Journal of Marketing Research
Article Postprint

Cumulative Timed Intent:

A New Predictive Tool for Technology Adoption

Koert van Ittersum
Assistant Professor of Marketing
College of Management
Georgia Institute of Technology
800 West Peachtree Street NW
Atlanta, GA 30308-0520
Phone: 404 385 4884
koert.vanittersum@mgt.gatech.edu.

Fred M. Feinberg*
Hallman Fellow and Professor of Marketing
Stephen M. Ross School of Business
University of Michigan
701 Tappan Street
Ann Arbor, MI 48109-1234
Phone: 734 764 4711
feinf@umich.edu.

* The authors acknowledge the financial support of Deere & Company and thank Jerry Duncan for his support and advice on this research. The authors further thank Wendy Rogers, Leonard J. Parsons, Muge Capar, Kelly E. Caine, Marita O'Brien, Sung Park, and Arthur D. Fisk for their support during the data collection process, Jerry Thursby and Laurie Garrow for their feedback on this research, and Eric Bradlow and Anocha Aribarg for clarifying certain issues in Bayesian

computation. Finally, the authors thank Joel Huber and Tülin Erdem, the Associate Editor, and the three anonymous reviewers for their constructive feedback.

Cumulative Timed Intent:

A New Predictive Tool for Technology Adoption

ABSTRACT

Despite multiple calls for the integration of time into behavioral intent measurement, surprisingly little academic research examines timed intent measures directly. In two empirical studies, the authors estimate individual-level cumulative adoption likelihood curves—curves calibrated on self-reported adoption likelihoods for cumulative time intervals across a fixed horizon—of 478 managerial decision makers, self-predicting whether and when they will adopt a relevant technology. A hierarchical Bayes formulation allows for a heterogeneous account of the individual-level adoption likelihood curves as a function of time and common antecedents of technology adoption. A third study generalizes these results among 354 consumer decision makers and, using behavioral data collected during a two-year longitudinal study involving a subsample of 143 consumer decision makers, provides empirical evidence for the accuracy of cumulative adoption likelihood curves for predicting whether and when a technology is adopted. Cumulative adoption likelihood curves are shown to outperform two single-intent measures as well as two widely-validated intent models in predicting individual-level adoption for a fixed time period of two years. The results suggest great promise for future research on using and optimizing cumulative timed intent measures across a variety of application domains.

Keywords: self-reported intentions, cumulative timed intent, technology adoption, predictive accuracy, hierarchical Bayes

Two of the most critical uncertainties associated with new-technology introductions are whether and when the target market will adopt them. Both uncertainties pose serious challenges for marketing managers planning a technology's production, pricing, distribution, and promotion (Morwitz, Steckel, and Gupta 2007). One way to reduce these uncertainties is to survey the target market about their intentions to adopt the technology (Silk and Urban 1978). Because self-reported adoption intentions can be collected prior to launch at relatively low costs, they are among the most widely applied proxy measures for actual adoption (Sun and Morwitz 2009).

The widespread use of intention measures to predict adoption behavior hinges on the belief that intentions are accurate indicators of individuals' behavior (Young, DeSarbo, and Morwitz 1998). Research in social psychology suggests that intention measures should be among the best predictors of behavior, because they allow each individual to incorporate and appropriately balance all relevant factors that may influence his or her actual behavior (Ajzen and Fishbein 2005). Intention measures have been shown to be a valuable input to predicting purchase behavior and sales forecasts (Armstrong, Morwitz, and Kumar 2000). Furthermore, multiple meta-analyses have shown that intention measures relate to behaviors, both for examinations of specific behaviors (e.g., Albarracin et al. 2001) and more generally across behaviors (e.g., Morwitz, Steckel, and Gupta 2007; Sheppard, Hartwick, and Warshaw 1988).

