

EDUCATION AND SOCIAL CAPITAL

John F. Helliwell

University of British Columbia

and

Canadian Institute for Advanced Research

and

Robert D. Putnam

Harvard University

and

The University of Manchester

Education is one of the most important predictors—usually, in fact, the most important predictor—of many forms of political and social engagement—from voting to chairing a local committee to hosting a dinner party to trusting others. Over the last half century (and more) educational levels in the United States have risen sharply. In 1960 only 41 percent of American adults had graduated from high school; in 1998 82 percent had. In 1960 only 8 percent of American adults had a college degree; in 1998 24 percent had. Yet levels of political and social participation have not risen *pari passu* with this dramatic increase in education, and by some accounts [Putnam, 1995a; 1995b; 2000] have even fallen. For at least two decades, political scientists have mused about this paradoxical “puzzle” [Brody, 1978].

Recently, however, Norman Nie, Jane Junn, and Kenneth Stehlik-Barry [1996, hereafter NJS-B] have offered an elegant and potentially powerful resolution to this paradox, beginning with a crucial distinction between the “relative” and “absolute” effects of education. If more people now have a college degree, they argue, perhaps the sociological significance of the credential has been devalued. Social status is, for example, associated with education, but we would not assume that just because more Americans are educated than ever before, America has a greater volume of social status than ever before. To the extent that education is merely about sorting people, not about adding to their skills and knowledge and civic values, then no puzzle remains to be explained. In fact, NJS-B conclude, participation is affected primarily by relative educational levels, and thus has not been (and should not have been expected to be) rising with aggregate educational levels.

The distinction that NJS-B have introduced is important. Education has external effects, as well as internal ones. In principle, my behavior can be affected not only by my education, but also by that of others around me. The core issue is whether (holding constant my own education), I am more likely or less likely to participate politically and socially if those around me become more educated. Besides its academic inter-

John F. Helliwell: Department of Economics, University of British Columbia, 997 - 1873 East Mall, Vancouver, BC, V6T 1Z1 . E-mail: john.helliwel@ubc.ca.

est, the NJS-B conclusion has practical significance. If the negative effects of average education match or exceed the positive effects of absolute education, then raising educational levels is a pointless or even counter-productive way to increase civic engagement. Another way of putting this point, more familiar to economists, would be that the coefficient on average education provides an estimate of the external effects of increases in education. If there are no external effects, then the average level of education would attract a zero coefficient, and all of the consequences of education would be revealed by the effect flowing through one's own education. If average education takes a negative coefficient, then the external effects of education are negative. If the negative effects on average education are as large as the positive effects on the own-education variable, as found by NJS-B for participation, then the negative externalities are so large as to fully offset the own-education effects, so that economy-wide increases in education would have no effect on participation. By contrast, if the coefficient of average education is positive, as NJS-B and we find for trust, and we find for most forms of participation, then education has positive externalities.

The possible existence of positive or negative externalities from education, with respect to the accumulation of social capital is parallel to the long-studied issue of the effects of education on the accumulation of human capital. The basic human capital model [Becker, 1964] argues that education is of value to individuals because of its impact on their knowledge and skills. By contrast, the signalling model [Spence, 1974] suggests that education is instead valuable to employers and employees as a sorting device, just as NJS-B argue is the case for education and social engagement. Recent progress in this debate has relied on the use of natural experiments and instrumental variables methods [Card, 2001] to assess the income effects of education, and to search for external effects. The results of this research tend to favour the human capital model over the sorting model, and to support the possibility of positive economic externalities from education [Moretti, 2004].

Since the working paper version of this paper was written [Helliwell and Putnam, 2001], some research has used instrumental variables to estimate the effects of own-education on social engagement (see, for example, Milligan, Moretti and Oreopoulos [2004] and Dee [2004]), and the results have shown, consistent with the results in NJS-B and our own work, that individuals with more education tend to be more engaged citizens. But to test the NJS-B contention that these individual-level effects are reversed by negative contextual-level effects requires direct estimation of a relative education model, as done by NSJ-B and in this paper¹.

It is important to search for evidence of positive or negative externalities from education, since the presumption of positive externalities provides important theoretical underpinning for public support of higher education. The NSJ-B analysis finds positive externalities for social trust, and negative ones for various types of social engagement, a mixed message. The latter finding drew our attention because it seemed to contradict what we and others had been finding for the individual and community-level effects of education on the accumulation of social capital. We decided to investigate the sources of these differing results, to attempt a reconciliation, and, most importantly, to try to settle the question of whether education has positive or negative externalities for the accumulation of social capital. This paper presents the

results of that reconciliation. When appropriate definitions for relative education are used, we find that the contextual effects of education on social participation are generally positive, and never significantly negative, even in using the same data and basic equations used by NSJ-B. We have also been able to confirm our findings by replication in the context of several large samples of survey data not available when the first version of this paper was written.

When we began our reconciliation, we first discovered that the NSJ-B finding of negative externalities from education, with respect to social participation, flowed entirely from their rather special, and we think theoretically inappropriate, definition of “relative education”. Both spatially and temporally, the operational standard adopted by NJS-B [1996, 119 and 227-233] is puzzling. Spatially, by using *national* standards, NJS-B in effect assume that my civic behavior is affected by educational levels in communities on the other side of the continent. Logically, this operationalization means that civic participation in Seattle—voting, group membership, and so on—should tend to fall if educational levels in rural North Carolina rise. Indeed, the operationalization adopted by NJS-B assumes that the effect of education in rural North Carolina on Seattle participation rates is fully as great as the effect of education levels in Seattle itself. In some domains—the job market for astronauts [NJS-B, 174], for example—educational externalities may be undiluted by distance, but whether participation in community affairs is like that is, we believe, worth exploring. So we propose to measure relative education relative to the respondent’s census region.²

Temporally, by comparing each respondent’s education to the level of education of all Americans who were between 25 and 50 years of age when the respondent reached the age of 25, NSJ-B assume a static, backward-looking measure of educational externalities. In some job markets, this may perhaps be a reasonable assumption, but in civic participation it seems implausible. For example, this operational measure of relative education means that the participation rate of a 55-year-old is influenced not at all by the educational credentials of her 54-year-old neighbors, but is influenced instead by the educational credentials of people long dead. In other words, in NSJ-B’s oddly asymmetric world of civic competition, no one ever competes against anyone younger, but everyone always competes against everyone older (including the dead). Here, instead, we propose to compare each respondent’s education to all other living adults, both older and younger.³

In short, while it seems to us well worth investigating whether the civic participation of a Seattle high school drop-out is influenced (positively or negatively) by the educational levels of his neighbors, it seems to us implausible to assume that his participation rate is equally or more influenced by the educational level of dead North Carolinians.

