Do affirmative action policies hinder economic development?

Graham K. Brown

University of Western Australia, Australia
graham.brown@uwa.edu.au

Christian Almer

University of Bath, UK

May 2, 2016

Abstract

In this paper we examine whether affirmative action policies systematically impact economic development in terms of growth and overall inequality. We use the synthetic control method to estimate the impact of the introduction of affirmative action policies on economic growth across 38 countries. Our results are largely a null finding, but an important one. We find a very small positive impact of affirmative action on growth, averaging a cumulative impact of 0.02 percentage points after five years and 0.04 percentage points after ten years, but also a small net increase in Gini coefficient of 0.42 point after five years and 0.77 points after ten years. These results are only robust to a randomised placebo test for national-level policies, however, where the effects are slightly stronger. We conclude that national-level affirmative action policies may promote long-run economic growth and certainly don't reduce it, but may also lead to slight deterioration in overall levels inequality, even if they are able to address the inter-group inequalities that they typically target.

1 Introduction

The aim of this paper is to provide the first cross-national estimates of the impact of affirmative action policies on growth and inequality conditioned on the type of policy and the socio-economic characteristic of the recipient group by utilising recent advances in econometric methods, including synthetic control methods and propensity score matching. Estimating the impact of affirmative action programmes on aggregate socio-economic outcomes, including growth, employment and inequality, is bedevilled by technical difficulties. Indeed, with the exception of the pioneering work of Jonathan Leonard in the 1980s, there have been few attempts to produce such estimates, and Leonard's own work is limited to the US context (see, e.g. Leonard, 1984, 1990).

Three main challenges face any attempt to estimate the impact of affirmative action internationally. The first is the range of different policies that come under the general rubric of 'affirmative action', from relatively common policies such as quotas in education to more extensive policy frameworks, such as with broad array of policies that have been implemented in Malaysia under the New Economic Policy. The second challenge is the range of demographic contexts in which such policies are applied, with beneficiary groups vary from relatively small population proportions, such as the Aboriginal population of Canada, which is estimated at 4.3 percent of the total population, to South Africa, where Black Africans constitute the vast majority of the population. Together, these two dimensions constitute a massive variation in policy scope and potential impact (see Brown and Langer, 2015). The third challenge is related to this variation but is more technical: determining an appropriate counterfactual. While subnational variation in the application of policies in places like the US has allowed for some estimation of the impact of educational quotas, estimating the impact of affirmative action for nationally-implemented policies, is much more problematic. In particular, countries tend to be rather heterogenous with substantial differences in various dimensions. That makes the identification of the appropriate counterfactual even more challenging than at the sub-national level. This project seeks to address these shortcoming through application of new econometric techniques applicable to such problematic counterfactuals: the synthetic control method developed by Abadie et al. (2010). In this paper we use this method to analyse the impact of affirmative action to test the impact of policy type and target group on growth and inequality outcomes. We find that affirmative action has some positive impact on economic growth when implemented at the national level, but may also lead to a slight increase in overall inequality. The findings on economic growth are robust to an alternative specification using propensity score matching.

2 Affirmative action and economic development

Given the conceptual and technical problems of policy variation and estimation of effects, it is perhaps unsurprising that there has been little attention in the literature to the relationship between affirmative action and macroeconomic outcomes, including growth and inequality. The micro-mechanisms through which affirmative action might impact overall economic growth and performance have been widely explored and tested in a range of contexts, however, with the US higher education system providing fertile empirical ground for such studies.

Broadly speaking, we can identify two theoretical hypotheses about the impact of affirmative action of economic development, which we term here the Misallo-

