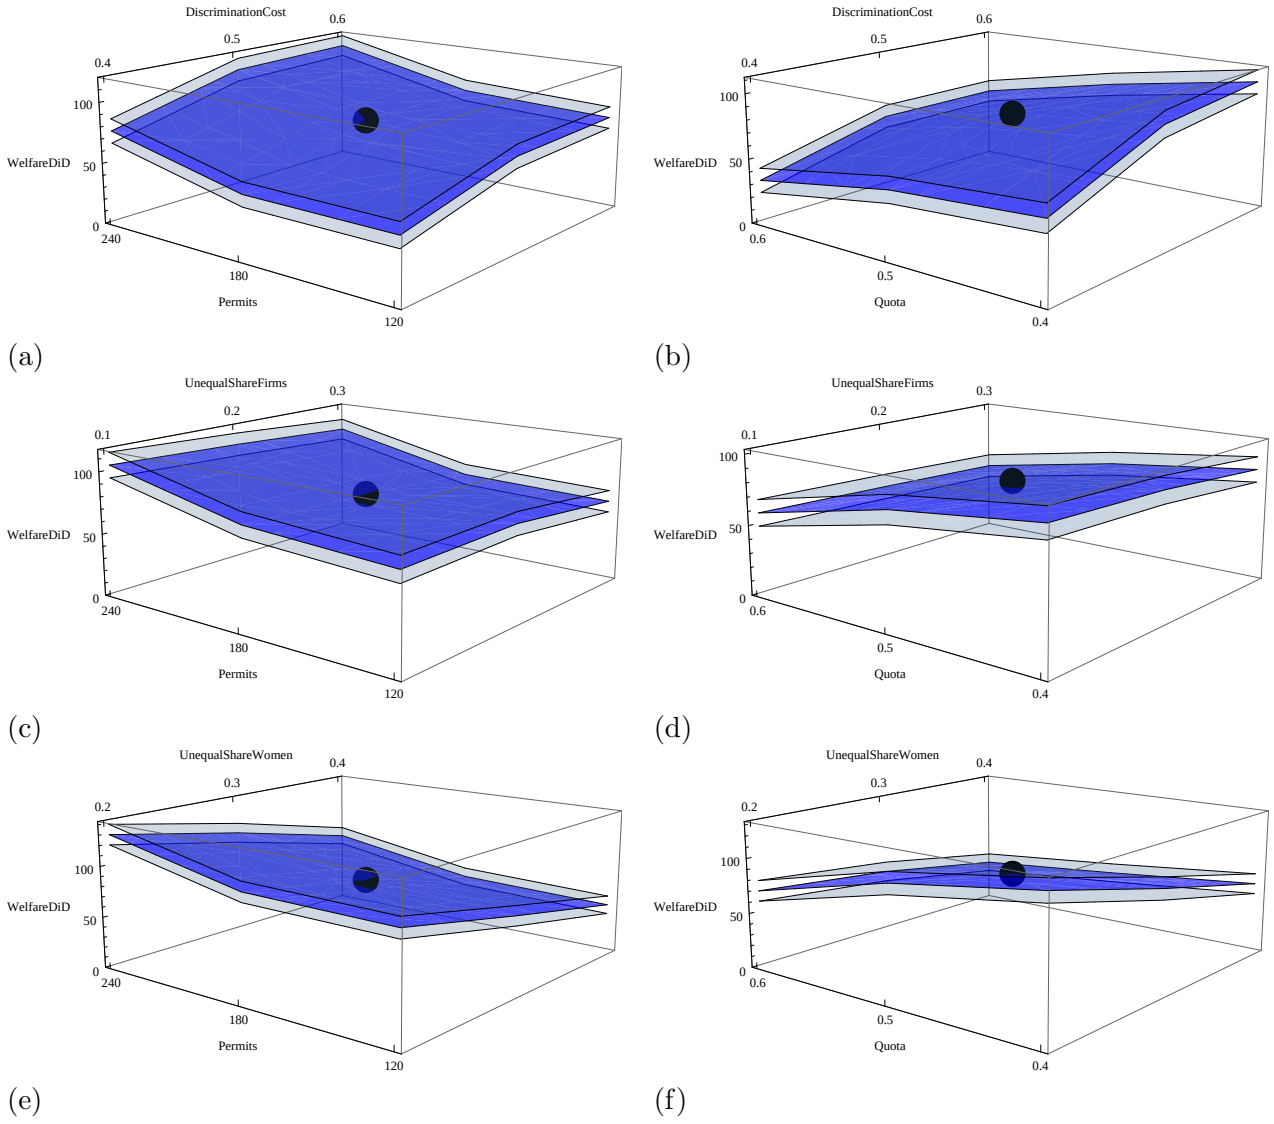

FIGURE 4.6: Robustness with respect to extent of discrimination

Appendix

Table 4.7: Women on boards

Country	Members of board			Board chairs [%]	CEOs [%]
	Ave. share [%]	Min [%]	Max [%]		
USA (a)	16				
Japan	1				
EU-27	16			3	2
Finland	29			4	0
Latvia	28			13	3
Sweden	26			0	4
...					
Germany	18	0 (b)	44 (b)	3	0
...					
Portugal	7			0	0
Hungary	7			0	0
Malta	4			0	5
Norway (c)	39				

Data source: European Commission Database on women and men in decision making for 2012 if not otherwise stated. (a) European Commission: Women on boards - Factsheet 2, Gender equality in member states; (b) Der Spiegel 48/2013, p.74; (c) Ahern and Dittmar (2012, p.143).

Chapter 5

Key Players in Labor Market Networks[1]

5.1 Introduction

The importance of social networks in shaping behavior and economic outcomes is well documented (for an overview see for instance Durlauf (2004), Ioannides and Loury (2004), Jackson (2014)). Typical channels through which peers affect the economic outcomes of individuals are the dissemination of job information and the way in which education and other human capital decisions are influenced through connections in the network. We know for instance that a large percentage of jobs are found through personal networks. Mark Granovetter (Granovetter (1973), Granovetter (1995)) using survey data finds that over 50 percent of jobs were obtained through social contacts. A similar study by Rees (1966) finds an estimate of over 60 percent.[2] Brock and Durlauf (2001) show in a theoretical model how individuals' behavior and decisions are affected by their desire to conform to the behavior of their peers, which allows to study a wide array of social phenomena such as entry or withdrawal from the labor force, job search, participating in criminal activities vs. participation in the above ground economy, staying in or dropping out of school. There is also a vast and growing empirical literature that attempts to provide empirical estimates of the Degree to which local interactions and neighborhood effects influence individual decisions (a comprehensive survey of the empirical neighborhood literature is given by Ioannides and Topa (2010).

These results have important implications regarding labor market outcomes of individuals who locate in ethnic "enclaves" within metropolitan areas. Calvo-Armengol and

[1] This Chapter is based on Akyol (2016b)
[2] see also Bewley (1999)

Jackson (2004) for instance show in a theoretical model in which individuals hear about new jobs through their peers, how the correlation of unemployment between connected individuals in a network can cause a rise in inequality. Topa (2001) shows using Census Tract Data for Chicago how neighborhoods exhibit positive spatial correlations, i.e. individuals are more likely to be employed if their neighbors are employed. This interdependence between group and individual behavior can precipitate the emergence of so called low level equilibria or poverty traps, from which individuals cannot escape easily. At the same time, these social interaction effects also may have important implications regarding the effectiveness and efficiency of policy interventions (such as educational training measures), as policies directed at individuals in a network can create "social multiplier" effects, i.e policies can have spill-over effects on non-treated individuals who are connected in the network. The extent to which such repercussion effects are present is subject to debate.

A number of empirical studies attempt to identify under which conditions such repercussion effects can be utilized for policy measures, e.g. Jackson and Lopez-Pintado (2013) study these in a general model of diffusion and show how certain behaviors spread within social networks that exhibit segregation or homophily (the tendency of agents to associate with others similar to themselves) and identify conditions under which a behavior diffuses. Banerjee et al. (2014) also examine in a theoretical analysis, how the central member of a network who is most apt to diffuse information can be identified. Banerjee et al. (2013) use detailed demographic and social network data from 43 villages in South India to estimate the effectiveness of various "injection points", i.e. the individuals that were first informed about the program. These individuals are chosen based on various degrees of centrality in the network (where centrality is determined based on multiple measures, e.g. Degree centrality, Eigenvector centrality, etc.).

While there is a vast array of studies that examine to what degree individuals are affected by the network they are embedded in, the research that looks into implications or applications of these findings for labor market policies is more scarce. In particular, it would be of interest how labor market policies could benefit by the knowledge of certain network structures and possible peer effects. Ballester et al. (2006) examine how crime rates could be reduced through a policy that consists of targeting the key player, that is, the player who, once removed, leads to the optimal change in aggregate crime activity. The approach of this paper looks into a related question of how a labor market policy such as educational subsidies could increase its effect on a group of individuals if the network structure is taken into account. To do so, I use a computational labor market model that allows to evaluate policy measures that are targeted towards central players. This not only takes network effects into account, but also shows how they can be utilized by a policy maker to increase the impact of a policy measure. Individuals

in the model are influenced in their educational decisions by their connections in the network. The model then simulates the effect certain key player policies such as providing an educational subsidy to individuals who have a high centrality measure (e.g. Katz-Bonacich centrality, Eigenvector centrality). The policies can then be evaluated with regard to their ability to maximize the effect of the policy on the entire network through diffusion but also regarding the impact on inequality. From a policy maker perspective it is naturally of special interest to increase these type of peer effects as much as possible for a given economic policy intervention as they allow for a 'social multiplier' that increases the original effect of the policy. Within this context, the goal of this paper is to investigate these so called spill-over effects on non-treated individuals, specifically which conditions affect the size of these types of effects. The central questions that are addressed in the paper are: 1) Which agents in the network should a labor market policy (e.g. educational measures) be targeted towards in order for the policy to be most effective? 2) How does such a central player policy affect the inequality within the group of agents? Does an overall increase of utility also reduce inequality?

Using a computational model extends the literature in a number of ways. It first allows to introduce large scale networks with various degrees of characteristic network properties that are observed empirically, such as the degree distributions and the amount of clustering and thereby gives a better understanding of the diffusion processes that take place. Meanwhile, the behaviour of the agents can be modelled in various ways and is not limited by simplifying assumptions in order to ensure tractability. Specifically, once a formal model of the heterogeneous behavioural dispositions of the agents is given, the computational approach simulates aggregate outcomes that will result from this individual behaviour and possible strategic interaction. Furthermore, it features heterogeneous agents and thereby give deeper insight into inequality within the group of agents.

The results indicate that overall utility can be increased the most by targeting the subsidies towards the most central individuals. Hereby, I find that there is no significant difference in the type of centrality that was used to find the most central individuals. This finding is in contrast to the results of previous work that compares centrality measures (e.g. Ballester et al. (2006) and Banerjee et al. (2013)). This result largely stems from that fact that unlike small scale networks that are highly highly stylized, on large scale network the various centrality measures strongly overlap in their choice of central players and thereby render the differences in overall results insignificant. A second result is that such a "targeted" approach comes at the cost of increased inequality. That is, targeting the individuals who are very central in the network, results in identifying individuals who are also the most affluent ones in terms of education and income (due to the correlation in centrality and educational investments). This precipitates a "Matthew effect" in the sense that, the individual who are already at an advantage due to their

position in the network, receive an additional benefit which widens the gap between them and the agents who are at the periphery of the network.

The remainder of this paper is organized as follows: Chapter 2 introduces the theoretical model and describes the labor market policies that will be implemented. Chapter 3 then discusses the computational results of a random transfer policy for the baseline case with a random graph. The robustness section relaxes the network creating process and allows for homophily and clustering. Chapter 4 then introduces the targeted approach and compares the welfare and inequality implications of this policy. Chapter 5 then concludes.

5.2 Model Specification

5.2.1 Model Specification

Following Ballester et al. (2006) and Benabou (2002) I use a simple network model which features strategic complementarities in the education level of individuals and includes inter-temporal educational investment. In particular, the network **g** consists of a finite set of agents N=1,..., n. Connections in this graph are represented by an n $\times$ n adjacency matrix **G**, in which entries g_{ij} represent connections between agents *i* and *j*, whereas if $g_{ij} = 1$ agent *i* is connected to *j*, and otherwise $g_{ij} = 0$. By convention $g_{ii} = 0$ (zeros on the diagonal) and $g_{ij} = g_{ij}$, i.e. the graph is undirected (links are reciprocal).