One special reason for caution regarding the NSJ-B implementation of their important theoretical insight is (as they fully recognize at [NSJ-B, 134]) that since national educational levels have risen monotonically throughout this century, operationally NSJ-B’s measure of relative education (defined as it is in national and static terms) is virtually a perfect linear transform of the respondent’s year of birth.⁴ Thus, there is the risk that this operational definition of relative education might falsely take credit for many other factors that have also been changing nationwide and affecting generations differently.

In this paper we confirm the NJS-B's results showing that education has positive externalities with respect to trust. However, we find that the NSJ-B's results showing average education to have negative effects on participation disappear if a more theoretically appropriate definition of the educational environment is employed.

ARE THE EFFECTS OF EDUCATION RELATIVE OR ABSOLUTE?

The basic theory underlying the NJS-B hypothesis is that education is for many activities a sorting mechanism used to distribute a fixed number of opportunities for gainful participation. Thus more education for an individual makes him or her better able and more likely to compete for a place of influence and activity, while increases in general education have no such effect. They argue that this relative education hypothesis has no claim to universal applicability, and recognize that, at least for some activities and attitudes, increases in average education levels may well have positive effects. They distinguish three main alternatives: a purely *relative* effect, as described above; an *absolute* effect, whereby the education effect depends only on an individual's own education and not on the education levels of others; and a *cumulative* effect, whereby there are positive feedbacks from general education levels to the individual's own actions or attitudes.

They test among these possibilities by defining separate variables for an individual's own education, and for the average level of education in the individual's "educational environment," as defined above. If the positive effects of one's own education are offset by equal and opposite effects from the average education level, then the relative education model dominates. If one's own education has a positive effect, with no significant effect from the average level of education, then the absolute or additive model dominates. Positive effects from both own education and average education would provide evidence favouring the cumulative model, where the effects of education are super-additive, since education has both direct and indirect positive effects. Mixed cases are of course possible.

Their reasoning and empirical results provide them with examples of each type of effect in operation. For political and social engagement they argue that the relative education model should well dominate, because of a relatively fixed amount of total benefits from participation. For cognitive abilities, they find a large positive effect of own education, combined with a relatively small negative feedback from average education levels.⁵ They see the predominance of the own-education effect as reasonable, because education can increase one person's cognitive skills without lowering those of anyone else. This, they argue, is in contrast with political and social participation, where increased participation by one is expected to reduce the gains available to others, because of competition for a fixed pool of benefits. They find, for organizational memberships, that each year of own-education has a substantial positive effect, but that this is offset by an even larger negative effect from each year of increase in average education levels [NSJ-B, 1996, 163]. For tolerance, they find that one's tolerance is increased not only by one's own education but also by the average education level of those in the surrounding community. They find that both own-education and average education have significant positive effects, with the effects of average education being even larger than those for own education [NSJ-B, 1996, 148].

The theoretical argument supporting the possible finding of cumulative or super-additive effects of education can also be made when social trust is the dependent variable. Higher average education levels may help to create a climate of trust that is self-reinforcing. If individuals know that higher education levels make others more likely to be trusting (and perhaps also more trustworthy), then they are in turn more likely to trust others. Hence the returns to trusting behaviour are increased where there are increases in average levels of education, so that it should be expected that people of any level of education are in fact more trusting of others in an environment marked by higher average education levels.

Going further, we argue that the same theoretical reasoning can be used equally well to support a finding of cumulative or super-additive effects of education for many types of social interaction, including political and community engagement. When deciding whether we want to participate in clubs or community life, are we not more likely to find such activities personally and socially rewarding if there is a climate of trust and tolerance among those with whom we are working? It is fairly well established that high levels of trust reduce the costs of getting things done, since in the absence of trust it is necessary to have rules and enforcement that provide expensive and cumbersome insurance that the agreed purposes of the organization will be pursued as hoped and expected. This will be true even if the number of leadership positions in such organizations does not increase proportionately with membership, and even if some organizations may have to compete harder for opportunities to make their way if the field becomes more crowded.

Moreover, it is likely that for some types of organization the opportunities for beneficial involvement by individuals may actually increase with the extent of involvement by others. For instance, a reading group is more likely to be of interest if there are other members with educated interests in the same issues. Likewise, participation in a community sports team will be more attractive if there are enough clubs and teams to make up a good schedule of games with comparable teams. Even the leadership point may cut both ways. Some organizations have jobs to do that are widely regarded as important, while individuals would in general prefer that someone else does the work. One example may be provided by home and school organizations. These may not even exist without a sufficiently large and committed group of interested parents, so in this case we might expect that I am actually more likely to participate in such organizations if educational levels around me increase.

But beyond some point, especially where there is a given task to be done, one might expect the relative education model to come into play. After some point, any increase in the number of interested and able parents may diminish the need for any particular parent to be involved without thereby threatening the ability of the organization to do its job to general satisfaction. This might lead to a negative effect, after some point, in the effects of average education on participation. This need not reflect increased competition for the number of available positions, but could instead signal the favourable effects of a larger pool of available volunteers.

Leadership roles in community organizations may not be prizes for which there is competition, but jobs that are taken by those willing to take their turn or do their share to keep valuable activities moving along. Thus one may be grateful that someone

else is able and public-spirited enough to do the organizing of the local sports club, and be more inclined to join if the club has a larger group of potential helpers with whom to share the job of arranging the schedules and teams. This would provide another reason for anticipating that average education levels might have positive rather than negative effects for at least some types of group memberships.

