Habitual Voting and Behavioral Turnout

James H. Fowler University of California, Davis

Bendor, Diermeier, and Ting (2003) develop a behavioral alternative to rational choice models of turnout. However, the assumption they make about the way individuals adjust their probability of voting biases their model towards their main result of significant turnout in large populations. Moreover, the assumption causes individuals to engage in casual voting (sometimes people vote and sometimes they abstain). This result is at odds with a substantial literature that indicates most people engage in habitual voting (they either always vote or always abstain). I develop an alternative model to show how feedback in the probability adjustment mechanism affects the behavioral model. The version of this model without feedback yields both high turnout and habitual voting.

Fiorina (1990) has called voter turnout "the paradox that ate rational choice theory." Standard assumptions about rationality typically yield models with vanishing turnout in large electorates (Myerson 1998; Palfrey and Rosenthal 1985). This is because a single vote becomes less and less likely to have an impact on the election as the size of the population increases. If the cost of voting is significant (e.g., the cost of learning about the candidates, going to the polls, and so on), then it is likely to dominate any benefits derived from the infinitesimal probability of affecting the outcome. Unless we assume collateral benefits like the rewarding feeling of doing one's civic duty, rational choice models yield predictions that are at odds with the reality that millions of people vote in large elections.

The paradox of voting has recently caused formal theorists to move away from a rational model of choice towards a *behavioral* model of choice. In particular, Bendor, Diermeier, and Ting (2003; hereafter BDT) explore the possibility that reinforcement learning can explain voter turnout. Their behavioral model of turnout discards any notion that individuals are *prospective optimizers*. Instead, individuals are *adaptive satisficers*. Each person's well-being is affected by the choice to vote or abstain and the outcome of the election. If a person achieves a satisfactory level of well-being then the turnout choice is reinforced and becomes more likely in the next election. If not, the choice is inhibited and becomes less likely in the next election.

The behavioral model is innovative and promising. Unlike the rational model, the BDT model generates significant turnout in large electorates, even when the cost of voting is relatively high. However, this article shows that significant turnout is due in part to *moderating feedback* in the BDT model. As an individual's probability of voting decreases from .5, moderating feedback decreases the strength of downward adjustment and increases the strength of upward adjustment in the probability of voting. This method of updating the individual probability of voting after each election inherently biases the model towards its main result of high turnout in large populations.

This article also shows that moderating feedback causes unrealistic turnout behavior. Most individuals in the BDT model are *casual voters*. In other words, sometimes people make it to the polls and sometimes they do not, but hardly anyone in the model makes it a habit always to vote or always to stay home. This result is at odds with a substantial literature that indicates most people are *habitual voters*—they either always vote or always abstain (Gerber, Green, and Shachar 2003; Green and Shachar 2000; Miller and Shanks 1996; Plutzer 2002; Verba and Nie 1972).

I offer an alternative model that posits a different method of reinforcement and inhibition for adjusting turnout probabilities. I show that the alternative model usually does not yield moderating feedback and then compare its behavior to the behavior of the BDT model. Both models generate significant turnout in large populations, but the model without feedback yields much more plausible levels of habitual voting among individuals. Thus, the model without feedback appears to correspond better to empirical data at both the individual and aggregate levels.

The Journal of Politics, Vol. 68, No. 2, May 2006, pp. 335–344

 ISSN 0022-3816

The BDT Behavioral Model of Turnout

Bendor, Diermeier, and Ting (2003) lay out general conditions for a behavioral model of turnout, but I will focus on the particular computational model from which they derive most of their results. The model they use can be briefly summarized as follows.

As in Palfrey and Rosenthal (1985), a finite electorate of size N is composed of $n_D > 0$ Democrats and $n_R > 0$ Republicans such that $n_D + n_R = N$. In each time period t an election is held in which each citizen i chooses whether to vote (V) or abstain (A). If a citizen chooses to vote, she votes for her own party. Thus the winner of the election is the party with the most turnout (with ties decided by a fair coin toss). All members of the winning party receive a fixed payoff b (regardless of whether or not they voted), and all citizens who choose to vote pay a fixed cost c. Winning abstainers get b, winning voters get b–c, losing abstainers get 0, and losing voters get $-c$. To incorporate uncertainty, a random shock $\theta_{i,t}$ is added to each payoff. This shock is i.i.d. across citizens and time periods and is drawn from a mean 0 uniform distribution with support ω.

Each citizen i in each period t has a *propensity* that defines the probability she will vote $p_{i,t}(V) \in [0,1]$. The probability of abstention is simply $p_{i,t}(A) = 1 - p_{i,t}(V)$. For simplicity of presentation, the propensity to vote will be denoted by a propensity without an associated action: $p_{i,t} = p_{i,t}(V)$. Each citizen also has an *aspiration* level $a_{i,t}$ that specifies the payoff she hopes to achieve. Depending on the propensity, each citizen realizes an action $I \in \{V,A\}$. This determines the election winner and the resulting payoff $\pi_{i,t}$ for each citizen.

