Habitual Voting and Behavioral Turnout *

James H. Fowler[†]
University of California, Davis

February 3, 2004

Abstract

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.

*I would like to thank Jonathan Bendor, Eric Dickson, Ben Highton, Bob Jackman, and Oleg Smirnov for helpful comments. A copy of the most recent version and supporting programs can be found at http://jhfowler.ucdavis.edu.
[†] Department of Political Science, University of California, Davis, One Shields Avenue, Davis, CA 95616. (530) 752-1649. E-mail: jhfowler@ucdavis.edu

Morris 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 significant. 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 0.5, 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 (Miller and Shanks 1996; Plutzer 2002; Verba and Nie 1972).

To study the effect of moderating feedback I describe a model with an alternative method of reinforcement and inhibition for adjusting turnout probabilities. This method allows me to control the degree of feedback in the model. I then use the model to compare the behavior of the BDT model to the behavior of a model without feedback. Both models generate significant turnout in large populations, but the results indicate that feedback in the behavioral model has a significant effect on the aggregate turnout rate. Moreover, the model without feedback yields much more realistic 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 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 $\pi_{i,t} \geq a_{i,t}$). A successful outcome reinforces an action, making it more likely in the next period:

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

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

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

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:

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

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

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: n_D =5,000, n_R =5,000, b=1, c=0.25, α=0.1, λ=0.95, ω=0.2, $\varepsilon_p = \varepsilon_a$=.01, and $p_{i,t=0} = a_{i,t=0}$ =0.5 for all i. To maintain comparability, the simulations in this article always 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. This is especially puzzling since Bush-Mosteller learning rules were largely abandoned by psychologists in the 1970s for their inability to predict individual level behavior (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 0.25 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 *montonicity*. 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 0.5, and inhibition is stronger than reinforcement for propensities above 0.5. Consequently, the strongest vector of change is always towards propensities of 0.5. 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 0.5 and the magnitude of the change due to inhibition is greater than the magnitude of the change due to reinforcement for propensities greater than 0.5.*

Notice also that the strength of the feedback is increasing as propensities move away from 0.5. The solid line in Figure 1 shows the ratio of change towards 0.5 vs. 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 0.5. For example, suppose $\alpha = 0.1$ and the previous propensity to vote is $p_{i,t} = 0.1$. If the propensity is reinforced, then the new propensity will increase by 0.09. However, if it is inhibited then the new propensity will decrease by a mere 0.01. This means that in order for 0.1 to be a stable

Figure 1. Moderating Feedback in the BDT Model of Turnout

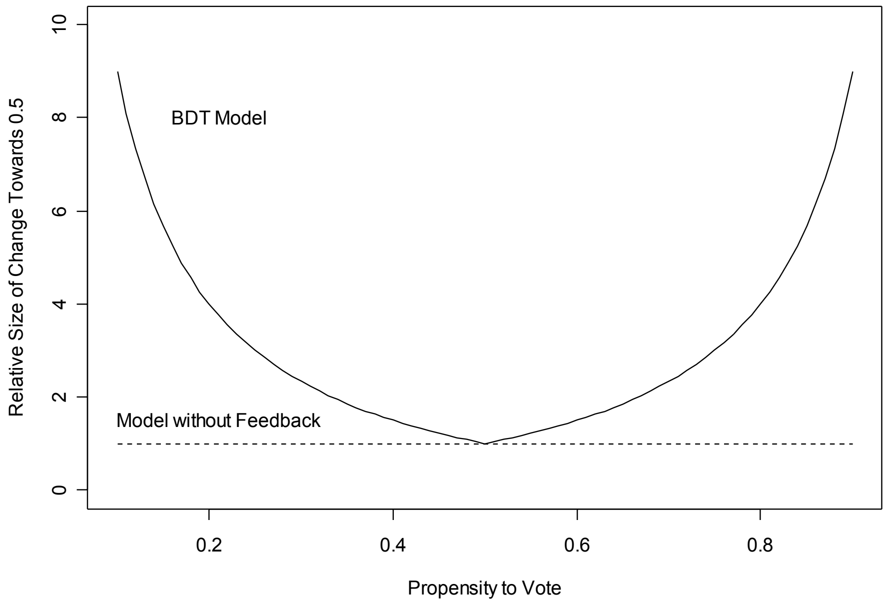

probability of turnout, *every* reinforcement must be matched by *nine* inhibitions.[2]

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 $\alpha = 0.1$ and the previous propensity to vote is $p_{i,t} = 0.1$. If the probability of success is $\Pr(\pi_{i,t} \geq a_{i,t}) = 0.5$, then there is a 50% chance that the propensity will be reinforced and go

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

up by 0.09 and a 50% chance it will be inhibited and go down by 0.01. 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 $E[p_{i,t+1}] = 0.14$.