For us, it is thus an open question, in terms of theory, whether education effects should be relative, additive, or super-additive, for different types of trust and participation. We agree with NJS-B that the case for expecting cumulative or super-additive effects may be greater for trust and tolerance than for some types of participation. Thus we think it is important to consider a number of different types of participation, as well as to use multiple sources of primary data. Our results, presented in the next section, are much more optimistic than those found by NJS-B, because we find much less evidence of relative education effects. We find large and pervasive positive effects of general education increases on levels of trust and participation. While NJS-B find that increases in average education have no net effect on social engagement (i.e., the positive own-effects are offset by negative effects from the rising educational environment), we find that social engagement increases with average levels of education. We will therefore be able to make the more optimistic conclusion that education can be seen as increasing rather than merely redistributing social capital.

EMPIRICAL RESULTS

In this section we present results on the effects of own and average education levels on measures of both trust and participation, using pooled time-series and cross-sectional data from the US General Social Survey (GSS) from 1972 through 1996⁶, and from the DDB-Needham Life Style survey data from 1975 through 1997. The advantage of both data sources is that they cover a large enough span of years, with enough individual observations in each year that there is some possibility of disentangling the complex interplay of generational, cohort and individual factors affecting trends in trust and participation. The GSS employs interviews with randomly selected samples, while DDB uses annual written mail questionnaires of a recruited panel of participants. Although the demographic makeup, questions, and survey procedures differ between the two surveys, the trends evident in the two bodies of data are strikingly similar, thus increasing the assurance that our conclusions are not a result of specific features of the methods and questions used in a particular survey. The GSS data have also been used by NJS-B and many other researchers, while the DDB data provide a fresh set of observations.

We look first at the effects of education on measures of social trust, and then deal with measures of social engagement. The social trust findings are in Table 1, which shows the results from equations with fully specified controls as well as from stripped-down equations based only on individual and average levels of education, controlling only for a possible time trend. As described above, the measure used for average education differs in two important ways from that used by NJS-B, first by restricting the comparison group to those in the census region in which the respondent lives, and second by permitting the comparison group to change as time progresses. We do this by defining the educational environment as the average education levels of all adult respondents to the same survey in the respondent's census region. As we

have noted, an important statistical benefit of making the comparison specific to the census region is that the variable then has more variation among respondents in the same time period, and is therefore likely to be free of contamination from correlation with other national time trends that may influence trust and participation.⁷

TABLE 1
Comparing GSS and DDB Evidence on Social Trust

Equation	(i)	(ii)	(iii)	(iv)	(v)	(vi)
	GSS	DDB	DDB	GSS	DDB	DDB
	1972-96	1975-97	1975-97	1972-96	1975-97	1975-97
Sample	22445	76156	76156	22445	76156	76156
Dependent variable: Trust		Honest	Honest	Trust	Honest	Honest
			(binary)			(binary)
	Equations with full control variables			Simple equations with year effects		
Education in years	.0439	.0093	.0175	.0391	.0062	.0133
	(40.3)	(24.4)	(24.8)	(37.0)	(16.0)	(18.7)
Average education in region (yrs)	.0244	.0057	.0096	.0602	.0130	.0206
	(3.2)	(1.6)	(1.4)	(11.0)	(3.6)	(3.1)

The GSS and DDB ask slightly different questions about social trust, and scale their results differently, raising difficulties for exact comparisons. The GSS asks a standard question used in many social science surveys: “Generally speaking, would you say that most people can be trusted, or that you can’t be too careful in dealing with people?” The DDB’s simpler question asks respondents to record their agreement or disagreement, on a six-point scale, to the statement “Most people are honest”. We present two sets of results using the DDB data, one based on the six point scale, converted so that 1.0 represents full agreement and 0 complete disagreement, and the second converted to a binary equivalent, with any form of agreement being given the value 1.0, and any form of disagreement the value of zero. The second form is more directly comparable with the GSS binary coding, which uses the value 1.0 for agreement that most people can be trusted and 0 for agreement with the alternative that you can’t be too careful. As might be expected, there are more affirmative answers to the DDB question than to the GSS, since the GSS offers an alternative that many may agree with more strongly than they accept general trustworthiness.⁸

The equations in Table 1 show separately the effects of own-education and average education, both measured in years, on the GSS and two DDB measures of social trust. Equations (i) through (iii) show the education effects based on a fully specified model attempting to account for a number of other factors that have been found to influence social trust, while equations (iv) through (vi) present results from a much cruder model including only education, average education, and a linear time trend.⁹ Both models show positive effects from both own and average education levels for all three measures of social trust. These positive effects are very large and highly significant in the simple model, and smaller and weaker in the fully specified models, presumably because the larger models contain a number of other variables that share some of the same regional and time series variation as average education.¹⁰ In the simple models, the effects of average education on trust are about twice as large as those from own-education, echoing the earlier results of NJS-B for tolerance. In the fuller models they are about half as large. Using the GSS results shown in the first

equation, each additional year of education increases one's likelihood of being trusting by .044, or just over 10 percent of the average likelihood of .38. A one-year increase in the average level of education in one's region increases one's trust by .024. Thus increases in average education levels, acting both through own-effects and education levels in the surrounding community, have had strong positive effects on social trust over the past twenty years. With average education increasing by almost 1.5 years from the mid-1970s to the late 1990s, the implied impact on trust is about .1, an increase equal to almost one-quarter of the average level of trust on the mid-1970s, and almost as large as the entire drop since that time. Thus the effects of changes over time in average education levels double the size of the decline in trust to be explained by other factors.

We turn now to consider the results for participation. First we shall compare the GSS and DDB results for some key forms of participation, and then present the GSS results for sixteen different types of organization. Table 2 shows the education results from both fully specified and simple models for total memberships from the GSS, plus three key types of social engagement as measured by the DDB surveys. These DDB participation measures represent the number of times that the respondent reports participating in each of these activities in the previous year: (1) attend a club meeting, (2) work on a community project, (3) give or attend a dinner party. For three out of four measures of social engagement, own-education has a strongly positive effect, and average education levels, where significant, are positive. The only estimated negative effect of average education levels is for total memberships from the GSS survey. Even there the estimated negative effect of average education levels is small and insignificant. This is in sharp contrast to the NJS-B [1996, 247] results for GSS total memberships. They report an own-education effect of +.20, fairly close to our estimate of +.23. However, for average education, they find a significant negative effect of -.24, while our negative effect is only one-tenth as large, and is insignificant. Since the explanation of this difference is a primary aim of this paper, it will be treated separately, in the next section.

