Observations on the Massachusetts Special Election for U.S. Senate

Charles Stewart III
Department of Political Science, MIT
Draft of January 20, 2010

Here are some thoughts and empirical observations about the special election for the U.S. Senate from Massachusetts.

After drafting these notes, I discovered some analysis by Charles Franklin, at pollster.com, that is consistent (I believe) with what I've written, and as always, in a more sophisticated fashion. (We southern-born election geeks are like that.) See http://www.pollster.com/blogs/how_massachusetts_votes_shifte.php

In addition, the New York Times provides some nice graphics about the election. See http://www.nytimes.com/interactive/2010/01/19/us/politics/massachusetts-election-map.html.

Boston.com also has some interesting data plots that provide helpful insights. See http://www.boston.com/interactive/graphics/20100119_turnout/

Now, my own take.

Commentary on the race breaks down the results into two effects, which I call the (1) vote shift and (2) turnout effects. Because most commentators are interested in importing meaning into the race, these effects usually come bundled with causal arguments. The vote shift effect, for instance, is usually associated with an argument that the results were caused by an angry electorate out to "throw the bums out," "send Obama a message," etc. The turnout effect is usually associated with an argument that Coakley failed to mobilize a base of Democratic voters who have recently provided big majorities to Governor Patrick and President Obama. I am more interested in understanding the trends than in providing normative commentary. However, I can't always resist, as you'll see.

Vote shift

Brown won 52% of all votes cast. This compares with 36% for John McCain for president in 2008 and 35% for the Republican gubernatorial candidate in 2006, Kerry Healey, in 2006. This represents a 17% pro-Republican shift over recent elections.

Because of third party candidates, it is sometimes helpful to calculate vote shifts in terms of vote margins. Using that method, Brown's victory margin over Coakley was 5% points, compared to a deficit of 26% points for McCain and 21% points for Healey. This leads to a calculation that the pro-Republican shift was between 26% and 31% in the victory margin.

Because I think it is more intuitive, I will analyze the vote shift using the first measure, which is simply the percentage of the vote received by the candidates. Also, because 2008 presidential

election is more recent, I use that election as the baseline, rather than the 2006 gubernatorial election. I suspect (but I haven't confirmed) that the analysis doesn't change.

The first question is whether the shift, or swing, was uniform throughout the state. The answer is illustrated in Figure 1. The *y*-axis shows the percentage of the vote received by Brown in each of the 351 cities and towns; the *x*-axis is the percentage of the vote received by McCain. The tokens are scaled to be proportionate to the number of votes cast in each municipality. The solid line is the 45-degree line, where the scatterplot would line up if Brown had received the same vote share in each town as McCain. The dashed line is simply the 45-degree line shifted up 17% points, to show where Brown would have been if his statewide vote swing were evenly distributed across the municipalities.