cation Hypothesis and the Resource Mobilisation Hypothesis.
Derived from neo-classical macroeconomic assumptions, the Misallocation Hypothesis asserts in the broadest sense that affirmative action policies are, by definition, causing market distortions in the labour market and will hence lead to sub-optimal outcomes. Sophisticated economic models of affirmative action in competitive economies typically point to precisely such outcomes (Moro and Norman, 2003). The ethical corollary of this argument is that affirmative action policies undermine the principle of 'merit' because they place less-qualified candidates from target groups into positions—employment, college places, etc.—that would otherwise be taken by better-qualified candidates from non-target groups. Feinberg (1985) uses simple mathematical models to contrast strategies for achieving ethnic equality that through affirmative action on the one hand; and through economic growth on the other hand, suggesting that while affirmative action may provide a faster route to equality, growth strategies may be better at providing employment opportunities for poorer groups in a kind of ethnic Rawlsian 'maximin' strategy (Rawls, 1999).
A specific instantiation of the Misallocation Hypothesis is the well-known 'mismatch' hypothesis, which asserts that university entry affirmative action in the US is counterproductive because it places weaker students from disadvantaged backgrounds into high-performing universities, leading to higher drop-out rates among those students with obvious negative consequences for their own careers and, by extension, the broader economy. The mismatch hypothesis was developed in the context of US higher education, where it has also been extensively empirically tested with some studies confirming the hypothesis (see, e.g. Hinrichs, 2014) and other rejecting it (see, e.g., Alon and Tienda, 2005).
Other micro-level studies in a range of contexts have sought to test misallocation-style hypotheses on more directly economic impacts of affirmative action. Deshpande and Weisskopf (2014), for instance, test the impact of affirmative action on labour productivity on the Indian railways, finding no evidence of a negative impact and some evidence of a positive impact, against the Misallocation Hypothesis.
The alternative broad hypothesis is one we term the Resource Mobilisation Hypothesis. The underlying assumptions undergirding this broad hypothesis is that affirmative action can have a beneficial long-term effect on economic development by providing marginalized groups the opportunity to contribute more effectively to the economy. The implicit assumption here is precisely the opposite of the Misallocation Hypothesis—that the status quo ante is a non-competitive economy with severe ethnic (gender, etc.) discrimination.
This broad hypothesis has been developed more theoretically to explain the persistence of group inequalities over time. John Goldthorpe's classic sociological model of the persistence of intergenerational inequality, for instance, is based on the observations that groups with lower educational attainment typically also have lower returns to education and hence have less incentive to invest in their own children's education, creating an intergenerational persistence of educational disparity (see Breen and Goldthorpe, 1997, for a formal statement of this model). Goldthorpe is primarily concerned with 'vertically' aligned social

classes, but Stewart and Langer (2008) develop a similar model for explaining the persistence of 'horizontal' inequalities between different ethnic groups. The motivating assumptions in their model is that marginalised groups typically suffer from deprivation across multiple dimensions of capital—human, financial, and social—and if the returns to these forms of capital are interactive, then improved access to one form of capital may not be sufficient to move that group out of poverty. Stewart and Langer point in particular to the role of social capital for the individual, arguing that even when marginalised groups have improved access to education or cash transfers though the removal of discriminatory barriers, they still face a challenge 'catching-up' with non-marginalised groups if they lack the social networks that those groups have. By extrapolation from this kind of model, affirmative action policies provides the route for these groups to establish better social capital and hence to be able to use their other resources more effectively, realising their full potential to contribute to economic development. The key point of the resource mobilisation hypothesis is that eradicating any existing discrimination is insufficient to maximise the potential for formerly marginalised groups to contribute economically.

3 Method and data

3.1 The Synthetic Control Method applied to multiple treatment units

As outlined above, the major challenge with estimating the impact of affirmative action policies on macro-economic outcomes is the lack of a clear counterfactual against which to judge impact. The Synthetic Control Method (SCM), developed by Abadie et al. (2010), is designed to estimate a 'synthetic' counterfactual for precisely such circumstances. The SCM estimates a counterfactual for the case of interest by constructing a synthetic version of that case as a weighted combination of control units from a 'donor' pool of cases unaffected by the intervention. Hence, for instance, in their expository study, Abadie et al. (2010) estimates the impact of Proposition 99 on cigarette consumption in California where the synthetic control was established using other states' cigarette consumption patterns.

In order to establish the weighting of 'donor' cases in the synthetic control, a set of predictor variables independent from the 'treatment' variable are specified and appropriate weights are established through a linear panel regression on the pre-treatment period. One way to motivate the SCM is hence as a generalization of the linear panel difference-in-differences model. Mathematically, the model assumes that the impact of the intervention can be estimated additively such that for outcome Y in country i at time t,

