PS 343 Political Economy of Developed Democracies
Notes on Surviving Hibbs Ch 5

In the first part of chapter 5, Hibbs develops a 'formal model' of first, why an individual chooses to support the president (approve of the way he's doing his job), and second, how to go from that individual-level decision to talk about aggregate popularity levels across the population. We do not have data on what individual voters/respondents think about presidential behavior over time, as if we had followed the same 1,000 people around and asked them every 3 months, "Do you approve or disapprove of the way President Flobbergap is doing his job?" Instead, we have results from randomly selected respondents in Gallop's regular national surveys, with a different group of individuals polled each time. Since we have to pick a new sample every time, and the individuals we poll are not always very well informed, or not always paying attention to the question, etc., we have to account for some level of error in how closely the survey's results about presidential support reflect the true underlying population value.

Yes, the math looks rather intimidating. (I'll admit, even I wasn't too excited when I saw it.) In actuality, it's not bad at all. It's all common sense and very "Well, duh." The math he uses isn't even very challenging – Hibbs just uses incredibly strange notation that makes it look a lot worse than it is. The hardest math in it is a single stage where he's talking about simple first derivatives (not even *doing* them, just talking about them), a place where he does a trick with statistical distributions, and finally a spot where he takes an inverse. IT'S NOT HARD. It's just moderately ugly. But I strongly urge you to take the time *now* to work through this, because it only gets uglier from here (take a flip through Franzese and you'll see what I mean).

To survive this chapter, and actually get something from it, you should plan to devote a couple of hours to reading it and working through the logic. Pick a time and place where you are comfortable and know you will not be disturbed, somewhere you can focus and go from start to finish without interruption. (This is really not coffeehouse reading for most people.) I won't lie to you; this is time-consuming and tiring, and it can be very frustrating if you are not in the right mood for it. I also strongly recommend that you use pencil and paper to work through as much of the math as you can alongside Hibbs. There's something about moving the terms around yourself that makes most people feel more comfortable, or more satisfied, that all the algebra really means something. The good news is, though, once you've gone through the logic process of this once, future models will be a lot less intimidating and also much easier to read. I suggest you use this guide by reading the paragraph here, then reading Hibbs' version and trying to work through it. Others, though, may prefer to read all of this, then read all of Hibbs. No matter your choice, feel free to curse Hibbs frequently as you read.

Utility is the benefit or satisfaction an actor gets from something. Utility theory argues that actors select the option that gives them the most utility (denoted U). So when Chris, our mythical "random individual," only has two choices A and B, Chris will do whichever of those has the greater utility. Formally, Chris will pick A if $U^A > U^B$. (Satisfy yourself by doing the algebra that $U^A > U^B$ means the same thing as $[U^A - U^B] > 0$.) What the first part of Hibbs' model does is say that an individual will support the president – pick *i* – if he or she thinks U^i is greater than U^0, where U^0 means "the utility the individual thinks s/he'd get if the other party were in office." If U^i is less than U^0, then the individual will choose to 'not support' the president. [We're now down to equation 5.1.]

Now Hibbs defines what individuals use to determine utility. They use what they can observe about each administration's (*i*) macroeconomic performance, which he refers to as *x*, and they use what they can infer about the probable macroeconomic performance of the out-party (aka the 'shadow' government), referred to as *x*-hat (*x* with a ^ over it). Rather than write out both every time, though, Hibbs lumps these two pieces of information together into a unit he calls *x**. *x** now becomes the key explainer/predictor in an equation. [We're at equation 5.2.] When trying to determine U^i, the individual places some weight B' on that lumped-together-set-of-information *x**. B' here is a coefficient, a slope term, that suggests the change in *y* (the thing to be explained or predicted, here U^i) for every unit change in *x* (here, *x**). It tells us about how much weight is placed on economic evaluations in the decision to support the president. Since individuals aren't perfect and don't know everything, though, and macroeconomic performance isn't the only thing that they use to determine whether to support the president, Hibbs adds an error term that reflects all these other things (the little swirly E is epsilon, and the superscripts, as usual, refer to the incumbent (*i*) or the shadow government (0)).