The DDB equations provide further and independent confirmation of the general lack of relative education effects, since all of the four types of activity show positive effects from average education levels, two of them with statistical significance. The DDB participation measures are in some ways preferable to the GSS data, since the latter show only the number of types of organization to which the individual belongs, and not the extent of involvement. The DDB data, by contrast, ask for the frequency of involvement in each of the types of activity. The sample size is also much larger for DDB, although in both surveys the range of years covered, and the number of respondents in each year is large enough, that small sample size is not likely to be a problem.

Comparing the results for the GSS and DDB equations is made more complicated by the different form of the questions and coding. The coefficients are generally higher for the DDB responses, but then so are the average values for the participation measures. To get a rough measure of the extent to which the proportionate responses are similar, the coefficients shown in the table should be divided by the sample averages for the dependent variables, as shown at the bottom of Table 2. With these adjustments made, the effects of own education are quite similar across the types

of involvement. A one-year increase in education leads to an estimated increase in participation ranging from 9 percent for number of dinner parties to 14 percent for number of involvements in community projects. The effects of average education are more varied, and much less precisely estimated. An increase of one year in the average level of education in one's region increases the number of dinner parties given or attended by 20 percent, increases the involvement in community projects by 5 percent, has no effect on club meetings attended, and reduces the number of memberships (the GSS variable) by about 1 percent.

TABLE 2
GSS and DDB Evidence on Social Engagement

Equation	(i) GSS 1974-94	(ii) DDB 1975-97	(iii) DDB 1975-97	(iv) DDB 1975-97
Sample	19214	71246	71246	71246
Dependent variable:	Number of memberships	Club meetings (times/year)	Community projects	Dinner parties
A. Results from Full Equations with control variables.				
Education in years	.2283 (51.7)	.7141 (36.1)	.3504 (33.5)	.4808 (41.5)
Average education in region (years)	-.0201 (0.7)	.0057 (0.1)	.1710 (1.7)	.9600 (8.7)
B. Results from Simple Equations including only education variables and a time trend.				
Education in years	.2286 (51.8)	.6167 (31.4)	.3273 (31.6)	.4313 (37.4)
Average education in region (years)	-.0362 (1.1)	.0979 (0.5)	.1344 (1.4)	1.227 (11.4)
Mean of dep variable	1.78	7.46	2.25	5.32

We turn to Table 3 to consider the effects of education on memberships in the sixteen different types of organization covered by the GSS memberships question. The dependent variable for each of the sixteen equations is equal to the fraction of the respondents saying they were a member of an organization of the type mentioned. The figure for total memberships is the sum of the answers for the sixteen types. Given the linear regressions used, it follows that the coefficients for total memberships are simply the sums of the coefficients on the same variable in each of the sixteen categories. Thus the equations for the different types of membership show where the aggregate effects are coming from.

Inspection of the results for specific groups shows how hard it is to generalize across groups with very different purposes, recruiting methods, advantages to members, and degree of commitment required of members. Own-education effects are positive for all forms of membership except for unions. The proportionate own-education effects are largest, unsurprisingly, for professional organizations, since such organizations generally have advanced education as prerequisites for membership. The negative own-education effect for union membership follows from similar reasoning, since unionization rates are generally higher on the shop floor than in management positions, while educational qualifications follow a reverse pattern.

Average education effects are generally small and insignificant. Significant positive effects for average education are estimated for literary groups, a finding foreshadowed by our earlier discussion. A book group is more likely to be formed, and to offer a links to other readers with similar interests, where average education levels are higher, and there seems to be no reason to think that the social structure sets any upper limit to the number of book groups. Average education effects are also positive for sports groups, although significantly so only in the simple form of the equation.

TABLE 3
Effects of Absolute and Relative Education on Types of Membership

Column Membership type	Full Equation			Simple Equation		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
	Educ	Aved	P-value	Educ	Aved	P-value
Fraternal	.0155 (21.6)	.0003 (0.7)	.0002	.0110 (15.7)	.0017 (0.5)	.0003
Service	.0187 (25.5)	-.0031 (0.6)	.0026	.0170 (24.3)	-.0066 (1.9)	.0026
Veterans	.0034 (5.3)	-.0054 (1.2)	.6407	.0004 (0.6)	-.0057 (1.8)	.0799
Political	.0092 (18.9)	.0006 (0.2)	.0000	.0079 (17.1)	.0028 (1.2)	.0000
Union	-.0054 (6.4)	-.0217 (3.6)	.0000	-.0031 (3.7)	.0337 (8.1)	.0000
Sport	.0171 (17.6)	.0128 (1.8)	.0000	.0230 (24.5)	.0179 (3.8)	.0000
Youth	.0093 (12.3)	.0075 (1.4)	.0017	.0127 (17.5)	-.0009 (0.2)	.0010
School	.0178 (21.1)	-.0021 (0.3)	.0083	.0218 (26.7)	-.0196 (4.8)	.5825
Hobby	.0113 (15.5)	.0014 (0.3)	.0137	.0105 (15.2)	-.0056 (1.6)	.1533
Greek	.0174 (33.3)	-.0052 (1.4)	.0009	.0166 (35.5)	-.0125 (5.0)	.0957
National	.0059 (13.1)	.0044 (1.3)	.0013	.0051 (11.9)	.0047 (2.2)	.0000
Farm	.0018 (3.8)	-.0081 (2.3)	.0719	.0011 (2.3)	-.0097 (4.1)	.0002
Literary	.0217 (31.6)	.0154 (3.2)	.0000	.0195 (29.8)	.0094 (2.8)	.0000
Professional	.0515 (65.3)	.0070 (1.2)	.0000	.0500 (66.7)	-.0078 (2.1)	.0000
Church	.0233 (20.1)	-.0415 (5.0)	.0278	.0155 (13.6)	-.0844 (14.8)	.0000
Other	.0100 (13.5)	.0013 (0.3)	.0300	.0075 (10.7)	.0111 (3.1)	.0000
Total memberships	.2286	-.0362	.0000	.2165	-.0716	.0003
Total	(51.8)	(1.1)		(51.1)	(3.3)	

Notes: The full equations (i) to (iii) include the same independent variables as the equations shown in Table 2, except that the variable for year is excluded, and fixed effects are added for each year. The simple equations (iv) to (vi) include only education and average education, plus fixed effects for each year.