$$Y_{it}^{I} = Y_{it}^{N} + \alpha_{it} D_{it} \quad (1)$$

where Y_{it}^I is the outcome if the observation is exposed to the treatment; Y_{it}^N is the outcome in the absence of the intervention; D_{it} is 1 in the presence of the treatment and 0 otherwise, and α_{it} is the effect of the treatment. In the presence of the treatment, however, we are not directly able to observe Y_{it}^N. The SCM estimates this by finding the best approximation of a set of weights on the 'donor' pool in the pre-treatment period whereby the weighted average of the donor pool in the pre-treatment period replicates both the vector of covariates Z of the treated unit and the pre-treatment outcome in the treated unit. In the ideal case, for the treatment unit j_1, we would establish weights $(w_2^*, ..., w_{J+1}^*)$ on donor units $j_2, ..., j_{J+1}$ such that

$$Z_1 = \sum_{j=2}^{J+1} w_j^* Z_j \tag{2}$$

and

$$Y_1 t = \sum_{j=2}^{J+1} w_j^* Y_{jt} \forall \in \{T_1, ..., T\} \tag{3}$$

In reality, we are unlikely to find such an exact match. The Stata implementation of the SCM in the `synth` command[1] implements a data-driven approach to identifying the best approximation to 2 and 3 by minimising $(X_1 - X_0W)'V(X_1 - X_0W)$, where X_0 and X_1 are, respectively, the pretreatment vector of covariates and outcomes for the donor pool and the treated unit.

While the SCM hence provides the 'best' estimate of the non-treated counterfactual Y_{1t}^N, it does not easily provide any 'significance' test of this estimate. A typical approach adopted by Abadie et al. (2010) and others (see, e.g. Singh and Nilakantan, 2012) is to estimate a 'placebo' treatment on a control unit that is not exposed to the treatment in reality. With multiple cases, we generalise this into a set of placebos. Abadie et al. (2015) show that such 'placebos' are, at least in specific applications, robust to small variations in the donor pool.

The treatment variable we are interested in in this paper is the implementation of affirmative action policies. We identify 44 countries that have implemented some degree of affirmative action policy (see below for details) and create a synthetic control for each of these cases across two outcome variables: economic growth; and income inequality. Six of these countries drop out of the analysis due to data- or time-period constraints, leaving 38 countries included in the analysis.

As we are dealing with multiple cases, we are less concerned with the year-by-year impact within a particular country, and more with the overall impact across countries. We hence estimate the impact of the policies at two time periods—five years after implementation and ten years after implementation. We use a standard ten-year pre-treatment period to establish the weights in each individual case. In some cases, the pre-treatment period is shortened due to data paucity and in other (recent) cases we are only able to estimate a five-year post-treatment effect.

[1]Available online at http://web.stanford.edu/ jhain/synthpage.html

With multiple observations, rather than create a single 'placebo' to test significance, we create 100 random placebos based on a random identification of the placebo country at a random point in time. We exclude from this all countries that have ever had an affirmative action policy, e.g. if a country implemented affirmative action in 2002, we would still exclude a random placebo for that country in 1980. With this random set of placebos, we are able to more robustly test the significance of our pool of actual treatments through simple t-tests on the average differences at five years and ten years.

3.2 Identifying the treatment: Affirmative action policies

The treatment we are interested in is the implementation of affirmative action policies, but, as noted above, this general terminology covers a wide range of different policies that affect differing proportions of the population. Moses et al. (2014) have systematically scoped and identified affirmative action policies across six continent, which they then classify according to target groups and level of implementation—Local or National/Federal. We collapse the target groups into three categories: Gender, Ethnicity, and Class by, for instance, assigning the 'caste' basis of Indian affirmative action policies to the 'ethnic' category (although India drops out of the analysis due to its early start date and concomitant lack of data). We do not include policies targeted at disabled groups as these will typically affect a very small proportion of the population and are hence unlikely to have a systematic macroeconomic impacts. In some cases where there is lack of clarity over exactly when policies were initiated, we have had to make a judgment call based on available literature. In these cases, we also estimated policy impact with alternative start dates, none of which had any significant impact on the estimates. [2] In addition to these classifications derived from the Moses et al. (2014) dataset, we also estimate the proportion of the total population targeted by the policies. For gender policies, we have assumed this to be 50 percent; for ethnically-based policies we have used commonly available population estimates, including the CIA World Factbook, to estimate the target population. We have not provided any such estimates for class-based policies as the numerical basis of these policies were not clear. These different classification of policies are not used in estimating the treatment effect, but are used to analyse variation in the treatment effect between different types of policy. Table 1 lists the countries and policy types according to these criteria.

