

Stichnoth, Holger; van der Straeten, Karine

Working Paper

Ethnic diversity and attitudes towards redistribution: a review of the literature

ZEW Discussion Papers, No. 09-036

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Stichnoth, Holger; van der Straeten, Karine (2009) : Ethnic diversity and attitudes towards redistribution: a review of the literature, ZEW Discussion Papers, No. 09-036, Zentrum für Europäische Wirtschaftsforschung (ZEW), Mannheim

This Version is available at:

<https://hdl.handle.net/10419/27717>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Discussion Paper No. 09-036

**Ethnic Diversity and
Attitudes Towards Redistribution:
A Review of the Literature**

Holger Stichnoth and Karine Van der Straeten

ZEW

Zentrum für Europäische
Wirtschaftsforschung GmbH

Centre for European
Economic Research

Discussion Paper No. 09-036

**Ethnic Diversity and
Attitudes Towards Redistribution:
A Review of the Literature**

Holger Stichnoth and Karine Van der Straeten

Download this ZEW Discussion Paper from our ftp server:

<ftp://ftp.zew.de/pub/zew-docs/dp/dp09036.pdf>

Die Discussion Papers dienen einer möglichst schnellen Verbreitung von neueren Forschungsarbeiten des ZEW. Die Beiträge liegen in alleiniger Verantwortung der Autoren und stellen nicht notwendigerweise die Meinung des ZEW dar.

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.

Non-technical summary

There is a large literature on the costs and benefits of immigration within a *given* system of social security. More recently, economists have begun to address a related question: does immigration and, more generally, ethnic diversity *change* this system of social security in turn?

A number of empirical studies suggest that ethnic diversity does indeed matter for the extent of redistribution. First, there is evidence that actual public spending is associated with the degree of ethnic diversity. Second, studies that attempt to explore the mechanisms behind this aggregate relationship have found that individual attitudes and behaviour are affected by ethnic diversity.

The purpose of this paper is to survey this empirical literature. We cover the studies that have appeared since the survey by Alesina and La Ferrara (2005) in the Journal of Economic Literature. In particular, we review the fast-growing literature that uses controlled experiments to study the effects of ethnic diversity on redistribution.

Our main conclusion from this survey is that although numerous studies document a negative and statistically significant relationship, most of these studies do *not* point to a *quantitatively important* role for ethnic diversity in shaping natives' preferences for redistribution. In most studies, the association is much weaker than for other factors such as own income (current or expected) or beliefs about the role of effort versus luck in determining this income.

Moreover, it seems that the sizeable negative association between ethnic diversity and support for redistribution that is sometimes found in U.S. studies does not generalize to Canada or Europe. However, the evidence for countries other than the U.S. is scarce so far, and there is certainly need for further research.

Das Wichtigste in Kürze (German summary)

Zahlreiche Studien haben die Kosten und den Nutzen der Einwanderung innerhalb eines *gegebenen* Systems der sozialen Sicherung untersucht. Seit einigen Jahren gerät aber zunehmend eine verwandte Frage in den Blickpunkt des Interesses: Beeinflusst die Einwanderung (und allgemeiner, die ethnische Vielfalt) ihrerseits die Gestaltung der sozialen Sicherung?

Eine Reihe von empirischen Untersuchungen deutet darauf hin, dass tatsächlich ein solcher Einfluss besteht. Zum einen lässt sich ein Zusammenhang zwischen der ethnischen Vielfalt und den öffentlichen Ausgaben im Bereich der sozialen Sicherung nachweisen; zum anderen zeigen Studien, die die Mechanismen hinter diesem Zusammenhang untersuchen, dass auch die Einstellungen und das Verhalten (etwa die Spendenbereitschaft) der Einzelnen von der ethnischen Vielfalt beeinflusst werden.

Das Ziel dieses Artikels ist es, einen Überblick über diese empirischen Studien zu geben. Insbesondere werden die Studien berücksichtigt, die seit dem von Alesina und La Ferrara (2005) im *Journal of Economic Literature* veröffentlichten Übersichtsartikel erschienen sind. Darunter sind zunehmend Studien, die den Einfluss der ethnischen Vielfalt auf das Verhalten im Rahmen von Experimenten untersuchen.

Aus dem Überblick über die Literatur ergibt sich der Schluss, dass es zwar zahlreiche Belege für einen statistisch signifikanten Zusammenhang zwischen der ethnischen Vielfalt und der Umverteilung oder Umverteilungsbereitschaft gibt, dass aber dieser Zusammenhang in den meisten Fällen eher schwach ist. Insbesondere hat die ethnische Vielfalt einen deutlich schwächeren Einfluss auf die Umverteilungsbereitschaft als Faktoren wie das eigene Einkommen oder die Auffassung darüber, ob dieses Einkommen eher der eigenen Anstrengung oder äußeren Umständen zu verdanken ist.

Außerdem zeigt sich, dass der negative Zusammenhang zwischen der ethnischen Vielfalt und der Umverteilungsbereitschaft, der für die USA nachgewiesen werden konnte, in Kanada und in europäischen Ländern schwächer ausgeprägt ist. Allerdings gibt es erst relativ wenige Studien, die mit anderen als US-amerikanischen Daten arbeiten; hier besteht noch Bedarf an weiteren Untersuchungen.

Ethnic Diversity and Attitudes towards Redistribution: A Review of the Literature

Holger Stichnoth*

ZEW Centre for European Economic Research, Mannheim
Paris School of Economics

Karine Van der Straeten

Toulouse School of Economics and Paris School of Economics

July 7, 2009

Abstract

We review the empirical literature that studies the effect of ethnic diversity on the welfare state and on individual attitudes. The outcome variables that we cover in the survey are on the one hand public spending, and on the other hand individual attitudes and behaviour, including charity spending. We also review the fast-growing literature that uses experiments to study the effects of ethnic diversity. Many of these studies have appeared since the pioneering survey by Alesina and La Ferrara (2005a), and have not been covered by a survey before.

Keywords: redistribution, social security, welfare state, immigration, ethnic diversity, survey

JEL codes: H53, H55, I38, J15, J61

*Corresponding author. ZEW Mannheim, Postfach 103443, 68034 Mannheim, Germany. E-Mail: stichnoth@zew.de.

1 Introduction

There is a large literature on the costs and benefits of immigration within a *given* system of social security. More recently, economists have begun to address a related question: do immigration and, more generally, ethnic diversity *change* this system of social security in turn?

One mechanism through which ethnic diversity is thought to affect the social security system, and redistribution in particular, is through individuals' preferences. If citizens tend to be more supportive of redistribution when people from their own ethnic group benefit from it, ethnic diversity will reduce the support for public spending, which in turn will decrease the actual level of public spending.¹

The question is delicate and important. Over the last four decades, Europe has experienced large-scale immigration, often from countries with cultural, religious, or ethnic backgrounds that are quite different from those of the native population. Alesina and Glaeser (2004, 11) argue that this inflow will affect the European welfare state: "one natural implication of our conclusion that fractionalization reduces redistribution is that if Europe becomes more heterogeneous due to immigration, ethnic divisions will be used to challenge the generous welfare state." The possibility of such a trade-off is increasingly recognized and discussed within the political Left, which traditionally tends to be both pro-redistribution and pro-immigration. (See for instance Gitlin (1996), Goodhart (2006), and the debate in Van Parijs (2004).)

A number of empirical studies suggest that ethnic diversity *does* indeed matter for the extent of redistribution. There is evidence that both aggregate public spending and individual attitudes and behaviour are affected by ethnic diversity. The purpose of this paper is to survey this empirical literature. We cover the studies that have appeared since the survey by Alesina and La Ferrara (2005a). In particular, we review the fast-growing literature that uses controlled experiments to study the effects of ethnic diversity on redistribution.

¹This argument is part of a broader attempt to enrich the seminal paper by Meltzer and Richard (1981) about the political economy of redistribution. In Meltzer and Richard's median voter model, individuals' demand for redistribution is led by (static) self-interest. At least since the mid-1990s this larger literature on the determinants of individual support for redistribution has studied factors beyond narrow self-interest. In particular, the importance of expectations about future income or future dependence on transfers, and hence of social mobility, has been recognized. (See for instance Piketty 1995; Ravallion and Lokshin 2000; Bénabou and Ok 2001; Corneo 2001; Corneo and Grüner 2000, 2002; Senik 2004, 2006; Alesina and La Ferrara 2005b). A second important extension has been to take into account beliefs about the deservingness of recipients, which is in turn influenced by beliefs about the determinants of social position and about the extent of social mobility. There is by now a huge amount of evidence, both from surveys and from experiments, that considerations of fairness and reciprocity matter for individual behaviour and attitudes. (See for instance the survey by Fong et al. 2006).