Average education effects are significantly negative only for unions, church groups, and farm organizations in the full equations, and for school groups in the

simple equations. The relative education hypothesis does not apply easily to the union and church group cases. For unions, both the own and the average education effects are negative, presumably for the same reason, with lower average education providing the greater critical mass needed for successful working-class organization. For church groups, the negative effect from average education exceeds the positive effect from own education, suggesting that regions (and time periods) with higher average education levels also have lower average ratios of church membership, though we doubt that this is attributable to the social competition emphasized by NJS-B. For farm and school groups, the negative effect of average education may reflect the sorting effect emphasized by NJS-B, or perhaps some burden sharing of the sort we have described earlier. Some farm groups may represent professional organizations in which membership is simply part of being in the business, with a negative average education effect reflecting in part differences in average education levels in census regions with different concentrations in agriculture.

UNRAVELLING THE PUZZLING CONTRAST

Why do our results for average education differ so much from those of NJS-B? The differences could be due to estimation method, sample size, choice of control variables, or differences in the definition of average education. We shall show below that the differences are entirely due to differences in the definitions of average education. In any equation using our definition of average education, the NJS-B negative effect of average education disappears. This happens whether we use their control variables, our alternative control variables, or simple equations with no control variables. Although we use the same corrected GSS data for total memberships used by NJS-B, our control variables differ somewhat from theirs, and their sample size is about 15 percent smaller. Their smaller sample size (15887 compared to 19214) is due to the unavailability of data for some of their control variables. We have tried to include only control variables that are available for almost the full sample, while testing to see if allowing for other control variables available only for some of the data (e.g. parents' education) makes any difference to the fundamental results. There are some slight differences in the statistical methods used. We include, in the Table 2 equation for GSS memberships, a linear time trend, while NJS-B use a more general method by inserting separate dummy variables for all but one of the individual years. To see if the different method of allowing for year effects has any impact on the education results, we have re-estimated equation (i) of Table 2 using year fixed effects instead of a linear time trend. The resulting coefficients for education and average education are shown in the last row of Table 3. As can be seen, the education coefficients are unaffected. If the results did differ, the NJS-B allowance for year fixed effects would be preferred to what we do in Table 2, since it is a more general statistical method. We include both sets of results to show they are identical, and use the linear time trend elsewhere in Tables 1 and 2, since we are interested in showing the extent to which trust and social engagement are changing through time, with and without controlling for other variables.

To provide a more definitive answer to the puzzle, we have drawn a sample of data as close as possible to that used by NJS-B, and then tested average education effects

in the context of three different equations: first using their control variables, then using ours, and finally using a simple model containing only education, average education, and fixed effects for each survey year.¹¹ We have then estimated these equations using the NJS-B measure of average education, our measure, and two intermediate measures designed to show which of the differences play the most important roles in explaining the puzzle. All these results are shown in Table 4.

TABLE 4
Total GSS Memberships with Differing Measures of Average Education

	(i)	(ii)	(iii)	(iv)
	Static peer group National average (NJS-B measure)	Static peer group Regional average	Dynamic peer group National average	Dynamic peers Regional average (Our measure)
Using the NJS-B control variables				
Education	.2058 (37.6)	.2068 (37.5)	.2007 (36.8)	.2006 (36.5)
Average education	-.2001 (8.6)	-.1188 (6.6)	.0174 (0.2)	.0021 (0.1)
Using the Helliwell and Putnam control variables				
Education	.2333 (45.8)	.2343 (45.7)	.2323 (45.7)	.2335 (45.5)
Average education	-.2648 (2.3)	-.0941 (2.5)	.4699 (2.6)	-.0478 (1.3)
Simple equation with year effects but no other controls				
Education	.2404 (48.0)	.2446 (48.1)	.2216 (45.1)	.2239 (44.8)
Average education	-.2118 (16.4)	-.1911 (16.2)	-.1949 (1.8)	-.0826 (3.2)

As noted above, NJS-B define the educational environment as the national average of education years of those who reached the age of 25 in the same year as the respondent, or in any of the 24 preceding years. In our view, this is too broad in being national rather than regional, and too unresponsive to the changing environment, by excluding all those who are younger than the respondent while continuing to include many who are dead and gone. We define average education as the average number of years of education of currently surveyed adults in the same census regions, arguing that this best represents the social environment which the individual is deciding whether to trust and join. Our measure thus uses a changing regional peer group while NJS-B use a static national reference group. The results for education and average education based on the NJS-B definition are shown in column (i) of Table 4. To help show the influence of the two differences between our definition and that of NJS-B, we have also defined two mid-way variables, one using a static regional peer group and the other a dynamic national peer group. These results are shown in columns (ii) and (iii) of Table 4, while the results from our variable, which employs a dynamic regional peer group, are shown in column (iv).

The three panels of results, moving from top to bottom in Table 4, make use of the NJS-B control variables, our control variables, and no control variables. All equations include year fixed effects. The first important conclusion from the results is that the negative effect of average education falls sharply moving from left to right in the table, whatever control variables are used in the equation. Looking first at the top panel, using the NJS-B control variables, the negative average education effect falls almost in half using a static regional peer group, and disappears entirely with either a dynamic national or a dynamic regional peer group. Using our control variables, as shown in the middle panel, the negative effect of average education falls by more than half with a static regional peer group, becomes positive with a dynamic national peer group and is insignificantly negative with the dynamic regional peer group. In the simple equation, with no controls, the negative effect of average education remains strong with the static regional peer group. It is still large, but statistically weak, with the dynamic national peer group, and drops to one-third the size of the own-education effect when our dynamic regional measure is used. As may be recalled from Table 3, the negative average education effect in the simple equation for total memberships can be traced to church membership. This explains why the negative average education effect disappears in the fully specified NJS-B equation, but not in the simple equation, since the NJS-B control variables include church attendance with a strong positive coefficient.