The predictive accuracy of intention measures, however, often remains modest at best (Abraham et al. 1999). One suggested approach to improve predictive accuracy is to collect timed intent measures—those collected for a pre-specified time frame (Morwitz 1994). However, despite multiple calls for the integration of time into behavioral intent measurement (Morwitz and Schmittlein 1992; Venkatesh, Maruping, and Brown 2006), academic research on timed intent measures remains scarce (Morwitz 1994). To that end, we introduce a novel measure, *cumulative timed intent*, for assessing whether and when people intend to adopt a new technology. We demonstrate the measure using three unique data sets, collected as part of a longitudinal research collaboration with a global Fortune 100 Company operating in the heavy equipment industry, and establish how cumulative timed intent compares to two single-intent measures (Juster 1966; Morwitz 1994) and two classic intent models (Bemmaor 1995; Morrison 1979).

Our focus in model development, and in all three empirical studies, is on estimating *individual-level cumulative adoption likelihood curves*: curves calibrated (via hierarchical Bayes techniques) using self-reported adoption likelihoods for cumulative time intervals across a fixed time horizon. To estimate these curves, we ask decision makers about the expected likelihood that they will have purchased a technology at several points in the future (e.g., within one month, 6 months from now, etc.), conditional on their not having already adopted. This allows for heterogeneous, individual-level adoption likelihood curves as a function of time and common antecedents of adoption, as well as an assessment of their accuracy in predicting “whether and when” people actually adopt a given technology.

SELF-REPORTED INTENTION MEASURES

Intentions appear to frequently provide biased measures of adoption propensity, sometimes underestimating actual adoption and other times overestimating it (Young, DeSarbo and Morwitz 1998). The most obvious problems with these biased intentions are that they may result in inaccurate predictions and may also bias estimates of the relationship between correlates of intentions and adoption (Sun and Morwitz 2009). As one case in point, Alexander, Lynch, and Wang (2008) report that the estimated probability of people following through on their intentions to acquire incrementally new products is a modest .32, and that this probability is barely half so much for ‘really new’ products.

Various sources for these biased intentions have been suggested: intention measures themselves may be biased (Hsiao and Sun 1999); intentions may change over time (Morrison 1979); and the assumed relationship between intentions and behavior may be imperfect (Gollwitzer 1999). A common antecedent of these three sources of bias is the temporal distance between intentions and behavior (Salisbury and Feinberg 2008). As the temporal distance between reporting an intention and the actual behavior increases, it becomes more difficult to anticipate and account for everything that determines final adoption. As a result, decision makers may be more likely to provide biased intentions (Trope and Liberman 2003). Furthermore, their intention may change in time as they are more likely to encounter the unanticipated, which, combined with the increased likelihood of providing biased intentions, may erode the presumably strong relationship between intentions and eventual behavior (Sutton 1998).

One approach to address this inconsistency has been to develop models that account for biases in self-reported intentions, customer heterogeneity, changes in intentions over time, and the stochastic and nonlinear nature of the intention-behavior relationship, among others (Bemmaor 1995; Hsiao, Sun, and Morwitz 2002; Kalwani and Silk 1982; Morrison 1979; Sun and Morwitz 2009). Instead of *modeling* these inconsistencies, other research has tried to *reduce* them by developing response scales that lessen the biases associated with intention measures. Juster (1966) concludes that subjective purchase probabilities (e.g., "what is the likelihood that you will buy a car in the next 12 months?") explain roughly twice the cross-sectional variance in automobile purchase rates as do buying intentions (e.g., "is it your intention to buy a car in the next 12 months?"). Sheppard, Hartwick, and Warshaw (1988) refer to these probability measures as a behavioral expectation—an individual's self-reported subjective probability of his or her performing a specified behavior, based on his or her cognitive appraisal of volitional and non-volitional behavioral determinants—and confirm its increased accuracy for predicting future behavior. Measuring subjective probabilities encourages respondents to consider 1) their attitudes, subjective norms, and intentions toward the action or outcome of interest; 2) alternative actions or outcomes; as well as 3) various factors that could cause them to be unsuccessful in their attempt to carry out such intentions (Davis and Warshaw 1992). Compared to the formation of intention measures, the formation of subjective probabilities is thus more heavily based on respondents' evaluation of the motivational drivers and environmental facilitators/inhibitors of behavior, anticipated changes in these determinants, and how they will influence the probability of behavioral performance. In sum, asking respondents for their subjective probabilities encourages and allows them to better account for future uncertainty than asking them for their intentions alone. As a consequence, subjective probabilities should more accurately predict future behavior than intentions when the behavior becomes distal in time (Venkatesh, Maruping, and Brown 2006).