Following Bush and Mosteller (1955), propensities are then adjusted according to whether or not the outcome is deemed successful (i.e., whether the resulting payoffs exceeded or equaled aspirations for each citizen $\pi_{i,t} \leq a_{i,t}$). A successful outcome reinforces an action, making it more likely in the next period:

$$p_{i,t+1}(I) = p_{i,t}(I) + \alpha(1 - p_{i,t}(I)) \qquad (1)$$

An unsuccessful outcome ($\pi_{i,t} < a_{i,t}$) inhibits the action by making it less likely in the next period:

$$p_{i,t+1}(I) = p_{i,t}(I) - \alpha p_{i,t}(I) \qquad (2)$$

The parameter $\alpha \in (0,1]$ determines how quickly propensities change in response to reinforcement and inhibition.[1] This parameter has no effect on the limiting distribution that determines how much turnout the model produces, but it does affect how quickly this distribution is reached. Thus we can think of α as representing the speed of learning in the model.

Aspirations are also adjusted in each time period. As citizens experience higher payoffs they become more accustomed to them and raise their aspirations. Similarly, lower payoffs cause citizens to lower their aspirations. Following Cyert and March (1963), each citizen's aspiration is assumed to be a weighted average of the previous aspiration and payoff:

$$a_{i,t+1} = \lambda a_{i,t} + (1-\lambda)\pi_{i,t}$$

where $\lambda \in (0,1)$.

Finally, there are two technical details that need to be mentioned. First, some individuals are *inertial*, meaning they do not update their propensities or their aspirations in a given period t. The probability of not updating propensities is denoted ε_p and the probability of not updating aspirations is denoted ε_a. Second, because BDT assume a finite state space, all propensities and aspirations are rounded to three digits (reinforcement is rounded up and inhibition is rounded down).

The dynamic aspect of the model makes it quite complicated and difficult to solve in closed form. Therefore, BDT use simulation to study the behavior of the model. For comparison, they refer to a set of base model assumptions as follows: a population of 10,000, equally divided between Democrats and Republicans (n_D = 5,000, n_R = 5,000), costs equal to one fourth the benefits of winning ($b = 1$, $c = .25$), a moderate pace of learning ($\alpha = .1$), a moderate pace of aspiration adjustment (λ = .95), a moderate amount of noise in the payoffs ($\omega = .2$), a low proportion of nonresponsive "inertial" individuals ($\varepsilon_p = \varepsilon_a = .01$), and moderate initial turnout propenstites and aspirations ($p_{i,t=0} = a_{i,t=0} = .5$) for all i. To maintain comparability, the simulations in this article use these assumptions unless otherwise noted.

Moderating Feedback in the BDT Model

Bendor, Diermeier, and Ting (2003) do not explain why they choose the Bush-Mosteller reinforcement rule shown in (1) and (2) instead of some alternative. Although this set of rules has been used successfully to explain *aggregate* behavior in pigeons, goldfish, and in some situations in humans, it was abandoned by psychologists in the 1970s in part for its inability to predict individual-level behavior (Camerer 2003;

[1]BDT denote a separate parameter for inhibition, β, but all their results assume that the reinforcement and inhibition rules are symmetric, $\alpha = \beta$.

FIGURE 1 Moderating Feedback in the BDT Model of Turnout

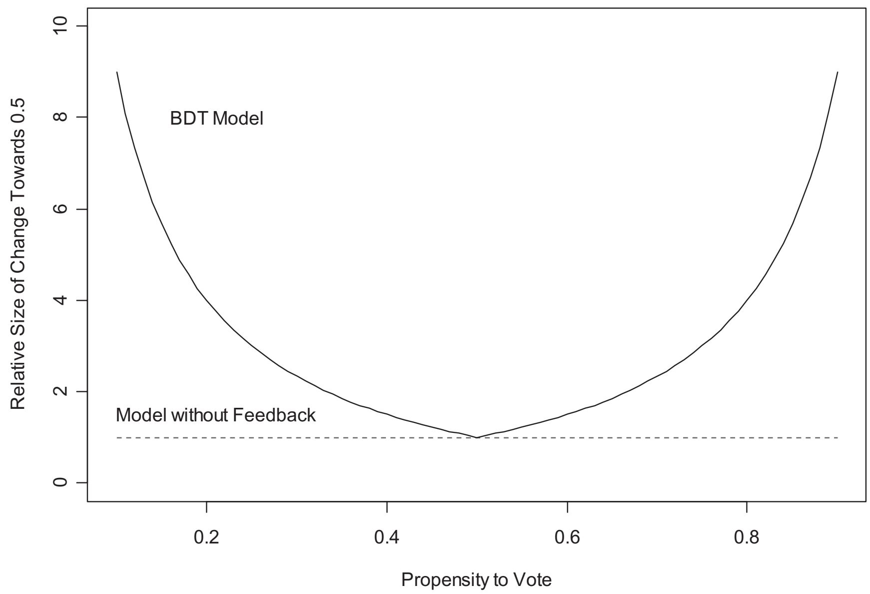

Diaconis and Lehmann 1987). Nonetheless, one might argue that this rule is simple with few parameters. One might also argue that it is elegant in the sense that it eliminates the need to introduce a mechanism to ensure that propensities remain between 0 and 1. However, this elegance has a substantial cost—it biases the model towards their main result.