The outline of this paper is as follows. In section 2 we survey the literature on the effects of ethnic diversity on public spending. Section 3 then looks at one of the building blocks behind this aggregate relationship, namely the effect of ethnic diversity on natives' *attitudes* towards the welfare state. As noted, we will consider both studies using observational data and studies that use data from experiments. Section 4 concludes.

2 Ethnic diversity and public spending

A large number of studies have documented a negative association between ethnic diversity and public spending, both across countries (section 2.1) and across regions within the U.S. (section 2.2).

2.1 Cross-country evidence

Alesina, Glaeser, and Sacerdote (2001) argue that ethnic diversity is a key factor in explaining why the U.S. does not have a European-style welfare state (see also Alesina and Glaeser (2004) for a book-length treatment). They show that some prominent economic explanations alone cannot explain the difference and point to the influence that political institutions have on the size and the design of the welfare state. Alesina et al. argue that ethnic diversity not only helps explain why the U.S. has more property-friendly institutions, but that its negative influence on redistribution holds even for a *given* set of institutions.

To support this argument, they regress social spending as a share of GDP on an indicator of ethno-linguistic fractionalization, while controlling for a number of variables from Persson and Tabellini (2000): GDP per capita, the percent of the population aged 15–64, a dummy for a majoritarian election system, and dummies for Caribbean, Asian, and Latin American countries. The regression is cross-sectional: social spending over GDP is the average for 1960–98, and ethno-linguistic fractionalization was measured in the 1960s. Alesina et al. do not find a significant relationship between the share of social spending over GDP and ethno-linguistic fractionalization. However, they show that a variable of “racial” fractionalization (measured over the 1990s) *is* negatively related to social spending as a share of GDP. The coefficient on racial fractionalization is now the only coefficient in the regression that is statistically significant. The association seems rather strong: if the probability that two people drawn at random from the population will belong to two different racial groups increases by one percentage point, then the share of social spending over GDP is estimated to be lower by 7.5 percentage points.

In an early study, Mueller and Murrell (1986) find for a sample of OECD countries and a few other countries with “sufficiently democratic institutions” (p. 136) that the total amount of public spending is negatively related to a country’s ethnic diversity. James (1993) shows that ethno-linguistic and especially religious diversity are associated with a greater share of private schools in total enrollments. McCarty (1993) finds that in countries with greater ethnic and religious diversity, the (central) government spends less on transfers; other government expenditure is unaffected by ethnic and religious diversity. In a study by Kuijs (2000), ethnic diversity, as measured by ethno-linguistic fractionalization, is negatively related to health spending and to public spending on education. He also finds a negative correlation between ethnic diversity and indicators of health and education even when the amount of money spent on these areas is controlled for. Kuijs interprets this finding as suggesting that what he calls the “technical efficiency” of public spending may be lower in ethnically diverse countries. More recently, Soroka et al. (2006) have studied the relationship between immigration and the share of social spending over GDP. They find that countries with a high share of immigrants in the population in 1970 did not experience a smaller increase in social spending between 1970 and 1998. However, this increase is negatively related to the *change* in the population share of immigrants over the period.

Banting and Kymlicka (2004) study whether there is a trade-off between multiculturalism policies (as opposed to immigration or ethnic diversity per se) and the welfare state. The possibility of such a trade-off is increasingly recognized and discussed within the political Left, which traditionally tends to be both pro-redistribution and pro-immigration. Critics worry that “a multicultural agenda [...] crowds out redistributive issues from the policy agenda, corrodes trust among vulnerable groups who would otherwise coalesce in a pro-redistribution lobby, or misdiagnoses the real problems facing minorities, leading them to believe that their problems lie in cultural misunderstandings rather than economic barriers that they confront along with vulnerable members of many other cultural groups” (Banting 2005, 8). On the possibility of such a trade-off, see for instance Gitlin (1996), Goodhart (2006), and the debate in Van Parijs (2004). To test whether there is evidence for such a trade-off, Banting and Kymlicka construct an index of multiculturalism policies, and classify the countries of the OECD into three groups of strong (Australia and Canada), modest (e.g., the U.S. and the U.K.), or weak policies (e.g., France, Germany). For each group, they calculate the average change in social expenditure between 1980 and 1998 (from OECD data), and the average change in the redistributive impact of taxes and transfers between the early 1980s and late 1990s (from the Luxemburg Income Study). A cross-tabulation shows no evidence of a negative relationship between multiculturalism policies and the change in social spending and redistribution (Banting

and Kymlicka 2004), and this finding is confirmed in a multivariate analysis (Banting et al. 2007).

These cross-country studies have been criticized for a number of reasons. A first criticism was directed at the cross-country *growth* literature, but applies to the present case as well. Brock and Durlauf (2000) argue that the uncertainty in the specification of the model should be explicitly recognized. This uncertainty concerns the variables that should be included in the model and the homogeneity of parameters across observations. Brock and Durlauf criticize the cross-country growth literature for *assuming* (without much justification and testing) that the observed data points are draws from the same statistical distribution. Drawing lessons (and policy conclusions) for a particular country from observing other countries has to rest on such a homogeneity assumption, but Brock and Durlauf argue that this assumption is particularly tenuous when the units of observations are countries, which are more heterogeneous than smaller units of observations (and given the small number of observations in a cross-section of countries, the number of control variables that can be included is limited).

Another criticism concerns general issues of measurement and of comparability across countries. For instance, in his discussion of Alesina et al. (2001), Durlauf (2001) points to a possible measurement problem in the dependent variable. He argues that their dependent variable (social spending over GDP) underestimates the true amount of welfare spending in the U.S. First, the Earned Income Tax Credit should be included under welfare spending. Moreover, Durlauf agrees with Alesina et al. that Americans tend to believe in equality of opportunities, not of outcomes; for him, the policies of affirmative action should therefore be taken into account when studying the link between ethnic diversity and redistribution.

The measurement of ethnic diversity has sparked even more criticism. Most of the early cross-country studies measure ethnic diversity by the indicator of ethno-linguistic fractionalization (ELF) from the Atlas Narodov Mira (Bruk and Apenchenko 1964). The ethnologists who compiled the data in the early 1960s relied mostly on language to delimit ethnic groups. ELF gives the probability that, at the time when the data were collected, two individuals drawn at random from the population belong to different linguistic groups.

Chandra (2001), Laitin and Posner (2001), and Posner (2002, 2004a,b) criticize the use of this variable in the cross-country growth literature for two reasons. First, most studies use data on ethno-linguistic groups from the early 1960s to test hypotheses about the consequences of ethnic diversity in much later periods; that is, the studies assume, sometimes implicitly, that ethnic diversity does not change or changes slowly. This assumption

is questionable. Migration and different rates of fertility and mortality may have changed the ethnic make-up of the population since the early 1960s. Moreover, most sociologists and anthropologists now agree that ethnic identity (and, by implication, ethnic diversity) is socially constructed and may change over time. In fact, causality may run from the dependent variable such as the growth of GDP per head to the degree of ethnic diversity, and third factors such as political stability that are not always controlled for in the studies may influence both ethnic diversity and the growth of GDP. (See the references in Fedderke et al. 2008, 262–63). One mechanism may be that in countries that are politically unstable or that grow slowly, ethnic diversity increases as people fall back on their ethnic groups. The possible endogeneity of many regressors is of course a general problem in cross-country studies on growth and in observational studies more generally.

Laitin and Posner (2001)'s second criticism is that ethnic identity has more than one dimension and that an exclusive focus on ethno-linguistic groups is therefore misleading: the salient dimension of ethnicity may not be language but some other cultural trait or the phenotype (e.g., skin colour). In other words, even if ethnic identity were static, the data set would be ill-suited to test the hypotheses of the empirical studies.

Another criticism concerns not so much the classification of ethnic groups, but the relevant measure of diversity. Collier (2001) argues that a single index of fractionalization misses some important dimensions. In particular, he calls for a distinction between ethnic fragmentation and ethnic dominance, since political economy models suggest that the two do not have the same effect on policies and on growth. Arcand et al. (2000) also argue that polarization may be more important than fractionalization.

In reaction to these criticisms, better data sets have been constructed by Annett (2001), Yeoh (2001), Posner (2002), Alesina et al. (2003), and Fearon (2003). These studies capture more dimensions of ethnicity than just language and try to better identify the salient ethnic cleavages in each country. Montalvo and Reynal-Querol (2002, 2005a,b) propose an index of ethnic polarization; Bossert et al. (2006) develop a generalized index of fractionalization that takes into account similarities (for instance, with respect to income or education) among individuals.