3.3 Baseline pre-treatment models

In order to operationalise the SCM, we need to establish the set of predictor variables that will be used to identify the factor weights among the donor cases. To do this, we use two simple models of, respectively, economic growth and

[2] Results available from authors on request

Table 1: Affirmative action policies

Country	Year of Implementation	Policy Type	Level	Data Issues
AUS	2001	Class; Ethnicity	Local	
AUT	2002	Gender;	National	
BEL	1999	Class;	Local	
BRA	2001	Gender; Class; Ethnicity	National	
CAN	1996	Gender; Ethnicity	Local	
CHN	1999	Class; Ethnicity	National	
COL	1993	Class;	Local	
ESP	2006	Ethnicity	Local	(a)
ETH	1993	Gender;	National	(b) (c)
FIN	2007	Gender; Ethnicity	Local	(a)
FJI	1997	Gender; Class; Ethnicity	Local (d) (e)	
FRA	2001	Class;	Local	
GBR	2004	Class;	Local	(a)
GHA	2001	Gender; Class;	Local	
GRC	1996	Ethnicity	National	
GRD	2001	Ethnicity	Local	(c) (d) (e) (h)
HUN	2005	Gender; Class; Ethnicity	National	(a)
IRL	1999	Class;	Local	
ISR	2002	Gender; Class; Ethnicity	Local	
JAM	1999	Gender;	Local	(h)
JPN	1969	Ethnicity	National	(f) (g)
KEN	2001	Gender;	Local	
LKA	1970	Class; Ethnicity	National	(f) (g)
MEX	2002	Gender; Class; Ethnicity	Local	
MKD	2001	Gender; Ethnicity	National	(c)
MYS	1970	Ethnicity	National	
NGA	1979	Ethnicity	National	(c) (f) (h)
NOR	1978	Gender;	National	
NZL	2001	Gender; Class; Ethnicity	Local	
ROM	2007	Gender; Class; Ethnicity	Local	(a)
SRB	2002	Ethnicity	National	(c) (g)
TJK	2001	Gender;	National	
TTO	2004	Class;	Local	(a) (h)
UGA	1990	Gender;	National	
UKR	2001	Gender; Class; Ethnicity	Local	
ZAF	1999	Gender; Ethnicity	National	
ZMB	2001	Gender;	National	(f)
ZWE	1995	Gender;	Local	

Data issues:
(a) Recent implementation: Only five-year estimates possible
(b) Missing trade data: Reduced model for GDP growth
(c) Missing human capital data: Reduced model for GDP growth
(d) Tropic missing from CID data: Climate manually set to tropical
(e) Coast missing from CID data: Small island nation; set at first percentile
(f) Missing gross capital formation data: Reduced model for GDP growth
(g) Early implementation: Shortened pre-treatment period
(h) Missing Gini estimates: No inequality model possible

income inequality. For economic growth, our predictor variables are based on a simple model of long-run economic growth in the tradition of Barro and Sala-i-Martin (see, e.g. Barro, 2003; Sala-i-Martin, 1997). Specifically, we model economic growth with the following variables, all lagged by one year unless otherwise specified:

- *GDP per capita.* We use a natural log transformation of GDP per capita, derived from the World Bank's World Development Indicators (WDI). We take the 'initial' value of GDP per capita, that is the value in the first year of the pre-treatment period;

- *Human capital.* We proxy human capital with the average years of education in the population aged 15 and above, using the estimates from the Barro-Lee dataset (Barro and Lee, 2013). The Barro-Lee dataset provide five-yearly estimates; we interpolate intervening years. As with GDP per capita, we take the 'initial' value at the start of the pre-treatment period. For some countries this data is missing, and in these cases a reduced model is used with human capital;

- *Gross capital formation.* Derived from the WDI database. In some cases, this data is missing for some or all of the relevant time periods, and in these cases a reduced model is used without GCF;

- *Trade.* The sum of exports and imports, reported as a percentage of GDP. Derived from the WDI database;

- *Coast.* The average distance to any coastline within a country, derived from the Harvard University Center for International Development geography datasets (Gallup et al., 2001). In two cases—Fiji and Grenada—this data was missing from the CID data. Both small island nations, they were manually set near the first percentile (10km);

- *Tropics.* The geographical proportion of the country classified as a tropical climated, derived from the CID dataset. Again, Fiji and Grenada were missing. Both are classified as tropical (Fiji Marine Tropical) and were manually set at the maximum value 1.0; and,

- *Ethnic fractionalisation.* Derived from the Alesina et al. (2003) dataset.