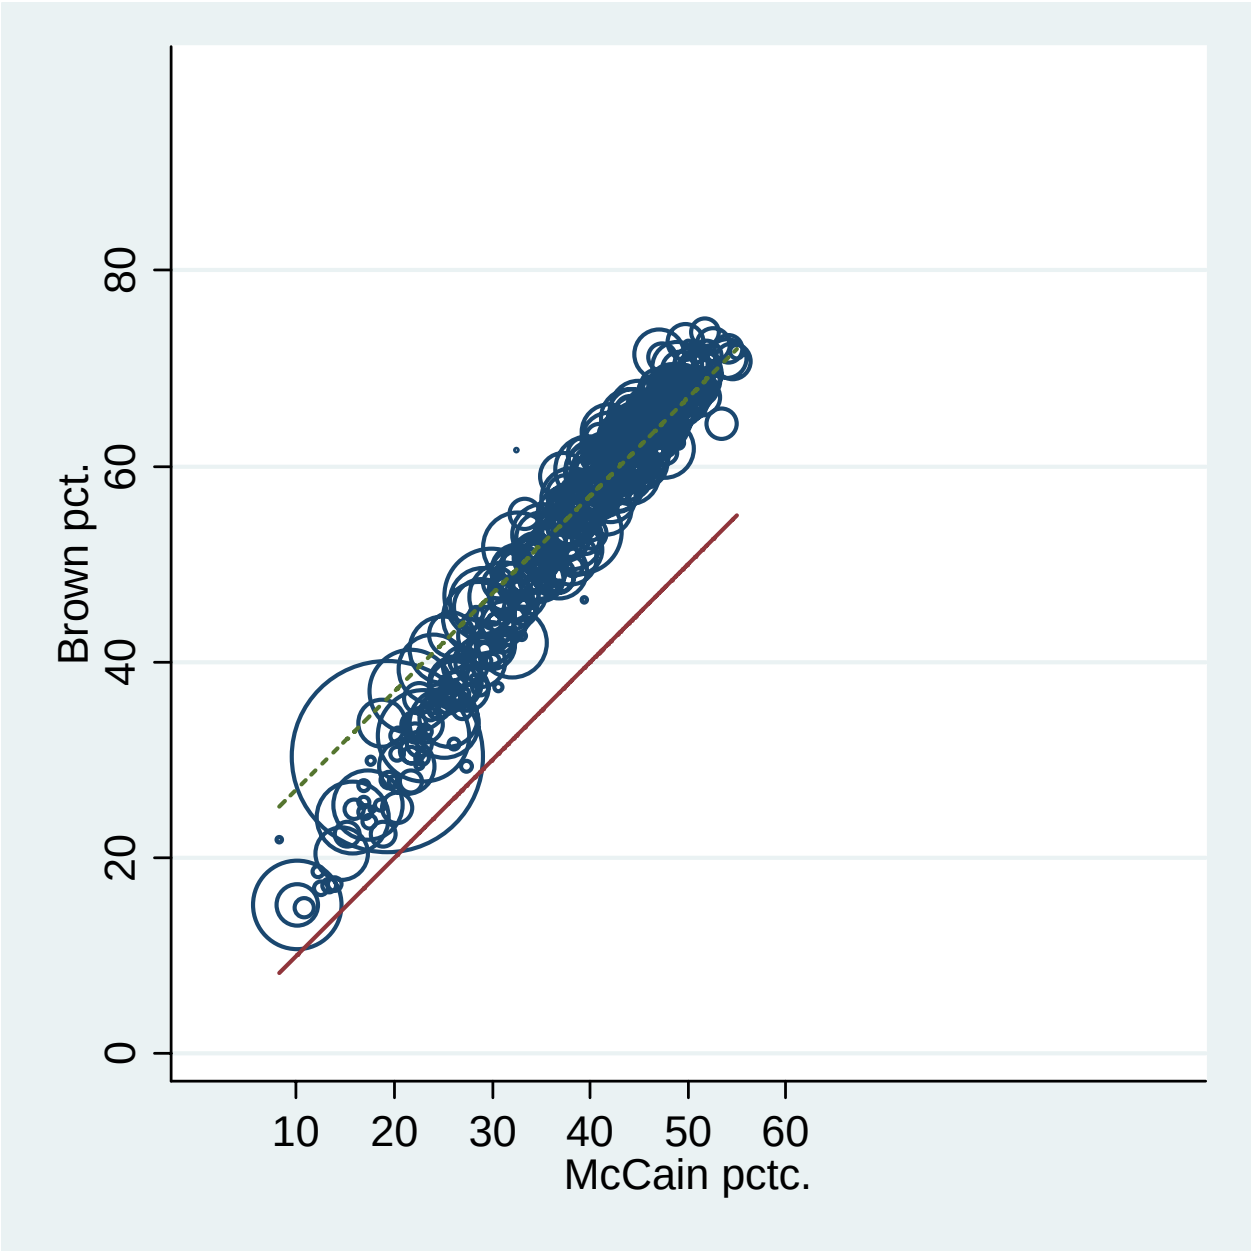

The plot shows that the swing wasn't uniform across municipalities. Brown's biggest gains came in the cities and towns that provided McCain his largest vote margins — which I will refer to as the "most Republican towns," as a shorthand. It is important to note that Brown also gained in the more Democratic cities, too, but by a smaller amount.

Looking at the 100 most Republican towns (by the McCain vote), Brown's swing was 18% points; in the 100 least Republican towns, Brown's swing was 12% points.

So, more of a shift in the more Republican areas, but still a shift statewide — Brown did better in every city and town than McCain did.

The same can be said within the city of Boston, where I've been able to get the precinct returns. Figure 2 is that picture. There are two differences from Figure 1. First, the data tokens aren't proportional to turnout because the precincts, to a first approximation, are equal in population. Second, the dashed line reflects a uniform swing in Boston of 11% points, rather than the statewide swing of 17% points.

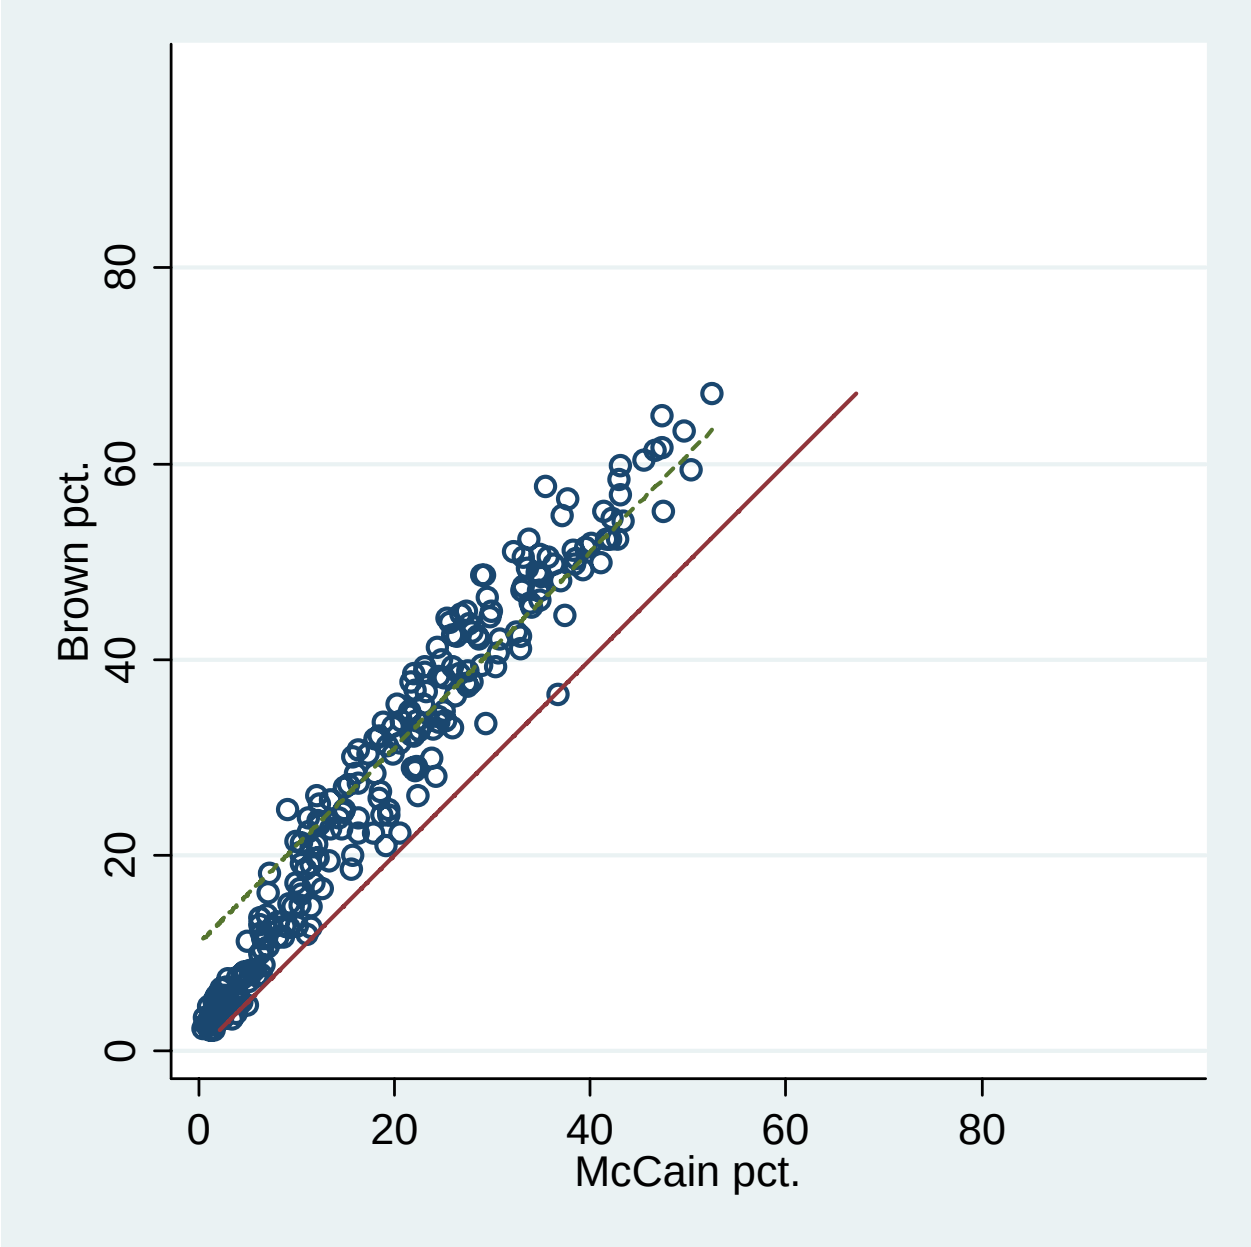

The statewide pattern is evident in Boston, though to a smaller degree.. The 100 most Republican precincts in Boston showed a 13% point shift to Brown; the 100 least Republican precincts showed a 4% shift.