In sum, there are two general approaches to address biases that limit the predictive accuracy of intention measures: (1) to develop a formal statistical model that accounts for these biases or (2) to focus on developing a response scale that aims to attenuate these biases at the point of collection. While correcting for biases via dedicated statistical models has been shown to improve predictive accuracy of self-reported intentions, reducing them in the first place decreases the need for such complex approaches. Furthermore, actual sales data are often required for calibration purposes, reducing the practical value of many of these models. Since self-reported adoption intentions can be collected *prior* to launch, this research aims to develop and validate a response scale that not only can help provide more accurate predictions of *whether* a future behavior will occur, but also *when* this behavior is most likely to.

CUMULATIVE TIMED INTENT

It has been suggested that the accuracy of intention measures can be further improved by the integration of time in behavioral intent measurement, for example, by measuring timed intent (Morwitz 1994). Timed intent measures allow for the assessment of both whether *and* when a technology is most likely to be adopted, which is among the reasons timed intent measurement is commonly used by practitioners (Morwitz and Schmittlein 1992). Academic research on timed intent measures remains scarce. One notable exception is Morwitz (1994), who measures timed intent by asking respondents: "When do you intend to purchase the product?" with response options "within 6 months," "7-12 months", or "not within a year." She reports that these measures accurately predict the proportion of purchasers for short time horizons, but provide biased estimates of longer term purchasing.[1]

Following Morwitz, we propose to measure timed intent by presenting respondents with multiple time intervals for a specified time horizon. However, our approach differs in several ways. First, we systematize the time horizon by making it technology-specific. The cumulative time intervals are selected based on a trade-off between granularity (desirable) and burden on the respondents (undesirable). Second, we use a probability response scale that circumvents respondents having to translate their internal probability estimates to a multinomial scale and assess which time interval seems most likely. Third, by measuring intent for each respective time interval, we attempt to recover the underlying timed probability adoption curve respondents rely upon when using more traditional intention measures (yes/no). As such, prior methods of intention elicitation represent *reduced* forms of what we propose, in the sense that the information collected

[1] Using a similar timed-intent measure, Sun and Morwitz (2009) develop a unified model of the relationship between intentions and purchasing that 1) takes into account possible sources of discrepancies between intentions and purchasing; 2) forecasts purchasing probability at the individual level by explanatory variables and intentions with actual purchasing; 3) considers multiple levels of purchase decisions rather than the simple purchase/no purchase decision. Their model, however, differs substantially from ours, in that it relies on a single timed-intent measure and is calibrated using actual purchase data.

for (and modeled by) our technique can be readily funneled down to suitable inputs for others. Finally, asking for *cumulative* probabilities avoids people having to do their own math, so long as those probabilities rise and stay between zero and one. In sum, we measure timed intent by 1) providing a clear and sufficient time horizon (e.g., 36 months); 2) dividing the entire time horizon into shorter, cumulative time intervals (e.g., 0-3 months, 0-6 months, …, 0-36 months); and 3) asking respondents to provide adoption probabilities for each cumulative time interval (e.g., 0-100%). As mentioned previously, we refer to any measure of this type as *cumulative timed intent*.