The dependent variable for this stage of the analysis is GDP per capita growth, derived again from the WDI dataset. For the inequality model, we base the predictors on the Kuznet's hypothesis of an inverse-U relationship between economic development and inequality (see Barro, 2000, for a recent evaluation of Kuznets' hypothesis). The predictor variables, derived from the same datasets as above, are:

- *GDP per capita.* We use the same log transformation as above combined with a squared version to capture the non-linearity in the Kuznets hypothesis;

- *Coast*;
- *Tropics*; and,
- *Fractionalisation.*

The inequality data is derived from the Standardized World Income Inequality Dataset (SWIID) (Solt, 2009, 2014). The SWIID is a standardisation of the UN-WIDER World Income Inequality Database that uses a missing-data algorithm to impute data from other sources. The SWIID provides multiple imputation points (100) for each observation; for our purposes, we collapse this into a single estimate by averaging across the imputations. Table 2 provides the summary statistics on this dataset; for reference, Tables 3 and 4 provide the results of this dataset run as a simple panel regression across the entire dataset.

Table 2: Descriptive Statistics

Variable	Obs	Mean	Std. Dev.
GDP per capita (USD, log)	8,755	7.50	1.71
Human capital (y.o.e.)	8,116	5.96	3.13
Gini coefficient	4,499	37.07	9.77
Gross capital formation	7,587	22.99	10.54
Trade (% of GDP)	7,991	78.14	51.37
Mean distance to coast (km)	8,935	404.03	467.61
Tropics (% of area)	8,935	0.48	0.48
Ethnic fractionalization	10,885	0.43	0.26

Table 3: Reference Model: GDP per capita growth

	Coef.	Std. Err.	z-stat
GDP per capita (log)	-1.158	0.104	-11.10
Human capital (y.o.e.)	0.368	0.057	6.43
Gross capital formation	0.036	0.010	3.62
Trade (Percent of GDP)	0.011	0.002	4.15
Mean distance to coast	-0.001	0.0003	-3.02
Tropics (Percent of area)	-1.186 0.332	-3.58	
Ethnic fractionalization	-1.301	0.600	-2.17
Constant	8.158	0.680	11.99

4 Results and discussion

5 provides the country-by-country estimates of the impacts of affirmative action policies on growth and inequality after five year and after ten years. Figures 1

Table 4: Reference Model: Inequality

	Coef.	Std. Err.	z-stat
GDP per capita (log)	-1.939	0.438	-4.42
GDP per capita (log, squared)	0.136	0.027	5.11
Mean distance to coast	0.001	0.001	0.42
Tropics (Percent of area)	9.982	1.334	7.48
Ethnic fractionalization	4.434	2.542	1.74
Constant	38.195	2.135	17.89

and 2 show the distribution of the five year estimates for growth and inequality respectively. The Average Treatment Effect of affirmative action policies is estimated as a cumulative positive impact on economic growth of 0.02 percent after five years and 0.04 percent after 10 years, but accompanied by a small deterioration in the Gini coefficient by 0.42 percentage points after five years and 0.77 percentage points after ten years.
These raw results broadly confirm the resource mobilisation hypothesis that affirmative action has long-term beneficial consequences for economic growth by enabling marginalised groups to realise their potential and, in doing so, contribute towards overall economic growth. While the average impact on overall inequality is detrimental, the impact is not substantial, suggesting that this may be a tolerable consequence given the positive impact on economic growth which, following the Kuznets hypothesis, should lead to longer-term reductions in inequality in any case.

These raw results, however, are not robust to comparison with the random synthetic placebo group. Table 6 reports the results of a series of t-tests on the treatment outcomes against the randomly generated set of placebo treatments. In all tests, the average effect in the placebo group is—as we would expect for a placebo—closer to zero than in the treatment group, but the differences between the effects on the treatments and on the placebos are not statistically significant at any reasonable level. Without disaggregating policy type, then, the best interpretations of our findings are a null result–affirmative action policies do not systematically affect growth or inequality.