Finally, Vigdor (2002) and Rushton (2008) argue that studies using ethnic *diversity* as a regressor should better justify why they believe that it is only diversity that matters, and not the share of the groups. Rushton also points out that it is often of interest *which* groups are represented and in which proportions. The index of fractionalization has the same value regardless of whether 30% of the population are black and the remaining 70% white, or if it is the other way around, but Rushton argues that in the context of the U.S. which group is in the majority may well make a differ

2.2 Studies for the U.S.

Ethnic diversity has also been found to be associated with differences in public spending across states, counties, cities, and metropolitan areas within the U.S. Exploiting within-country variation offers the advantages of better data quality and of greater comparability of jurisdictions, mitigating the risk of bias from omitted variables.

However, many spending decisions are made at a local level, and “state average spending levels therefore conceal substantial heterogeneity within states” (Poterba 1997, 52); in other words, there is not always an “exact match between the unit of observation and the relevant jurisdiction for voting on the public good” (Alesina et al. 1999, 1255–56). In these cases, it is unclear how to interpret the correlations one observes. Some studies deal with this problem by running regressions at different geographical levels. Although a useful robustness check, the problem remains that sometimes *none* of these levels is the relevant jurisdiction. Another potentially serious problem is that the demographic composition of a jurisdiction is likely to be endogenous: it is affected by decisions on public spending (reverse causality) and possibly correlated with unobserved influences on public spending (omitted variables bias). This is true also at the country-level, but the endogeneity bias likely gets worse at more disaggregated geographical levels.

With these caveats in mind, the three main results from the studies using variation in ethnic diversity within the U.S. are as follows: First, the relationship between ethnic diversity and the *overall* level of public spending is ambiguous. Cutler et al. (1993) show that the relationship seems to be positive at the local level and negative at the state level.

Second, for a given level of public spending, ethnic diversity seems to lead to a shift from public spending on public goods to the public provision of private goods, arguably because the latter can be targeted to particular ethnic groups (“ethnic patronage”). In two related papers, Alesina et al. (1999, 2000) show evidence that greater racial fractionalization in a jurisdiction is associated with less spending on public (non-excludable, non-rival) goods. By contrast, in ethnically diverse jurisdictions more tends to be spent on the public provision of private goods, and public employment tends to be higher.

An area of public spending that has received particular attention in this context is spending on public education. James (1987) finds that across U.S. states the percentage of blacks in the population is correlated with the percentage of schools that are private. Goldin and Katz (1999) study the rapid rise public secondary education in the U.S. between 1910 and 1940. Among other things, they show that this “high school movement” was stronger the more homogenous a community was in terms of ethnicity, religion, and income. Using panel data on U.S. states for the period 1961–91, Poterba (1997) finds that

public-school spending per pupil tends to be lower the greater the share of people aged 65 or over in a state's population.² This negative relationship is stronger "in states with a substantial population of older individuals who are from a different ethnic or racial group than the school-aged population" (Poterba 1998, 317). As for the size of the association, "a 1 percentage point increase in the share of nonwhites in the 0 to 17 population, holding the share of nonwhites in the 65 and over population constant, reduces the log of per-child school spending by -0.006, or approximately one-half of 1 percent." (Poterba 1997, 60). However, Poterba bases this calculation on a point estimate that is not statistically significant at conventional levels. Moreover, as he himself points out, the point estimate becomes smaller (and even less significant) once the percentage of the population living in urban areas or under the poverty line are included.

Finally, greater ethnic diversity has been found to be associated with less generous Aid to Family with Dependent Children (AFDC), the main program of social assistance until the mid-1990s. There is considerable variation in the level of these benefits: Alesina and Glaeser (2004, 147) report that in 1990, the state of Alaska paid over \$800 dollars per family per month, whereas the maximum benefit in Alabama and Mississippi was less than \$150. Alesina et al. (2001, 235) point out that in 1990 the raw correlation between the maximum level of AFDC benefits and the percentage of blacks in a state was around -0.5. In a regression controlling for median income in the state, they find that "a 1 percentage point change in the share of the population that is black reduces the maximum monthly AFDC payment by \$6.92, and a \$1,000 increase in median income increases the maximum payment by \$17" (p.236). Of course, there are many other differences between states that may affect both the share of blacks and the generosity of AFDC spending and that are not controlled for in this regression. Many of these differences are controlled for in an article by Ribar and Wilhelm (1999). Using data for 1982–1992 (from 1969 onwards in some models), they show that the share of African-Americans is negatively correlated with AFDC levels. One can calculate the magnitude of the association from their table 2 on page 101. In their preferred specification, an increase of 1 percentage point in the share of African-Americans (the mean share was around 9 percent in the period of study) is associated with average monthly AFDC benefits that are lower by about 9 percent. The estimated coefficient is highly statistically significant. This estimate is for AFDC benefits for a family of four with no other income; for other measures of the generosity of AFDC or welfare spending more generally, the coefficients are considerably lower and not always statistically significant.

²Note that Poterba measures spending *per pupil* and not per head of the state's population. The latter measure would tend to go down mechanically, since the demand for education is lower in an older population.

Lind (2007) constructs a panel of U.S. states in six years between 1969 and 2000 (mostly from the 1990s). Using data from the Luxembourg Income Study, he calculates for each state a measure of the overall income inequality (before taxes and transfers), and then decomposes this overall measure into inequality between and within racial groups. Lind shows that income inequality between racial groups tends to reduce welfare spending by the state, whereas income inequality within groups increases it. These results are consistent with the predictions from his model; that is, with the median-voter model by Meltzer and Richard (1981) extended to include social preferences with higher weights for members of one's own racial group. Since inequality may be endogenous in these regressions, Lind instruments current inequality by lagged values and finds that the results change little, but he admits some doubt as to the validity of the instruments. Closer to the studies above, Lind finds that racial fractionalization is negatively related to welfare spending. However, he finds that the association "is not very strong and it is not robust to the introduction of state fixed effects. As fractionalization changes little over time, this is not surprising" (p. 68).

3 Individual attitudes and behaviour

In the previous section we have reviewed the evidence for a negative (although sometimes weak and not always robust) association between ethnic diversity and public spending. To shed some light on the forces behind this relationship, a second strand of the empirical literature has used data on individual attitudes and behaviour. Most of the studies use observational data from surveys (section 3.1); more recently, the effect of ethnic diversity on attitudes and behaviour has also been studied in experiments (section 3.2).

3.1 Observational studies

The observational strand of the literature can be further subdivided into studies that focus on the effect of ethnic diversity on social capital and studies that directly measure the effect on the (self-reported) support for the welfare state.

3.1.1 Social capital

During the 1990s, social scientists have become increasingly interested in social capital and in particular in trust. Trust is seen to facilitate economic transactions in a world in which not all eventualities can be dealt with through formal contracts. Cross-country

studies, such as the one by Knack and Keefer (1997), show a positive correlation between the average level of trust in a country and the country's economic performance. Furthermore, it has been remarked that countries with high average levels of trust, such as the Nordic countries, tend to have homogenous populations. In this subsection, we focus on observational studies that test whether there is such a link between ethnic diversity and trust *within* countries. The effect of ethnic diversity on trust has also been studied in experiments. We will come back to these studies in subsection 3.2 below.

Using individual-level data from the U.S. General Social Survey for the years 1974–1994, Alesina and La Ferrara (2000) regress (self-reported) participation “in a variety of social activities including recreational, religious, civic, and educational groups” (p. 850) on individual and community characteristics (at the level of metropolitan sampling areas and primary metropolitan sampling areas). Ethnic and racial fractionalization, the variables of interest, are calculated from the 1990 Census, and this value is used for the whole period. In probit models, which include state- and time- fixed effects, Alesina and La Ferrara “find that, after controlling for many individual characteristics, participation in social activities is significantly lower in more unequal [in terms of income] and in more racially or ethnically fragmented localities.” (p. 847). The negative association is rather strong: an increase of one standard deviation in a community's racial fractionalization is associated with a probability of participating that is lower by eight percentage points. This is twice as large as the reduction in the participation probability associated with having a child below the age of five. The negative partial correlation is strongest for groups in which interaction is typically more frequent (youth clubs and churches, as opposed to, say, professional associations). Finally, Alesina and La Ferrara find “that those individuals who express views against racial mixing are less prone to participate in groups the more racially heterogeneous their community is” (p. 847). However, the study cannot identify to what degree the lack of participation is *caused* by racial attitudes, and to what degree it is the other way around.