CRITICAL SIDEBAR: It's important to understand here that what *x** stands for is a comparison between how well the individual thinks current administration is doing on economic matters, and how well the individual thinks the other party would have done in the same circumstances. What this ultimately captures is a preference for one party's macroeconomic management over the other's. (Think about it: No matter how well a Republican president is doing, someone who's a strong Democratic partisan would almost always express an opinion that a Democrat would have done better.) These evaluations of the 'shadow' government can't actually be measured, though, since the other party can't be given a chance to manage the economy at the same exact time, so that gives Hibbs a couple of problems later on when he tries to test his model.

All that happens in equation 5.3 is the parts of 5.2 are substituted into a statement of probability. Hibbs says that the probability of an individual in the population supporting the president – P(Y = 1) – should just be written as "P," and the probability this is true is just the probability that the individual finds $U^i > U^0$. Line 2 has the substitution. In line 3, he factors out the B' and rearranges the terms. In line 4, "diff" means simply 'the difference in performance between the incumbent and the opposition' – it's shorthand for $(x^{*i} - x^{*0})$. Finally, line 5 contains a sneaky bit of statistics. We start out with the term inside the brackets, which is just some weight placed on the difference in performance. But now, epsilon and >0 have disappeared, as has the P (probability) that's been carried down at the front of the equations; instead, we see F in front of the brackets. Epsilon is error. The value it takes for each individual is selected from some distribution. (Think of picking a number 1-5 out of a hat, but not knowing how many of each digit are in there—the 'how many of each digit are in there' is the *distribution* of the probability of picking each different digit.) What the 'F' out front means is that once we have the value for the term inside the brackets, we draw some error value. The actual probability, then, that an individual will support the president is that individual's weighted value of performance, plus or minus some error value drawn from the distribution of epsilon. (And since the probability of all possible outcomes has to add up to 1, the probability of supporting the president [known as P] + probability of *not* supporting the president = 1. So the probability of not-supporting is just 1-P.)

Stop here and think for a moment about the terms that are left at the bottom of 5.3. Does it make sense that those terms are things that matter when an individual chooses to support the president? What does B' stand for, in practical terms? Think about what would happen to P if B' took on a large positive value, a smaller positive value, a small negative value, and a large negative value. (You might find sketching a coordinate axis and some prospective lines with the suggested slopes useful here. Remember that P has to fall between 0 and 1.)

Now, the ugliest part of all: Figure 1. First, the diagonal line represents "B'*x**diff." In plain English, that's the weighted evaluation of differing macroeconomic performance by the parties. You'll notice that this diagonal line crosses a solid horizontal line labeled 0, suggesting Chris can think overall performance by the incumbents was good (evaluated above 0), or bad (evaluated below 0). B' represents a slope coefficient, a weight, on this evaluation of performance: if B' is a high positive number, the line would be steeper than shown, suggesting that macroeconomic performance is very important to Chris when deciding to support the president. A negative slope would mean that the better performance is, the less Chris supports the president; this is not very probable (we know from evidence that most people don't think this way), but in theory it's possible.

Second, the vertical axis represents (net) utility. Remember, this is [U^i - U^0], and that U^i and U^0 are determined by x^{*i} and x^{*0} - just each side's performance evaluation, plus some error. If the expression [U^i - U^0] is positive (greater than 0), then Chris supports the incumbent; if it's negative (less than 0), Chris dis-supports the incumbent. The top solid horizontal line represents this critical tipping point of U=0. When U=0, Chris is indifferent (i.e., has no preference one way or the other) for either the incumbent or the opposition. What determines, then, whether Chris 'supports the president'? [Think about it: what else is in that bottom line of equation 5.3?]