5.2.2 Household Characteristics

I consider the educational investment and consumption choices of individual *i* in the job market (e.g. investment into a college education). In period *0* agent *i* chooses consumption c_0^i and c_1^i to maximize intertemporal utility U_i subject to a lifetime budget constraint that that features an intertemporal investment into education (which subsequently affects next periods income). These choices then determine the available education investment into human capital. Any income y_t is determined by the amount of time worked l_t and the human capital that is available h_t in that period. Income is then used to finance consumption in the same period, while any income that is not used for consumption in period $t = 0$ can then be invested into educational attainment e_t, i.e. agents choose in today's period ($t = 0$) how much they are willing to forgo in today's consumption (c_0) in order to invest into their educational attainment which increases

tomorrow's income and subsequently tomorrow's consumption (c_1). Investment in education is predicated on the assumption of a return in future periods, given by ψ. This basic model is then extended by introducing a government.

A formal description of the household preferences is given by:

$$\max_{c_t^i} \mathcal{U} = lnc_0^i + \beta lnc_1^i$$

$$s.t.$$

$$y_t^i = h_t^i * \overline{l}_t^i \qquad \forall \quad t = 0, 1 \tag{5.1}$$

$$\hat{y}_t^i = y_t^i - \tau_t^i \qquad \forall \quad t = 0, 1 \tag{5.2}$$

$$\hat{y}_t^i = c_t^i + e_t^i - \alpha_t^i \qquad \forall \quad t = 0 \tag{5.3}$$

$$\hat{y}_t^i = c_t^i \qquad \forall \quad t = 1 \tag{5.4}$$

$$h_{t+1}^i = h_t^i + e_t^i(1+\psi) + \gamma \sum_{i \in \nu, j \neq i} g_{ij} e_t^i e_t^j \qquad \forall \quad t = 0 \tag{5.5}$$

where $0 < \beta < 1$ is a discount factor ($\beta = \frac{1}{1+r}$ with r being the discount rate), and where $0 \leq \gamma < 1$. For all values of $\gamma \neq 0$ we thus have an element of complementarity, in the sense that agents benefit of the educational attainment of their connections or peers through getting better job referrals etc. The reasons to assume these type of complementaries abound (see for instance Ballester et al. (2006), Ballester et al. (2010), and Durlauf (2004)). Complementarity is a fundamental concept of the social interactions literature and characterizes how an individual i are affected by individuals to which i is connected (e.g. friends, neighbors, colleagues). The term comprises two channels for such an influence, namely the attributes of the connected agents and their behaviour, whereas the former is treated as an exogenous predetermined characteristic while the latter often is endogenously determined. In the setting of this model, the characteristic item would be the position of the individual in the network (which is exogenously determined and does not change throughout the analysis) while the choice between consumption and educational investment encompasses the behavioural choice. In the context of this model, one conceivable reason for such a complementarity could be as follows: the increased educational investment of the connections of agent i will increase their ability to find a job and thereby increase the probability for job referrals they can relay back to agent i.[3]

[3]Note that the model abstracts from a possible global substitution effect, i.e. given a limited demand for workers the increasing education of one individual agent could then also have negative effect on the job opportunity of all the agents in the model, see for instance Ballester et al. (2010)

5.2.3 Government Characteristics

The government redistributes income using taxes, τ_t and transfers, α_t. The government furthermore keeps a balanced budget in every period, i.e. any expenses in form of subsidies in one period have to be offset by taxes in the very same period (a "pay-as-you-go" rule).[4] This balanced budget constraint is given by:

$$\sum_{i \in \nu} \tau_t^i = \sum_{i \in \nu} \alpha_t^i \qquad \forall \quad t = 0, 1 \tag{5.6}$$

$$h_0^i = \underline{h}^0$$

$$\forall i \in \nu$$

The effects of this redistribution scheme on overall welfare and inequality will be analyzed within this class of policies, or its analogue for education finance.

5.2.4 Steady State Analysis

In order to determine the behavior of the agents in the model, the maximization problem of households is extended in order to include the government budget constraint. Put differently, households are now taking the balanced budget constraint of the government into account when solving their maximization problem. The problem of the household then becomes:

[4]The robustness section relaxes this assumptions and provides the government with an option to balance the budget inter temporally

$$\max_{c_t^i} \mathcal{U} = lnc_0^i + \beta lnc_1^i$$

$$s.t.$$

$$y_t^i = h_t^i * \overline{l}_t^i \qquad \forall \quad t = 0,1 \tag{5.7}$$

$$\hat{y}_t^i = y_t^i - \tau_t^i \qquad \forall \quad t = 0,1 \tag{5.8}$$

$$\hat{y}_t^i = c_t^i + e_t^i - \alpha_t^i \qquad \forall \quad t = 0 \tag{5.9}$$

$$\hat{y}_t^i = c_t^i \qquad \forall \quad t = 1 \tag{5.10}$$

$$h_{t+1}^i = h_t^i + e_t^i(1+\psi) + \gamma \sum_{i\in\nu, j\neq i} g_{ij} e_t^i e_t^j \qquad \forall \quad t = 0 \tag{5.11}$$

$$\sum_{i\in\nu} \tau_t^i = \sum_{i\in\nu} \alpha_t^i \qquad \forall \quad t = 0,1 \tag{5.12}$$

$$h_0^i = \underline{h}^0$$

$$\forall i \in \nu$$

The first order condition for this maximization problem is then given by:

$$h_1^i = h_0^i(2+\psi+\theta)\frac{\beta}{1+\beta} - \mu_0^j[1+\psi+\theta]\frac{2+\beta}{1+\beta} + \frac{\kappa}{1+\beta} \tag{5.13}$$

where:

$$\rho_0^j = [h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$

$$\mu_0^j = [\sum_{i\in\nu\setminus i} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i] = \tau_0^i - \alpha_0^i$$

$$\theta = \gamma \sum_{i\in\nu, j\neq i} g_{ij}\rho_0^j$$

$$\kappa = \sum_{i\in\nu} \alpha_1^i - \sum_{i\in\nu\setminus i} \tau_1^i$$

$$\zeta_0^j = \gamma \sum_{i\in\nu, j\neq i} g_{ij}(h_0^j - c_0^j)$$

That is, each individual acts according to this 'best-response'-function in each period in order to maximize their utility given the exogenous variables and the choices of the other agents.

In equation 5.13, if we set taxes equal to zero we obtain:

$$h_1^i = h_0^i(2 + \psi + \zeta_0^j)[\frac{\beta}{(1+\beta)}]$$

For this simplified version, h_1^i increases for a higher initially given h_0^i, a higher return on investment, and a higher β. A higher discount factor β implies that the individual is less impatient and thus willing to shift more income into the next period.

Note that this is a simultaneous move n-player game with payoffs according to the utility function and strategy spaces $\mathbb{R}_+$. Thus, solving the model for $n \geq 2$ will require a computational approach in order to determine possible outcomes. To do so, I use an agent based model that consists of n-agents which are initialized with with random values for h_1 and previously determined values of the exogenous variables (which will depend on the type of policy experiment). Afterwards, the agents sequentially begin to update their value for h_1 based on their connections in the network and continue to do so each round. Put differently, the agents are put in a random order and then the first agent determines whether the value of h_1 is in optimal based on the best response function and the values for h_1 of the connected agents. If this is not the case, the value is corrected. Then the next agent goes through the same process and updates h_1 if necessary. This process is repeated until no individual changes their value for h_1 any more, i.e. a steady state is achieved in which no agent has an incentive to change their behaviour unless another agent does so (comparable to a Nash Equilibrium). Below the

Algorithm 2 conveys the method of implementation:

```
create agents;
create network;
for t = 0 to T = 50 do
    for all agents i = 0 to n do
        set h_1 according to best response function
    end
end
if transfer policy then redistribute according to specific policy;
for t = 0 to T = 50 do
    for all agents i = 0 to n do
        set h_1 according to best response function
    end
end
```

Algorithm 2: Pseudocode of model implementation

Figure 5.1 illustrates how the model converges towards such a steady state for a network of 10 agents. After a an initial burn in phase in which agents observe the behavior of their peers and react based on their best-response function. As this process is repeated, the adjustments become smaller and converge towards zero where all agents reach a value of h_1 at which they do not have an incentive to change their behavior unless other agents do so. For $n = 10$ this steady state is typically reached within 10 time steps.

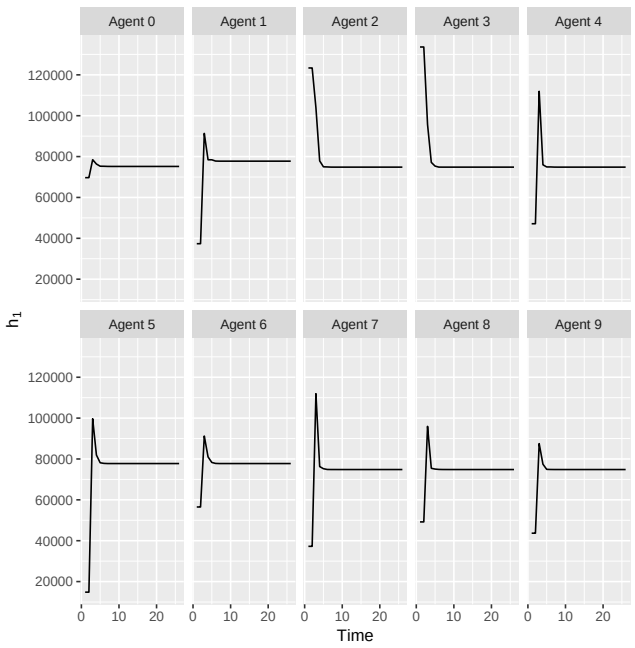

FIGURE 5.1: Equilibrium Values of h_1

In order to determine the uniqueness and stability of this steady state I use two approaches. I initiate each agent with a random value for h_1 and ensure that for the same network, the same equilibrium is reached. Additionally, once the equilibrium is achieved, I perturb the agents by readjusting their h_1 to a random value. The appendix contains further details on the robustness of the steady state. Once the group of agents arrives at a steady state, the policy experiments are introduced. The policy experiments set in motion a similar adjustment process in which the agents sequentially re-adjust their values for h_1. I report the results from the final time step when the steady state is reached again.

5.3 Experiments and Results

In this section, the computational results are presented. I first explain the basic results from a simplified model that consists only of 2 agents. This serves the purpose of describing the underlying model in more detail and also helps to unfold the underlying mechanisms of the model. When the model is extended to 2 agents agents with unequal endowments of initial human capital, the severity of inequality can be alleviated very simply with a transfer from the high endowment agent to the low endowment agent. The following section then examines the effects of such a transfer policy on overall utility and inequality as the network agents increases.

5.3.1 Baseline Model

5.3.1.1 Baseline Model with 2 Agents

In order to get an intuition for the basic model, we look into the outcome of a simplified model consisting of two identical agents with income in period 1 of $y_0 = \$50.000$,[5] which implies human capital $h_0 = 50,000$. The parameters β and ψ are set to $\beta = 0.99$, $\psi = 0.5$, and we dispense with the peer effect initially, that is $= \gamma = 0$. We then obtain the values for the endogenous variables in the steady state as summarized in the first line of Table 5.1.

γ	h_0	c_0	e_0	h_1	c_1	α_0	τ_0	Mean U	Aggr. U
0	50,000	41,876	8,124	62,186	62,186	0	0	21.68	43.36
0.00001	50,000	40,920	9,080	64,445	64,445	0	0	21.69	43.39

Table 5.1: Peer Effects with 2 Agents

[5]Comparable to the median household income in the US for 2014 \$53,657, Source: Federal Reserve Bank of St Louis

We can see that the individual consumes \$41,876 of her income in the first period and invests \$8,124 into her education. She then obtains total consumption equal to $c_0 + c_1 =\$104,062$. Thus, the investment into education yields a 'return on investment' for education of \$4,062 compared to the counterfactual event in which the individual would not have invested any of its income in period $t = 0$ and consumed \$50,000 in both periods (notice that this 'return on investment' is equal to the value indicated by $\psi = 0.5$). The last two columns contain the mean utility and aggregate utility respectively, and this will be used to compare the overall welfare effects of policies further below.