In a related study and using the same data from the General Social Survey, Alesina and La Ferrara (2002) find that the probability that respondents declare that “most people can be trusted” is lower in more racially heterogeneous communities. They also find that income inequality in a community is negatively related with inter-personal trust. By contrast, they do not find a significant association between trust and ethnic (as opposed to racial) diversity. “The magnitude of the coefficient on racial fragmentation is substantial. Moving from the most homogeneous MSA where racial fragmentation assumes the value of 0.06 to the most heterogeneous where it is 0.61, the likelihood of trusting others would fall by 12 percentage points, i.e. about 30 per cent of the mean. Starting from the sample mean, an increase by one standard deviation in racial fragmentation decreases

the probability of trusting others by 3 percentage points, i.e. almost 10 per cent of the mean value. This is larger than the effect of having had a traumatic experience in the last year, and almost the same size as the effect of being divorced or separated” (p. 222). Interestingly, Alesina and La Ferrara find that only trust in other people is lower in more racially heterogeneous communities, not trust in institutions. As in their 2000 article, they find that the negative association is strongest for people who declare themselves to be against racial integration.

Costa and Kahn (2003) show additional evidence on the links between ethnic diversity and social capital in the U.S. They measure social capital using information from several data sets on volunteer activity, organizational membership and activity, and entertaining and visits with friends, relatives, and neighbours. In a pooled cross-section with year-specific intercepts, and at the level of metropolitan areas, they find significant negative association with ethnic diversity only for the probability of volunteering (DDB Lifestyle Survey, 1975–98) and for membership in non-church organizations (American National Election Survey, ANES, 1952–72). As for the size of the association, “the increase of 0.092 in the racial fragmentation index from 1950 to 1970 predicts a decline in nonchurch group membership in the ANES of 0.04.” (p. 105). Interestingly, even though they use the same data from the General Social Survey as Alesina and La Ferrara, they do not arrive at the same results. Costa and Kahn find a positive (but insignificant) relationship between racial fractionalization and the probability of being a member of an organizations; as for trust, they do find a negative point estimate, but the estimated coefficient is not statistically significant.—Finally, Costa and Kahn study whether *changes* in ethnic diversity can predict changes in social capital in the period between 1952 and 1998. They find that the increase in community heterogeneity can predict between 6 and 13 percent of the decline in social capital from the mid-1970s to the end of the 1980s (or the end of the 1990s in some models). Much more of the decline can be predicted by the increase in wage inequality.

In two recent studies, Hungerman (2007, 2008) examines the link between ethnic diversity and charitable activity by church congregations in the U.S. In a first article, Hungerman (2008) finds that congregations that are all-white tend to be less active (relative to congregations that are not all-white) in communities with a higher share of blacks in the population. Hungerman (2008) acknowledges the possibility of omitted variables bias (e.g., white congregations in black communities may be poorer on average), but his results survive a number of robustness checks: he controls for the income of the congregation members, includes community fixed effects, and instruments the current share of blacks in the community by the lagged share. Moreover, it is only charitable activity that tends to be lower in communities with a higher share of blacks; other spending by the

congregations is if anything higher.

In a second study on charitable church activity, Hungerman (2007) studies how the crowd-out of private charity by government spending differs by the ethnic diversity of the community. He uses the natural experiment of an expansion of a federal welfare program (Supplemental Security Income, SSI), mandated by the U.S. Supreme Court in 1991. His outcome variable is charitable activity by the United Methodist church for the years 1984 to 2000. Hungerman confirms that government spending crowds out charitable activity, but he finds that this average effect is driven almost exclusively by ethnically homogeneous communities. In ethnically diverse communities, by contrast, there is only a modest effect, which suggests that charity in diverse communities is driven more by “warm glow” motive (pleasure derived from the act of giving *per se*) than by pure altruism.

3.1.2 Attitudes towards the welfare state

A number of other studies have examined the effect of ethnic diversity on people’s support for the welfare state, based on answers to survey questions.³

The earliest studies use data from the U.S., mostly from the General Social Survey (section 3.1.2). More recently, the scope of countries covered has widened. Soroka et al. (2004) study Canada (section 3.1.2). We (the authors of the present article), have also contributed to this literature. In a recent book, Roemer et al. (2007) study the U.S., but also three European countries: Denmark (see also Roemer and van der Straeten 2006), France, and the United Kingdom. We choose not to review these studies here: for a short review of their main results, see Lee et al. (2006), and for a non-technical review (in French) summing up their results for France, see Roemer et al. (2005). We rather choose to present here two more recent contributions for Europe, using data from the European Social Survey and from the German Socio-Economic Panel (section 3.1.2).

United States Using data from the 1986 National Election Study, the 1994 General Social Survey, and the 1991 National Race and Politics Study, Gilens (1995, 1996, 1999)

³A methodological point on studying individual *attitudes* may be in order. What is observed in these studies are self-reported answers to survey questions about trust or the welfare state. While it is straightforward to estimate how these self-reported answers covary with other variables of interest such as ethnic diversity in a respondent’s region, the ultimate interest is not in answers to survey questions but in the respondent’s true attitude towards the welfare state. The true attitude is unobserved, however, and so statements about it have to rely on untestable identifying assumptions. If additional data are available, cross-validation can be used to justify these assumptions; in the literature on the determinants of life satisfaction, such cross-validation has relied on asking other people to independently judge a person’s life satisfaction, or on physiological and neurological evidence. This cross-validation has increased people’s confidence that something meaningful is estimated when using “subjective variables,” at least for the literature studying life satisfaction (see the survey by Senik 2005).

shows that racial stereotypes are important predictors of the support for welfare spending among white Americans.⁴ Moreover, Gilens (1999) documents that black people are over-represented in the media coverage of poverty, and in particular in the least sympathetic stories. Consistent with this overrepresentation in the media, white Americans tend to overestimate the proportion of poor people who are black.

With data from the General Social Survey for the period from 1972 to 1993, Luttmer (2001) shows that in the U.S. people are more likely to express support for welfare spending if they live in a neighborhood where the share of people of their own race among welfare recipients is high. This is true whatever the economic situation of the respondents, even among wealthy people who have only a very small risk of being welfare recipients themselves. Luttmer therefore interprets this relationship as evidence for group loyalty, that is, of solidarity being stronger towards members of one's own ethnic group than towards members of other ethnic groups. He links his results to the findings on aggregate public spending: "I find that over 30 percent of the variation in levels of welfare benefits across states can be explained by applying my estimates of interpersonal preferences to the differences in the demographic composition of states. Hence, interpersonal preferences seem to transform difference in racial composition into differences in redistribution within the United States" (Luttmer 2001, 502).

Alesina et al. (2001), in their study discussed at greater length above, have a short section on individual attitudes as well; like Gilens and Luttmer, they use data from the General Social Survey. Alesina et al. find that blacks are more supportive of welfare spending. For whites, they confirm Gilens' findings that racial attitudes are correlated with support for welfare spending: whites who believe that blacks are lazy tend to be less, and whites who have had a black person over for dinner more supportive of welfare spending. By contrast, support among whites is not significantly associated with the share of blacks in the population of the respondent's state. This difference with Luttmer's results may be due to the fact that Luttmer looks at shares of blacks among welfare recipients and not in the population as a whole.

Lind (2007), whose study has already been reviewed above in the section on aggregate outcomes, also shows evidence from the General Social Survey. For the period 1972–2002 he confirms that blacks tend to be much more supportive of welfare spending than whites, even controlling for a number of observable characteristics. He also confirms that white respondents who do not mind having or actually have had an African-American at home tend to be more supportive of welfare spending. On a shorter sample from 1996 to 2002,

⁴Welfare is understood in the narrow sense of means-tested, in-cash assistance; the biggest program of this kind was Aid to Family with Dependent Children (AFDC).

Lind shows that blacks who feel closer to blacks tend to be more, and blacks who feel closer to whites less supportive of welfare spending. For whites the interaction terms have the expected (opposite) signs, but are insignificant.

Finally, a recent article by Keely and Tan (2008) uses classification and regression trees to detect homogenous subgroups with respect to attitudes towards welfare spending and income redistribution. These exploratory techniques are a useful complement to the models reviewed so far, which focus on the testing of hypotheses and typically assume away the problem of model uncertainty (that is, they explore at best a few interactions in an ad-hoc fashion; see the critique by Brock and Durlauf (2000) mentioned above in the context of studies on aggregate spending). Applying these techniques on data from the General Social Survey for the period 1978–2000, Keely and Tan confirm that the race of the respondent is important for classifying people with respect to attitudes towards the welfare state.