BDT are primarily interested in whether or not their model produces significant turnout in large populations when the cost of voting is high relative to the benefit of winning the election. Several of their results indicate that the model yields turnout at or near 50% when the cost of voting is as high as .25 and the benefit of winning is 1. This is much higher than predicted by a variety of formal models, most notably Palfrey and Rosenthal (1985).

Yet close examination of the BDT computational model reveals why it consistently produces turnout near 50%. Notice that reinforcement in equation (1) takes place most quickly when propensities are *low*. When the previous propensity is 0, reinforcement causes the new propensity to *increase* by α. However, for propensities near 1, the effect of reinforcement diminishes to 0. Conversely, inhibition in equation (2) takes place most quickly when propensities are *high*. When the previous propensity is 1, inhibition causes it to *decrease* by α. But for propensities near 0, the effect of inhibition diminishes to 0. BDT refer to this property of the reinforcement and inhibition rules as *monotonicity*. In fact, weak monotonicity is a requirement for most of the analytical results in the BDT general model.

However, monotonicity has a very important effect on the behavior of the model. It means that reinforcement is stronger than inhibition for propensities below .5, and inhibition is stronger than reinforcement for propensities above .5. Consequently, the strongest vector of change is always towards propensities of .5. I call this *moderating feedback* and define it as follows:

> *Definition. Moderating feedback occurs when the magnitude of the change due to reinforcement is greater than the magnitude of the change due to inhibition for propensities less than .5 and the magnitude of the change due to inhibition is greater than the magnitude of the change due to reinforcement for propensities greater than .5.*

Notice also that the strength of the feedback is increasing as propensities move away from .5. The solid line in Figure 1 shows the ratio of change towards .5 versus change towards 0 or 1 in the BDT computational model. For comparison, the dotted line shows what this ratio would be in a model without feedback. In the BDT model very high and very low propensities are subject to the strongest adjustment towards .5. For example, suppose $\alpha = .1$ and the previous propensity to vote is $p_{i,t} = .1$. If the propensity is reinforced, then the new propensity will increase by .09. However, if it is inhibited then the new propensity will decrease by a mere .01. This means that in order for .1 to be a stable probability of turnout, *every* reinforcement must be matched by *nine* inhibitions.[2]

[2]Note that the reasoning is symmetric whether we are thinking of the propensity to turn out or the propensity to abstain. It would be difficult to sustain either very high or very low turnout in a model with moderating feedback.

While it is not impossible for the BDT model to produce such a sequence of reinforcements and inhibitions, it is unlikely because of adaptive aspirations. A successful action yields not only an increase in the propensity but an increase in the aspiration level. The higher aspiration makes it *less* likely that the next action will be successful. Similarly, unsuccessful actions yield lower aspiration levels and make it *more* likely that the next action will be a success. Thus negative reinforcement in the aspiration level tends to equalize the number of successes and failures.

In turn, an equal number of successes and failures drives up the propensity to vote. Consider the above example where $\alpha = .1$ and the previous propensity to vote is $p_{i,t} = .1$. If the probability of success is $\Pr(\pi_{i,t} \geq a_{i,t}) = .5$, then there is a 50% chance that the propensity will be reinforced and go up by .09 and a 50% chance it will be inhibited and go down by .01. The expected change in the propensity to turnout will be the previous propensity plus the changes due to reinforcement and inhibition weighted by the probabilities of success and failure:

$$E(p_{i,t+1}) = p_{i,t} + \Pr(\pi_{i,t} \geq a_{i,t})\alpha(1-p_{i,t}) + \Pr(\pi_{i,t} < a_{i,t})(-\alpha p_{i,t})$$

In the previous example the propensity to vote is therefore expected to increase to $E[p_{i,t+1}] = .14$.

This process continues driving up the expected propensity until it reaches a point where it equals the previous propensity $E[p_{i,t+1}] = p_{i,t}$. Rearranging the above equation, it is easy to see that this occurs if and only if the propensity to vote equals the probability of success: $p_{i,t} = \Pr(\pi_{i,t} \geq a_{i,t})$. Hence, a success rate of 50% tends to drive the turnout propensity towards 50% in the BDT computational model. This reasoning also applies to the BDT general model. As long as the reinforcement and inhibition rules are monotonic, they will also yield expected values that tilt propensities towards .5.

Casual Voting in the BDT Model

Moderating feedback in the propensity adjustment causes nearly everyone in the BDT model to engage in *casual voting*.[3] In other words, sometimes people make it to the polls and sometimes they do not, but hardly anyone in the model makes it a habit always to vote or always to stay home. This is inconsistent with the well-known empirical phenomenon of *habitual voting*. A number of studies have demonstrated that most people either vote all the time or not at all.

[3]For example, in 1,000 simulations using BDT's base model assumptions, 98% of the individual propensities end up closer to .5 than to 0 or 1.