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 0.5.

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* (Miller and Shanks 1996; Plutzer 2002; Verba and Nie 1972). A number of studies have demonstrated that most people either vote all the time or not at all.

To illustrate the difference between the model and empirical reality, I draw on data from the South Bend Election Survey (Huckfeldt and Sprague 1985). This survey is useful because it includes validated turnout information from a series of 6 general elections and 7 sets of primary elections for

[3] For example, in 1000 simulations using BDT's base model assumptions, 98% of the individual propensities end up closer to 0.5 than to 0 or 1.

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

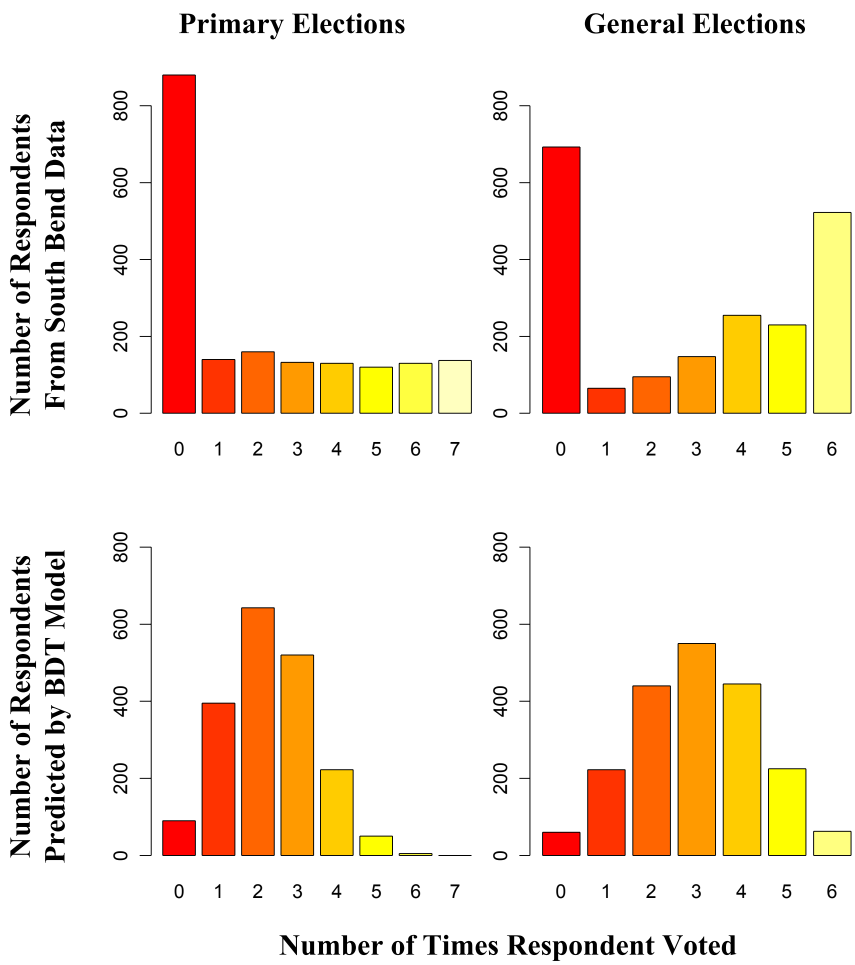

residents who lived in South Bend for the years 1976-1984. Figure 2 shows the distribution of turnout frequency—that is, how many individuals never voted, voted once, voted twice, and so on. The upper left graph shows the frequency of voting in primary elections and the upper right shows the frequency 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 6 periods for general elections and 7 periods for primaries. The number of individuals sampled is equal to the number 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 0.5, 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 0.5 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.

[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.

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

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

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

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

The α parameter is the same, representing speed of adjustment, but there are two new features to consider. First, the min and max conditions ensure that propensities stay within [0,1]. Second, there is an additional parameter $\tau \in [0,1]$ that determines feedback in the propensity adjustment.