Canada Contrary to Luttmer (2001)’s result for the U.S., Soroka et al. (2004) find for Canada that “the link [between regional ethnic diversity and support for social programs] is weak at best” (p. 50); “moving from 100% majority to 50% majority leads to a decrease in aggregate support for unemployment and welfare of about .0025%” (p. 51). Their data are from the first wave of the Canadian “Equality, Security, and Community” Study; regional demographic data are from the 1996 Canadian Census, and ethnic diversity is measured at the level of census tracts or census subdivisions. They measure support for the welfare state in three domains: employment insurance and welfare (here, several items are combined into a single index “by a combination of factor analysis and reliability tests” (p.44), details of which are not provided); health; and pensions. As noted, ethnic diversity in the respondent’s region is not correlated with support for the welfare state in any of these three domains. By contrast, Soroka et al. confirm the finding of Alesina and La Ferrara (2002) that members of the ethnic majority tend to report less trust in other *people* (here measured by a question on how likely it would be to get back a wallet lost in the neighbourhood) the greater the share of ethnic minorities in their region. As in Alesina and La Ferrara’s study, trust in government *institutions* is found to be unrelated to ethnic diversity.

Europe Senik, Stichnoth, and Van der Straeten (2009) use survey data for 22 European countries from the 2002/2003 round of the European Social Survey to investigate the link between immigration and support for the welfare state. They want to find out (1) how the perceived *presence* of immigrants is related to natives’ support for the welfare state, and

how this relationship varies (2) with natives' *attitudes* towards immigrants and (3) across countries. The specificity of the study lies in its emphasis on parameter heterogeneity and in the use of European instead of U.S. or Canadian data.

Their main result is that, for Europe as a whole, the association between the perceived *presence* of immigrants and natives' support for the welfare state is weak at best. They do find some evidence that natives who perceive the share of immigrants in the population to be high tend to be less supportive of the welfare state. But the association is very weak and not even statistically significant for two of their three dependent variables. An increase in the perceived share of immigrants of one standard deviation (about 16 percentage points) is associated with a decrease in the probability of supporting the welfare state of about one percentage point. This is small compared with the associations that they find for other covariates such as income or education.

However, although they find only weak evidence for a negative association between immigration and natives' support for the welfare state in general, it may well be that a negative association does exist for certain sub-groups of natives. In a second step they therefore interact the variables measuring the actual or perceived presence of immigrants with variables measuring natives' *attitudes* towards immigrants. Exploring this parameter heterogeneity is made possible by the rich set of such attitudinal variables in the European Social Survey.

As expected, Senik, Stichnoth, and Van der Straeten (2009) find that the association between the perceived share of immigrants and support for the welfare state is most negative for natives who both dislike immigrants and express concern about the economic consequences of immigration. By contrast, the association is positive for those natives who view immigration positively along both dimensions. There is no clear evidence concerning the relative importance of the two channels: often, the parameter estimates for the two intermediate types are not significantly different from another.

When Senik, Stichnoth, and Van der Straeten (2009) include the two channels also as main effects (instead of interacting them with the perceived share of immigrants), they find that attitudes towards immigration and attitudes towards the welfare state are strongly associated, but that this association is little affected by the (perceived) *presence* of immigrants.

A third contribution of their article is to look at differences across countries within Europe. They find that the small average effect masks considerable heterogeneity across countries. However, the general result is preserved: the practical significance of the association between the perceived presence of immigrants and natives' support for the welfare state is small even in the countries for which the estimated coefficients are largest.

Based on data from the 1997 and 2002 waves of the German Socio-Economic Panel and from official statistics, Stichnoth (2009) studies whether natives are less supportive of state help for the unemployed in regions where the share of foreigners among the unemployed is high. Unlike in previous studies using (repeated) cross-sections, the models are estimated using individual-level panel data, which allows more convincing identification of a causal effect.

The main result is that there is at best weak evidence that German natives' support for the unemployed is influenced by the regional share of foreigners among the unemployed. Mean comparisons show that people in areas with a high share of foreigners among the unemployed do tend to be less supportive of state help for the unemployed. However, much of this difference is driven by common influences such as income or East German origin. Once these individual characteristics are controlled for, the share of foreigners among the unemployed is still negatively associated with natives' support for the unemployed, but the association is rather weak compared to other variables such as income, self-employment, or East German origin. Stichnoth (2009) finds that a one standard deviation increase in the share of foreigners among the unemployed is associated with a reduction in support for financially helping the unemployed of about two percent of the standard deviation of the dependent variable.

3.2 Experiments

The studies using survey data reviewed in the preceding section exploit regional variation in ethnic diversity and study whether respondents are more or less supportive of the welfare state in regions with greater ethnic diversity. Ethnic diversity is measured either in the population as a whole, or in some sub-group such as the recipients of AFDC transfers (Luttmer 2001) or the registered unemployed (Stichnoth 2009). Since individuals self-select into regions, the estimated partial correlation between ethnic diversity and self-reported support for the welfare state can be interpreted as a causal effect only under the strong assumption that there is no bias from omitted variables. reverse causality? Perhaps not an issue here since dependent variable is at individual level This strong assumption can be relaxed somewhat by using only temporal (that is, within-individual) variation, which requires individual-level panel data as in Stichnoth (2009).

As a reaction to these limitations of observational studies, the effects of ethnic diversity on trust and on altruism have been increasingly studied in experiments in the last few years, mainly in trust, ultimatum, and dictator games.⁵

⁵In an *ultimatum game*, two players interact to divide a sum of money. The first player proposes

Experiments offer the advantage that the experimenter can create exogenous variation in the experimental conditions through randomization; as a result, experiments are the gold standard for the identification of causal effects. A second, related advantage is that because of this control of the experimental conditions, experiments can be used to explore the channels through which an effect operates. In the case of ethnic diversity, experiments can shed some light on whether people behave differently towards members of another ethnic group because of preference-based or statistical discrimination. Third, experiments allow the observation of behaviour, as opposed to self-reported attitudes in a survey. These reasons explain why experiments have become increasingly popular in the literature on the effects of ethnic diversity on support for redistribution.

However, a major drawback of experiments is the loss of external validity: because the subjects are often a non-random sample of some specific population (typically, undergraduate students), and because their behaviour is observed in the artificial setting of the experiment, it is unclear to what extent the results can be generalized. See Levitt and List (2007) for a discussion.

In this subsection we briefly review this very active strand of the literature. We focus on experiments in *economics* and leave aside the much older and larger literature in social psychology on the effects that group affiliation in general and race and ethnicity in particular have on individual behaviour. Moreover, we review only experiments on the effects of *inter-ethnic* relations and not on comparisons between groups of different but more or less homogenous cultural or ethnic backgrounds. A well-known study of this second type is by Henrich et al. (2001), who study behaviour in games in fifteen ethnic groups around the world. Furthermore, we focus not on groups in general, but on *ethnic* groups.⁶ Finally, because of this focus on ethnic *relations* we do not report results on whether particular ethnic or racial groups are generally more or less trusting or altruistic in these games (e.g., Eckel and Grossman 2001); what we are interested in is whether people *condition* their strategies on the (supposed) ethnic identity of the other player.

how to split this sum, and the second player either accepts or rejects the proposal. If the second player rejects the proposal, neither player receives anything. In a *dictator game*, the first player receives a sum of money and can then decide on how much of this money to keep and on how much of it to pass on to the second player. The second player's role in this game is passive, that is, he or she cannot reject the proposed division of money. Finally, in a *trust game* the first player receives a sum of money, of which he or she can send some share to the second player. The money sent is typically increased by the experimenter, and the second player then decides on how much of this money he or she wants to send back to the first player.

⁶The effects of group-affiliation per se have often been studied using Tajfel et al. (1971)'s minimal group paradigm; that is, the experimenters create otherwise meaningless groups for the purpose of the experiment, for example by random assignment. Studies of this type are numerous in social psychology; two recent examples in economics are Güth et al. (2007) and Chen and Li (2009). Goette et al. (2006) use random assignment to a group that *is* meaningful (platoons in the Swiss army).

In economics, the experimental literature on the effects of ethnic group affiliation on individual behaviour begins with a study by Glaeser et al. (2000). They show that in a trust game played with Harvard undergraduates the recipient returns more when he or she is of the same race or nationality as the sender. By contrast, the amount sent does not differ by the race or nationality of the recipient.