As noted above, however, the policies we have coded as 'affirmative action' cover a wide range of target groups and levels of implementation. In order to identify whether particular types of policy have more systematic impacts, we perform a further series of t-tests against the placebo group with the treatment group restricted to cases with different types of policy. These reduced-sample t-tests produce no significant results across the different 'target' population groups (i.e. whether 'ethnic', 'class', or 'gender'), or for locally-implemented policies [3], but significant results do emerge for nationally- or federally-implemented poli-

[3] Test results omitted for space consideration; available from authors on request

Table 5: Country results

Country	GDP - 5 Years	GDP - 10 Years	Gini - 5 Years	Gini - 10 Years
AUS	-0.008	0.033	-0.924	1.333
AUT	0.003	0.027	0.140	0.730
BEL	–0.029	-0.044	-0.735	-3.1223
BRA	-0.097	-0.005	0.399	-0.252
CAN	0.046	-0.031	2.418	1.790
CHN	0.330	1.078	8.249	15.202
COL	-0.076	-0.059	-3.138	-2.341
ESP	-0.050		2.115	
ETH	0.062	-0.052	-1.624	0.856
FIN	-0.043		0.047	
FJI	-0.027	-0.436	-2.631	-2.717
FRA	-0.012	0.053	-0.240	2.071
GBR	-0.088		2.030	
GHA	-0.051	0.222	4.392	
GRC	0.054	0.134	-1.838	-4.358
GRD	0.075	-0.005		
HUN	-0.119		-1.894	
IRL	0.083	-0.058	-1.272	-3.267
ISR	0.088	0.146	3.112	
JAM	-0.032	-0.199		
JPN	-0.090	-0.085	-3.078	-3.109
KEN	0.009	-0.012	-2.171	
LKA	-0.072	-0.021	3.410	9.677
MEX	-0.035	-0.041	0.686	-0.446
MKD	-0.098	-0.065	7.587	
MYS	0.224	0.425	-2.624	0.588
NGA	-0.219	-0.268		
NOR	0.081	0.056	1.018	-1.455
NZL	-0.048	-0.141	-0.183	-1.487
ROM	-0.051		-2.815	
SRB	0.173	0.149	-0.815	
TJK	0.075	0.263	2.167	
TTO	0.121			
UGA	0.299	0.380	-6.180	1.0131
UKR	0.227	0.120	-2.210	
ZAF	0.011	0.047	4.509	3.661
ZMB	0.103	0.276	7.277	
ZWE	-0.110	-0.397	-0.756	1.830
ATE	0.018	0.036	0.424	0.771
St.Dev.	0.119	0.294	3.289	4.514

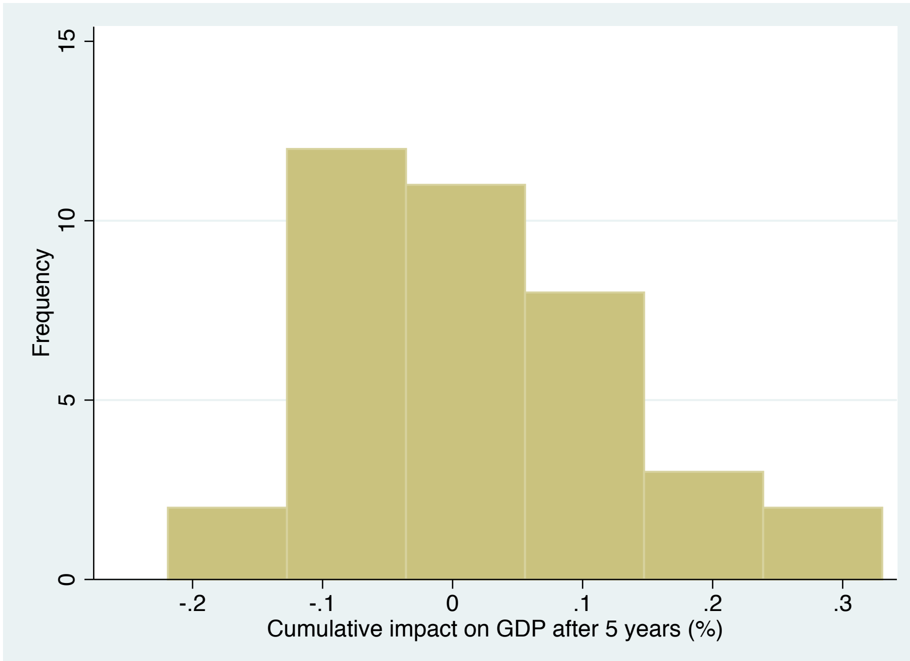

Figure 1: Distribution of estimated cumulative impact on economic growth after 5 years

cies. This is intuitively unsurprising, as local policies—such as the affirmative action measures for Catholics in Northern Ireland—are less likely to have an effect on aggregate economic performance at the national level. Table 7 provides the t-test results comparing national-level policies only to the placebo group. The size of the estimated treatment effect among this group of countries is considerably larger across both outcome variables at both time points than among the entire pool of exposed countries. On the economic growth variables, these estimates are now significantly different from those in the placebo group; the inequality estimates are at best marginally significant after ten years compared with the placebo group.