Next, we examine the effect of peer effects, i.e. we set $\gamma = 0.00001$ in the second line of table 5.1. Both agents benefit from being connected to each other now, as both benefit from the others educational investment. This can be explained by the effect of having $\gamma \neq 0$ in Equation 5.11. This *complementarity effect* causes an additional increase in the 'return on investment' (in addition to ψ), rendering an increase in e_0 more attractive (and subsequently making up for the foregone consumption today). Consumption in period $t = 0$, c_0 is reduced to \$40,920 (a 2.28% decrease compared to the model without peer effects) while educational investment, e_0 thus increases to \$9080 (a 11.77% increase compared to the model without peer effects). Notice that this causes an increase in aggregate utility from 43.36 to 43.39. The *complementarity effect* that is demonstrated here will be explored further in the following subsections.

5.3.1.2 Inequality

We now examine the effect of inequality by changing the initial endowment of both agents to $h_0 = 10,000$ for agent 1 and $h_0 = 90,000$ for agent 2 (and leaving $\gamma = 0.00001$ from here on), which leaves the mean endowment unchanged:

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U
50,000	40,920	9,080	64,445	64,445	0	0	21.69	43.39
10,000	8,071	1,929	13,184	13,184	0	0	18.48	41.34

Table 5.2: Inequality with 2 Agents- Low Endowment Agent

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U
50,000	40,920	9,080	64,445	64,445	0	0	21.69	43.39
90,000	74,994	15,006	112,799	112,799	0	0	22.86	41.34

Table 5.3: Inequality with 2 Agents- High Endowment Agent

We can see that the overall utility decreases from 43.39 to 41.34, which is caused by the decreasing marginal utility feature of the utility function. The re-distribution of

initial endowment causes an increase in utility for the high endowment agent which does not offset the simultaneous decrease in utility for the low endowment agent, i.e. we have a 5.4% increase in utility for the former and a 14.8% decrease in utility for the latter (put differently, while the increase and decrease in endowments offset each other with the total endowment remaining constant, the overall utility decreases). Given this decrease in overall utility introduced by inequality we next look into a possible remedy through transfer system, i.e. a tax on the high endowment agent and a subsidy to the low endowment agent.

5.3.1.3 Subsidies and Taxes

We now examine what happens if we introduce a transfer policy that provides the low endowment agent with a subsidy $\alpha_0 = 20,000$ and finances that transfer via taxing the high endowment agent, i.e. $\tau_0 = 20,000$.

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U
10,000	8,071	1,929	13,184	13,184	0	0	18.48	41.34
10,000	18,279	11,721	91,129	91,129	20,000	0	21.23	42.79

Table 5.4: Subsidides with 2 Agents- Low Endowment Agent

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U
90,000	74,994	15,006	112,799	112,799	0	0	22.86	41.34
90,000	63,141	6,859	36,403	36,403	0	20,000	21.56	42.79

Table 5.5: Subsidides with 2 Agents- High Endowment Agent

When we introduce a subsidy/tax, we observe an increase in aggregate utility, due to the decreasing marginal utility of the utility functions (the decrease in utility for the high endowment agents is outweighed by the increase in utility of the low endowment agents). Thus, the transfer reverses the process from the previous section to a small degree and alleviates the burden on the low endowment agent. The transfer policy for a two agent model is comparatively straightforward in the sense that the intricacies introduced through a network or connections are dispensed with. In the next sections, these simplifying assumptions will be gradually relaxed to allow for a more complex network structure and thereby allow for an analysis of peer effects of such a transfer policy.

5.3.2 Networks with 10 Agents

We now increase the size of the network to 10 agents (with $k = 2$ connections for each agent). In order to demonstrate the functioning of the model, Figure 5.2 depicts a random network with 10 agents, i.e. each agent randomly selects 2 connections from array of available connections.

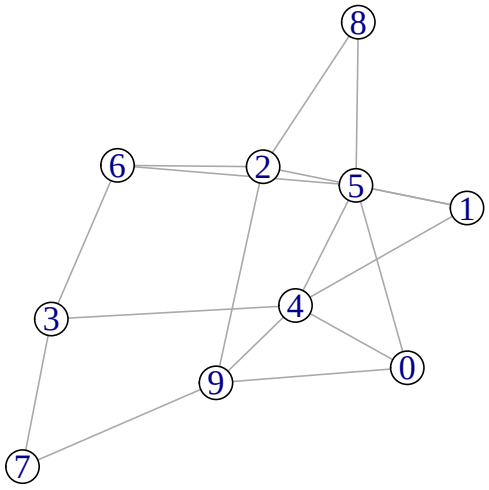

FIGURE 5.2: Random Network Graph with 10 agents

All of the agents in the network are identical regarding their initial endowment of h_0. The only difference in outcomes for h_1 is thus caused by their position in the network and the resulting measure of centrality. Due to the fact that connections are reciprocal we observe a strong variation in the actual number of connections agents end up with. For instance, while agent 7 has indeed a degree of 2, we find agent 4 on the other end of the degree spectrum with 5 connections. Table 5.6 contains a number of centrality measures for each agent for the network displayed above. These centrality measures are defined as follows:

- Degree centrality is defined as the number of connections an agent i has:

$$Degree(i) = \sum_{j=1}^{n} g_{ij}$$

- Closeness centrality measures the distance between an agent and all other agents in the network, whereas the distance is defined by the length of the shortest path between any two agents. Closeness centrality for an agent i then simply is the reciprocal of the sum of all those distances:

$$closeness(i) = \frac{1}{\sum_{j=1}^{n} dist(i,j)}$$

- Betweenness centrality quantifies the number of times that an agent lies on the shortest path between two other agents. For agent i (with $i \neq j$, $i \neq k$, and $j \neq k$) betweenness is then defined as:

$$betweenness(i) = \sum_{i,j,k \in \nu} \frac{p_{jk}(i)}{p_{jk}}$$

  where $p_{jk}(i)$ is the number of shortest paths between j and k which go through i and p_{jk} is the total number of shortest paths between j and k.

- Eigenvector centrality measures the influence of an agent in a network through measures how well that agent is connected to other well connected agents. For a given adjacency matrix $\mathbf{G}$, the Eigenvector centrality of an agent i is defined as:

$$ev(i) = \frac{1}{\lambda} \sum_{j=1}^{n} g_{ij} ev(j)$$

  or in Matrix form:

$$\lambda x = xG$$

  where x is the left-hand Eigenvector of the adjacency Matrix $\mathbf{G}$ and $\lambda \neq 0$ is the largest eigenvalue in absolute value of the adjacency matrix $\mathbf{G}$.[6]

- Bonacich centrality[7] is an extension of the Eigenvector centrality that includes an "attenuation factor" which allows to indicate the type of influence an individual exerts on others. An individual that is connected to other well connected agents is very central but not very influential, since those connections are themselves very well connected. On the other hand an individual that is well connected to other individuals that are not well connected themselves, is less central than the previous

[6] By virtue of the Perron-Frobenius theorem, choosing the largest eigenvalue ensures the Eigenvector solution is both unique and positive for a a given square matrix $\mathbf{G}$ that is irreducible and non-negative.

[7] see Bonacich (1987)

example, but more influential. This centrality measure can be expressed as:

$$bonacich(i) = \frac{1}{\lambda} \sum_{k=0}^{\infty} \beta^k G_{tr}^k$$

where k is the walk length and G_{tr}^k transformed adjacency matrix.

- Intercentrality as proposed by Ballester et al. (2006) extends Bonacich centrality by measuring each agent i's centrality and their contribution to every other agents' Bonacich centrality. It is formally defined as:

$$intercentrality(i) = \frac{bonacich(i)^2}{m_{ii}(G)}$$

where $m_{ii}(G)$ is an element of the matrix $M(G) = [I - G]^{-1}$.

Depending on the specific position in the network, even agents with the identical Degree, could still differ with regard to their other centrality measures. For instance, Agents 1 and 6 both have a Degree of 3, however their betweenness centrality is very different (Agent 1 has a betweenness of 0.05 and Agent 6 has a betweenness of 0.17.

Table 5.6 also displays the endogenous value of h_1 for this specific network. A comparison of the value for h_1 reveals that the agent 4 who has the highest value of \$76,892 is also the one with the highest Degree centrality. Figure 5.2 also displays this central role that agent 4 plays in this network. The fact that a higher Degree causes a higher h_1 follows directly from Equation 5.11, where h_1 for each agent is a function of the sum of the respective connections of the agent and their h_1. Thus, those agents with more connections will necessarily end up with a higher value for h_1.

	A 0	A 1	A 2	A 3	A 4	A 5	A 6	A 7	A 8	A 9
H_1	71,574	71,553	73,494	70,795	76,892	76,540	71,156	67,999	68,354	73,815
Degree	3.00	3.00	4.00	3.00	5.00	5.00	3.00	2.00	2.00	4.00
Closeness	0.20	0.18	0.21	0.18	0.23	0.21	0.20	0.15	0.16	0.21
Betweenness	0.05	0.05	0.35	0.18	0.37	0.38	0.17	0.03	0.02	0.38
Eigenvector	0.20	0.20	0.19	0.14	0.27	0.26	0.16	0.10	0.12	0.21
Bonacich	0.19	0.19	0.20	0.19	0.23	0.22	0.19	0.17	0.17	0.21
Intercentrality	0.18	0.18	0.20	0.16	0.27	0.26	0.17	0.11	0.12	0.21

Table 5.6: Centralities of a 10 Agent Network

While the Degree centrality is able to explain the increase in h_1 relative to agents with lower Degrees, it remains incumbent to explain the differences between agents with similar Degrees. Table 5.6 gives some indication of the high correlation between other centrality measures and the level of h_1.

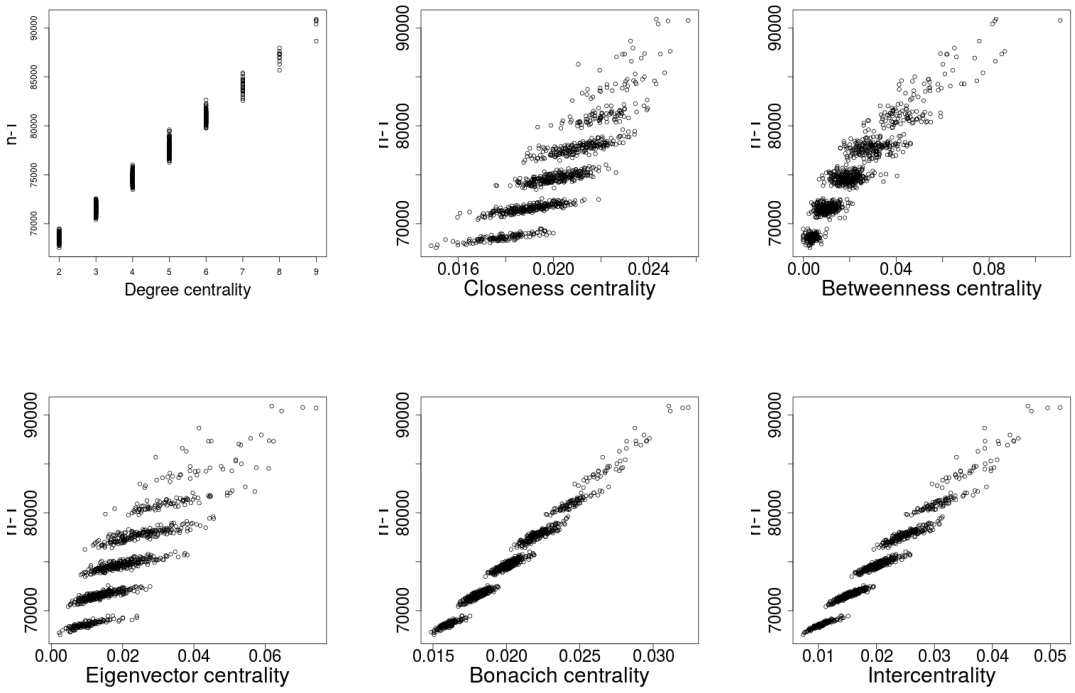

FIGURE 5.3: Correlation of centrality Measures with h_1

In order to examine this characteristic further, Figure 5.3 plots the correlation between h_1 and the various centrality measures for a larger network consisting of 100 agents. While the relationship is not perfectly linear a positive trend becomes visible and demonstrates the role of the other centrality measures in determining the value of h_1. Given that the emergence of the network position of each agent in the model is determined entirely by a random process in the beginning, we can exclude any influence of h_1 on the network centrality of the agent. We thus know that, not only can the variation in h_1 be attributed to a large degree to the position of the individual in the network, but we can also infer that a high centrality by every measure used here is indicative of a high value of h_1.