Table 1 Validated Turnout in the 1972–1974–1976 NES Panel Survey

	Voted in '76		**Abstained in '76**	
	Voted in '74	**Abstained in '74**	**Voted in '74**	**Abstained in '74**
Voted in '72	1,169	376	27	158
Abstained in '72	67	188	60	782

Plutzer reviews the turnout literature and argues that "there is a longstanding agreement that voting behavior is habitual" (2002, 42). For example, Milbrath (1965, 31) describes the important role of reinforcement learning in the development of political participation habits. Studies of cohort and age effects on turnout also suggest that voting behavior is persistent over time and related to past behavior (Miller and Shanks 1996; Verba and Nie 1972). In particular, Miller and Shanks note that the U.S. population is mostly made up of "regular voters" and "persistent non-voters" (1996, 17). Notice, for example, the turnout behavior of respondents in the 1972, 1974, 1976 Panel of the National Election Study presented in Table 1. About 70% of the respondents in this sample either voted in all three elections or abstained in all three of them. By comparison, if behavior were not habitual and completely independent across elections, the turnout rates in these three elections would imply that only 25% of the respondents would always vote or always abstain. A chi-square test suggests this difference is statistically significant ($P < 10^{-16}$).

Of course, what appears to be habitual behavior might also be explained by persistence in the underlying variables that influence the behavior. For example, if income has an impact on turnout and individual incomes usually do not change, then what seems like habitual turnout might be persistence in income levels over time. To that end, several studies have controlled for both known and unknown factors that might influence turnout using standard analysis of National Election Studies panel data (Brody and Sniderman 1977), instrumental variables analysis of the same data (Green and Shachar 2000), and evidence from a randomized field experiment (Gerber, Green, and Shachar 2003). These studies all find that past voting behavior is positively and strongly correlated with future voting behavior—people who voted

FIGURE 2 **Distribution of Individual Turnout Frequency in South Bend (1976–1984) vs. Turnout Frequency Predicted by the BDT Behavioral Model of Turnout**

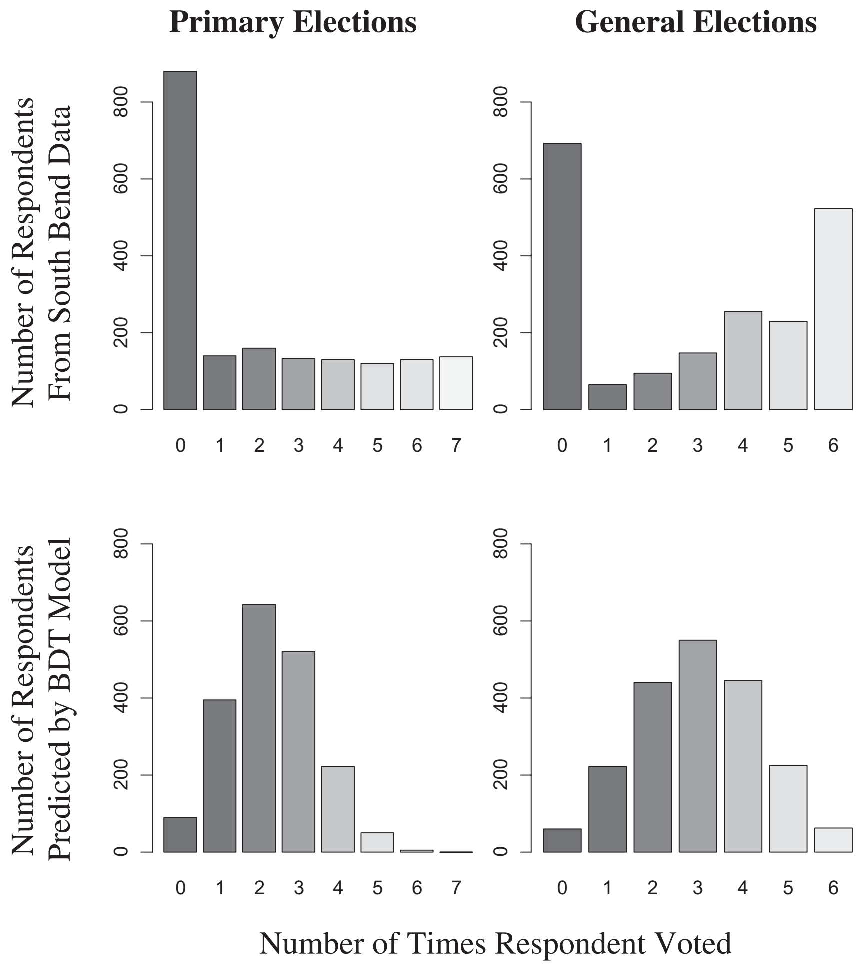

in the previous election turn out at a rate of about 50 percentage points higher than those who do not.

To illustrate more sharply the difference between the BDT model and empirical reality, I draw on data from the South Bend Election Survey (Huckfeldt and Sprague 1985). This survey can help us examine the habitual behavior of the average voter because it includes validated turnout information from a series of six general elections and seven sets of primary elections for residents who lived in South Bend for the years 1976–1984. Figure 2 shows the distribution of turnout frequency—that is, how many individuals never voted, voted once, voted twice, and so on. The upper-left graph shows the frequency of voting in primary elections and the upper-right shows the frequency of voting in general elections. Notice the mode at 0 in both graphs—the plurality of people stay home all the time. Notice also that a substantial group always votes in the general election. Habitual voting and nonvoting dominates casual voting. More than half of the respondents always vote or always abstain.