These statewide and Boston patterns can be summarized in a simple regression that uses the Brown percentage as the dependent variable and the McCain percentage as the independent variable, weighting each observation by the number of votes cast. (Standard errors in parentheses.)

	Stateside	Boston
McCain pct.	1.28 (0.01)	1.27 (0.02)
Intercept	5.25 (0.47)	3.79 (0.45)
N	351	253
R^2	.97	.95

There are (at least) two explanations for this un-uniform swing pattern. The first is individual: lots of Republican self-identifiers who voted for Obama switched to vote for Brown. The second is purely contextual: voters living in districts with lots of Republicans, regardless of their own partisan leanings, were more drawn to Brown than those living in areas with lots of Democrats. This could be a friends-and-neighbors effect — when you talk to your neighbors about politics, you discover more of them supporting Brown, and you are convinced to do so, too. Or, this could be due to common social characteristics shared by neighbors — for instance, people living in Republican neighborhoods, regardless of their own party affiliation, may already have health insurance and view Brown as a vehicle for helping to retain the value of the health insurance coverage they already have.

Unfortunately, without exit polls (ugh) or panel studies of identical voters across time, we can't distinguish among these various explanations. It's probably a bit of both.

The statewide data also provide evidence that the shift toward Brown, compared to McCain, was *not* driven by income differences. This is illustrated in Figure 3, which plots the Brown share of the town-level vote against median household income.

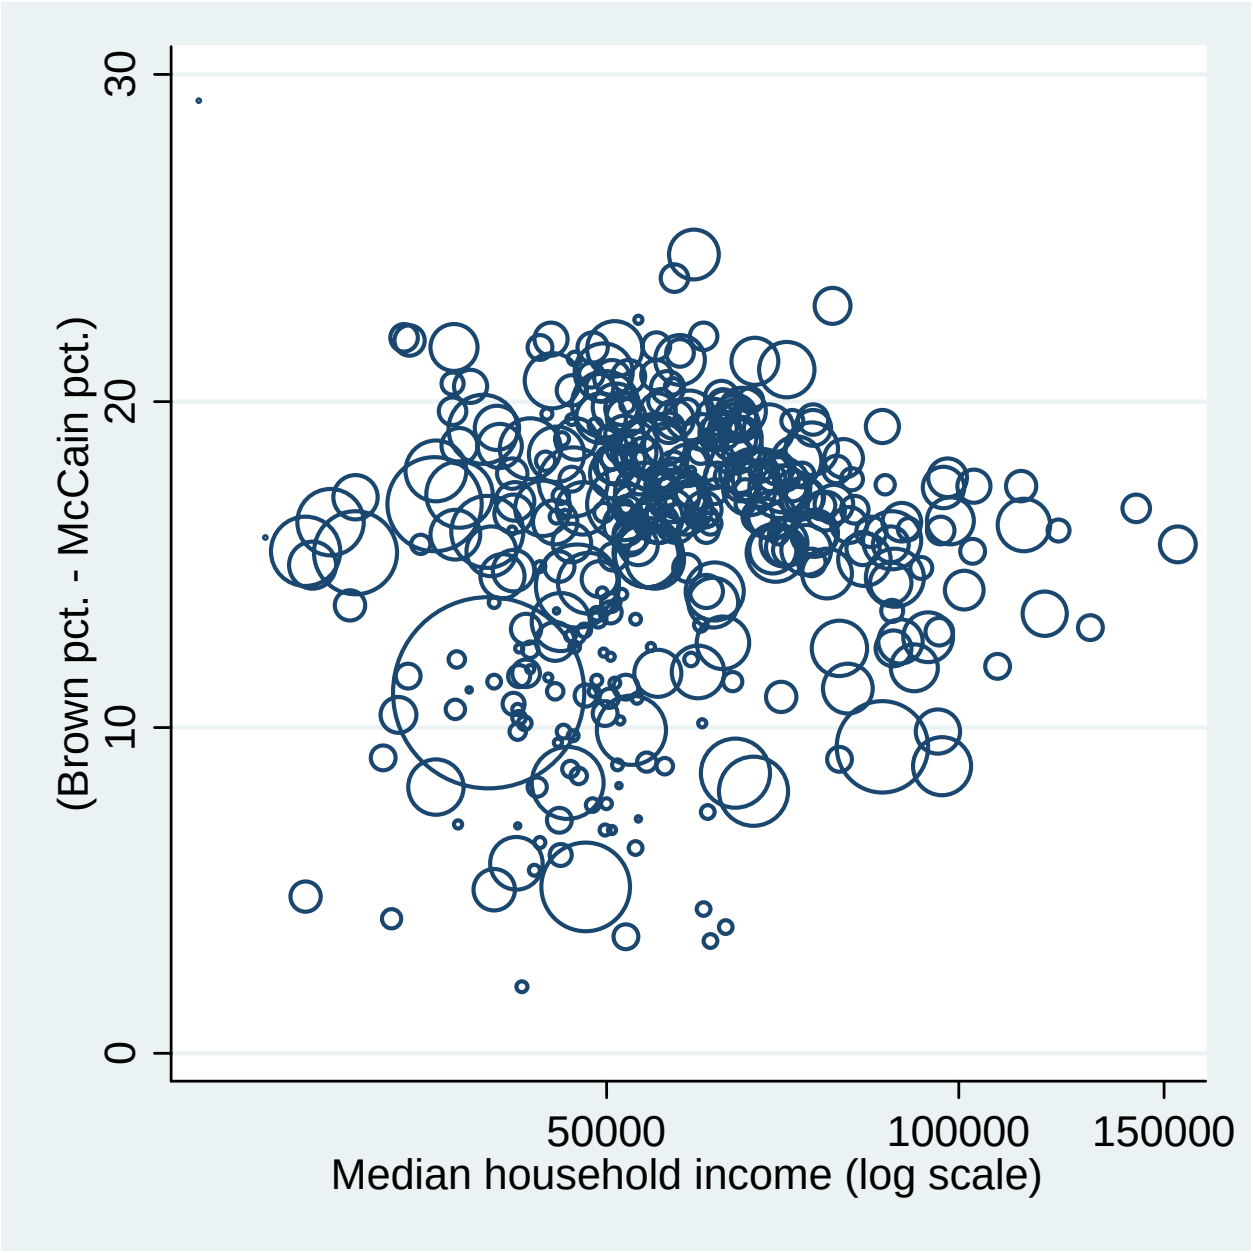

The correlation is very weak (r = .01). Nothing to look at. Move along. (The same won't be true when we look at turnout.)

The Boston data allow us to look at race. (Alas, we don't have income statistics for Boston precincts and I haven't been able to put together the comparable race figures for Massachusetts towns.)