These results are indicative of the role of centrality in explaining the activity level of an individual in the network and are in line with Ballester et al. (2006) and Ballester et al. (2010) who use a crime network model where each agent's effort level depends on the effort level of their contacts in the network. They demonstrate that the effort level is proportional to the Katz-Bonacich centrality of each agent. Calvo-Armengol et al. (2009b) studies how centrality in friendship networks affects educational outcomes and

find a similar mapping between Katz-Bonacich centrality of each agent and the effort levels in education. Using the unique dataset of friendship networks from the National Longitudinal Survey of Adolescent Health (Add Health) they find that after controlling for observable individual characteristics and unobservable network specific factors, the individual's position in a network (as measured by her Katz-Bonacich centrality) is a key determinant of her level of activity. Each of these approaches purports to estimate the ability to predict the activity level of individual based on their centrality in their network.

We now return to the previous case of two separate groups of agents, i.e. the group of agents is now divided into two groups of agents of equal size ($n_A = n_B = 5$) with agents within each group being identical except for their position in the network. The human capital distribution from before remains, i.e. we have:

- $h_0^A = 10,000$
- $h_0^B = 90,000$

Apart from the initial endowment the agents across groups are identical and their equilibrium values of endogenous variables are entirely determined by their position in the network. The networks are created entirely randomly as before, i.e. each agent chooses k connections out of the nine possible agents to connect with (connections are undirected).[8] Again, some agents then end up with more connections than k, as they are chosen randomly by more agents to connect with. Figure 5.11 illustrates a representative network consisting of these two groups.

Given the random nature of network creation there is a large number of possible network realizations for a size of 10.[9] Given that the individual outcomes of h_1 will entirely depend on this network structure there a great variety of possible aggregates of h_1. In order to evaluate policies, I thus compare average outcomes for 100 networks for each setting.

The policy from the previous section is now slightly adjusted in the sense that the subsidy is provided to one agent from the low endowment group and then equally paid for via taxing the high endowment agents by an equal share. This "shock" or "treatment" is then propagated through the network as agents who are connected to this particular agent have to re-adjust their initial educational investments. The recipient of the subsidy is chosen randomly initially. Later on, in order to determine the importance of

[8] In the robustness section alternative ways of creating the network such as homophily and clustering are discussed

[9] for a network size of $n = 10$, there could be $2^{(n(n-1)/2)} = 2^{45}$ possible networks.

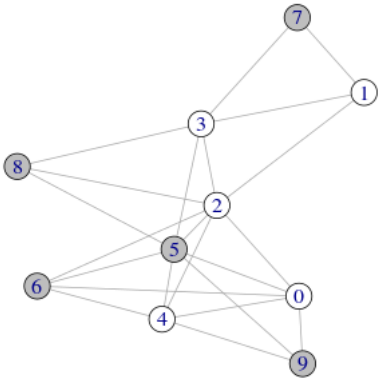

FIGURE 5.4: Random Network Graph with 50 Agents of Each Type

various centrality measures when choosing an initial 'key player', I compare the resulting aggregate utility levels and inequality before and after the introduction of the policy.

For a given network, I introduce the policy after 50 time steps when the equilibrium is reached and then allow for another 50 time steps for all agents to reach their equilibrium state. This allows for ample time to reach a stable equilibrium. I reiterate this process for 100 networks for each parameter setting and report averages.

5.3.2.1 Inequality

The distribution under inequality is similar to the previous case, i.e. the low endowment group has an initial human capital endowment of \$10,000 while the high endowment group starts with an initial human capital, h_0, endowment of \$50,000. This yields a Gini-Coefficient of $Gini = 0.41$ based on c_1. Notice that in the case of the equal distribution, the Gini is not equal to 0 because even though all agents start at the same human capital endowment, small differences arise due to their position in the network which yields different outcomes for c_1.

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U	Gini
50,000	38,230	11,770	72,783	72,783	0	0	108.73	217.45	0.02
10,000	7,657	2,343	14,569	14,569	0	0	92.64	207.23	0.41

TABLE 5.7: Inequality with 10 Agents- Low Endowment

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U	Gini
50,000	38,230	11,770	72,783	72,783	0	0	108.73	217.45	0.02
90,000	69,199	20,801	130,206	130,206	0	0	114.59	207.23	0.41

TABLE 5.8: Inequality with 10 Agents- High Endowment

As before, we can see that the overall utility decreases, caused by the decreasing marginal utility. We will examine in the next section analogously to the previous section the effect of a redistribution process.

5.3.2.2 Subsidies with 10 Agents

We can see that again as in the case with two agents aggregate utility increases through the subsidy. Not only is the average utility higher, but utility is actually higher in each one of the simulation runs. We can also observe that the inequality decreases to $Gini = 0.33$.

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U	Gini
10,000	7,657	2,343	14,569	14,569	0	0	92.64	207.23	0.41
10,000	9,663	4,337	33,501	33,501	4,000	0	95.49	209.15	0.33

TABLE 5.9: Subsidies with 10 Agents Low Endowment

h_0	c_0	e_0	h_1	c_1	α_0	τ_0	U	Aggr. U	Gini
90,000	69,199	20,801	130,206	130,206	0	0	114.59	207.23	0.41
90,000	67,115	18,885	111,209	111,209	0	4,000	113.66	209.15	0.33

TABLE 5.10: Subsidies with 10 Agents High Endowment

As in the case of two agents, it was possible to alleviate the effects of inequality through the transfer system. However, the increased size of the network raises additional questions regarding a more efficient way to provide the transfer. Until now the recipient of the subsidy was randomly selected and then depending on the position of the agent, this caused stronger or weaker peer effects as the 'shock' was propagated through the network. The next section will explore whether a policy maker could take advantage of the knowledge of the network position of individuals and what effects this will have on overall utility and inequality.

5.3.3 Key Player Policy

Until now, the educational subsidy was provided to a randomly chosen agent from the low endowment group. While this choice is interpreted as random given the limited characteristics that are available in the model, it could be interpreted as the real life equivalent of a choice based on alternative characteristics of the agents. Put differently, from the perspective of the position of agents in the networks the allocation appears random as it does not take centrality into account. This does not obviate however, that the choice was based on other factors that are outside of the scope of this model. While this increased overall utility and reduced inequality, the policy did not take into account the position of individuals in the network of agents and the resulting differences in peer effects. In this section, I examine how the result of such a policy can change if the underlying network of individuals is taken into account. The motivation for this is twofold. On the one hand, it is obviously in the interest of the policy maker to increase the social multiplier effect as the overall benefit of the group might be increased while keeping the costs constant. For this reason, it might appear straightforward to simply make the most central individuals the recipients of the subsidy and thereby increase and potential peer effect (since more central individuals would appear to be more influential and thereby affect more agents in the network compared to less central recipients). On the other hand, since the inequality is also a concern to the policy maker (which is why the transfer is introduced in the first place), it is not quite clear, which policy should be chosen to reduce the inequality between individuals. Figure 5.11 from the previous section displays the nature of the problem of focusing on central agents. We exacerbate the concentration of activity levels that occur at the very core of the network while neglecting the activity levels of agents in the periphery of the network. Under the goal of inequality it might appear more sensible to choose less central individuals to avoid a concentration of the educational subsidies within a very central group while leaving behind the peripheral agents in the network. For this reason, we now compare the effectiveness of the policy when we use different characteristics to provide the subsidy such as the centrality of potential candidates.

5.3.3.1 No Policy Vs Random Policy with 100 Agents

For ease of comparison I first report the same comparisons as in the previous Section 5.3.2.2. I first compare the random allocation of the subsidy to 10 agents to the case of no subsidy. Introducing the policy increases utility significantly, since any redistributive policy will yield an overall increase of utility due to the fact that the utility function features decreasing marginal utility. We also see in the right panel of Figure 5.5 how inequality decreases. These are the same results as in Tables 5.9 and 5.10.

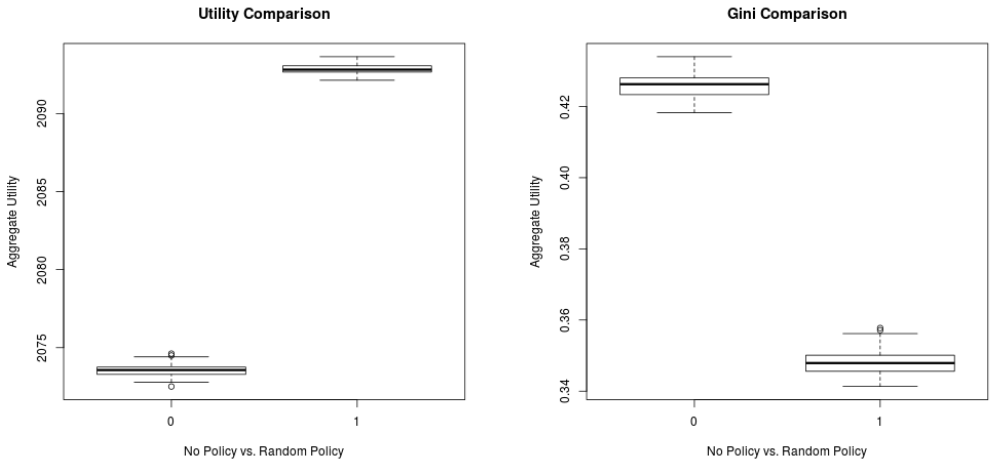

FIGURE 5.5: Overall Utility and Gini- Comparing no policy to random policy

Figure 5.6 then shows how utilities of both groups develop differently, as the group A observes an increase in utility while utility for group B decreases through the policy.

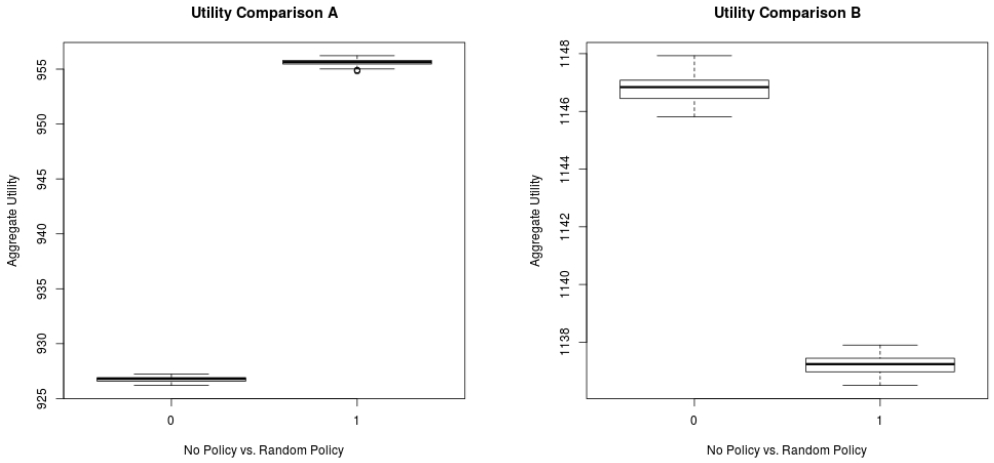

FIGURE 5.6: Utility Group A and Group B- Comparing no policy to random policy

5.3.3.2 Random Policy Vs Central Policy with 100 Agents

I next compare the utility levels from before to a policy where agents are chosen based on their centrality in the network. Specifically, all potential recipients of the subsidy from group A are ranked based on their centrality, and then the agent with the respective highest measure is chosen. I employ the same 6 centrality measures as before.

The results obtained from the various centrality measures are very similar as the rankings based on centrality only differ slightly. This is caused by the network structure which results in an overlap in the ordering according the centrality measures (e.g. the agent who has the highest Degree centrality also happens to have the highest Closeness centrality). For this reason, I do not distinguish between results among the various measures but rather use averages over all measures instead.

I create 100 different networks and for each network I choose the recipients based on each of the centrality measures listed above and randomly. Then, I take the average utility values that result from the centrality measures and compare those to the utility levels that result from the random choice.