Notice that the BDT computational model is a special case of the alternative model when $\tau = 1$. Like the BDT model, the alternative model is monotonic and produces moderating feedback for all positive values of τ. Unlike the BDT model, however, the alternative model can also yield propensity adjustment *without* feedback. The relationship between the parameter $\tau \in [0,1]$ and feedback in the alternative model can be formally characterized with the following statement:

Proposition 1. *If the speed of adjustment* (α) *is not too fast then there exists a range of propensities* $p_{i,t} \in [p^{\min}, p^{\max}]$ *such that for* $\tau > 0$ *there is moderating feedback and for* $\tau = 0$ *there is no feedback.*

Proof: See appendix.

Returning to our previous example, suppose $\alpha = 0.1$ and the propensity to vote is $p_{i,t} = 0.1$. When voting satisfices, the propensity to vote will be reinforced and increase by $0.1 - 0.01\tau$. When voting does not satisfice, the propensity to vote will be inhibited and decrease by $0.1 - 0.09\tau$. Note that reinforcement yields a larger change than inhibition for all $\tau > 0$. This will tend to drive the propensity to vote towards 0.5. Only when $\tau = 0$ are reinforcement and inhibition in balance at 0.1.

Before comparing the BDT model to a version of the alternative model without feedback, I characterize feedback with the following two corollaries to proposition 1:

Corollary 1.1 (BDT computational model). *If* $\tau = 1$, *then all propensities* $p_{i,t}(I) \in [0,1]$ *are subject to moderating feedback.*

Corollary 1.2 (model without feedback). *If* $\tau = 0$, *then propensities in the range* $p_{i,t}(I) \in [\alpha, 1-\alpha]$ *are not subject to moderating feedback.*

Proof: See appendix.

It is important to note that the model without feedback 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%. However, letting $\tau = 0$ removes feedback from the model for a wide range of propensities when the speed of adjustment is not too high. For example, if we assume as BDT do that $\alpha = 0.1$, then propensities between 0.1 and 0.9 are not subject to feedback in the alternative model when $\tau = 0$.

Recall that in the BDT model when the success rate is fixed, individual propensities will tend to approach the success rate $p_{i,t} = \Pr(\pi_{i,t} \geq a_{i,t})$. Thus a 50% success rate tends to yield a 50% turnout rate. In contrast, individual propensities in the model without feedback are not subject to such a tendency. In the alternative model, the expected change in the propensity to vote is:

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

When there is no feedback, this simplifies to $\alpha(2\Pr(\pi_{i,t} \geq a_{i,t}) - 1)$. Notice that a 50% success rate $\Pr(\pi_{i,t} \geq a_{i,t}) = 0.5$ 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 can be stable. However, it is not clear how this will affect the aggregate behavior of the model.

I use simulation to analyze the behavior of the alternative model. Note that this 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 the model will converge to a unique limiting distribution from any initial set of propensities and aspirations (Bendor, et al. 2003). 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.

RESULTS

Population and Turnout

The main concern raised by BDT is the effect of population size on turnout. In rational choice models such as Palfrey and Rosenthal (1985), turnout approaches 0 as the population increases. However, BDT show that this is not necessarily true in a behavioral model. Figure 3 shows the effect of population size on aggregate turnout for both the BDT model and the model without feedback. Notice that turnout in both models declines as the population size increases in accord with empirical findings (Hansen, et al. 1987). The main difference is that the model without feedback yields much lower turnout. This suggests that when reinforcement and inhibition pressures are in balance, reinforcement learning yields a smaller but still significant amount of turnout relative to that predicted by rational choice models.