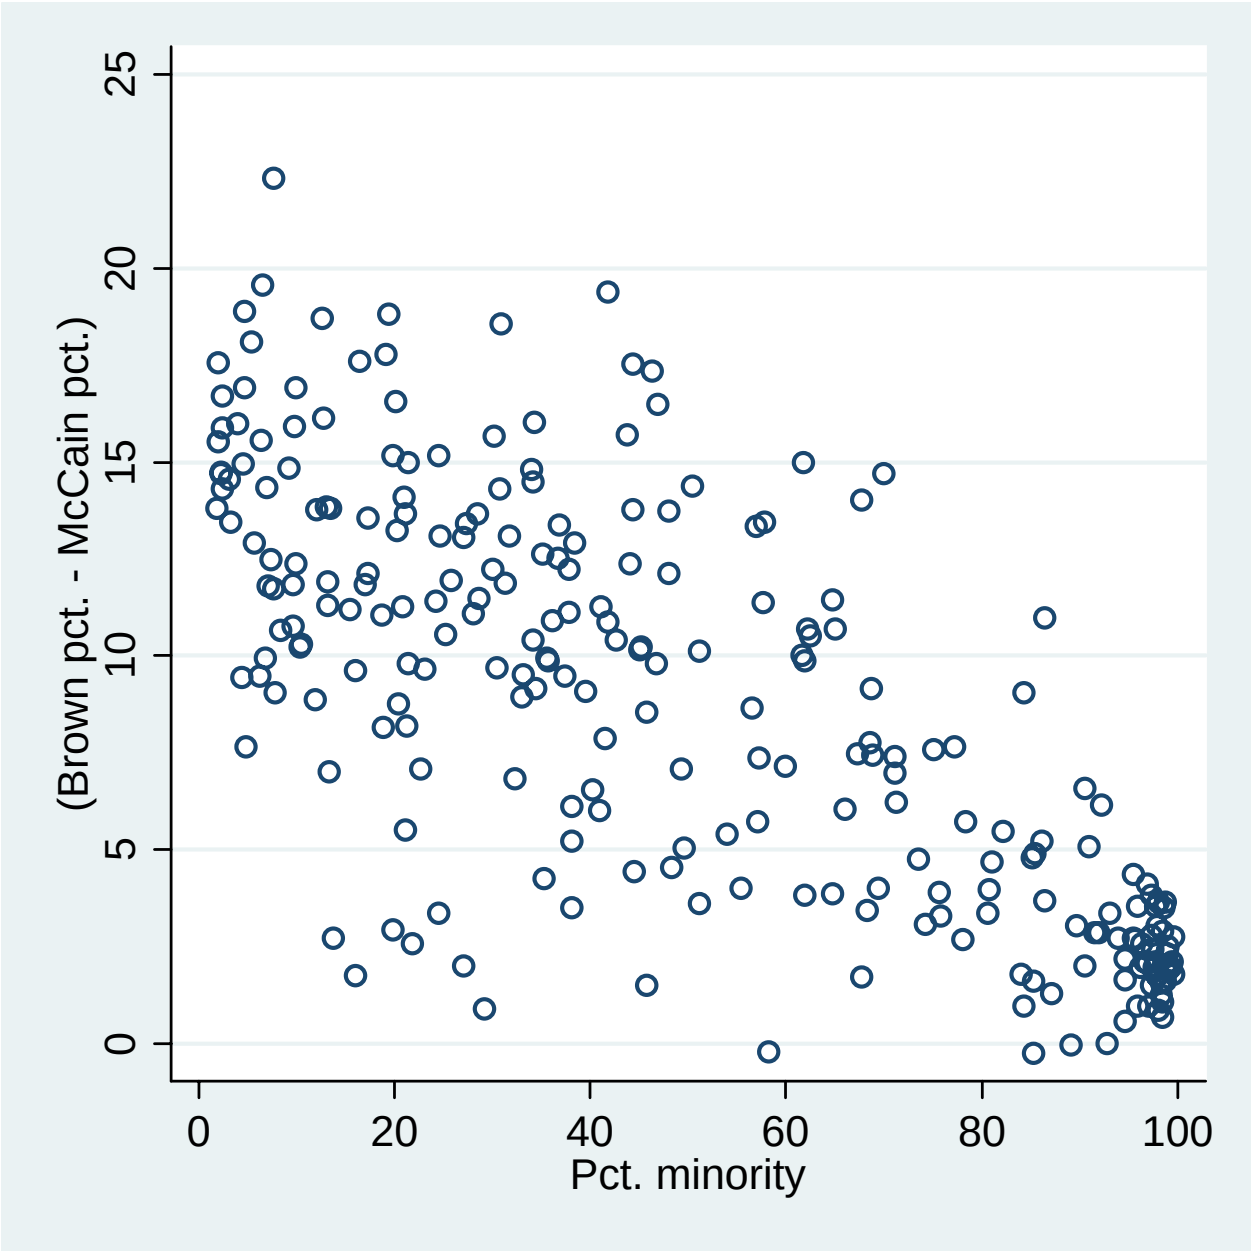

This plot is quite a stunner (though not necessarily a shocker). The shift toward Brown was in the "whiter" parts of the city. In fairness, even the heavily minority precincts shifted toward Brown, but by a much smaller margin. In the 100 precincts with the lowest minority population, the shift toward Brown was 13% points; in the 100 precincts with the largest minority population, the shift was an average 5% points.

Here is the regression describing the Figure 3 scatterplot. (Here, I weight by turnout.)

Pct. minority	-0.12 (0.01)
Intercept	14.54 (0.39)
N	253
R^2	.51

This looks like an increase in racial polarization in the vote in Boston, compared to the 2008 presidential election. To test this, I do the standard regression in which the dependent variable is the percentage of the vote for the Democratic candidate and the independent variable is the percentage of the precinct that is minority. Here are the regressions for 2008 and 2010:

	2008	2010
Pct. minority	0.35 (0.02)	0.46 (0.02)
Intercept	62.9 (0.86)	49.2 (1.0)
N	253	253
R^2	.67	.68

Using these estimates, we can come up with quick-and-dirty estimates of the racial voting patterns in these two races. The intercepts are the estimate of the percentage of white voters who supported the Democratic candidate. These estimates suggest that 62.9% of White Boston voters voted for Obama and that 49.2% of White Boston voters voted for Brown, for a 13.3% shift among white voters. Combining the intercepts and slope terms give us estimates of how minority voters voted. They also suggest that minority Boston voters voted 98% for Obama and 95% for Brown, for a 3% pro-Brown shift among minority Boston voters.[1]

Stated another way, the racial polarization in these two elections increased by about 11% points between the two elections.

As always, there are problems of ecological inference here, but they point to a racial dimension of the vote that has gone mostly unremarked in commentary about the election.

---

[1]The calculation is done this way. For 2008, the minority estimate is derived by setting minority population to 100% and then doing this calculation: 62.9 + .35*100 = 97.9.