Figure 5.7 shows how utility is actually increasing more under the centrally chosen policy (all reported changes are significant at the 1% level unless reported otherwise). However, this comes at the cost of a higher inequality as shown by the higher Gini coefficient in the right panel of Figure 5.7. Therefore, we can see that providing the subsidy to the most central agent in the network causes an increase in overall utility as this increase in educational investment disseminates through the network and causes a stronger effect on the peers of the receiving individual. However, this focus on the central agents in the network also causes an increase in the inequality of the network. This results appears counter-intuitive as the overall goal of the subsidy policy was to reduce inequality through a transfer from the high endowment group to the low endowment group. It would thus appear to be serving both goals to choose a more central recipient of the transfer, as this would cause a higher impact of the policy. However, as more central agents have a higher educational investment level other things being equal, they already are ahead of their peers in terms of income. Choosing them as recipients then reinforces this status quo and meanwhile neglects agents in the periphery of the network with lower incomes.

Figure 5.8 shows how the central policy yields higher utility levels for the receiving agents while yielding slighty higher (but not significantly higher) results for group B. These findings indicate that a policy that is targeted at the most central agents can indeed yield a higher dissemination of the intended policy measure but will also neglect individuals who are not very prone to be influential on their peers. The possibility for

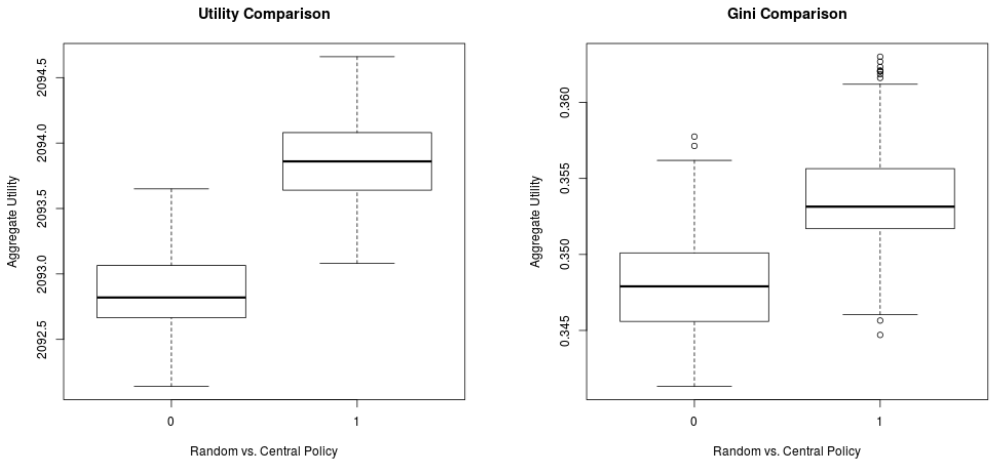

FIGURE 5.7: Utility and Gini Coefficient

such a detrimental effect suggests that such a policy would require a very careful design in order to avoid negative outcomes.

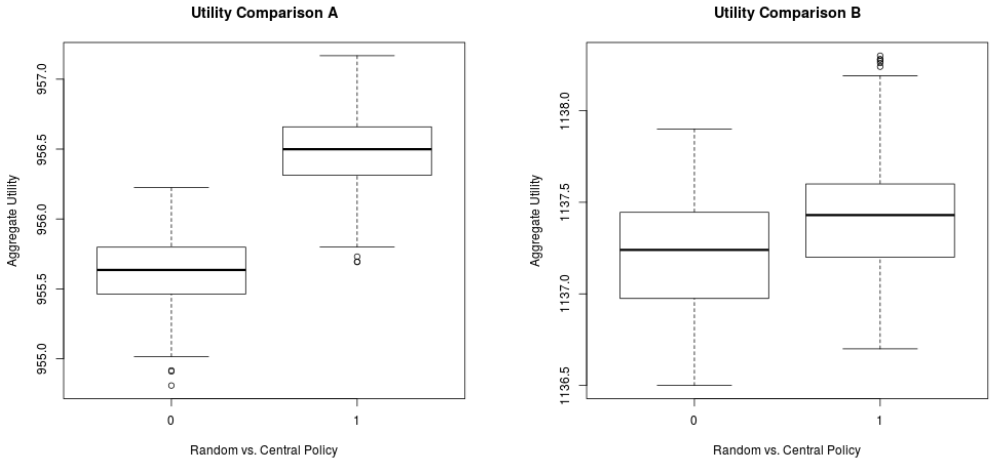

FIGURE 5.8: Utility for Group A and Group B

5.4 Sensitivity Analysis

In order to test for the robustness of the obtained results, the main findings are subjected to a sensitivity analysis. In particular, the parameters γ, β, and ψ are varied systematically within meaningful boundaries (e.g. $\beta > 1$ is excluded from the analysis), by first increasing the parameters either by 10% and 20% and then decreasing it by the same increments. Furthermore, I vary the number of connections each agent is allowed to have, that is parameter k, by allowing up to 10 connections in increments of 2. Then, instead of allowing each agent to have the same number of connections, I relax this assumption and introduced a degree distribution (such as Power-law degree distributions) that match more closely what can be observed empirically. Specifically, for each individual, k_i is drawn from a lognormal distribution that allows to create a degree distribution with very few agents having a very high Degree while the majority of agents has only a small number of connections. Figure 5.9 depicts the normalized Degree densities for the case $k = k_i \ \forall i$ and the lognormal distribution of k_i.

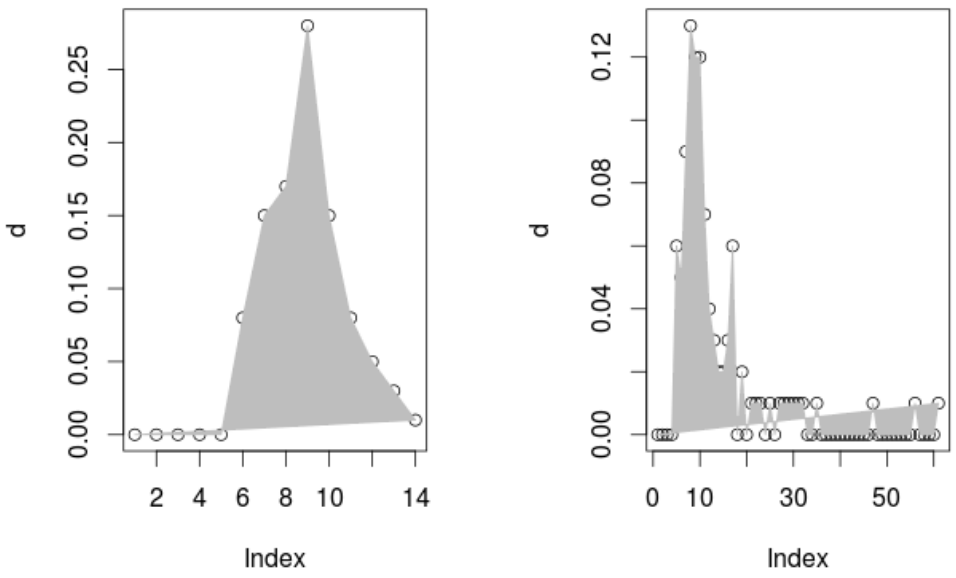

FIGURE 5.9: Kernel Density of Degree Distribution

In order to emulate empirically observed networks[10], I allow for individuals to choose their connections based on preferences such as homophily, where individuals prefer connections to agents of their own type and clustering, where agents choose a certain fraction of their friends randomly and a certain fraction based on their friend's connections. Both

[10]The National Longitudinal Study of Adolescent to Adult Health (Add Health) dataset for instance provides a nationally representative sample of high school student networks

of these mechanisms allow to create networks that match more closely the empirically observed ones. Alternatively, individuals have a certain preference to connect to their own group, i.e. homophily is introduced. A third specification allows individuals to choose a certain fraction of friends randomly and then lets them pick additional connections from the list of their friends' friends (this allows for clustering). These modifications of the initial random graph, allow to introduce more realistic graphs as we observe them in reality, specifically with regard to the degree distributions, as it allows for fat tails. All two graphs are plotted below for one simulation run.

FIGURE 5.10: Network Graph with 50 Agents and Homophily

Additionally, some of the assumptions made in the model regarding the government budget are relaxed, such that the government budget only has to be balanced over both periods, i.e. restriction 5.8 is replaced with:

$$\sum_{i \in \nu} \tau_0^i + \tau_1^i = \sum_{i \in \nu} \alpha_0^i + \alpha_1^i \tag{5.14}$$

$$\tag{5.15}$$

This allows to raise the income to finance subsidies in the next period after the subsidies are distributed.

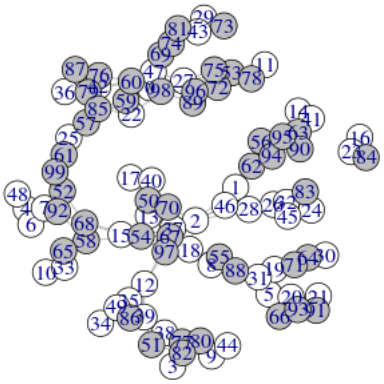

FIGURE 5.11: Network Graph with 50 Agents with Clustering

In all of the performed variations the main findings did not change qualitatively. Another variation included targeting the least central agents. In this setting the overall utility that resulted from the policy was lower compared to the random allocation and the 'centralized' allocation. This finding confirms the previous findings that a centralized allocation can increase overall welfare while an allocation to less central agents yields lower spill over effects. In terms of the effect inequality, the effect was not significantly different from the random allocation. Furthermore, a policy targeted at the median agent yielded results very close to the random allocation in terms of overall welfare and inequality.

In order to ensure the stability of the equilibrium, for each simulation run I initiate the endogenous variable h_1 with random numbers to determine the uniqueness of equilibrium values. In addition, to determine whether the system was stable, I perturbed the model after a stable value was reached with random vales for h_1. Figure 5.12 displays such a process for one run and how the h_1 for each agent returned to their equilibrium again. The model returned to the initial steady state in all of these cases.

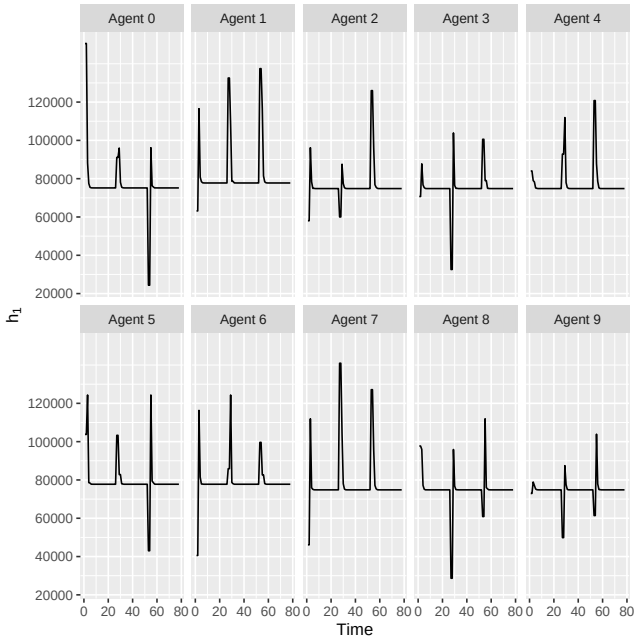

FIGURE 5.12: Perturbation of Equilibrium

5.5 Conclusion

The labor economics literature has by now established the paramount role of social networks in determining the labor market outcomes of individuals. To a large extent the empirical literature provides a number of natural experiments that identify significant network effects on labor market outcomes (see for instance Bayer et al. (2008), Beaman (2012), Kramarz and Skans (2010), Munshi (2003), Munshi and Rosenzweig (2006)).

While these analyses have provided a great deal of evidence, these type of interaction effects are not very prominent when it comes to labor market policies. From a policy maker's perspective the mere existence of such neighborhood effects merit some attention as these regional shocks or policy interventions leads to spill-over effects on non-treated individuals, i.e. social multipliers arise.

The dissemination of job information is often investigated as one of the primary aspects of social networks. Other aspects such as the peer influence play a similarly important role. In this paper I demonstrate how the knowledge of the existence of such effects can be taken into account when deciding about the provision of educational trainings in the hope to increase the potential multiplier effects to other non-treated individuals in the network. I investigate the neighborhood effects that emerge by providing individuals within a network with such a labor market training subsidy that allows them to invest into their human capital.