Table 6: T-tests: Treatment vs placebo, all cases

	Placebo	Treatments	Diff.	t-stat	P
Economic Growth - Five Years	0.002	0.018	-0.016	-0.518	0.303
Economic Growth - Ten Years	-0.008	0.036	-0.045	-0.769	0.222
Gini coefficient - Five Years	-0.149	0.424	-0.572	-0.823	0.206
Gini Coefficient - Ten Years	-0.502	0.771	-1.274	-0.991	0.162

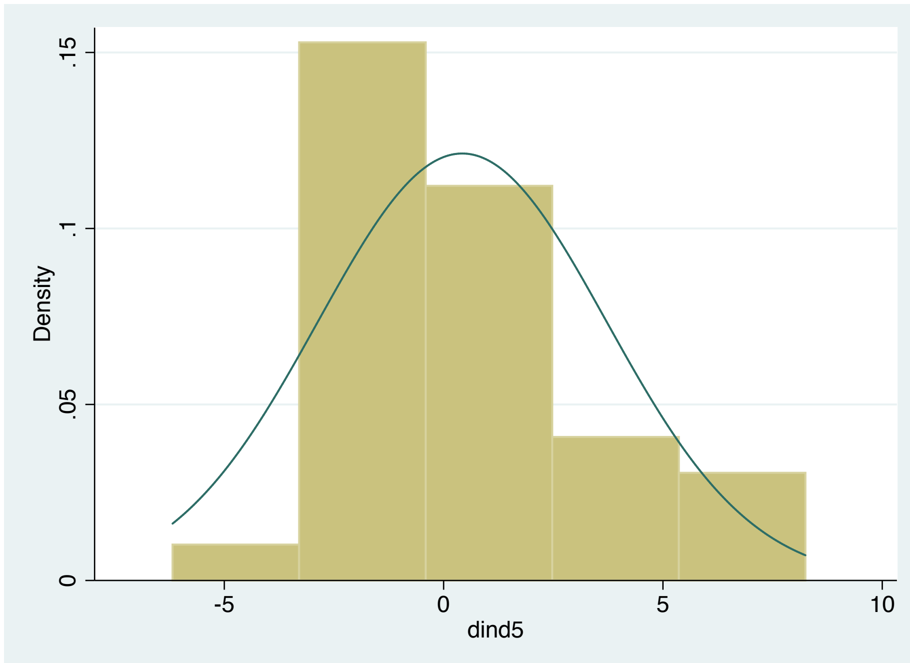

Figure 2: Distribution of estimated impact on Gini coefficient after 5 years

5 Conclusions

Synthetic control estimates of the impact of affirmative action policies on economic development in 38 countries suggests that national-level policies have a small but significant positive impact on economic growth after ten years, although this may be at the price of a small increase in overall inequality. The effect on growth is very small although significantly different from a 'placebo' control group. While we have provided some systematic evidence that nationally-implemented affirmative action policies *may* promote economic growth, we can certainly more robustly claim that there is no evidence that they systematically retard economic growth or increase overall inequality.
Relating this back to our broad stylised hypotheses about the economic impacts

Table 7: T-tests: Treatment vs placebo, all cases

	Placebo	Treatments	Diff.	t-stat	P
Economic Growth - Five Years	-0.008	0.095	–0.103	-1.610	0.055
Economic Growth - Ten Years	-0.007	0.220	-0.228	-1.709	0.045
Gini coefficient - Five Years	-0.149	1.043	-1.192	-1.209	0.115
Gini Coefficient - Ten Years	-0.502	2.050	-2.553	-1.454	0.074