Turnout

These vote swings can come about through two mechanisms: (1) Some members of a fixed electorate change how they vote or (2) the electorate itself shifts, favoring Brown. The data clearly suggest the latter was the dominant factor here.

First, here is the most striking statistic to me: In 2008, Barak Obama received 1,904,098 votes; Coakley received 1,058,682, for a decline of 845,416 votes. Also in 2008, John McCain received 1,108,854; Brown received 1,168,107, for a *pick-up* of 59,253 votes. We will never know for sure, but this looks like some evidence that Brown's "ground game" was in a different league from Coakley's. (Some might say that Democrats were just demoralized, but I consider that just to be the other side of the coin. Why shouldn't Republicans, in a heavily Democratic state, have been demoralized after the primary?)

At first blush, it appears that overall turnout, compared to 2008, declined fairly uniformly throughout the state. Here is the plot of 2010 turnout at the town level against 2008 turnout: (I apologize for changing graph formats, but Stata was not cooperating when I tried to do the log plots.)

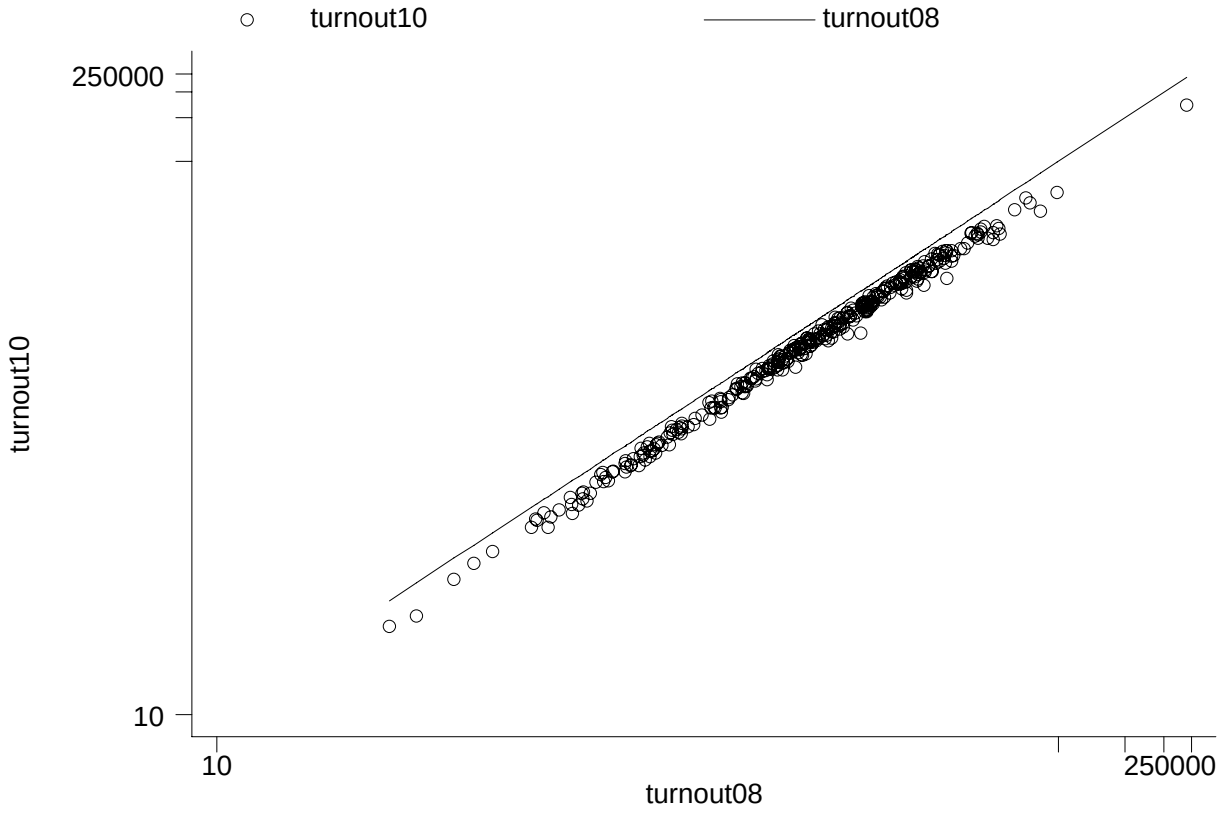

The regression confirms this: (The dependent variable is logged turnout for the special election against logged turnout for the 2008 election. Observations are not weighted.)

log(Turnout 2008)	1.00 (0.004)
Intercept	-0.31 (0.03)
N	351
R^2	.995

The same is not true in Boston, however:

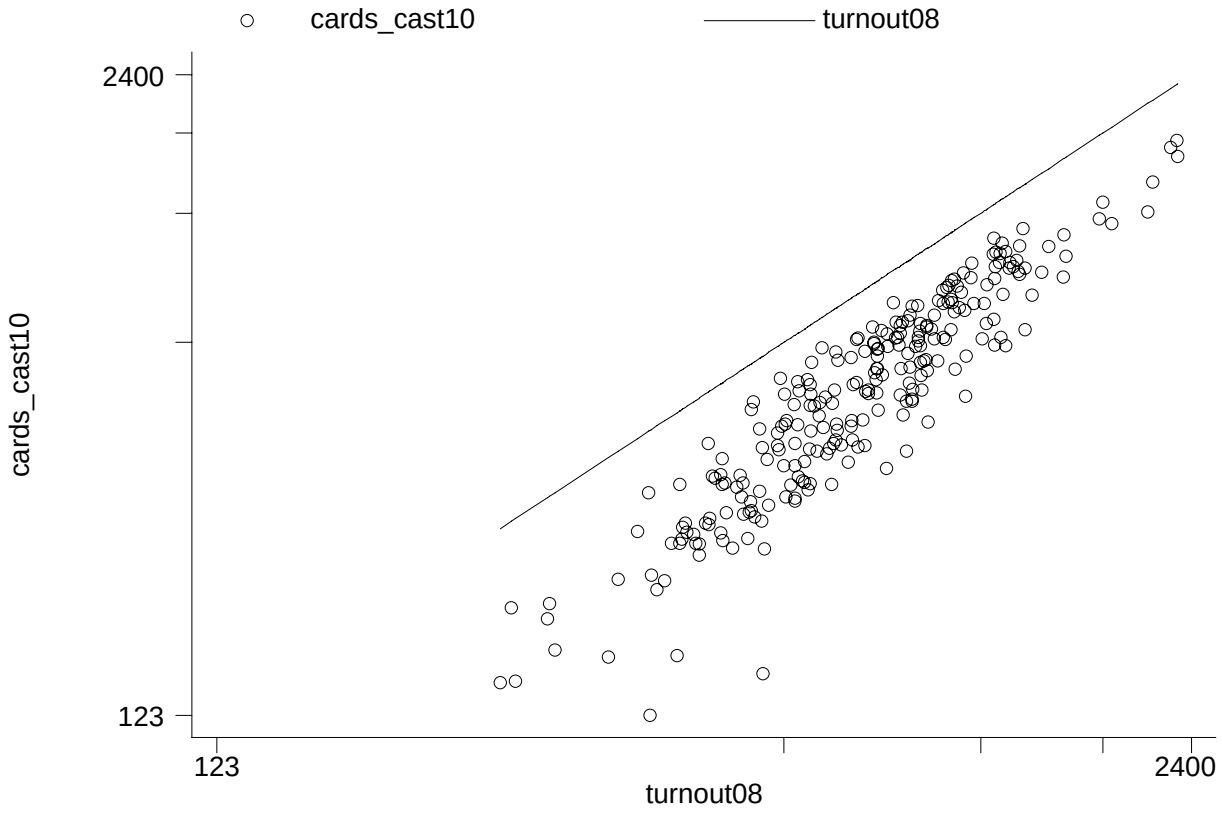

log(Turnout 2008)	1.20 (0.03)
Intercept	-1.86 (0.21)
N	253
R^2	.86

Whether this is politically relevant is unclear to me.

Returning to the statewide results, although the initial look suggests that turnout change was uniform across cities and towns, further analysis reveals that it was greater in Republican areas and in high-income areas.

Here is the plot of *change*[2] in turnout from 2008 to 2010 against the vote share for McCain:

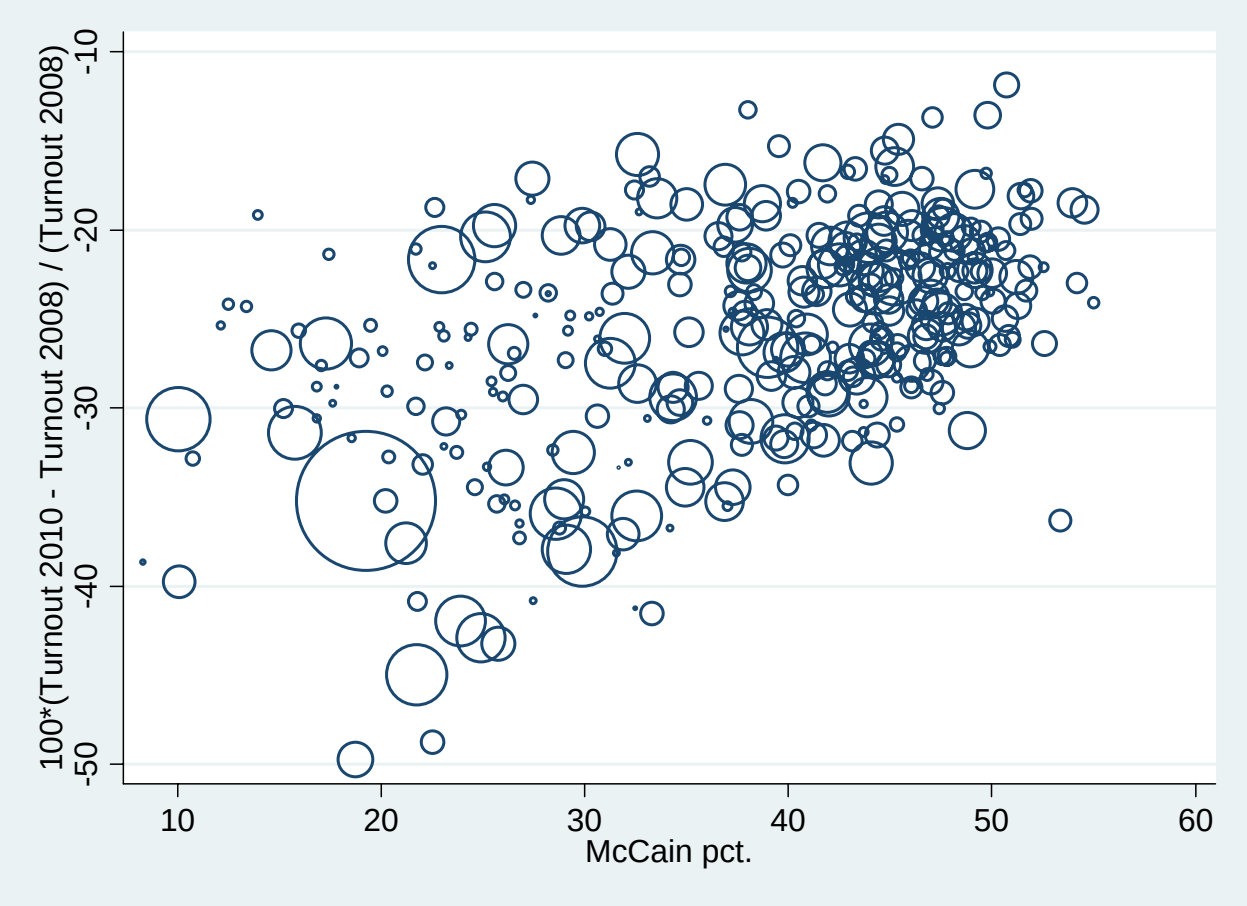

(As before, the token sizes are proportionate to turnout.)

Here is the plot of change in turnout against (logged) per capita income:

[2]Turnout change is defined as the percentage change in turnout, compared to 2008: 100 x (Turnout 2010 - Turnout 2008) / (Turnout 2008)