Figure 3. The Effect of Population Size on Aggregate Turnout

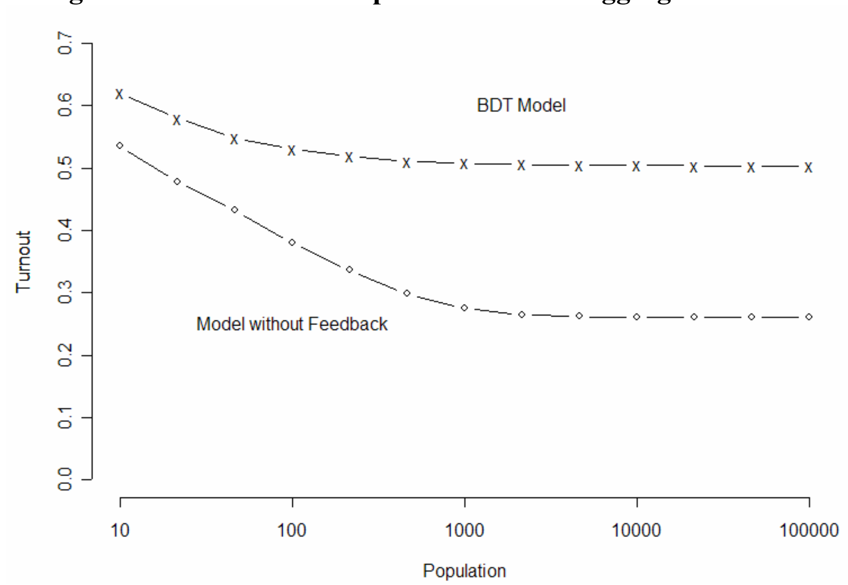

Notice also that turnout tends to approach a lower limit as the population increases. The difference in behavior between a system with 10,000 voters and 1,000,000 voters is negligible. This suggests that simulations based on smaller population sizes can be used to approximate the behavior of larger electorates.

Feedback and Aggregate Turnout

To see how moderating feedback causes the model to generate higher levels of aggregate turnout, we can hold all aspects of the model fixed and gradually increase the parameter τ from 0 to 1. Using the BDT base model assumptions, Figure 4 shows the results of this exercise. Notice that the model without feedback ($\tau = 0$) yields the lowest turnout (about 26%), and turnout increases as the feedback increases. The last point in the graph replicates the BDT finding that their model yields a turnout rate of about 50%. Clearly, the choice of the propensity adjustment mechanism has an important impact on the outcome of the model, and the one chosen by BDT biases the model towards their main result of high aggregate turnout in large electorates.

Figure 4. Effect of Moderating Feedback on Turnout

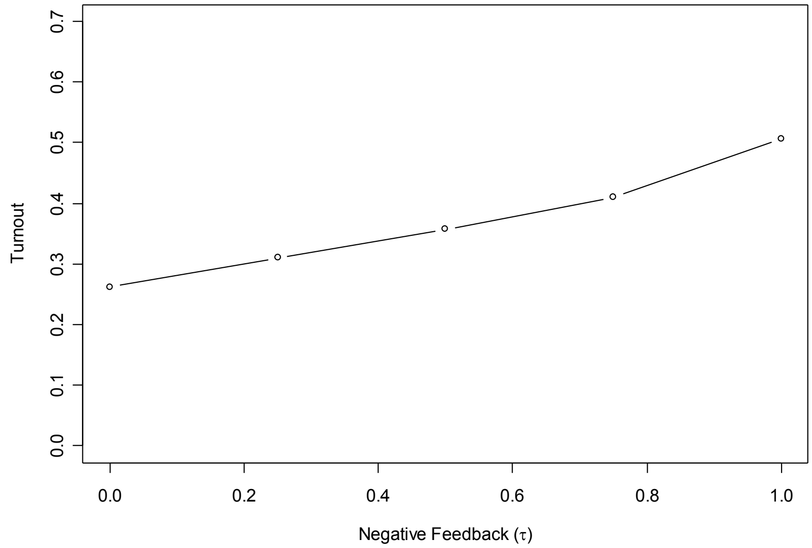

On its face, this might seem to imply that we should reject the model without feedback because it yields low aggregate turnout. However, it is important to remember that these results are based on the arbitrary choice of unknown values like the cost of voting. For the sake of comparison, all assumptions have been kept the same as they were in BDT. However, lower (but still positive) costs of voting yield aggregate turnout in the model without feedback that is similar to turnout in the BDT model. For example, the base model assumptions generate turnout of 49% when the cost of voting is changed to 0.025. By comparison, models based on rational choice would yield near-zero turnout at this level. Thus, like the BDT model, the model without feedback can generate high levels of aggregate turnout.

Feedback and Habitual Voting

Feedback in the behavioral model has an additional effect at the individual level. The tendency to drive propensities towards 0.5 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).