Fershtman and Gneezy (2001) use a series of experiments to study trust and altruism between Ashkenazi and non-Ashkenazi Jews in Israel (the latter group tends to fare less well in economic and social terms). The study is interesting because they use different games to shed some light on whether the behaviour towards a particular group is driven by statistical or by preference-based discrimination. The experiment is conducted with university students, and group affiliation is inferred by the surname of the players (the players do not actually meet).

Fershtman and Gneezy find that non-Ashkenazi Jews receive smaller amounts of money in a trust game. However, two pieces of evidence suggest that this discrimination does not seem to be based on preferences: first, even non-Ashkenazi Jews themselves tend to send less money when the receiver is a non-Ashkenazi Jew. Second, in a dictator game, there is no evidence that non-Ashkenazi Jews are discriminated against. Fershtman and Gneezy therefore conclude that the smaller amounts sent in the trust game reflect statistical discrimination—which, in the experiment, is based on *mistaken* beliefs, for non-Ashkenazi Jews are found to be no less trustworthy than Ashkenazi Jews: the amount that the receiver returns is not significantly different between the two groups.

Finally, the study confirms that gender matters for trust and altruism. More to our point, the importance of group affiliation also differs by gender. Fershtman and Gneezy find that there is no statistically significant evidence of discrimination in the trust game when the receiver is a woman. Moreover, women did not discriminate between Ashkenazic and non-Ashkenazi male players.

In a trust game with small businessmen of Belgian and Turkish ethnic origin in the city of Ghent in Belgium, Bouckaert and Dhaene (2004) find that “the average levels of trust and reciprocity are independent of ethnic origin and, moreover, independent of the ethnic origin of the opposite party” (p. 869). In their study, ethnicity is inferred from the first names of the participants. Studying small businessmen (instead of students as in the studies by Glaeser et al. and by Fershtman and Gneezy) is interesting because it helps gauge whether the previous results were specific to students or valid more generally; moreover, trust is arguably more important in their professional lives for small businessmen than for students. Bouckaert and Dhaene speculate that the absence of discrimination in their study may be due to the fact that the participants “were equal with respect to

gender, socio-professional status, and place of residence, *and they knew this*. [...] This leads to the tentative conclusion that ethnic discrimination between interacting parties vanishes when enough other characteristics are equal.” (p.872; italics in the original).

Habyarimana et al. (2006) use a number of different games to identify the *mechanisms* through which ethnic group affiliation affects behaviour. They focus on the provision of public goods; observational studies have documented a negative correlation between ethnic diversity and the provision of (certain types of) such goods (see the references in the introduction). The experiments are conducted with 300 subjects in several slums of Kampala, Uganda, a city that is ethnically diverse and in which ethnicity is highly salient. Participants are drawn at random from the local population, inference from the sample to this population are possible. In some of the games, ethnic group affiliation is ascertained through face-to-face interaction; in other games, subjects are shown photos and videos.

Habyarimana et al. confirm that the degree of ethnic fragmentation of an area is negatively correlated with “whether, during the last six months, residents of the LC1 [an administrative district in Kampala] had organized community efforts in the area of crime prevention and security” (p.9).

Their main finding is that this negative correlation probably results from strategic concerns: “co-ethnics cooperate because they adhere to within-group norms and institutions that facilitate the sanctioning of individuals who fail to contribute to collective endeavors.” (p.3). By contrast, the link from ethnic group affiliation to the provision does *not* seem to pass through preferences or technology.

Concerning preferences, they ask subjects about their priorities for public goods provision and on how these public goods should be provided. It turns out that answers do not differ significantly across ethnic groups, which (at least in this particular context) casts doubt on the assumption that public good provision is lower because ethnic groups differ in their preferences over the type of public good that should be provided—an assumption that is made by Alesina et al. (1999), for instance. The second preference-related explanation is that people are less altruistic towards members of other ethnic groups; this argument of group-specific weights in an inter-dependent utility function is formalized in Lind (2007), among others. However, consistent with the earlier findings by Fershtman and Gneezy (2001), Habyarimana et al. (2006) do *not* find evidence that in a dictator game people are less altruistic towards members of other ethnic groups.

Concerning technology they conclude that “Nor are technological explanations central. Co-ethnics do not appear to cooperate with one another simply because cooperation is easier with people that share a common language or modes of interacting.” (p. 3).

However, this conclusion is based on evidence from two games in which two players have to work together to open a lock or to put together a jigsaw puzzle—that is, two rather specific forms of cooperation. In any case, although a common language and strong social ties do not seem to facilitate cooperation directly, they underlie what Habyarimana et al. call the “strategic explanation”: “co-ethnics benefit from stronger network linkages among them, linkages which make it possible for individuals to sanction co-ethnics that defect” (p.3).

Haile et al. (2008) conduct a trust game with students of two universities in South Africa. Their contribution is to disentangle the effects of ethnicity and of income on trust. They note that in the studies by Glaeser et al. and by Fershtman and Gneezy, the group that is discriminated against (blacks in the former and non-Ashkenazi Jews in the latter case) is weaker economically, so that it is not clear whether the senders in the trust game condition their strategies on the responders’ economic position or on their ethnicity. Somewhat surprisingly perhaps given the history of race relations in South Africa, Haile et al. find no evidence that either income or race by themselves affect the amounts transferred. However, there is evidence that whites with low income tend to transfer less money to high-income blacks, and that low-income blacks transfer less to high-income whites—Haile et al. term this “cross-racial envy”. Low-income blacks do not expect lower returns from high-income whites, so in their case the conditioning of strategies appears to reflect preference-based and not statistical discrimination. Consistent with this, low-income blacks tend to return less money to high-income whites; this is the only statistically significant evidence for discrimination at the stage of the responder. As for low-income whites, they expect less reciprocity from high-income blacks, but this expectation turns out to be wrong in the experiment.

Whereas the studies reviewed so far have used one-shot versions of the trust game, Bornhorst et al. (2006) conduct repeated trust games; moreover, in their version of the trust game senders choose the players to whom they transfers. The subjects of the study are Ph.D. students at the European University Institute in Florence. Bornhorst et al. find evidence that participants condition their strategies on whether the other player comes from a Northern or a Southern European country. Southern Europeans tend to receive less money; more importantly, they are chosen less frequently as receivers. These differences tend to increase in the course of the game; they are mainly driven by the behaviour of Northern Europeans. Bornhorst et al. argue that Southern Europeans tend to be contacted less frequently and tend to receive less money because they exhibit lower levels of trust themselves. Since generous transfers tend to be rewarded by other generous transfers, the smaller transfers that Southern Europeans tend to make hamper the build-up of trust towards them.

Falk and Zehnder (2007) study whether people condition their strategies in a trust game on the city district that the recipient lives in. They work with about 1000 subjects, drawn at random from the population of Zurich. As a result, Falk and Zehnder (2007) have a more heterogeneous subject pool than studies working with university students; moreover, the large number of participants allows them to study the individual correlates of trust and discrimination. Falk and Zehnder find that about 55% of their subjects discriminate by district, whereas 45% send the same amount to all districts. The district that receives the highest average amount gets about 11% more than the district with the lowest average amount. They find that districts with higher social status (measured by median income and share of people with at least high school education) tend to receive more. There is also evidence for in-group favoritism at the level of residential districts: controlling for other factors, people tend to send about ten percent more to recipients in their own district. Most important for the focus of this survey, Falk and Zehnder find that recipients living in districts with a higher share of foreigners tend to receive less.

In an additional experiment conducted with the help of a local newspaper, Falk and Zehnder show that people (who did not participate in the original experiment) can predict well which districts would receive the highest average amounts; this suggests that beliefs about the trustworthiness associated with certain districts (and, arguably, with socio-economic and immigration status) are widely shared among the residents of Zurich.

Interestingly, Falk and Zehnder find evidence that the discrimination is not purely preference-based. First, the amount of money invested in a district is positively correlated with the return expected from that district. Moreover, unlike in the study by Fershtman and Gneezy (2001), differences in trust are at least partly justified by actual differences in trustworthiness and are not just mistaken stereotypes.

Finally, as for the individual correlates of discrimination, Falk and Zehnder find that older people tend to discriminate more, whereas more educated people and people with children tend to discriminate less. Falk and Zehnder do not find significant differences between men and women, between foreigners and Swiss nationals, or for participants that grew up as an only child.

Fong and Luttmer (2009) study whether race affects actual charitable giving and self-reported attitudes towards private and public spending to help the victims of hurricane Katrina. They present slide shows and audio stories to 1300 respondents, roughly representative of the U.S. adult population. Through the images and the audio commentary, Fong and Luttmer manipulate respondents' perceptions of the income, deservingness, and race of victims.