Using a unique approach of re-distribution of funds within a group of agents, I analyze the effects of transfer policy that chooses agents not randomly but based on their centrality within the social network. The goal of this approach is to determine how the overall utility and the inequality within the network would be affected by the policy. Focusing on both aspects allows us to gauge effects that cannot be trivially determined a priori. The computational results indicate that it is indeed possible for policy interventions to exert social multiplier effects, i.e. to achieve a significant increase in the overall welfare by making very central individuals the recipients of the policy. However, this type of targeting of central individuals introduces a trade-off regarding the inequality, i.e. a centrally allocated subsidy increases inequality significantly compared to a random allocation of the subsidy. This is caused by the fact that central agents are already characterized by a very high level of educational attainment and by a targeting them, this relative advantage is exacerbated. These findings suggest that any policy measure that targets central individuals does not unequivocally yield positive aggregate outcomes. Rather, it would need to be carefully designed in order to avoid such negative outcomes regarding inequality.

5.6 Appendix

5.6.1 Deriving the Best-Response Function

The model is reproduced for convenience:

$$\max_{c_t^i} \mathcal{U} = lnc_0^i + \beta lnc_1^i$$

$$s.t.$$

$$y_t^i = h_t^i * \bar{l}_t^i \qquad \forall \quad t = 0, 1 \tag{5.16}$$

$$\hat{y}_t^i = y_t^i - \tau_t^i \qquad \forall \quad t = 0, 1 \tag{5.17}$$

$$\hat{y}_t^i = c_t^i + e_t^i - \alpha_t^i \qquad \forall \quad t = 0 \tag{5.18}$$

$$\hat{y}_t^i = c_t^i \qquad \forall \quad t = 1 \tag{5.19}$$

$$h_{t+1}^i = h_t^i + e_t^i(1+\psi) + \gamma \sum_{j \in \nu, j \neq i} g_{ij} e_t^i e_t^j \qquad \forall \quad t = 0 \tag{5.20}$$

$$\sum_{i \in \nu} \tau_t^i = \sum_{i \in \nu} \alpha_t^i \qquad \forall \quad t = 0, 1 \tag{5.21}$$

$$h_0^i = \underline{h}^0$$

$$\forall i \in \nu$$

Setting $\bar{l}_t^i = 1$, we substitute constraint 5.16 into constraint 5.17:

$$\hat{y}_t^i = h_t^i - \tau_t^i \qquad \forall \quad t = 0, 1 \tag{5.22}$$

Then we substitute this expression into constraint 5.18:

$$h_t^i - \tau_t^i = c_t^i + e_t^i - \alpha_t^i \qquad \forall \quad t = 0 \tag{5.23}$$

Now we solve constraint 5.23 for e_t^i in $t = 0$:

$$e_0^i = h_0^i - \tau_0^i - c_0^i + \alpha_0^i \tag{5.24}$$

and in $t = 1$

$$e_1^i = 0 \tag{5.25}$$

Substituting this into the transition equation in order to replace e_0^i (and e_0^j) from the original problem yields:

$$h_1^i = h_0^i + e_0^i(1+\psi) + \gamma \sum_{j\in\nu, j\neq i} g_{ij} e_0^i e_0^j$$
$$h_1^i = h_0^i + [h_0^i - \tau_0^i - c_0^i + \alpha_0^i](1+\psi) + \gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^i - \tau_0^i - c_0^i + \alpha_0^i][h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$

Now the problem becomes:[11]

$$\max_{c_t^i} \mathcal{U} = lnc_0^i + \beta ln[h_1^i - \tau_1^i]$$
$$s.t.$$
$$h_1^i = h_0^i + [h_0^i - \tau_0^i - c_0^i + \alpha_0^i](1+\psi) + \gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^i - \tau_0^i - c_0^i + \alpha_0^i][h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$
$$\sum_{i\in\nu} \tau_t^i = \sum_{i\in\nu} \alpha_t^i \qquad \forall \quad t = 0, 1$$
$$h_0^i = \underline{h}^0$$
$$\forall i = 1, ..., n.$$

We now consider the constraint that requires that the government budget constraint to be balanced at all times, i.e.:

$$\sum_{i\in\nu} \tau_t^i = \sum_{i\in\nu} \alpha_t^i \qquad \forall \quad t = 0, 1$$

[11] To avoid any lack of generality the superscript i is not omitted at this point from any parameter or variable

We now solve the constraint for τ_t^i (in the subsequent analysis we are going to solve the problem from the perspective of agent $i \in \nu$

$$\tau_t^i = \sum_{i\in\nu} \alpha_t^i - \sum_{i\in\nu\setminus i} \tau_t^i \qquad \forall \quad t = 0, 1$$

We can now insert τ_t^iinto the transition equation of agent i, i.e.:

$$h_1^i = h_0^i + [h_0^i - \tau_0^i - c_0^i + \alpha_0^i](1+\psi) + \gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^i - \tau_0^i - c_0^i + \alpha_0^i][h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$

$$h_1^i = h_0^i + h_0^i(1+\psi) - [\sum_{i\in\nu} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i](1+\psi) - c_0^i(1+\psi) + \alpha_0^i(1+\psi)$$

$$+\gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^i - [\sum_{i\in\nu} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i] - c_0^i + \alpha_0^i][h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$

$$h_1^i = h_0^i(2+\psi) - [\sum_{i\in\nu} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i](1+\psi) - c_0^i(1+\psi) + \alpha_0^i(1+\psi)$$

$$+\gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^i - [\sum_{i\in\nu} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i] - c_0^i + \alpha_0^\eta][h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$

$$h_1^i = h_0^i(2+\psi) - [\sum_{i\in\nu\setminus i} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i](1+\psi) - c_0^i(1+\psi)$$

$$+\gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^i - [\sum_{i\in\nu\setminus i} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i] - c_0^i][h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$

In order to simplify notation we set:

$$\rho_0^j = [h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$
$$\mu_0^j = [\sum_{i\in\nu\setminus i} \alpha_0^i - \sum_{i\in\nu\setminus i} \tau_0^i] = \tau_0^i - \alpha_0^i$$
$$\theta = \gamma \sum_{j\in\nu, j\neq i} g_{ij}\rho_0^j$$
$$\kappa = \sum_{i\in\nu} \alpha_1^i - \sum_{i\in\nu\setminus i} \tau_1^i = \tau_1^i$$
$$\zeta_0^j = \gamma \sum_{j\in\nu, j\neq i} g_{ij}(h_0^j - c_0^j)$$

and obtain:

$$h_1^i = h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\gamma \sum_{j\in\nu, j\neq i} g_{ij}[h_0^j - \tau_0^j - c_0^j + \alpha_0^j]$$
$$h_1^i = h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\gamma \sum_{j\in\nu, j\neq i} g_{ij}\rho_0^j$$
$$h_1^i = h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta$$

Pluggin this back into the original problem then gives us the final maximization problem (remember that $c_1 = h_1 - \kappa$):

Final Maximization Problem

$$\max_{c_t^i} \mathcal{U} = ln[c_0^i] + \beta ln[h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta - \kappa]$$
$$s.t.$$
$$h_0^i = \underline{h}^i$$
$$\forall i \in \nu$$

Taking the derivative of $\mathcal{U}$ w.r.t. c_0^0 gives us:

$$\frac{\partial \mathcal{U}}{\partial c_o^i} = \frac{1}{c_0^i} - \beta \frac{1+\psi+\theta}{h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta - \kappa} \stackrel{!}{=} 0$$

Rearranging yields:

$$\begin{aligned}
&\frac{1}{c_0^i} = \beta \frac{1+\psi+\theta}{h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta - \kappa} \\
&\frac{h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta - \kappa}{1} = \beta c_0^i \frac{1+\psi+\theta}{1} \\
&h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta - \kappa = \beta c_0^i[1+\psi+\theta] \\
&h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - \kappa = \beta c_0^i[1+\psi+\theta] + c_0^i(1+\psi) + c_0^i\theta \\
&h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - \kappa = c_0^i[\beta[1+\psi+\theta] + 1 + \psi + \theta] \\
&h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - \kappa = c_0^i[(1+\beta)[1+\psi+\theta]] \\
&c_0^i = \frac{h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - \kappa}{(1+\beta)[1+\psi+\theta]}
\end{aligned}$$

setting taxes and subsidies zero we get:

$$c_0^i = \frac{h_0^i(2+\psi) + h_0^i\theta}{(1+\beta)[1+\psi+\theta]}$$

and this is the solution from before without the taxation (remember that for zero taxes and zero subsidies $\theta = \zeta$):

$$\begin{aligned}
c_0^i &= \frac{h_0^i(2+\psi) + h_0^i\zeta_0^j}{(1+\beta)(1+\psi+\zeta_0^j)} \\
c_0^i &= \frac{h_0^i[2+\psi+\zeta_0^j]}{(1+\beta)(1+\psi+\zeta_0^j)}
\end{aligned}$$

We now insert c_0^i into h_1^i:

$$h_1^i = h_0^i(2+\psi) - \mu_0^j(1+\psi) - c_0^i(1+\psi) + [h_0^i - \mu_0^j - c_0^i]\theta$$
$$h_1^i = h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - c_0^i(1+\psi) - c_0^i\theta$$
$$h_1^i = h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - c_0^i(1+\psi+\theta)$$
$$h_1^i = h_0^i(2+\psi) + h_0^i\theta - \mu_0^j(1+\psi) - \mu_0^j\theta - c_0^i(1+\psi+\theta)$$
$$h_1^i = h_0^i(2+\psi+\theta) - \mu_0^j(1+\psi+\theta) - c_0^i(1+\psi+\theta)$$
$$h_1^i = h_0^i(2+\psi+\theta) - \mu_0^j(1+\psi+\theta) - \frac{h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - \kappa}{(1+\beta)[1+\psi+\theta]}(1+\psi+\theta)$$
$$h_1^i = h_0^i(2+\psi+\theta) - \mu_0^j(1+\psi+\theta) - \frac{h_0^i(2+\psi) - \mu_0^j(1+\psi) + [h_0^i - \mu_0^j]\theta - \kappa}{(1+\beta)}$$
$$h_1^i = h_0^i(2+\psi+\theta) - \frac{h_0^i(2+\psi+\theta)}{1+\beta} - \frac{h_0^i\theta}{1+\beta} - \mu_0^j(1+\psi+\theta) - \frac{-\mu_0^j(1+\psi) - \mu_0^j\theta}{(1+\beta)} + \frac{\kappa}{(+\beta}$$
$$h_1^i = h_0^i[(2+\psi+\theta) - \frac{2+\psi+\theta}{1+\beta}] - \mu_0^j(1+\psi+\theta) - \mu_0^j\frac{-(1+\psi)-\theta}{1+\beta} + \frac{\kappa}{(1+\beta)}$$
$$h_1^i = h_0^i(2+\psi+\theta)[1 - \frac{1}{1+\beta}] - \mu_0^j(1+\psi+\theta) - \frac{-1-\psi-\theta}{(1+\beta)} + \frac{\kappa}{1+\beta}$$
$$h_1^i = h_0^i(2+\psi+\theta)\frac{\beta}{1+\beta} - \mu_0^j[1+\psi+\theta+\frac{1+\psi+\theta}{(1+\beta)}] + \frac{\kappa}{1+\beta}$$
$$h_1^i = h_0^i(2+\psi+\theta)\frac{\beta}{1+\beta} - \mu_0^j[1+\psi+\theta][1+\frac{1}{1+\beta}] + \frac{\kappa}{1+\beta}$$
$$h_1^i = h_0^i(2+\psi+\theta)\frac{\beta}{1+\beta} - \mu_0^j[1+\psi+\theta]\frac{2+\beta}{1+\beta} + \frac{\kappa}{1+\beta}$$

We have now solved the entire problem in the sense that, given parameters for taxation and subsidies by the government (and given an initial state of h_0^i) agent i has a best response function, i.e. she can choose optimally the value of c_0^i and h_1^i in order to maximize consumption.

Chapter 6

Conclusion

In this dissertation I provide an evaluation of three policy proposals for the areas of education, gender inequality, and social networks. All three areas are of high relevance in the context of labor market economics and the research in this area is an important tool in order to provide better economic outcomes for individuals. In order to account for the complexities of the possible effects I employed the methodological approach of Agent-based models for all three policies.