The lower graphs in Figure 2 show the individual turnout frequency predicted by the BDT computational model. To generate these predictions, I use BDT's base model assumptions and change the cost of voting until mean turnout in the model equals observed turnout (general election turnout is 49% and primary turnout is 27% in the South Bend data).[4] The model is then run for 1000 elections and individual-level data is collected for the last six periods for general elections and seven periods for primaries. The number of individuals sampled is equal to the number

[4]If the model is not adjusted to yield the same aggregate turnout as the empirical data, then differences in the means of the two distributions may yield other differences in those distributions. The question is whether or not the model can simultaneously yield *both* realistic aggregate turnout *and* a realistic distribution of individual turnout behavior when the cost of voting is positive. I want to maintain comparability with BDT's results, so to match aggregate turnout rates between the model and empirical data I change a single parameter, the cost of voting. Note that changing the benefit instead of the cost yields substantively identical results.

sampled in the South Bend study (1,921 for primaries and 1,999 for general elections).

Notice that the modal turnout frequency is 2 for primaries and 3 for general elections. Few individuals in the model habitually abstain and even fewer habitually vote. In fact, habitual behavior is extremely rare. For primaries only 6% of the individuals repeat the same action for each election and for general elections this drops to 4%. Thus the BDT computational model fails to generate a realistic level of habitual behavior. Moreover, *any* version of the BDT model that relies on a monotonic adjustment mechanism is likely to have the same problem—if individual propensities are driven towards .5, then the probability of habitual behavior will continue to remain low.

Can the base model assumptions in the BDT computational model be altered to yield habitual behavior? One possibility is to increase the probability of being inertial (ε_p). After all, inertial behavior is behavior that does not change. However, it is important to remember that this parameter governs change in the *probability* of voting, not change in voting itself. An individual who has a propensity to vote of .5 who becomes inertial will still have a 50–50 chance of voting or abstaining. Thus, unless individual probabilities are already near 0 or 1, even inertial voters will continue to engage in casual voting.

Another possibility is to increase the speed of adjustment (α). When this becomes sufficiently large, it causes individual propensities to bounce back and forth between values near 0 and values near 1. Thus, for a single election the distribution looks right—most voters either have a near 0% or a near 100% chance of going to the polls. However, after the first election many individuals will update their propensities dramatically in the opposite direction. Many abstainers will become voters and many voters will become abstainers. As a result, hardly anyone will behave the same way over a series of elections.

An Alternative Behavioral Model of Turnout

In order to explore the effect of feedback on aggregate turnout and habitual voting, I develop an alternative model. This model keeps all features of the BDT computational model the same except the propensity adjustment rule in equations (1) and (2). In the alternative model, a successful outcome ($\pi_{i,t} \geq a_{i,t}$) reinforces an action with

$$p_{i,t+1}(I) = \min(1, p_{i,t}(I) + \alpha) \tag{3}$$

and an unsuccessful outcome ($\pi_{i,t} < a_{i,t}$) inhibits the action with:

$$p_{i,t+1}(I) = \max(0, p_{i,t}(I) - \alpha) \tag{4}$$

Notice that the α parameter is the same, representing speed of adjustment, but we must now use a min and max condition to ensure that propensities stay within [0,1].

Substantively, this change means that successes and failures cause people to change their voting behavior in the same way, regardless of their prior propensity to vote. For example, a success that causes a habitual abstainer to change from a 10% to a 20% chance of voting would also cause a habitual voter to change from an 80% chance to a 90% chance of voting. Unlike the BDT model, this alternative model yields propensity adjustment *without* feedback as I characterize formally with the following statement:

> *Proposition 1. If the speed of adjustment* (α) *is not too fast then there exists a range of propensities* $p_{i,t}(I) \in [\alpha, 1-\alpha]$ *such that there is no moderating feedback.*
>
> Proof: See web appendix at http://www.journalofpolitics.org.

To see why this is true, let us return to our previous example in which $\alpha = .1$ and the propensity to vote is $p_{i,t} = .1$. When voting satisfies, the propensity to vote will be reinforced and increase by .1. When voting does not satisfy, the propensity to vote will be inhibited and decrease by .1. Thus reinforcement and inhibition are in balance.

It is important to note that the alternative model is not *completely* without feedback. The fact that probabilities are bounded means there must always be some moderating feedback at the boundaries—for example a probability of 0 cannot be adjusted lower but it *can* be adjusted higher. BDT call this a "ceiling effect" and note that it is partially responsible for ensuring that turnout neither falls to 0% nor rises to 100%. In fact, given sufficient alternation in election outcomes, any stationary adjustment process will prevent individual-level voting behavior from remaining fixed at the extremes. In this sense, the alternative adjustment rule is *weakly* monotonic just like the adjustment rules studied by BDT.

However, the alternative model removes (strictly) monotonic feedback for a wide range of propensities when the speed of adjustment is not too high. For example, if we assume as BDT do that $\alpha = .1$, then propensities between .1 and .9 are not subject to feedback in the alternative model. Although individual voters cannot have a fixed 100% or 0% chance of voting, many of them will have very high and very low propensities that cause them to make the same turnout choice for a long series of elections. Thus, the alternative model is more likely to generate habitual behavior than the BDT model.