Figure 5. Distribution of Individual Turnout Frequency in South Bend (1976-1984) vs. Turnout Frequency Predicted by Behavioral Models of Turnout

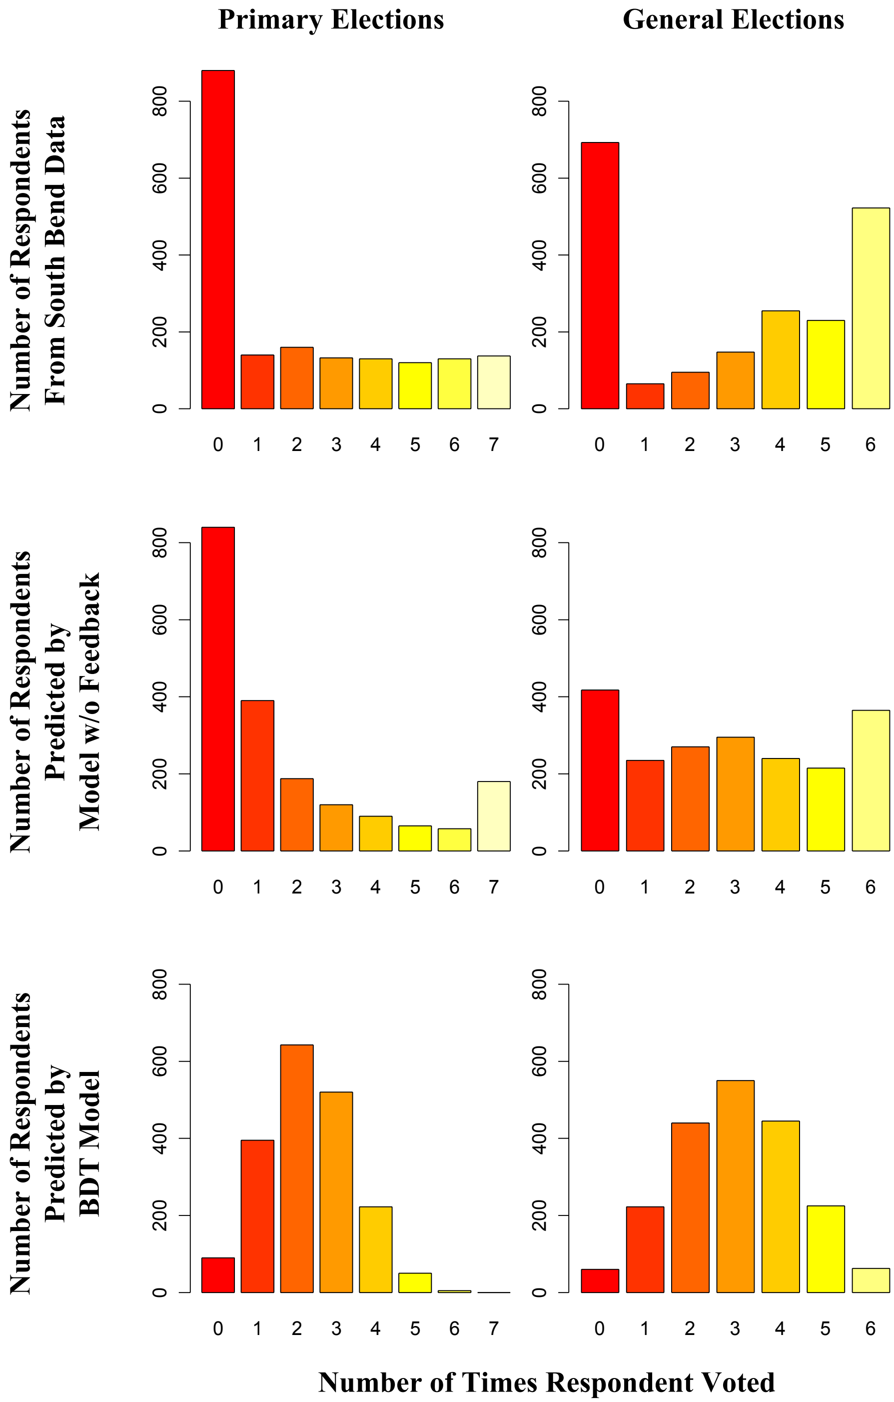

When we eliminate feedback from the model, there is a large shift away from casual voting towards habitual voting. Figure 5 compares actual turnout frequencies from the South Bend Election Survey (top) to turnout frequencies predicted by the model without feedback (middle) and the BDT model (bottom). The same procedure used to generate individual turnout frequencies in the BDT model is used again here.

Notice that the model without feedback fits 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 is more realistic because it yields substantially more habitual behavior.

Summary and Conclusion

Controlling feedback in a behavioral model of turnout allows us to see how it 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 less 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 is more empirically realistic 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 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 when it became clear that these learning rules could not explain the sequential behavior of individual subjects (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.