Fong and Luttmer study two types of outcome variables. First, to measure how the race

of the recipients affects actual behaviour in a dictator game, they ask respondents to split 100 dollars between themselves and a charity helping victims. The second type of outcome variables are answers to survey questions on whether government and charities should spend more or less helping the victims of the hurricane.

Concerning actual behaviour, Fong and Luttmer find that *on average*, there is no influence of race or deservingness on the amount that people give: “in the overall sample [...] the response is -\$2.2. Thus, the point estimate suggests that there is little effect of victims’ race on giving, but given that the 95% confidence interval on this estimate ranges from about -\$10 to \$5, we cannot rule out a moderately large racial bias in giving in the overall sample.” (p. 11). However, for sub-groups there *is* evidence of a racial bias: whites who strongly identify with their own racial group tend to give less when the victims are portrayed as predominantly black; blacks who strongly identify with their own racial group tend to give more in the same treatment. Fong and Luttmer conclude that “social identity is an important predictor of racial bias, but that race by itself is not a good proxy for social identity.” (p. 26)

As for *attitudes* towards public and private spending to help the victims, race and deservingness are important predictors. Unlike for actual behaviour, this holds for everybody, not just for the subgroup of people who strongly identify with their own racial group. Fong and Luttmer find that whites tend to be less supportive of helping the victims if these are portrayed as black.

To explain why race matters more for attitudes towards public assistance than for own giving to a private charity, Fong and Luttmer argue that the latter may be seen to better pre-select individuals who receive assistance: “respondents [may] believe that Habitat for Humanity [the charity to which respondents can contribute in the dictator game] only provides assistance to deserving individuals, while government programs may not. If this is true, then there may not be enough variation in perceptions of deservingness of recipients in our charity experiment. This could also explain why race does not matter in our charity experiment, if racial bias is mediated by perceptions of relative deservingness of white and black recipients, as some have argued (Gilens, 1999)” (p.3).

4 Conclusion

This paper has surveyed the empirical literature on the effects of ethnic diversity on natives’ attitudes towards redistribution. As mentioned in the introduction, there is growing concern about a trade-off between a generous immigration policy and a generous welfare state. For instance, in a well-known contribution Alesina and Glaeser (2004, 11)

have made the case that immigration and ethnic diversity help explain why the U.S. does not have a European-style welfare state.

However, our main conclusion from this survey is that most studies do not point to a quantitatively important role for ethnic diversity in shaping natives' preferences for redistribution. In most studies, the association is much weaker than for other factors such as own income (current or expected) or beliefs about the role of effort versus luck in determining this income.

Moreover, it seems that the sizeable negative association between ethnic diversity and support for redistribution that is sometimes found in U.S. studies does not generalize to Canada or Europe. However, the evidence for countries other than the U.S. is scarce so far, and there is certainly need for further research.

References

- Alesina, A., R. Baqir, and W. Easterly (1999). Public goods and ethnic divisions. *Quarterly Journal of Economics* 114(4), 1243–84.
- Alesina, A., R. Baqir, and W. Easterly (2000). Redistributive public employment. *Journal of Urban Economics* 48, 219–241.
- Alesina, A., A. Devleschawuer, W. Easterly, S. Kurlat, and R. Wacziarg (2003). Fractionalization. *Journal of Economic Growth* 8(2), 155–94.
- Alesina, A. and E. Glaeser (2004). *Fighting Poverty in the U.S. and in Europe: A World of Difference*. New York: Oxford University Press.
- Alesina, A., E. Glaeser, and B. Sacerdote (2001). Why doesn't the U.S. have a European-style welfare system? *Brookings Papers on Economic Activity* 2, 187–278.
- Alesina, A. and E. La Ferrara (2000). Participation in heterogeneous communities. *Quarterly Journal of Economics* 115(3), 847–904.
- Alesina, A. and E. La Ferrara (2002). Who trusts others? *Journal of Public Economics* 85, 207–234.
- Alesina, A. and E. La Ferrara (2005a). Ethnic diversity and economic performance. *Journal of Economic Literature* 43, 762–800.
- Alesina, A. and E. La Ferrara (2005b). Preferences for redistribution in the land of opportunities. *Journal of Public Economics* 89(5-6), 897–931.
- Annett, A. (2001). Social fractionalization, political instability, and the size of government. *IMF Staff Papers* 48(3), 561–92.
- Arcand, J.-L., P. Guillaumont, and S. Guillaumont Jeanneney (2000). How to make a tragedy: on the alleged effect of ethnicity on growth. *Journal of International Development* 12, 925–38.
- Banting, K., R. Johnston, W. Kymlicka, and S. Soroka (2007). Do multiculturalism policies erode the welfare state? An empirical analysis. In W. Kymlicka and K. Banting (Eds.), *Multiculturalism and the Welfare State: Recognition and Redistribution in Contemporary Democracies*. Oxford University Press.
- Banting, K. and W. Kymlicka (2004). Do multiculturalism policies erode the welfare state? In P. Van Parijs (Ed.), *Cultural Diversity versus Economic Solidarity*, Brussels. De Boeck.

- Banting, K. G. (2005). Multiculturalism and the welfare state: Recent evidence about ethnic diversity, multiculturalism policies and redistribution. Unpublished, Queen's University.
- Bénabou, R. and E. A. Ok (2001). Social mobility and the demand for redistribution: the POUM hypothesis. *Quarterly Journal of Economics* 116(2), 447–87.
- Bornhorst, F., A. Ichino, K. Schlag, and E. Winter (2006). Trust and trustworthiness among Europeans: South - North comparison. Unpublished working paper, European University Institute, Florence.
- Bossert, W., C. D'Ambrosio, and E. La Ferrara (2006). A generalized index of fractionalization. Technical report, IGIER - Università Bocconi.
- Bouckaert, J. and G. Dhaene (2004). Inter-ethnic trust and reciprocity: results of an experiment with small businessmen. *European Journal of Political Economy* 20, 869–886.
- Brock, W. A. and S. N. Durlauf (2000). Growth economics and reality. Working paper 8041, NBER.
- Bruk, S. I. and V. S. Apenchenko (Eds.) (1964). *Atlas Narodov Mira*. Moscow: Glavnoe upravlenie geodezii i kartografii gosudarstvennogo geologicheskogo komiteta SSSR and Institut etnografii im. H. H. Miklukho-Maklaia, Akademiia nauk SSSR.
- Chandra, K. (2001). Cumulative findings in the study of ethnic politics. *apsa-cp (Newsletter of the Organized Section in Comparative Politics of the American Political Science Association)* 12(1), 7–11.
- Chen, Y. and S. X. Li (2009). Group identity and social preferences. *American Economic Review* 99, 431–57.
- Collier, P. (2001). Implications of ethnic diversity. *Economic Policy* 32, 129–66.
- Corneo, G. (2001). Inequality and the state: comparing U.S. and German preferences. *Annales d'Economie et de Statistique* 63-64, 283–96.
- Corneo, G. and H. P. Grüner (2000). Social limits to redistribution. *American Economic Review* 90(5), 1491–1507.
- Corneo, G. and H. P. Grüner (2002). Individual preferences for political redistribution. *Journal of Public Economics* 83(1), 83–107.

- Costa, D. L. and M. E. Kahn (2003). Understanding the American decline in social capital, 1952–1998. *Kyklos* 56, 17–46.
- Cutler, D. M., D. W. Elmendorf, and R. Zeckhauser (1993). Demographic characteristics and the public bundle. *Public Finance* 48, 178–198.
- Durlauf, S. N. (2001). [Why doesn't the United States have a European-style welfare state?]. Comments and discussion. *Brookings Papers on Economic Activity*, 255–277.
- Eckel, C. C. and P. J. Grossman (2001). Chivalry and solidarity in ultimatum games. *Economic Inquiry* 39, 171–188.
- Falk, A. and C. Zehnder (2007). Discrimination and in-group favoritism in a citywide trust experiment. Discussion paper 2765, IZA.
- Fearon, J. (2003). Ethnic and cultural diversity by country. *Journal of Economic Growth* 8(2), 195–222.
- Fedderke, J., J. Luiz, and R. de Kadt (2008). Using fractionalization indexes: deriving methodological principles for growth studies from time series evidence. *Social Indicators Research* 85, 257–278.
- Fershtman, C. and U. Gneezy (2001). Discrimination in a segmented society: an experimental approach. *Quarterly Journal of Economics* 116, 351–377.
- Fong, C. M., S. Bowles, and H. Gintis (2006). Strong reciprocity and the welfare state. In S. Kolm and J. Mercier-Ythier (Eds.), *Handbook on the Economics of Giving, Altruism, and Reciprocity*, Vol. 2, pp. 1439–1464. Amsterdam: Elsevier.
- Fong, C. M. and E. F. P. Luttmer (2009). What determines giving to hurricane Katrina victims? Experimental evidence on racial group loyalty. *American Economic Journal: Applied Economics* 1, 64–87.
- Gilens, M. (1995). Racial attitudes and opposition to welfare. *Journal of Politics* 57(4), 994–1014.
- Gilens, M. (1996). "Race coding" and white opposition to welfare. *American Political Science Review* 90(3), 593–604.
- Gilens, M. (1999). *Why Americans Hate Welfare: Race, Media, and the Politics of Anti-Poverty Policy*. Chicago: University of Chicago Press.
- Gitlin, T. (1996). *The Twilight of Common Dreams: Why America Is Wracked by Culture Wars*. Owl Books.