To see why, recall that in the BDT model when the success rate is fixed, individual propensities will tend to approach the success rate $p_{i,t} = \Pr(\pi_{i,t} \geq a_{i,t})$. A 50% success rate tends to yield a 50% turnout rate. In contrast, individual propensities in the model without feedback between α and $1 - \alpha$ are not subject to such a tendency. In this interval, the expected change in the propensity to vote is $\Pr(\pi_{i,t} \geq a_{i,t})\alpha + \Pr(\pi_{i,t} < a_{i,t})(-\alpha)$, which simplifies to $\alpha(2\Pr(\pi_{i,t} \leq a_{i,t}) - 1)$. Notice that a 50% success rate $\Pr(\pi_{i,t} \geq a_{i,t}) = .5$ implies the expected change is 0, *regardless of the value of the prior propensity*. When reinforcement and inhibition pressures are in balance, *any* individual propensity to turnout between α and $1 - \alpha$ can be stable. However, it is not immediately clear how this will affect the aggregate behavior.

I use simulation to analyze the behavior of the alternative model. Note that the alternative model meets the same criteria as the BDT model for ergodicity since the new propensity adjustment continues to satisfy the condition of being a stationary aspiration-based adjustment rule. This means that voting propensities will converge to a unique limiting distribution from any initial set of propensities and aspirations (Proposition 1 in BDT). Therefore I use the same procedure to analyze the model that they do. Each simulation starts with an initial set of assumptions and then runs for 1,000 periods. Measurements are taken in the 1,000th period and then the simulation is repeated 1,000 times.

BDT rely heavily on simulation to illustrate properties of the Bush-Mosteller adjustment rule, but they also generate a number of analytical propositions that apply broadly to symmetric and weakly monotonic adjustment rules like Bush-Mosteller and the one presented here. Because these propositions cover a large class of adaptive rules, BDT argue that the computational results could not be generated by any particular features of the Bush-Mosteller rule but only by more general properties which are shared by the alternative model. However, closer inspection of their propositions shows why they do not apply to the alternative model. Proposition 2 in BDT establishes bounds on how many individuals will change voting propensities in the next time period, but this information is not sufficient to make aggregate-level predictions about expected turnout—without knowing the functional form of the reinforcement mechanism there is no way to tell *how much* the propensities will change.

Propositions 3 and 5 in BDT suggest that if voters use an aspiration-based adjustment rule like Bush-Mosteller or the alternative rule presented here, the average propensity to vote will increase when *all voters* have propensities less than or equal to .5 and decrease when *all voters* have propensities greater than or equal to .5. However, since any $p_{i,t}(I) \in [\alpha, 1 - \alpha]$ can be stable when using the alternative rule, there is almost always at least one voter with a propensity above .5 and one with a propensity below .5. *In fact, in 100,000 simulations using the alternative adjustment rule with randomly drawn parameters, the conditions for proposition 3 and 5 were not met even once* (see the web appendix at http://www.journalofpolitics.org for program code). The limiting distribution always has some individual propensities above and some below .5. All of the remaining propositions in BDT that address expected change in turnout rely on the same assumption about propensities or explicitly on the Bush-Mosteller adjustment rule, so their general analytical results do not apply to the alternative model.

Results

Aggregate Turnout and the Cost of Voting

The main concern raised by BDT is the ability to generate high-aggregate turnout in large populations when voting is costly. In the behavioral model, a higher cost of voting c reduces turnout because an individual choice to vote is more likely to be unsuccessful and yield an inhibition in the following period. However, moderating feedback puts the brakes on this process. As propensities decrease in the model with feedback, the relative size of the change due to inhibition also decreases. This is not true in the model without feedback—reinforcement and inhibition remain equal for nearly all propensities. Therefore voters in the model without feedback should be more sensitive to costs.

Table 2 compares the results of the two models when we make the base model assumptions and change the cost of voting, c. Notice that lower (but still positive) costs of voting yield aggregate turnout in the model without feedback that is similar to turnout in the BDT model. Thus both models perform well compared to the low turnout predictions of game theoretic models at these cost levels. However, notice also that

Table 2 The Effect of Cost on Aggregate Turnout

	Average Turnout (t = 1,000)			
	Model without Feedback		BDT Model	
C	Democrats	Republicans	Democrats	Republicans
.05	.471	.471	.498	.498
.25	.259	.261	.481	.483
.80	.058	.056	.416	.415

individuals in the BDT model seem to be unusually inured to ever-increasing costs of voting—in fact, about one-third of the voters in the BDT model continue to vote even when $c > b$! In other words, a substantial number of individuals in the BDT model choose to vote even when they think that they are paying more to vote than they could get if they alone chose the winner of the election. By comparison, individuals in the model without feedback more plausibly drop out of the political process when the cost of voting is extremely high.

Feedback and Habitual Voting

Feedback in the behavioral model has an additional effect at the individual level. The tendency to drive propensities towards .5 means that nearly everyone in the BDT model engages in *casual voting*. That is, most individuals vote part of the time and abstain part of the time. As shown in Figure 2 this feature of the BDT behavioral model is inconsistent with the phenomenon of *habitual voting*. Empirically, most people either vote all the time or abstain all the time (Miller and Shanks 1996; Plutzer 2002; Verba and Nie 1972).