Thus we can conclude that in either of the fully specified models the negative effects of average education on GSS memberships disappear when the static national peer group is replaced by a dynamic regional one, with the shift from a static to a dynamic reference group being the most important part of the story.

CONCLUSION

We have presented a range of data designed to show the effects of education on trust and social engagement, two key variables often used as measures of social capital. We have paid special attention to the relative education hypothesis of NJS-B, who report evidence suggesting that only relative education is likely to influence political and social engagement. For trust, our results support their findings for tolerance, with increases in own-education and average education both leading to significant increases in social trust.

For several measures of social engagement, our results differ significantly from those of NJS-B, and are much more optimistic about the social benefits of increases in the width and depth of education. We find no systematic evidence that increases in average education have any negative effects on participation, let alone of the size required to offset the large positive effects of own-education. Since we prefer our definition on theoretical grounds in both respects in which it differs from theirs, we are inclined to put more weight on our results. Our tables show confirming results from two large US surveys. Since the preparation of the first version of this paper, we have been involved with several new surveys measuring the nature and consequences of social capital in the United States, Canada, and other countries. We have recently used these new surveys (employing the same data samples used by Helliwell and Putnam

[2005]) to confirm the same pattern of generally positive contextual effects reported in this paper. We also now have even finer contextual measures of education, even to the level of the census tract, and we usually find, as we had previously surmised, that these are even more relevant than the broader contextual measures. We are thus able to be even more confident that rising general levels of education are likely to be accompanied by higher general levels of political and social engagement. This means that the answer to the Brody [1978] puzzle about declining political participation in America must lie elsewhere¹².

APPENDIX

Although the focus of this paper is on the effects of education and the educational environment, it may be helpful to provide some explanation of the control variables used in the fuller models for trust and participation. The control variables used by NJS-B for their membership equation are fully described by them, and the coefficients are reported in NJS-B [1996, 247]. The variables and coefficients from our alternative control variables are shown in Tables A1 and A2. The following paragraphs explain the variables and their effects on trust and participation.

The TV generation variable is a measure of the likely exposure to television during an individual's formative years, designed to enable a specific test of the negative link between television and social capital put forward by Putnam [1995b; 2000]. For each individual, the value taken by the variable is equal to the fraction of a person's pre-adult years (from birth to twenty) during which they were potentially exposed to television, weighted by a measure of the availability of television derived from data showing the increases in U.S. television ownership from almost zero in 1950 to over 90 percent in 1959. The TV generation variable takes a value of 1.0 for all those born since the late 1950s, and zero for all those born before 1930, with intermediate values for those born in the intervening period. The GSS results suggest that someone brought up before television has a trust level that is higher by .07, or almost 20 percent, compared to someone exposed to television throughout his or her formative years.

The TV era variable is the average value of the TV generation variable for all other current respondents in the same census region. It thus relates to the individual TV generation variable exactly as the average education variable relates to the individual measure of education. Exposure to television, like educational attainment, may have relative, additive or super-additive effects. The results for trust provide some evidence of super-additive effects, with the explanation presumably being the same as for education. If watching television makes individuals less trusting, then it is less rational to be trusting in an environment where many others have also been subject to the same influences. Whether this is in fact the appropriate explanation for the correlation between the TV-related variables and social trust is an issue that we leave aside here, since our primary purpose now is to explore the effects of relative and absolute education.

Lifecycle is a variable equal to the absolute value of the difference between one's age and 40. It is intended to capture a possible cohort effect in memberships, with presumed increases in participation up to the age of forty and decreases thereafter.

It is to be interpreted jointly with the age effect, as the two coefficients together determine the relative sizes of the participation changes before and after the age of 40. It is included in the equation to ensure that the TV generation variable is not taking credit or blame for the combination of a cohort effect and a population bulge working its way through the demographic structure. The lifecycle and age variables combine to provide a two-part age effect, with separate slopes, and possibly separate signs, before and after the age of forty. In the GSS equation for trust, the coefficient on lifecycle is negative and much larger than the positive coefficient on age. Putting the two together implies that trust rises with each year of age (by $.0283 + .0037 = .0320$) until the age of 40, and then falls with each year of age thereafter (by $.0283 - .0037 = .0245$).

The year variable is the year during which the survey took place. A negative coefficient implies that trust was falling through time after accounting for any effects of changing demographic structure, education and any other included variables, such as the extent of likely exposure to television. As already noted, it takes a significant negative coefficient in all the trust equations, with larger and more significant effects in the simpler equations.

The variable male takes the value 1.0 if the respondent is male, and zero otherwise. The GSS results suggest that males are significantly more trusting than females, while the DDB results show the reverse. How can this discrepancy be explained? Fortunately, there is additional evidence that helps to unravel this puzzle. In 1983 the GSS sample was split in two, with half asked the usual question, and the other half asked the simpler question "Generally, would you say that most people can be trusted?" Both males and females were more likely to answer yes to the simpler question, but the difference was far greater for females than for males. The difference is so great that females are significantly more trusting than males if asked the simpler question, but significantly less so when asked the question with the alternative "You can't be too careful". The implication would appear to be that females are more cautious than males, but are also more inclined to think others to be trustworthy. Thus it would appear that the difference in gender effects between the GSS and DDB surveys is sufficiently explained by the difference between the forms of the question, without settling the question of whether there are gender differences in perceptions of honesty and trustworthiness.

The variable divorced takes the value of 1.0 for every respondent reporting their marital status as divorced. The results show that divorced persons, in either the GSS or DDB results, are significantly less trusting. An alternative family status variable, used by NJS-B, takes the value of 1.0 for persons so describing their marital status. The two variables give very similar results, with a slight empirical preference for the married version. This suggests that single and widowed individuals have trust levels that are closer to those of divorced persons than of those who are currently married, a presumption that is supported by more specific tests. The choice of which variable to use for family status has no effect on the estimates of the effects of education. The same is true for various variables reflecting ethnic differences in trust and participation.