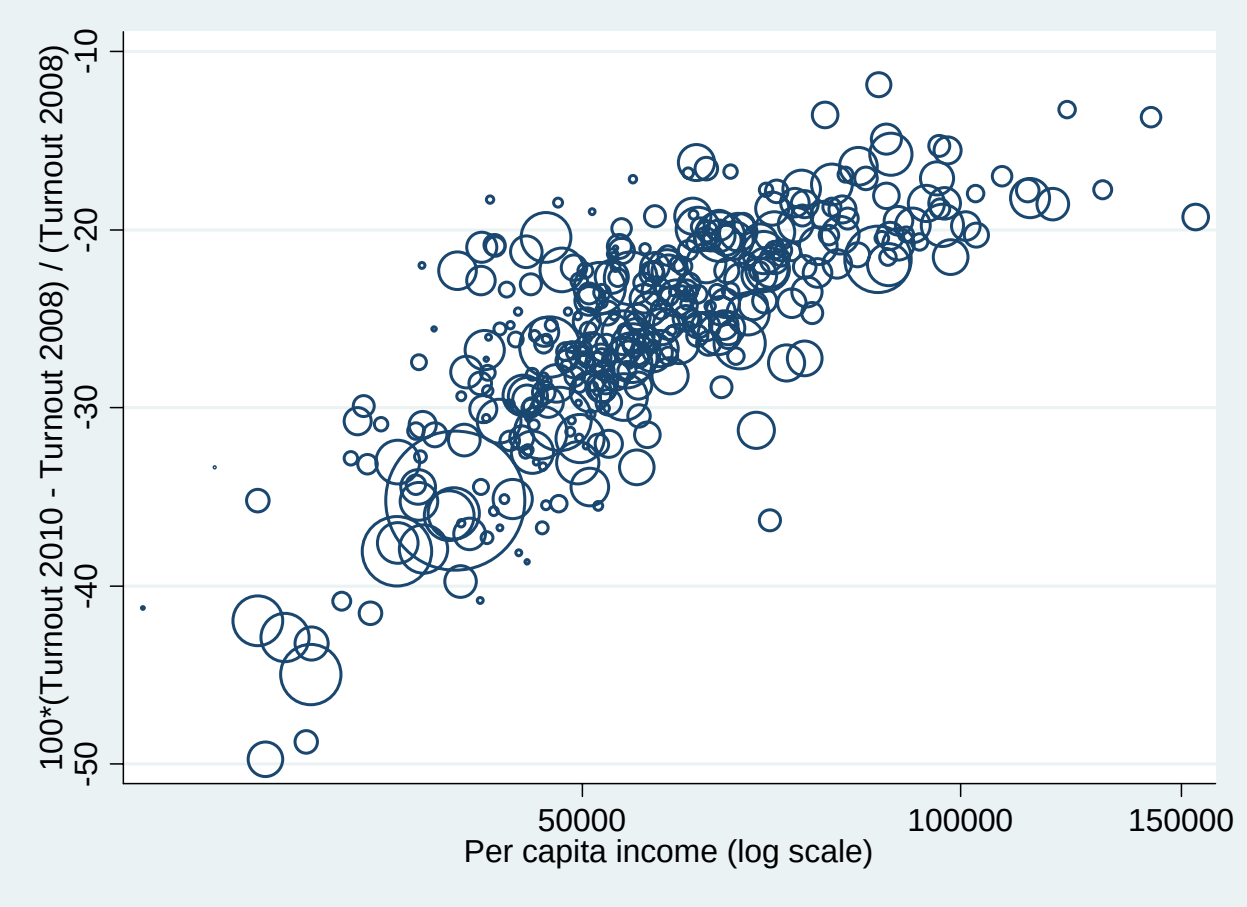

Turnout went up in Republican and high-income areas of the state. Of course the two are related, but the correlation is sufficiently low (r = .38), that both partisanship and income are likely separate factors.

Here are the supporting regressions, first doing separate bivariate regressions for each independent variable, and then combining them into a multivariate regression:

McCain pct.	0.36 (0.03)	—	0.16 (0.02)
log(Per capita income)	—	18.59 (0.66)	16.3 (0.66)
Intercept	-40.1 (01.1)	-229.8 (7.2)	-210.1 (7.0)
N	351	351	351
R^2	.31	.70	.75

The bivariate McCain coefficient is cut in (more than) half in the multivariate analysis, which suggests to me that while partisanship certainly had something to do with turnout differentials, the effect of income appears to be much more robust.

What do we learn from Boston? Again, we unfortunately don't have income statistics for precincts, but we do have racial composition. (Future drafts will add race to the statewide analysis.)

Here is the plot of the change in turnout against the vote share for McCain in Boston precincts:

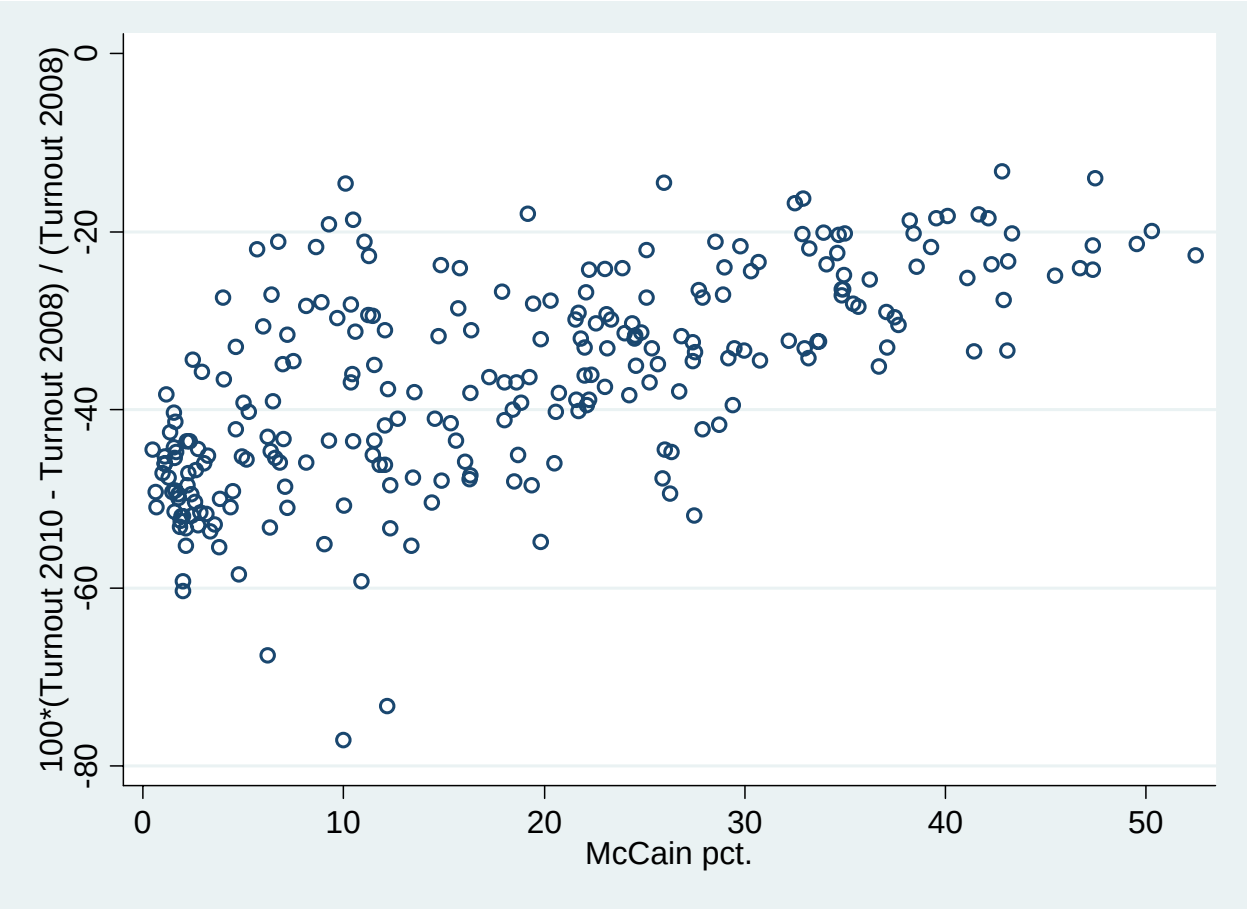

Here is the plot of the change in turnout against the percentage minority population:

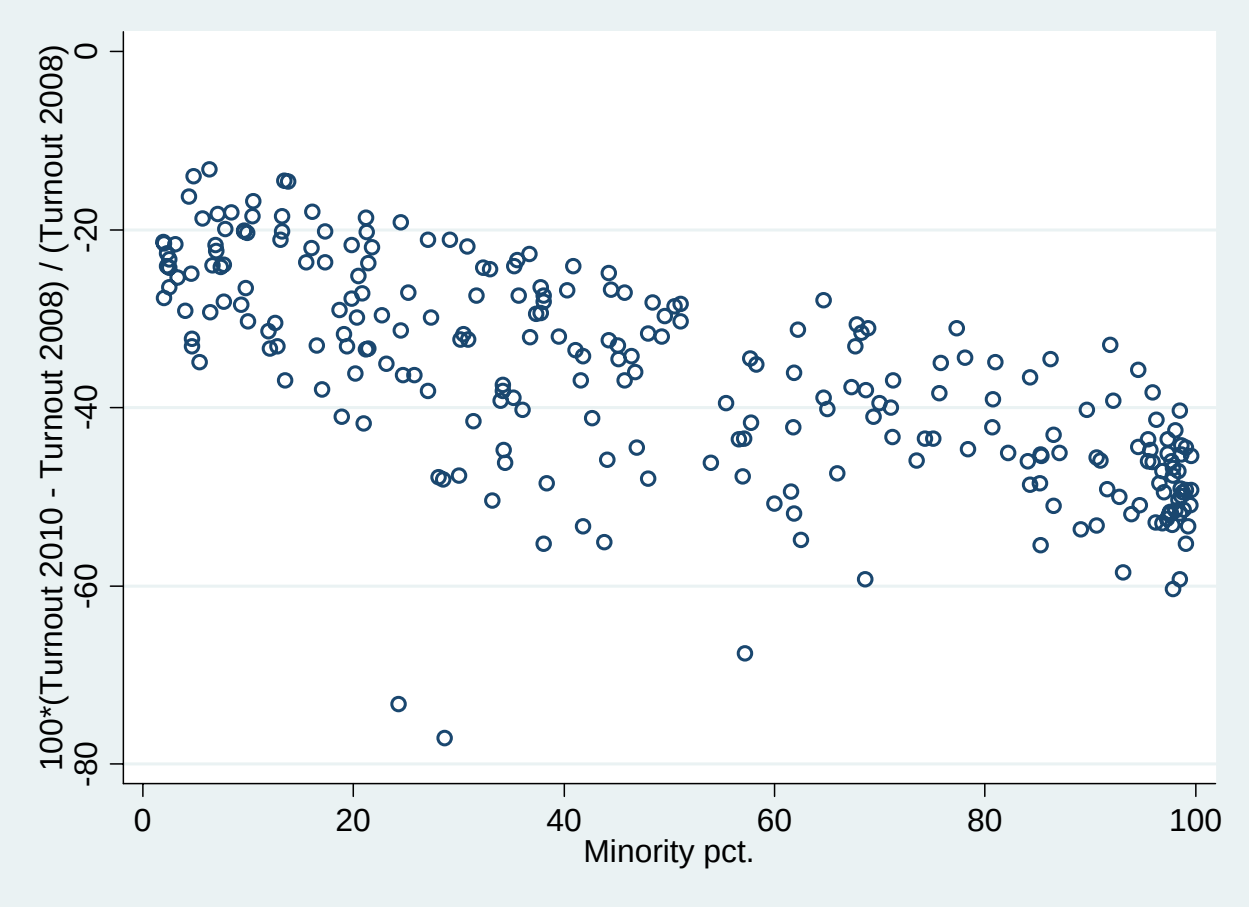

In Boston, turnout went up in the Republican parts of the city and went down more in the minority parts. The two independent variables are correlated pretty strongly (r = -.80), so this is where the multivariate analysis is particularly important.

Here are the regressions:

McCain pct.	0.51 (0.04	—	0.12 (0.06)
Minority pct.	—	-0.25 (0.01)	-0.21 (0.02)
Intercept	-43.8 (0.95)	-22.7 (0.76)	-26.6 (2.1)
N	253	253	253
R^2	.41	.55	.55

The partisan effect takes the big hit in the multivariate analysis, suggesting that it is primarily race, rather than partisanship, that affected turnout.

Conclusions

These are very preliminary analyses, subject to revision as new facts come to light and my own errors have been pointed out.

This is what we seem to *know* from the correlations:

1. Statewide, the big Brown vote shift was concentrated in the Republican parts of the state. Within Boston, the shifts came in the Republican and white parts of the city.
2. Statewide, the big turnout shift was concentrated in the Republican and wealthy parts of the state. Within Boston, turnout held firmest in the Republican and white parts of the city.

The statewide and Boston analyses have a bit of an apples-to-oranges quality to it. I will soon add racial data to the town dataset, to see if the racial patterns discerned in Boston carry over to the entire state. I will also work to see if I can get income figures for Boston precincts, but that is likely to be harder.

It would be nice to get data about how many residents of the different towns are covered by employee medical insurance, to test more directly the notion that Brown's appeal was, in part, due to his opposition to the health insurance legislation in Congress.

One question that emerges is what difference it made that wealthier, more Republican parts of the state turned out to vote in greater numbers? One way to address this is to run this counterfactual: Shift the turnout in the poorest half of municipalities to match that of the richest half, apply the vote shares we observe in those communities, and see what we get.

Prior to doing that, it is important to note that Brown received 57.0% of the vote in the 176 wealthiest municipalities, but only 46.6% in the poorest.

Running through the calculations, adding 7.4% more voters to the poorest towns would have increased turnout by about 77,600. Allocating 47% to Brown and 53% to Coakley would only net Coakley about 4,700 votes against Brown, which is not nearly enough to overcome the 100,000 vote margin he won by. Thus, it appears that Coakley would certainly have been helped by greater turnout in poorer, Democratic areas of the state, but fixing the turnout gradient would not have been sufficient for her to win.

So, what does this all mean? Without individual-level data, I can't comment directly on the individual-level explanations for Brown's victory. However, there is no escaping the fact that wealth and race are both likely suspects here. The geographical pattern of the vote shifts (see the *New York Times* maps), showing that the shifts occurred in the Boston suburbs and exurbs,

argues against the notion that this was largely a matter of angry working-class Democrats abandoning their party. More likely, it was borderline independents punishing the incumbent party for poor economic performance, along with Republicans who had no interest in George W. Bush coming home to support a national Republican who could string sentences together to make paragraphs.

Contrary to my previous analysis about a “new” Boston, the evidence here is not so sanguine for those looking to the Bay State for signs of new racial harmony.

More to come.`