- Glaeser, E. L., D. I. Laibson, J. A. Scheinkman, and C. L. Soutter (2000). Measuring trust. *Quarterly Journal of Economics* 115, 811–846.
- Goette, L., D. Huffman, and S. Meier (2006). The impact of group membership on cooperation and norm enforcement: evidence using random assignment to real social groups. *American Economic Review* 96, 212–216.
- Goldin, C. and L. Katz (1999). Human capital and social capital: The rise of secondary schooling in America, 1910 to 1940. *Journal of Interdisciplinary History* 29(4), 683–723.
- Goodhart, D. (2006). *Progressive Nationalism*. Demos.
- Güth, W., M. V. Levati, and M. Ploner (2007). Social identity and trust: An experimental investigation. Discussion paper, Max Planck Institute of Economics, Jena.
- Habyarimana, J., M. Humphreys, D. N. Posner, and J. Weinstein (2006). Why does ethnic diversity undermine public goods provision? An experimental approach. Discussion paper 2272, IZA.
- Haile, D., A. Sadrieh, and H. A. A. Verbon (2008). Cross-racial envy and underinvestment in South African partnerships. *Cambridge Journal of Economics* 32, 703–24.
- Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath (2001). In search of Homo Economicus: behavioral experiments in 15 small-scale societies. *American Economic Review Papers and Proceedings* 91, 73–78.
- Hungerman, D. M. (2007). Diversity and crowd-out: a theory of cold-glow giving. Working paper 13348, NBER.
- Hungerman, D. M. (2008). Race and charitable church activity. *Economic Inquiry* 46, 380–400.
- James, E. (1987). The public/private division of responsibility for education: an international comparison. *Economics of Education Review* 6(1), 1–14.
- James, E. (1993). Why do different countries choose a different public-private mix of educational services? *Journal of Human Resources* 28(3), 571–92.
- Keely, L. C. and C. M. Tan (2008). Understanding preferences for income redistribution. *Journal of Public Economics* 92, 944–961.
- Knack, S. and P. Keefer (1997). Does social capital have an economic payoff? A cross-country investigation. *Quarterly Journal of Economics* 112, 1251–1288.

- Kuijs, L. (2000). The impact of ethnic heterogeneity on the quantity and quality of public spending. Working Paper 00/49, International Monetary Fund, Washington, D.C.
- Laitin, D. and D. Posner (2001). The implications of constructivism for constructing ethnic fractionalization indices. *apsa-cp (Newsletter of the Organized Section in Comparative Politics of the American Political Science Association)* 12(1), 13–17.
- Lee, W., J. Roemer, and K. Van der Straeten (2006). Racism, xenophobia and distribution. *Journal of the European Economic Association* 4, 446–454.
- Levitt, S. D. and J. A. List (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives* 21, 153–174.
- Lind, J. T. (2007). Fractionalization and the size of government. *Journal of Public Economics* 91, 51–76.
- Luttmer, E. F. P. (2001). Group loyalty and the taste for redistribution. *Journal of Political Economy* 109(3), 500–28.
- McCarty, T. (1993). Demographic diversity and the size of the public sector. *Kyklos* 46(2), 225–40.
- Meltzer, A. H. and S. F. Richard (1981). A rational theory of the size of government. *Journal of Political Economy* 89(5), 914–27.
- Montalvo, J. G. and M. Reynal-Querol (2002). Why ethnic fractionalization? Polarization, ethnic conflict and growth. Unpublished working paper, Universitat Pompeu Fabra, Barcelona.
- Montalvo, J. G. and M. Reynal-Querol (2005a). Ethnic polarization, potential conflict, and civil wars. *American Economic Review* 95, 796–816.
- Montalvo, J. G. and M. Reynal-Querol (2005b). Fractionalization, polarization and economic development. *Journal of Development Economics* 76, 293–323.
- Mueller, D. C. and P. Murrell (1986). Interest groups and the size of government. *Public Choice* 48, 125–45.
- Persson, T. and G. Tabellini (2000). *Political Economics: Explaining Economic Policy*. MIT Press.
- Piketty, T. (1995). Social mobility and redistributive politics. *Quarterly Journal of Economics* 110(3), 551–84.

- Posner, D. (2002). Ethnic fractionalization in Africa: How should it be measured? What does it explain about economic growth? Unpublished working paper, University of California, Los Angeles.
- Posner, D. (2004a). The implications of constructivism for studying the relationship between ethnic diversity and economic growth. Unpublished working paper, University of California, Los Angeles.
- Posner, D. (2004b). Measuring ethnic fractionalization in Africa. *American Journal of Political Science* 48(4), 849–863.
- Poterba, J. M. (1997). Demographic structure and the political economy of public education. *Journal of Policy Analysis and Management* 16(1), 48–66.
- Poterba, J. M. (1998). Demographic change, intergenerational linkages, and public education. *American Economic Review Papers and Proceedings* 88, 315–320.
- Ravallion, M. and M. Lokshin (2000). Who wants to redistribute? The tunnel effect in 1990s Russia. *Journal of Public Economics* 76, 87–104.
- Ribar, D. C. and M. O. Wilhelm (1999). The demand for welfare generosity. *The Review of Economics and Statistics* 81, 96–108.
- Roemer, J., W. Lee, and K. Van der Straeten (2005). Politique sociale et immigration en France : une analyse des opinions et des comportements électoraux. *Economie Publique* 16, 3–36.
- Roemer, J., W. Lee, and K. Van der Straeten (2007). *Racism, xenophobia and distribution: A study of multi-issue politics in advanced democracies*. Harvard University Press & Russell Sage Foundation Press.
- Roemer, J. E. and K. van der Straeten (2006). The political economy of xenophobia and distribution: the case of Denmark. *Scandinavian Journal of Economics* 108, 251–78.
- Rushton, M. (2008). A note on the use and misuse of the racial diversity index. *Policy Studies Journal* 36, 445–459.
- Senik, C. (2004). When information dominates comparison: learning from Russian subjective panel data. *Journal of Public Economics* 88, 2099–2123.
- Senik, C. (2005). Income distribution and well-being: what can we learn from subjective data? *Journal of Economic Surveys* 19, 43–63.

- Senik, C. (2006). Ambition et jalousie : la perception du revenu d'autrui dans la "vieille Europe", la "nouvelle Europe" et les Etats-Unis. *Revue économique* 57, 645–653.
- Senik, C., H. Stichnoth, and K. Van der Straeten (2009). Immigration and natives' attitudes towards the welfare state: evidence from the European Social Survey. *Social Indicators Research* 91, 345–70.
- Soroka, S., K. Banting, and R. Johnston (2006). Immigration and redistribution in a global era. In P. Bardhan, S. Bowles, and M. Wallerstein (Eds.), *Globalization and Egalitarian Redistribution*. Princeton, NJ: Princeton University Press.
- Soroka, S. N., R. Johnston, and K. Banting (2004). Ethnicity, trust, and the welfare state. In P. Van Parijs (Ed.), *Cultural Diversity versus Economic Solidarity*, Brussels. De Boeck.
- Stichnoth, H. (2009). Does immigration weaken natives' support for the welfare state? Evidence from Germany. Discussion paper, ZEW Mannheim.
- Tajfel, H., M. Billig, R. Bundy, and C. Flament (1971). Social categorization and inter-group behavior. *European Journal of Social Psychology* 1, 149–177.
- Van Parijs, P. (2004). *Cultural diversity versus economic solidarity*. De Boeck.
- Vigdor, J. L. (2002). Interpreting ethnic fragmentation effects. *Economics Letters* 75, 271–276.
- Yeoh, K. K. (2001). Towards an index of ethnic fractionalization. Working Paper 2001-3, University of Malaya, Faculty of Economics and Administration, Kuala Lumpur.