There are two regional dummy variables in the equation. West North Central takes the value 1.0 for all respondents in the WNC census region. The variable south takes the value of 1.0 for all respondents in the South Atlantic, East South Central and West South Central census regions. Trust levels are higher in the WNC region, and

lower in the south, than elsewhere in the country, even after allowing for differences in education, demography and exposure to television. These are the only regional effects that were found to be systematically present in the trust and membership data, with south being important only for trust and for some individual membership categories, but not for total memberships.

TABLE A1
Comparing GSS and DDB Evidence on Social Trust
Full results including control variables

Equation	(i)	(ii)	(iii)	(iv)	(v)	(vi)
	GSS	DDB	DDB	GSS	DDB	DDB
	1972-96	1975-97	1975-97	1972-96	1975-97	1975-97
Sample	22445	76156	76156	22445	76156	76156
Dependent variable: Trust		Honest	Honest (binary)	Trust	Honest	Honest (binary)
Education in years	.0439 (40.3)	.0093 (24.4)	.0175 (24.8)	.0391 (37.0)	.0062 (16.0)	.0133 (18.7)
Average education in region (yrs)	.0244 (3.2)	.0057 (1.6)	.0096 (1.4)	.0602 (11.0)	.0130 (3.6)	.0206 (3.1)
TV generation	-.0700 (3.3)	-.0168 (2.6)	-.0403 (3.4)			
TV era (average of TVgen in region)	-.0583 (0.7)	-.0646 (2.1)	-.0514 (0.9)			
Lifecycle	-.0282 (6.5)	-.00098 (7.1)	-.0022 (8.9)			
Age	.0037 (6.8)	.00398 (21.4)	.00576 (16.7)			
Year	-.0071 (4.4)	-.0033 (5.2)	-.0054 (4.6)	-.0120 (20.8)	-.0048 (22.0)	-.0073 (18.1)
Divorced	-.061 (5.0)	-.0415 (10.9)	-.0683 (9.7)			
WNC Region	.102 (8.4)	.020 (6.0)	.0358 (5.6)			
South	-.051 (6.4)	-.0209 (10.3)	-.0358 (9.5)			
Male	.033 (5.4)	-.013 (7.1)	-.025 (7.4)			
R ²	.1074	.0758	.0517	.0754	.0135	.0110
SEE	.466	.248	.460	.475	.256	.470

The previous discussion has related mainly to the GSS results. The pattern of the DDB results is very similar, except for the gender difference already analysed. Comparing the size of the effects is made more difficult by the differences in the form of the two questions and the gradation of the answers. The DDB trust variable is explained in two ways, once coded as a fraction, and again converted to a binary form similar to that used for the GSS question. The fractional version gives a more precise equation, suggesting that the shades of agreement or disagreement are explained by the same variables used in explaining the binary choice, and this information is suppressed in the binary coding. Comparing the binary form of the DDB equation with the GSS equation, the standard error of the estimate is very similar (about .46 in both cases), while the explained variance is almost twice as high for the GSS

equation. The implication of this is that there is more variance of the explanatory variables among the GSS than among the DDB respondents. This would be the case if, for example, the techniques used to select the DDB sample had the effect of producing more homogeneity than is found in the GSS sample or in the population as a whole. This appears to be so, as the DDB sample has fewer divorced persons, and less variance in its distribution of both age and education. Thus the GSS sample has more variance in its dependent variable, and in its key independent variables. This explains why the explained variance is higher for the GSS sample. The fact that the unexplained variance is the same suggests that whatever is missing from the model is found equally in the GSS and DDB samples.

TABLE A2
GSS and DDB Evidence on Social Engagement
Full results including control variables

Equation	(i)	(ii)	(iii)	(iv)
	GSS	DDB	DDB	DDB
	1975-97	1975-97	1975-97	1975-97
Sample	19214	71246	71246	71246
Dependent variable:	Number of	Club meets	Comm	Dinner
	memberships	(times/yr)	projects	parties
Education in years	.2283 (51.7)	.7141 (36.1)	.3504 (33.5)	.4808 (41.5)
Average education in region (yrs)	-.0201 (0.7)	.0057 (0.1)	.1710 (1.7)	.9600 (8.7)
TV generation	-.5646 (6.1)	-.3000 (0.9)	-.3243 (1.8)	-.5206 (2.7)
TV era (average of TVgen in region)	-1.067 (3.2)	-2.988 (2.0)	-1.565 (2.0)	-1.768 (2.0)
Lifecycle	-.0085 (4.4)	-.0475 (6.7)	-.0289 (7.7)	.0153 (3.7)
Age	.0031 (1.2)	.1404 (14.5)	.0422 (8.2)	.0305 (5.4)
Year	-.0138 (2.1)	-.2311 (7.4)	-.0428 (2.6)	-.1746 (9.5)
Divorced	-.2434 (4.8)	-.7877 (4.0)	-.3816 (3.7)	-1.039 (9.0)
WNC Region	.290 (5.8)	1.106 (6.2)	.361 (3.8)	-.077 (0.7)
South	-.034 (0.9)	-.531 (5.0)	.080 (1.4)	-.755 (12.2)
Male	.238 (9.4)	-1.934 (20.5)	-.508 (10.2)	-1.051 (19.0)
R ²	.1469	.0594	.0265	.0576

Turning to the participation results in Table A2, we see that the control variables tend to have significant effects, of the same sign, for all four measures of social involvement. The negative effects of TV exposure on participation, which are strong for the GSS membership variable, are less strong for the DDB measures of participation. The negative effects of average exposure to TV are fairly similar across the categories, and larger than the estimated effects of individual exposure. Males are more involved in some types of GSS memberships, but are significantly less intensively involved

in any of the three DDB measures of social engagement. Participation is greater in the WNC census region (except for dinner parties) and lower in the south (except for involvement in community projects).

NOTES

The authors wish to thank the editors for their patience and good advice, and Haifang Huang and Josh Bolian for their assistance with the new empirical results referred to in this updated version of the paper. The delays are entirely ours.