I first looked into the possible effects of an educational voucher that is meant to alleviate the inferior educational outcomes of students from low income families who are forced to default into low performing public schools while increasing the competitive pressure on public schools. I show that initial concerns that the policy might lead to "cream skimming" over the group of public schools are not unsubstantiated and the universal voucher would not lead to unequivocal welfare gains. The vouchers increase the number of students who are actually able to choose their preferred school, thereby leaving the public schools no other choice than increasing their educational expenditure to attract more students or prevent current students from leaving. Meanwhile, due to the positive correlation between ability and preference for higher quality schools, the group of students most likely to leave public schools for private schools are high ability students, thus decreasing the welfare gains from the increased public school expenditure due to a lower mean ability. As an alternative, I analyze the effect of so called target vouchers, which are a function of student ability and show that these would allow to reap the benefits of higher competition while avoiding the "cream skimming" effect. These results indicate that the effects of a voucher based school are highly dependant on the particular voucher design and any implementation needs to take potential detrimental effects into account.

Furthermore, I analyze an alternative policy instrument to the fixed gender quota for the boards of publicly listed companies as it was recently introduced in Germany (as of 2016)

and Norway (since 2003). As the case of Norway suggested, the fixed quota can lead to large negative effects for the targeted companies due the difficulties of filling positions with qualified women. Due to the heterogeneous labor supply of women across sector, some firms find it particularly difficult to fulfil the quota. For this reason, a more flexible policy is required. I evaluate such an alternative policy, namely a tradable quota, that would provide companies with certificates that entitle them to hire men. This would avoid exposing firms to a fixed gender ratio and rather offer them to fill their positions based on the number of certificates they own. Companies with more men, could then buy certificates from companies that hire more women. Using a computational model, the chapter demonstrates that a tradable quota is a superior alternative (in terms of welfare) that would allow to achieve the same aggregate ratio of women to men at much lower adjustment costs to firms.

In the last chapter, I demonstrate the effects of labor market subsidy that takes the network structure that potential beneficiaries of the policy are embedded in, into account. In the analysis, I employ a setting in which individuals maximize intertemporal utility by choosing between consumption today and investment into their education. Hereby, the model takes into account potential externalities introduced by the peers of each individual. That is, each agent profits from the educational investment of their peers and vice versa. Using this setting I demonstrate how a policy that takes the network structure into account by targeting the most central agents can lead to an ambiguous effect. While potential spill-over effects on non-treated individuals are maximized through this strategy, inequality increases as individuals in the periphery of the network fall further behind through this approach. These finding are in contrast to similar studies such as Ballester et al. (2006) who argue unequivocally in favor of targeting the central agents when it comes to breaking up crime networks. Instead, for such a "positive" labor market policy the effects on the overall group of agens are not as clear and indicate that a careful design of such a policy is necessary.

Bibliography

Ahern, K. R. and Dittmar, A. K. (2012). The changing of the boards: the impact on firm valuation of mandated female board representation. *Quarterly Journal of Economics*, 127:137–197.

Akyol, M. (2016a). Do educational vouchers reduce inequality and inefficiency in education? *Working Paper*.

Akyol, M. (2016b). Key players in labor market networks. *Working Paper*.

Akyol, M., Neugart, M., and Pichler, S. (2015). A tradable employment quota. *Labour Economics*, 36:48–63.

Alesina, A., Ichino, A., and Karabarbounis, L. (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy*, 3(2):1–40.

Altonji, J. G. and Blank, R. M. (1999a). Race and gender in the labor market. *Handbook of labor economics*, 3:3143–3259.

Altonji, J. G. and Blank, R. M. (1999b). Race and gender in the labor market. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3C, pages 3143–3259. Elsevier.

Altonji, J. G. and Pierret, C. R. (2001). Employer learning and statistical discrimination. *Quarterly Journal of Economics*, 464:313–350.

Arrow, K. (1973). The theory of discrimination. In Ashenfelter, O. and Rees, A., editors, *Discrimination in Labor Markets*. Princeton University Press, Princeton, NJ.

Ballester, C., Calvó-Armengol, A., and Zenou, Y. (2006). Who's who in networks. wanted: the key player. *Econometrica*, 74(5):1403–1417.

Ballester, C., Zenou, Y., and Calvó-Armengol, A. (2010). Delinquent networks. *Journal of the European Economic Association*, 8(1):34–61.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2013). The diffusion of microfinance. *Science*, 341(6144):1236498.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2014). Gossip: Identifying central individuals in a social network.

Barrow, L. and Rouse, C. E. (2008). School vouchers: recent findings and unanswered questions. *Economic Perspectives*, (Q III):2–16.

Bayer, P. and McMillan, R. (2010). Choice and competition in local education markets. *NBER Working Paper*, (11802).

Bayer, P., Ross, S. L., and Topa, G. (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy*, 116(6):1150–1196.

Beaman, L. A. (2012). Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the us. *The Review of Economic Studies*, 79(1):128–161.

Becker, G. (1957). *The Economics of Discrimination.* University of Chicago Press, Chicago.

Benabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica*, 70(2):481–517.

Bergmann, B. (1974). Occupational segregation, wages, and profits when employers discriminate by race and sex. *Eastern Economic Journal*, 1:103–110.

Bergmann, B. R. (1990). Micro-to-macro simulation: a primer with a labor market example. *Journal of Economic Perspectives*, 4:99–116.

Berson, C. (2014). Oligopsony and discrimination on the labor market. Discussion Paper - Paris School of Economics.

Bertrand, M., Goldin, C., and Katz, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2:228–255.

Bewley, T. F. (1999). *Why wages don't fall during a recession.* Harvard University Press.

Bhaskar, V., Manning, A., and To, T. (2002). Oligopsony and monopsonistic competition in labor markets. *The Journal of Economic Perspectives*, 16(2):155–174.

Bhaskar, V. and To, T. (2003). Oligopsony and the distribution of wages. *European Economic Review*, 47(2):371–399.

Black, D. A. (1995). Discrimination in an equilibrium search and matching model. *Journal of Labor Economics*, 13:309–334.

Blau, F. D. and Kahn, L. M. (2000). Gender differences in pay. *Journal of Economic Perspectives*, 14(4):75–99.

Bohren, O. and Staubo, S. (2014). Does mandatory gender balance work? changing organizational form to avoid board upheaval. *Journal of Corporate Finance*, 28:152–168.

Bonacich, P. (1987). Power and centrality: A family of measures. *American journal of sociology*, pages 1170–1182.

Book, T. N. D. (2012). Statistical abstract of the United States. *US Gov. Print. Off.*

Boskin, M. J. and Sheshinski, E. (1983). Optimal tax treatment of the family: Married couples. *Journal of Public Economics*, 20(3):281–297.

Bowlus, A. and Eckstein, Z. (2002). Discrimination and skill differences in an equilibrium search model. *International Economic Review*, 43:1309–1345.

Bréchet, T. and Picard, P. M. (2010). The price of silence: markets for noise licenses and airports. *International Economic Review*, 51(4):1097–1125.

Brenner, T. (2006). Agent learning representation: Advice on modelling economic learning. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics*, volume 2 of *Handbook of Computational Economics*, chapter 18, pages 895–947. Elsevier.

Brock, W. A. and Durlauf, S. N. (2001). Discrete choice with social interactions. *The Review of Economic Studies*, 68(2):235–260.

Burdett, K. and Mortensen, D. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39:257–273.

Calvo-Armengol, A. and Jackson, M. O. (2004). The effects of social networks on employment and inequality. *American economic review*, pages 426–454.

Calvo-Armengol, A., Patacchini, E., and Zenou, Y. (2009a). Peer effects and social networks in education. *The Review of Economic Studies*, 76(4):1239–1267.

Calvo-Armengol, A., Patacchini, E., and Zenou, Y. (2009b). Peer effects and social networks in education. *The Review of Economic Studies*, 76(4):1239–1267.

Camerer, C. and Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874.

Camerer, C. and Ho, T.-H. (2008). Individual differences in ewa learning with partial payoff information. *Economic Journal*, 118(525):37–59.

Carrell, S. E., Sacerdote, B. I., and West, J. E. (2013). From natural variation to optimal policy? The importance of endogenous peer group formation. *Econometrica*, 81(3):855–882.

Casella, A. (1999). Tradable deficit permits - efficient implementation of the stability pact. *Economic Policy*, 14(October):322.

Caucutt, E. M. (2002). Educational vouchers when there are peer group effects- size matters. *International Economic Review*, 43(1):195–222.

Chakrabarti, R. (2013). Impact of voucher design on public school performance: Evidence from Florida and Milwaukee voucher programs. *The B.E. Journal of Economic Analysis & Policy*, 14(1):349–394.

Chen, Z. and West, E. G. (2000). Selective versus universal vouchers: Modeling median voter preferences in education. *The American Economic Review*, 90(5):1520–1534.

Chiang, H. (2009). How accountability pressure on failing schools affects student achievement. *Journal of Public Economics*, 93(9):1045–1057.

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics*, pages 1–44.

Coate, S. and Loury, G. (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review*, 85:1220–1240.

Coleman, J. S., Campbell, E. Q., Hobson, C. J., McPartland, J., Mood, A. M., Weinfeld, F. D., and Robert, L. (1966). York, 1966: Equality of educational opportunity. *Washington: US Government Printing Office*.

Dagum, C. (1975). A model of income distribution and the conditions of existence of moments of finite order. In *Proceedings of the 40th session of the International Statistical Institute*, volume 46, pages 199–205.

Dawid, H. and Neugart, M. (2011). Agent-based models for economic policy design. *Eastern Economic Journal*, 37(1):44–50.

De la Croix, D. and Gosseries, A. (2009). Population policy through tradable procreation entitlements. *International Economic Review*, 50(2):507–542.

Duffy, J. (2006). Agent-based models and human subject experiments. *Handbook of computational economics*, 2:949–1011.

Durlauf, S. N. (2004). Neighborhood effects. *Handbook of regional and urban economics*, 4:2173–2242.

Epple, D., Figlio, D., and Romano, R. (2004). Competition between private and public schools: testing stratification and pricing predictions. *Journal of Public Economics*, 88(7):1215–1245.

Epple, D. and Romano, R. (1996). Ends against the middle: Determining public service provision when there are private alternatives. *Journal of Public Economics*, 62(3):297–325.

Epple, D. and Romano, R. (1998). Competition between private and public schools, vouchers, and peer-group effects. *American Economic Review*, pages 33–62.

Epple, D. and Romano, R. (2008). Educational vouchers and cream skimming. *International Economic Review*, 49(4):1395–1435.

Epple, D. and Romano, R. (2011). Peer effects in education: A survey of the theory and evidence. *Handbook of social economics*, 1(11):1053–1163.

Epple, D. and Romano, R. (2012). Economic modeling and analysis of educational vouchers. *Annual Review of Economics*, 4(1):159–183.

Epple, D. and Romano, R. (2014). On the political economy of educational vouchers. *Journal of Public Economics*, 120:62–73.

Epple, D., Romano, R., Sarpça, S., and Sieg, H. (2013). The US market for higher education: A general equilibrium analysis of state and private colleges and public funding policies.

Fabra, N. and Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, 104:2872–2899.

Ferreyra, M. M. (2007). Estimating the effects of private school vouchers in multidistrict economies. *The American Economic Review*, pages 789–817.

Figlio, D. N. and Rouse, C. E. (2006). Do accountability and voucher threats improve low-performing schools? *Journal of Public Economics*, 90(1):239–255.

Fowler, C. S. (2011). Finding equilibrium: how important is general equilibrium to the results of geographical economics? *Journal of Economic Geography*, 11(3):457–480.

Freeman, R. B. (1998). War of the models: Which labour market institutions for the 21st century? *Labour Economics*, 5:1–24.

Friedman, M. (1955). The role of government in education. In Solo, R. A., editor, *Economics and the Public Interest*. Trustees of Rutgers College, New Jersey.

Friedman, M. (1962). *Capitalism and Freedom.* Number 9780226264202 in University of Chicago Press Economics Books. University of Chicago Press.

Friedman, M. (1997). Public schools: Make them private. *Education Economics*, 5(3):341–344.

Fu, C. (2014). Equilibrium tuition, applications, admissions, and enrollment in the college market. *Journal of Political Economy*, 122(2):225–281.

Gaviria, A. and Raphael, S. (2001). School-based peer effects and juvenile behavior. *Review of Economics and Statistics*, 83(2):257–268.

Glomm, G. and Ravikumar, B. (1998). Opting out of publicly provided services: A majority voting result. *Social choice and welfare*, 15(2):187–199.