Appendix

Proof of Proposition 1:

Notice in equation (3) that the min condition applies to propensities above a certain cutpoint to keep them from exceeding 1. It is easy to show that this cutpoint is $p^{\max} = (1-\alpha)/(1-\alpha\tau)$. Thus, change in the propensity due to reinforcement is $\left|p_{i,t+1}(I) - p_{i,t}(I)\right| = \alpha(1-\tau p_{i,t}(I))$ if $p_{i,t}(I) \le p^{\max}$. Notice in equation (4) that the max condition applies to propensities below a certain cutpoint to keep them from falling below 0. It is easy to show that this cutpoint is $p^{\min} = 1-(1-\alpha)/(1-\alpha\tau)$. Thus, change in the propensity due to inhibition is $\left|p_{i,t+1}(I) - p_{i,t}(I)\right| = \alpha(1-\tau(1-p_{i,t}(I)))$ if $p_{i,t}(I) \ge p^{\min}$.

Notice that the interval $p_{i,t}(I) \in [p^{\min}, p^{\max}]$ is empty unless $p^{\max} \ge p^{\min}$ which is only true if the speed of adjustment is not too large: $\alpha \le 1/(2-\tau)$. Compare the magnitude of the change due to reinforcement and the magnitude of the change due to inhibition. If $\alpha(1-\tau p_{i,t}(I)) > \alpha(1-\tau(1-p_{i,t}(I)))$, then $\tau(1-2p_{i,t}(I)) > 0$ which is only true when $\tau > 0$ and $p_{i,t}(I) < 1/2$. By symmetry, the opposite inequality is true when $\tau > 0$ and $p_{i,t}(I) < 1/2$. Thus, for $\tau > 0$ there is moderating feedback—propensities below 0.5 are subject to stronger reinforcement and propensities above 0.5 are subject to stronger inhibition. Notice further that when $\tau = 0$ there is no feedback since the magnitudes of change due to reinforcement and inhibition are equal: $\alpha(1-\tau p_{i,t}(I)) = \alpha(1-\tau(1-p_{i,t}(I))) \Rightarrow \alpha = \alpha$.

Proof of Corollary 1.1:

From proposition 1, there is moderating feedback in the interval $p_{i,t}(I) \in [p^{\min}, p^{\max}]$ when $\tau > 0$. If $\tau = 1$ then $p^{\min} = 1-(1-\alpha)/(1-\alpha\tau) = 0$ and $p^{\max} = (1-\alpha)/(1-\alpha\tau) = 1$. Thus all propensities in the interval $p_{i,t}(I) \in [0,1]$ are subject to moderating feedback.

Proof of Corollary 1.2:

From proposition 1, there is no moderating feedback in the interval $p_{i,t}(I) \in [p^{\min}, p^{\max}]$ if $\tau = 0$. The relation $\tau = 0$ also implies $p^{\min} = 1-(1-\alpha)/(1-\alpha\tau) = \alpha$ and $p^{\max} = (1-\alpha)/(1-\alpha\tau) = 1-\alpha$. Thus propensities between α and $1-\alpha$ are not subject to feedback.

References

Bendor, J., D. Diermeier, and M. Ting. "A Behavioral Model of Turnout." *American Political Science Review* 97, no. 2 (2003): 261-80.

Bush, Robert R., and Frederick Mosteller. *Stochastic Models for Learning*, *Wiley Publications in Statistics*. New York, etc. John Wiley & Sons, Inc., 1955., 1955.

Cyert, Richard Michael, and James G. March. *A Behavioral Theory of the Firm*, *Prentice-Hall International Series in Management*. Englewood Cliffs, N.J.,: Prentice-Hall, 1963.

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

Fiorina, Morris. "Information and Rationality in Elections." In *Information and Democratic Processes*, edited by John Ferejohn and James Kuklinski. Urbana: University of Illinois Press, 1990.

Hansen, S., T. R. Palfrey, and H. Rosenthal. "The Downsian Model of Electoral-Participation - Formal Theory and Empirical-Analysis of the Constituency Size Effect." *Public Choice* 52, no. 1 (1987): 15-33.

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

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

Myerson, R. B. "Population Uncertainty and Poisson Games." *International Journal of Game Theory* 27, no. 3 (1998): 375-92.

Palfrey, T. R., and H. Rosenthal. "Voter Participation and Strategic Uncertainty." *American Political Science Review* 79, no. 1 (1985): 62-78.

Plutzer, E. "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood." *American Political Science Review* 96, no. 1 (2002): 41-56.

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