1. Campbell [2006] includes both own and contextual levels of education in his analysis of data for a number of European countries, with the contextual level of education defined by nation and age cohort. His results generally favour the absolute model over the relative and cumulative alternatives. There are two main exceptions: relative education is dominant for his 'zero sum' forms of political engagement, while the cumulative model is strongly preferred for general social trust.
2. Logically, of course, the appropriate standard of comparison might be even more local; indeed, we suspect that it is. Our purpose here, however, is merely to show that even slightly narrowing the spatial standard of comparison (moving from national to regional standards) can substantially affect one's conclusions. A fortiori defining more localized standards of comparison should improve estimates still further, although at some point, an increasingly localized definition might in fact become smaller than the real range of externalities. For example, although we considered using state- or county-level standards in the present analysis, we set that aside out of concern that some externalities might be carried by inter-county or even inter-state commuters. Interregional commuting is vanishingly small. In subsequent work using "relative education," we recommend that sensitivity testing be done to assess the most suitable level of aggregation, but however low that turns out to be, it will, we believe, be smaller than the nation as a whole.
3. Again, of course, this implementation could be tuned more finely to fit specific models of civic competition. Our claim is merely that it is more plausible than a purely static, backward-looking model.
4. A close approximation to the NJS-B measure of educational environment, measured in average years of education, is provided by $edenv = 9.0 + .07273 * yob - .03333 * (yob - 55) * d55$, where $d55$ is a variable that takes the value 1.0 where $yob > 55$, and zero elsewhere. The equation is thus piecewise linear with a kink at $yob = 55$. The national average level of education thus rises by .073 for each year of birth up to 1955, and at .04 (= .073 - .033) per year for years of birth after 1955, as shown in Table F3 of NJS-B [1996, 232]
5. The authors argue that the negative effect from average education levels "...represents decreasing marginal returns for additional education rather than indicating increased competition." NJS-B [1996, 162].
6. The GSS participation results cover 1974 through 1994 only, because the membership questions were not asked in the 1972 and 1996 surveys.
7. The use of a reference group based on the current adult population also has the advantage that the change in the sample average level of own-education from year to year is equal to the corresponding change in the educational environment.
8. The DDB results are very similar to the results of a 1983 GSS experimental form that offered only the trusting alternative.
9. The nature and estimated effects of the control variables are reported in the Appendix. The coefficients reported are from the linear probability model estimated by OLS. The normality of residuals assumed in OLS regression is false, since the dependent variable falls into the range between 0 and 1. We used logit estimation of the binary equations to see if non-normality of the residuals is affecting the results. Logit estimation of the equations with binary dependent variables, with coefficients renormalized to be comparable with those from the linear probability model, shows no change in the pattern of results. For example, the transformed logit coefficients for education and average education in equation (i) are .050 ($t=37.1$) and .027 ($t=3.1$), and in equation (iv) .044 ($t=34.9$) and .067 ($t=11.2$). These are almost identical to those shown in Table 1.
10. Similarly, the time trend coefficients are larger and much more significant in the simpler model than in the fully specified model. This is as we would hope, as one of the objectives of fuller specification

is to expose the underlying reasons for the downward trends in social trust. A fully specified model would leave little or nothing to be explained by time alone, in the absence of some theory suggesting why the passage of time should be sufficient to erode social trust. The simpler model includes the time trend to ensure that any effects found for average education levels are not simply due to the fact that it shares a common time trend with some important excluded variables.

11. Our replication of the NJS-B equation does not yet provide an exact match, although the results for education and average education are so close to theirs that full replication is not likely to change the pattern of results. We have approximated the NJS-B measure of relative education by modelling it as a piece-wise linear function of year of birth, with a kink at 1955, as shown in Table F-3 of NJS-B [1996, 232]. Our current equation includes all of the other variables they use, except that hours worked are not included, since the number of valid observations in the GSS database is several thousand less than the NJS-B sample of 15887. Our current equation has a sample size of 18037.
12. But see the Appendix for some results suggesting that the prevalence of television during one's childhood appears to have contributed importantly to the inter-generational drop in political and social engagement.

REFERENCES

- Becker, G. S.** *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. New York: Columbia University Press, 1964.
- Brody, R.A.** The Puzzle of Political Participation in America, in *The New American Political System* edited by A. King. Washington: American Enterprise Institute for Public Policy Research, 1978.
- Campbell, D. T.** Education's Impact on Civic and Social Engagement. In *Measuring the Effects of Education on Health and Civic / Social Engagement*, edited by R. Desjardins and T. Schuller. Paris: OECD Centre for Educational Research and Innovation, 2006, 25-126.
- Card, D.** Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems. *Econometrica*, September 2001, 1127-60.
- Dee, T. S.** Are There Civic Returns to Education? *Journal of Public Economics*, August 2004, 1197-1720.
- Helliwell, J. F. and Putnam, R. D.** Education and Social Capital. *NBER Working Paper* No. 7121, Cambridge: National Bureau of Economic Research, 2001.
- _____. The Social Context of Well-Being, in *The Science of Well-Being*, edited by F. A. Huppert, N. Baylis and B. Keverne. Oxford: Oxford University Press, 2005, 435-60.
- Milligan, K. Moretti, E. and Oreopoulos, P.** Does Education Improve Citizenship? Evidence from the United States and the United Kingdom. *Journal of Public Economics*, August 2004, 1167-96.
- Moretti, E.** Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data. *Journal of Econometrics*, July-August 2004, 175-212.
- Nie, N. H., Junn, J. and Stehlik-Barry, K.** *Education and Democratic Citizenship in America*. Chicago: University of Chicago Press, 1996.
- Putnam, R. D.** Bowling Alone: America's Declining Social Capital. *Journal of Democracy*, January 1995a, 65-78.
- _____. Tuning In, Tuning Out: The Strange Disappearance of Social Capital in America. *PS*, December 1995b, 664-83.
- _____. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster, 2000.
- Spence, A. M.** *Market Signaling: Informational Transfer in Hiring and Related Screening Processes*. Cambridge: Harvard University Press, 1974.