When we eliminate feedback from the model, there is a large shift away from casual voting towards habitual voting. Figure 3 compares actual turnout frequencies from the South Bend Election Survey (top) to turnout frequencies predicted by the model without feedback (middle) and the BDT model (bottom). The same procedure used to generate individual turnout frequencies in the BDT model is used again here.

Notice that the model without feedback appears to fit the data better than the BDT model. The modal turnout frequencies match for both primaries (mode at 0) and general elections (modes at 0 and 6). This means there is a tendency in the model without feedback for people always to abstain or always to vote. In particular, notice that the correspondence in the distribution for primaries is relatively close. The correspondence in the distribution for general elections is somewhat weaker since the model without feedback under-predicts the incidence of habitual behavior, but it still does a much better job than the BDT model. About 35% of the individuals in the model without feedback repeat the same action for each general election compared to 4% in the BDT model. Thus, overall the model without feedback appears to conform more closely to the empirical data because it yields substantially more habitual behavior.

Even though this comparison is based on the parameter values used by BDT to justify their model, one might argue that a single set of parameter values is insufficient to test the superiority of the alternative model. Therefore I test the robustness of these findings by generating 100,000 combinations of the parameters b, c, α, ε, λ, and ω each drawn from a [0,1] uniform distribution. For each combination I let the model run for 1000 elections and then collect individual-level data for the last six periods for general elections and seven periods for primaries. I then calculate the mean squared error between the model distribution and the actual distribution observed in South Bend. Out of 100,000 combinations, the best-fitting BDT model predictions generated a mean squared error of 1,084 for primaries and 4,072 for the general elections. By comparison, the best-fitting alternative model predictions yielded errors of 324 for primaries and 2901 for general elections. The lower numbers suggest the alternative model fits the South Bend data better than the BDT model.

Summary and Conclusion

The alternative behavioral model of turnout presented in this article allows us to see how moderating feedback affects voting at both the aggregate and individual level. At the aggregate level, feedback increases the amount of turnout. This means that the method of propensity adjustment chosen by BDT biases turnout towards their main result. However, if we assume that voting is not extremely costly then feedback has less of an effect and both the BDT model and the model without feedback produce high levels of aggregate turnout. At the individual level, feedback causes most individuals to be casual voters. In the BDT model hardly anyone consistently votes or abstains all the time. In contrast, a large number of individuals in the model without feedback are habitual voters. Thus, the model without feedback matches observed data better because it can generate *both* habitual voting *and* high levels of aggregate turnout.

There is a broader lesson in these results. This is obviously not the first effort by scholars to formalize

Figure 3 Distribution of Individual Turnout Frequency in South Bend (1976–1984) vs. Turnout Frequency Predicted by Behavioral Models of Turnout