Glomm, G., Ravikumar, B., Schiopu, I. C., et al. (2011). The political economy of education funding. *Handbook of the Economics of Education*, 4:615–680.

Goldberg, J. and Smith, J. (2008). The effects of education on labor market outcomes. *Handbook of research in education finance and policy (688-708). New York, NY: Routledge.*

Goldberg, P. K. and Hellerstein, R. (2013). A structural approach to identifying the sources of local currency price stability. *The Review of Economic Studies*, 80(1):175–210.

Granovetter, M. (1995). *Getting a job: A study of contacts and careers.* University of Chicago Press.

Granovetter, M. S. (1973). The strength of weak ties. *American journal of sociology*, pages 1360–1380.

Greene, J. P. (2001). An evaluation of the Florida A-Plus accountability and school choice program. *Center for Civic Innovation at the Manhattan Institute.*

Greene, J. P. and Winters, M. A. (2003). When schools compete: The effects of vouchers on florida public school achievement. *Education Working Paper.*

Guryan, J. and Charles, K. K. (2013). Taste-based or statistical discrimination: The economics of discrimination returns to its roots. *The Economic Journal*, 123:F417–F432.

Heilman, M. E., Block, C. J., and Lucas, J. A. (1992). Presumed incompetent? stigmatization and affirmative action efforts. *Journal of Applied Psychology*, 77(4):536.

Heilman, M. E., Block, C. J., and Stathatos, P. (1997). The affirmative action stigma of incompetence: Effects of performance information ambiguity. *Academy of Management Journal*, 40(3):603–625.

Hemphill, F. C. and Vanneman, A. (2011). Achievement gaps: How hispanic and white students in public schools perform in mathematics and reading on the national assessment of educational progress. statistical analysis report. nces 2011-459. *National Center for Education Statistics*.

Henderson, V., Mieszkowski, P., and Sauvageau, Y. (1978). Peer group effects and educational production functions. *Journal of Public Economics*, 10(1):97–106.

Hindriks, J. and Myles, G. D. (2013). *Intermediate Public Economics*. MIT Press, Cambridge, Massachusetts, 2nd edition.

Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of computational economics*, 2:1109–1186.

Hoxby, C. M. (2000). Peer effects in the classroom: Learning from gender and race variation. NBER Working Papers 7867, National Bureau of Economic Research, Inc.

Hoxby, C. M. (2003a). School choice and school competition: Evidence from the united states. *Swedish Economic Policy Review*, 10(2):9–66.

Hoxby, C. M. (2003b). School choice and school productivity. could school choice be a tide that lifts all boats? In *The economics of school choice*, pages 287–342. University of Chicago Press.

Hoyt, W. H. and Lee, K. (1998). Educational vouchers, welfare effects, and voting. *Journal of Public Economics*, 69(2):211–228.

Ioannides, Y. M. and Loury, L. D. (2004). Job information networks, neighborhood effects, and inequality. *Journal of economic literature*, pages 1056–1093.

Ioannides, Y. M. and Topa, G. (2010). Neighborhood effects: Accomplishments and looking beyond them*. *Journal of Regional Science*, 50(1):343–362.

Jackson, M. O. (2014). Networks in the understanding of economic behaviors. *The Journal of Economic Perspectives*, pages 3–22.

Jackson, M. O. and Lopez-Pintado, D. (2013). Diffusion and contagion in networks with heterogeneous agents and homophily. *Network Science*, 1(01):49–67.

Jennings, N. R. (2000). On agent-based software engineering. *Artificial intelligence*, 117(2):277–296.

Judd, K. L. (2006). Computationally intensive analyses in economics. *Handbook of computational economics*, 2:881–893.

Kaas, L. (2009). Does equal pay legislation reduce labour market inequality? *The Scandinavian Journal of Economics*, 111:51–71.

Kaas, L. and Lu, J. (2010). Equal-treatment policy in a random search model with taste discrimination. *Labour Economics*, 17:699–709.

Knowles, J., Persico, N., and Todd, P. (2001). Racial bias in motor vehicle searches: theory and evidence. *Journal of Political Economy*, 109:203–229.

Kramarz, F. and Skans, O. N. (2010). With a little help from my... parents? family networks and youth labor market entry.

Lang, K., Manove, M., and Dickens, W. (2005). Racial discrimination in labor markets with posted wage offers. *American Economic Review*, 95:1327–1340.

Lavecchia, A. M., Liu, H., and Oreopoulos, P. (2014). Behavioral economics of education: Progress and possibilities. *NBER Working Paper No. 20609*.

Levin, H. M. and Belfield, C. R. (2002). The effects of competition between schools on educational outcomes: A review for the united states. *Review of Educational Research*, 72(2):279–341.

MacLeod, W. B. and Urquiola, M. (2009). Anti-lemons: school reputation and educational quality.

MacLeod, W. B. and Urquiola, M. (2012). Competition and educational productivity: Incentives writ large. IZA Discussion Papers 7063, Institute for the Study of Labor (IZA).

Manski, C. F. (1992). Educational choice (vouchers) and social mobility. *Economics of Education Review*, 11(4):351–369.

McEwan, P. J. (2003). Peer effects on student achievement: Evidence from chile. *Economics of Education Review*, 22(2):131–141.

Mechtenberg, L. (2009). Cheap talk in the classroom: how biased grading at school explains gender differences in achievements, career choices and wages. *Review of Economic Studies*, 76:1431–1459.

Misra, K., Grimes, P. W., and Rogers, K. E. (2012). Does competition improve public school efficiency? a spatial analysis. *Economics of Education Review*, 31(6):1177–1190.

Moraga, J. F. and Rapoport, H. (2014). Tradable immigration quotas. *Journal of Public Economics*, 115:94–108.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the us labor market. *The Quarterly Journal of Economics*, pages 549–599.

Munshi, K. and Rosenzweig, M. (2006). Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy. *The American Economic Review*, 96(4):1225–1252.

Neal, D. (2002). How vouchers could change the market for education. *Journal of Economic Perspectives*, 16(4):25–44.

Nechyba, T. J. (1999). School finance induced migration and stratification patterns: the impact of private school vouchers. *Journal of Public Economic Theory*, 1(1):5–50.

Nechyba, T. J. (2000). Mobility, targeting, and private-school vouchers. *American Economic Review*, pages 130–146.

Nechyba, T. J. (2003). Centralization, fiscal federalism, and private school attendance*. *International Economic Review*, 44(1):179–204.

Neugart, M. (2008). Labor market policy evaluation with ace. *Journal of Economic Behavior and Organization*, 67:418–430.

Neugart, M. and Richiardi, M. (2015). Agent-based models of the labor market. In Chen, S.-H. and Kaboudan, M., editors, *The Handbook of Computational Economics and Finance*, page forthcoming. Oxford University Press, Oxford.

Neugart, M., Richiardi, M., et al. (2012). Agent-based models of the labor market. *LABORatorio R. Revelli working papers series*, 125.

Nygaard, K. (2011). Forced board changes: Evidence from norway. *Norwegian School of Economics and Business Administration Discussion Paper*.

OECD-Report (2012a). Pisa 2012 results: What makes schools successful? *Volume IV*.

OECD-Report (2012b). Successful reformers in education- lessons from pisa 2012 for the united states. *OECD*.

Olson, M. (1965). The logic of collective action: public goods and the theory of collective action.

Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, 39(2):390–431.

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

Piolatto, A. (2010). Education and selective vouchers. *Economics of Education Review*, 29(6):993–1004.

Rangazas, P. (1995). Vouchers and voting: An initial estimate based on the median voter model. *Public Choice*, 82(3-4):261–279.

Rees, A. (1966). Information networks in labor markets. *The American Economic Review*, pages 559–566.

Richiardi, M. et al. (2003). The promises and perils of agent-based computational economics. *LABORatorio R. Revelli, Centre for Employment Studies LABORatorio R. Revelli Working Papers Series*, 29.

Richiardi, M. G. (2012). Agent-based computational economics: a short introduction. *The Knowledge Engineering Review*, 27(02):137–149.

Richiardi, M. G., Leombruni, R., Saam, N. J., and Sonnessa, M. (2006). A common protocol for agent-based social simulation. *Journal of artificial societies and social simulation*, 9.

Rosen, A. (2003). Search, bargaining, and employer discrimination. *Journal of Labor Economics*, 21:807–829.

Rosen, H. S. (1977). Is it time to abandon joint filing? *National Tax Journal*, pages 423–428.

Roth, A. E. (2002). The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70:1341–1378.

Rothstein, J. M. (2006). Good principals or good peers? parental valuation of school characteristics, tiebout equilibrium, and the incentive effects of competition among jurisdictions. *American Economic Review*, 96(4):1333–1350.

Rouse, C. E., Hannaway, J., Goldhaber, D., and Figlio, D. (2013). Feeling the florida heat? how low-performing schools respond to voucher and accountability pressure. *American Economic Journal: Economic Policy*, 5(2):251–281.

Sacerdote, B. (2011). *Peer Effects in Education: How Might They Work, How Big Are They and How Much Do We Know Thus Far?*, volume 3 of *Handbook of the Economics of Education*, chapter 4, pages 249–277. Elsevier.

Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, pages 141–156.

Sandel, M. J. (1997). It's immoral to buy the right to pollute. New York Times, December 15th, p. A29.

Schelling, T. (1978). Schelling micromotives and macrobehavior (cloth)(fels lectures on public policy analysis).

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186.

Schmalensee, R. and Stavins, R. N. (2013). The so2 allowance trading system: The ironic history of a grand policy experiment. *Journal of Economic Perspectives*, 27:103–122.

Scotchmer, S. (1997). On price-taking equilibria in club economies with nonanonymous crowding. *Journal of Public Economics*, 65(1):75–88.

Scotchmer, S. and Wooders, M. H. (1987). Competitive equilibrium and the core in club economies with anonymous crowding. *Journal of Public Economics*, 34(2):159–173.

Spiro Maroulis, Eytan Bakshy, L. G. and Wilensky, U. (2014). Modeling the transition to public school choice. *Journal of Artificial Societies and Social Simulation*, 17(2).

Stark, O. and Hyll, W. (2014). Socially gainful gender quotas. *Journal of Economic Behavior and Organization*, 105:173–177.

Stavins, R. N. (1998). What can we learn from the grand policy experiment? lessons from so2 allowance trading. *Journal of Economic Perspectives*, 12:69–88.

Summers, A. A. and Wolfe, B. L. (1977). Do schools make a difference? *The American Economic Review*, pages 639–652.

Tesfatsion, L. (2001). Structure, behavior, and market power in an evolutionary labor market with adaptive search. *Journal of Economic Dynamics and Control*, 25:419–457.

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial life*, 8(1):55–82.

Tesfatsion, L. and Judd, K. L. (2006). *Handbook of computational economics: agent-based computational economics*, volume 2. Elsevier.

Thomas D. Snyder, C. d. B. and Dillow, S. A. (2014). Digest of education statistics.

Topa, G. (2001). Social interactions, local spillovers and unemployment. *The Review of Economic Studies*, 68(2):261–295.

Vanneman, A., Hamilton, L., Anderson, J. B., and Rahman, T. (2009). Achievement gaps: How black and white students in public schools perform in mathematics and

reading on the national assessment of educational progress. statistical analysis report. nces 2009-455. *National Center for Education Statistics*.

Vriend, N. J. (2006). Ace models of endogenous interactions. *Handbook of computational economics*, 2:1047–1079.

Walsh, P. (2009). Effects of school choice on the margin: The cream is already skimmed. *Economics of Education Review*, 28(2):227–236.

Welch, F. (1976). Employment quotas for minorities. *Journal of Political Economy*, 84:S105–S139.

West, E. G. (1967). Tom paine's voucher scheme for public education. *Southern Economic Journal*, 33:378–382.

West, M. R. and Peterson, P. E. (2006). The efficacy of choice threats within school accountability systems: Results from legislatively induced experiments. *The Economic Journal*, 116(510):C46–C62.

Wilhite, A. (2006). Economic activity on fixed networks. *Handbook of computational economics*, 2:1013–1045.

Winker, P. (2000). Efficient labour contracts: impediments and how to circumvent them. *Labour, Review of Labour Economics and Industrial Relations*, 14:373–392.

Zimmerman, D. J. (2003). Peer effects in academic outcomes: Evidence from a natural experiment. *Review of Economics and Statistics*, 85(1):9–23.