Chris's support for the president is determined by two things: the weighted value of performance, and the error term epsilon. Follow the 0 axis over to where it intersects the diagonal line. This is where Chris's utility for the incumbent is 0, since that's where the function representing Chris's net utility (B'*x**diff) crosses the 0 axis. If Chris is indifferent about performance, then other things like personality, publicity, scandals or exogenous shocks (outside events not caused by the actors in the model, but affecting them anyway) determine whether—right now—Chris supports the president. All those other things (personality, publicity, scandals, shocks) are lumped together here in the error term, epsilon. Epsilon has some distribution of how likely Chris is to draw each of its possible values; Hibbs places a graph of epsilon's distribution on the figure (vertically, it looks like a bump around where the diagonal and 0 axes cross).

Epsilon can be positive or negative, with 0 in its distribution falling on the 0 axis. You'll notice, then, that in Chris's case, epsilon's mode (the highest point on its distribution, the most probable value to draw) is centered right on 0. If 'nature' gave Chris a positive value of epsilon, that's enough to tip Chris's choice this time: Chris will support the president. If instead, Chris drew a negative value for epsilon, Chris will dis-support the president. Which will Chris pull? We don't know; for this distribution, positive and negative values are equally likely.

Look now at the top-right end of the diagonal line, at the point labeled B'*x**diff3. We'll say that this point represents another mythical 'random citizen' named Pat. Ask yourself:

1. Does Pat have a positive or negative net utility for the incumbent? (HINT: What axis should you consult?) Does Pat have a positive or negative assessment of performance?
2. Where is the distribution of epsilon centered, in relation to the point of interest (where the dashed lines meet)?
3. Is Pat more likely to draw a positive value of epsilon, or a negative value? (Remember, probability is determined by the area under the curve; think about where 0 lies in *this* distribution of epsilon.)
4. Finally, is Pat likely to support the president, or dis-support him/her?

Repeat questions 1-4 for B'*x**diff1, which we'll assign to Sam, a third mythical 'random citizen.' Which of our three random citizens, Chris, Pat, or Sam, is *most* likely to support the president under any circumstances? Which is least likely, under any circumstances? Which might support the president, depending on circumstances?

Now, imagine that you are the president, and you control (or can influence) some of the things that Hibbs has included in epsilon: publicity, personality, scandals, shocks. This can help you manipulate support levels. Which of our random citizens might be susceptible to supporting you, if the circumstances were right? What are some things you can do to make this citizen more likely to support you?

Congratulations, the worst is over. You've gotten the basic bits of it now. We're down to the middle of p. 147, just above equation 5.4. In 5.4 itself, all Hibbs does is provide a distribution for epsilon (he selects one that he thinks has a plausible shape and useful characteristics). The key idea you should take away there is that, overall, the more one of our mythical random citizen perceives a positive difference between the performance of the incumbent and of the opposition, the more likely that citizen is to support the president. As the weighted gap (which is B'*x**diff) continues to grow positively (that is to say, as we move from Chris towards Pat on the figure), the more likely the citizen is to support—so likely that by the time we reach someone with a gap the size of Pat's that we are virtually certain that person will support the president. The reverse logic is also true: the larger the *negative* gap, the negative weighted performance assessment, the more certain we become that the individual will *not* support the president. (This corresponds to moving from Chris's position to Sam's.)

People who have performance assessments as extreme as Sam's and Pat's are virtually certain not to hold an opinion besides what's predicted. Sam is very likely to dis-support the president (i.e., P is very small, nearly 0), and Pat is very likely to support the president (P is very large, nearly 1). Picture another individual just to the right of Sam, and another just to the left of Pat. Are either of these new individuals much more or less likely than their respective neighbor to support the president? No. The probability of support changes slowly out there, where individuals have strong opinions. People with assessments like Chris's, though, could change their positions quite easily by drawing a different value of epsilon in another round; right now they have about a 50-50 chance of going either way. We express this idea of changing probabilities of support across increasing values of performance assessment by referring to the slope of the function. With a straight line, slope is easy to calculate. Think, though, about what shape this function must take. Try to sketch it out: plot P on the y axis [0-1], and performance assessment on the x. Then add Sam, Pat, their two 'neighbors,' and Chris, and connect the dots. For functions other than a straight line, we use a single step of calculus to take the derivative of the function. (Derivatives allow the slope to change over different parts of the function; they tell us the slope *at* any given point, rather than *between* two given points as with the slope of a line.)