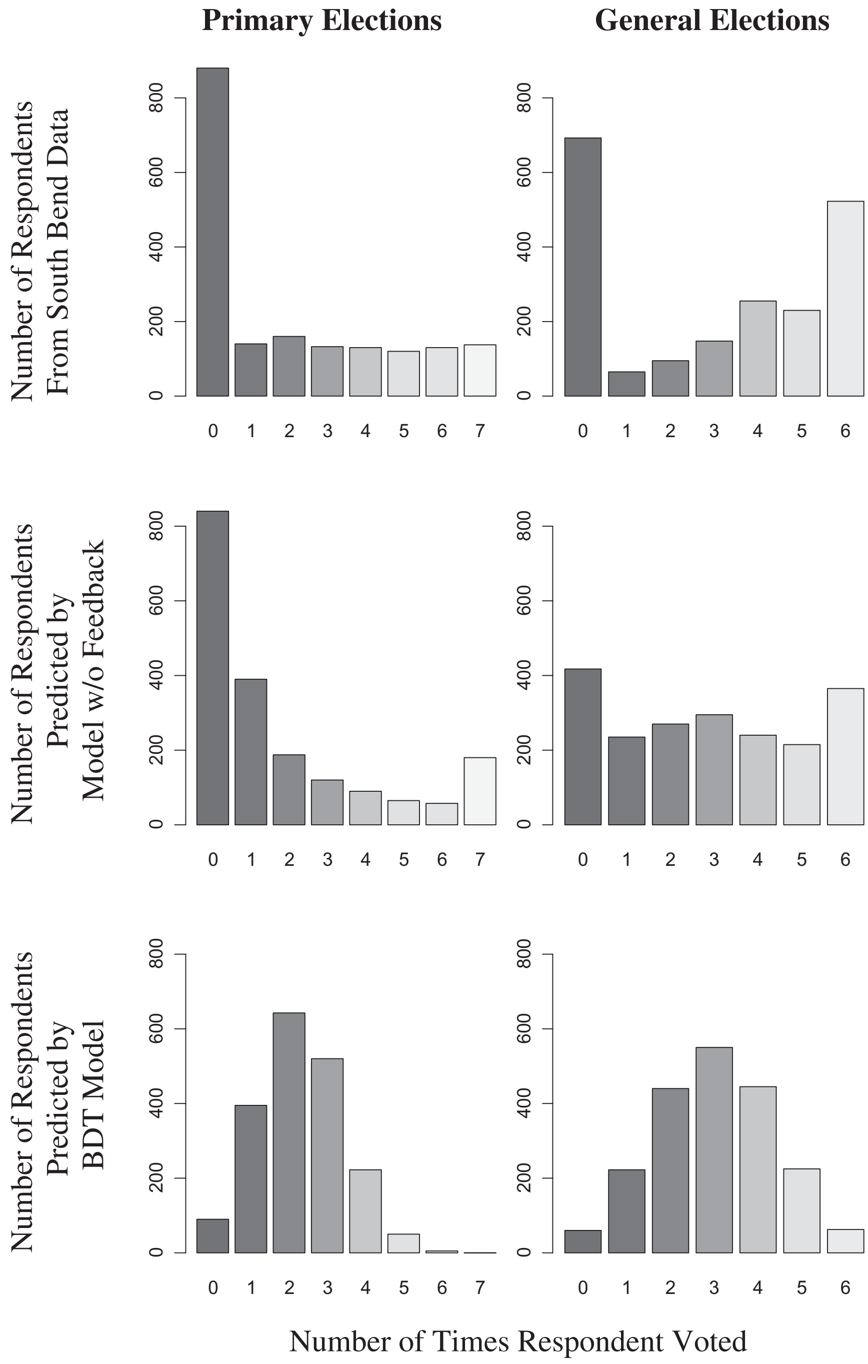

behavioral assumptions. In the 1950s and 1960s psychologists intensively studied stochastic learning rules like the one proposed by Bush and Mosteller (1955). However, much of this work was abandoned in the early 1970s in part because it became clear that these learning rules could not explain the sequential behavior of individual subjects (Camerer 2003; Diaconis and Lehmann 1987). It is precisely this weakness that affects the BDT computational model of turnout. Although it successfully predicts widespread turnout, it fails to account for the individual tendency to behave habitually. Thus, when we incorporate alternative behavioral assumptions into formal theories, it is very important that we analyze not only what happens at the population level but also what happens at the individual level. Otherwise we risk dooming our renewed interest in "formal behavioralism" at its outset.

Acknowledgments

I would like to thank Jonathan Bendor, Eric Dickson, Dan Diermeier, Ben Highton, Bob Jackman, Oleg Smirnov, and Michael Ting for helpful comments. A copy of the most recent version and supporting programs can be found at http://jhfowler.ucdavis.edu.

Manuscript submitted 25 October 2004
Manuscript accepted for publication 22 April 2005

References

Bendor, J., D. Diermeier, and M. Ting. 2003. "A Behavioral Model of Turnout." *American Political Science Review* 97 (2): 261–80.

Brody, R. A., and P. M. Sniderman. 1977. "Life Space to Polling Place—Relevance of Personal Concerns for Voting-Behavior." *British Journal of Political Science* 7 (3): 337–60.

Bush, Robert R., and Frederick Mosteller. 1955. *Stochastic Models for Learning*. New York: John Wiley & Sons.

Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press.

Cyert, Richard Michael, and James G. March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs: Prentice-Hall.

Diaconis, Persi, and Erich Lehmann. 1987. "Fred Mosteller as a Mathematical Statistician." In *A Statistical Model: Frederick Mosteller's Contributions to Statistics, Science, and Public Policy*, eds. S. Fienberg, D. Hoaglin, W. Kruskal and J. Tanur. New York: Springer-Verlag, pp. 59–80.

Fiorina, Morris. 1990. "Information and Rationality in Elections." In *Information and Democratic Processes*, eds. J. Ferejohn and J. Kuklinski. Urbana: University of Illinois Press, pp. 329–342.

Gerber, A. S., D. P. Green, and R. Shachar. 2003. "Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment." *American Journal of Political Science* 47 (3): 540–50.

Green, D. P., and R. Shachar. 2000. "Habit Formation and Political Behaviour: Evidence of Consuetude in Voter Turnout." *British Journal of Political Science* 30 (4): 561–73.

Huckfeldt, Robert, and John Sprague. 1985. *Presidential Election Campaign Study, 1984: South Bend, Indiana*. Indiana University, Center for Survey Research.

Milbrath, Lester W. 1965. *Political Participation: How and Why Do People Get Involved in Politics?* Chicago: Rand McNally.

Miller, Warren E., and J. Merrill Shanks. 1996. *The New American Voter*. Cambridge: Harvard University Press.

Myerson, R. B. 1998. "Population Uncertainty and Poisson Games." *International Journal of Game Theory* 27 (3): 375–92.

Palfrey, T. R., and H. Rosenthal. 1985. "Voter Participation and Strategic Uncertainty." *American Political Science Review* 79 (1): 62–78.

Plutzer, E. 2002. "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood." *American Political Science Review* 96 (1): 41–56.

Verba, Sidney, and Norman H. Nie. 1972. *Participation in America: Political Democracy and Social Equality*. New York: Harper & Row.

James H. Fowler is assistant professor of political science, University of California-Davis, Davis, CA 95616.