of affirmative action, our analysis suggests we can safely reject the 'Misallocation hypothesis'–there is no evidence that affirmative action policies, whether national or local, systematically retard economic growth. Our evidence can, however, be interpreted as providing weak support for the 'Resource Mobilisation' hypothesis that affirmative action policies reduce the unrealised potential for marginalised groups to contribute economically. Indeed, insofar as the Resource Mobilisation is a more 'slow-moving' hypothesis that the Misallocation hypothesis, the fact the significant positive effects we do find only emerge after ten years not five is intuitively in-line with this hypothesis. The safest interpretation of the analysis, however, may be a null finding–there is no clear evidence that affirmative action policies have any systematic effect on economic growth on inequality. Of course, this does not imply that such policies are useless: we have not been able to examine the impact on the 'horizontal' inequalities that are the main target of such policies. Rather, this evidence suggests that two common assumptions about possible negative externalities of affirmative action policies are not backed up by empirical analysis. Judged on their own terms, affirmative action policies may be more or less successful in achieving their goals, but argument that they have negative macroeconomic consequences that outweigh any potential benefits are firmly rejected by our analysis.

References

A. Abadie, A. Diamond, and J. Heinmuller. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association*, 105, 2010.

A. Abadie, A. Diamond, and J. Hainmueller. Comparative politics and the synthetic control method. *American Political Science Review*, 59(2):495–510, 2015.

A. Alesina, A. Devleeschauwer, W. Easterly, and S. Kurlat. Fractionalization. *Journal of Economic Growth*, 8(2):155–194, 2003.

S. Alon and M. Tienda. Assessing the 'mismatch' hypothesis: Differences in college graduation rates by institutional selectivity. *Sociology of Education*, 78:294–315, 2005.

R. J. Barro. Inequality and growth in a panel of countries. *Journal of Economic Growth*, 5(1):5–32, 2000.

R. J. Barro. Determinants of economic growth in a panel of countries. *Annals of Economics and Finance*, 4:231–274, 2003.

R. J. Barro and J.-W. Lee. A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics*, 104:184–198, 2013.

R. Breen and J. H. Goldthorpe. Explaining educaitonal differentials: Towards a formal rational action theory. *Rationality and Society*, 9(3):275–305, 1997.

G. K. Brown and A. Langer. Does affirmative action work? lessons from around the world. *Foreign Affairs*, 94:49–56, 2015.

A. Deshpande and T. E. Weisskopf. Does affirmative action reduce productivity? a case study of the indian railways. *World Development*, 64:169–180, 2014.

W. E. Feinberg. Are affirmative action and economic growth alternative paths to racial equality? *American Sociological Review*, 50:561–571, 1985.

J. L. Gallup, A. D. Mellinger, and J. D. Sachs. *General Measures of Geography Dataset*. Harvard University Center for International Development, http://www.cid.harvard.edu/ciddata/geographydata.htm, 2001.

P. Hinrichs. Affirmative action bans and college graduation rates. *Economics of Education Review*, 42:43–52, 2014.

J. S. Leonard. The impact of affirmative action on employment. *Journal of Labor Economics*, 2:439–463, 1984.

J. S. Leonard. The impact of affirmative action regulation and equal employment law on black employment. *Journal of Economic Perspectives*, 4:47–63, 1990.

A. Moro and P. Norman. Affirmative action in a competitive economy. *Journal of Public Economics*, 87:567–594, 2003.

M. S. Moses, L. D. Jenkins, C. Paguyo, and L. Wei. Assessing affirmative action programs on six continents. In L. D. Jenkins and M. S. Moses, editors, *Affirmative Action Matters: Creating Opportunities for Students around the World*. Routledge, London, 2014.

J. Rawls. *A Theory of Justice*. Oxford University Press, Oxford, 1999. Revised edition.

X. X. Sala-i-Martin. I just ran two million regressions. *The American Economic Review*, 87(2):178–183, 1997.

S. Singh and R. Nilakantan. Naxalite insurgency and the economic benefits of a unique robust security response. *HiCN Working Paper*, 127, 2012.

F. Solt. Standardizing the world income inequality database. *Social Science Quarterly*, 90(2):231–242, 2009.

F. Solt. *The Standardized World Income Inequality Database v. 5.0.* Harvard Dataverse, http://hdl.handle.net/1902.1/11992, 2014.

F. Stewart and A. Langer. Horizontal inequalities: Explaining persistence and change. In F. Stewart, editor, *Horizontal Inequalities and Conflict: Understanding Group Violence in Multiethnic Societies*. Palgrave, London, 2008.