Your graph should look a lot like Figure 2. If not, go back and work through it again before proceeding.

Middle of 149, now. Hibbs is about to confront that problem I mentioned earlier, where he ideally wants data from the same individuals in the population over time, but instead has to settle for generalizing from random sample surveys. This means he has to use his sample – which might not be perfectly random, and so might be biased one way or the other – to estimate the true population-wide level of support. He *wants* to talk about P, but this problem with the samples means he has to talk about P' instead (say "p-prime"). The large 'dented E' symbol with n=1 on the bottom and N_j on the top is a capital sigma, meaning "add up the values of (support for the president in each political group divided by the number of people in each group in this sample)' for all the groups and individuals in this sample." (*j* refers to groups, *i* refers to individuals)

Equation 5.6 gives us the last nasty bit of math. Remember that a statistical distribution says, "How likely are we to observe a given value?" What we have in survey data is the opposite: we know an estimated probability—the proportion, P', but we don't know what the true value of performance assessment is that we'd observe at that probability because there's all that error in the way. The distribution of error is what gives a statistical distribution its shape. Taking the inverse of the distribution (F) lets us go from the proportion to the value, since F is telling us about the distribution of epsilon across all the possible proportions we could have observed. (Just trust me on this one. I'd recommend that if you don't have advanced statistics training that you just skip to the next section at the top of page 150.) [Putting it a slightly different way, what we normally ask with a statistical distribution is, "Given that this is the Murglewump distribution, and we've drawn a certain value from the hat, how likely is it that we drew this value? What's the probability of pulling an *x* from a Murglewump distribution?" What the inverse says is, "Given that this is a Murglewump distribution, and I know I've gotten 26% of the same response, what's that response likely to be? What's the thing that makes up 26% of a Murglewump distribution?"]

At this point, you have enough of the basic model that you should be able to read through the rest of the bells and whistles that Hibbs adds. 'Dynamic' models allow change over time; the event in question (deciding to support the president) occurs repeatedly, but events may occur between evaluations which cause individuals to change their responses. *w* in 5.8 and 5.9 is a weight, some multiplicative factor. The subscript t means 'in any given time period.' 5.10 says that when individuals assess their preferred party, they discount (proportionally reduce) the weight they put on performance when this party is out of government (when it is the 'shadow' party and citizens assess its performance with *x*-hat) compared to when it is in the government and citizens actually observe its performance (producing assessments of *x*).

The other thing that gives Hibbs some grief here is that the 'incumbent' in question changes over time in two specific ways: in personal (individual) identity (Reagan, then Bush, then Clinton…) and in partisan identity (two Republicans, a Democrat, and a Republican). Both of these factors would be expected to influence levels of reported support, so Hibbs has to control for them in his statistical model. They're also of substantive interest in and of themselves, and to see the difference between different parties' and different individuals' levels of support. He's particularly interested in seeing what happens when there's a change of administrations, and when that change of administrations includes a change of party.

Don't stress over the remaining equations. We'll walk through some of them in class. You should at least look at them, though, and try to determine what the variables refer to and what the algebraic relations mean in normal English. In the end, we're not going to be testing you on your ability to understand algebra. We want you to have the concepts and arguments, the theories about political behavior, that are beneath the algebra. Focus on those: they're normally in the text rather than in the algebra. In short, the model is ***not*** worth losing sleep.

As always, feel free to visit my office hours if you're finding this too frustrating, or if it's making no sense at all. I'm happy to explain as much as I can.