There is ample evidence to suggest that measuring cumulative timed intent enhances the accuracy of the data. First, providing respondents with a specific and sufficient (in the sense of most adoptions taking place within it) time horizon improves contextual clarity and offers flexibility (Venkatesh, Maruping, and Brown 2006). Improving contextual clarity helps reduce uncertainty and thus improves the accuracy of the self-predictions (Armor and Sackett 2006). A specific, sufficient time horizon, divided into appropriately granular time intervals, also offers respondents flexibility. For example, if a respondent has a strong intent to adopt in 12 months, how might she respond when asked for her intent to adopt "in the next 6 months"? The respondent likely wants to communicate her actual intent despite the mismatch with the question's timing. A sufficient horizon and appropriate intervals should therefore lead to an enhanced ability to accurately communicate self-predictions.

Second, dividing the entire time horizon into shorter, cumulative time intervals and measuring adoption probabilities for each successive cumulative time interval, starting with the most immediate (e.g., zero to six months), should also contribute to the accuracy of self-stated intentions. A rationale is provided by temporal construal theory, which suggests that people, when evaluating products well in advance of buying them, focus on and consequently overweight the abstract benefits associated with the product, while largely ignoring and underweighting more concrete constraints that might prevent them from following through on their intentions (Trope and Liberman 2003). As a consequence, people are more likely to overstate their intentions when the temporal distance between the time of the measurement and the intended behavior increases. First measuring the adoption probability for the most immediate time-interval leads respondents to consider *both* the abstract benefits associated with the adoption of a new technology as well as the concrete constraints that might prevent them from following through on their intention to adopt. The increased salience of the abstract benefits and concrete constraints should serve to enhance the accuracy of the self-stated intentions for substantial temporal distances (Koehler and Poon 2006).

Finally, asking respondents to provide adoption probabilities for multiple cumulative time intervals stimulates cognitive effort (Osberg and Shrauger 1986). Respondents are thereby encouraged to 1) think extensively and in detail about factors that may hinder their ability to adopt the technology as intended (Venkatesh, Maruping, and Brown 2006); 2) take into consideration assumptions about future price and, for instance, quality developments with regards to the technology (Lipshitz and Strauss 1997); and 3) imagine themselves purchasing and using the technology (Klein and Crandall 1995). That is, "If someone's utilities at time t are to predict overt choices at $t + k$, the person's state-of-mind at t must match in important ways the state-of-mind during the critical preference judgment episode(s) at $t + k$" (Wright and Kriewall 1980, p. 278). While the task may be perceived as relatively complex and demanding, greater cognitive effort should enhance accuracy of the intention measures (Castano et al. 2008).

EXAMINING CUMULATIVE TIMED INTENT

To assess the performance of the newly proposed cumulative timed intent measure, we estimate individual-level cumulative adoption likelihood curves and examine them for two perpetual problems confronting intention measures, accuracy and predictive correlates (Sun and

Morwitz 2009). First, using behavioral data collected during a two-year longitudinal study involving 143 consumer decision makers, empirical evidence is provided for the accuracy of the estimated curves for predicting whether and when a technology is actually adopted. Second, using intention data from 478 managerial and 176 consumer decision makers, the relationship between common correlates of intentions and adoption is assessed. If asking people for their expected likelihood that they will have adopted a technology at several points in the future encourages greater cognitive effort, resulting in the provision of more accurate likelihoods, one would expect to see that reflected in the relationship between the slope of the adoption curves and common antecedents of adoption as for instance specified by Rogers' Innovation Diffusion Theory (2003).

Innovation Diffusion Theory. Consistent with the earlier observation that surprisingly little research integrates time into behavioral intent measurement, research on predicting the rate of adoption (based on self-reported behavioral intent measures) remains scarce. One exception is Rogers (2003), who suggests that individuals' perceptions of technology characteristics not only predict whether, but also when, they adopt a technology. Rogers (2003) defines five key characteristics of technologies: relative advantage, compatibility, complexity, trialability, and observability. Concretely, and concisely: *relative advantage* is the degree to which a technology is perceived to be better than the one it supersedes; *compatibility* is the degree to which a technology is perceived as consistent with the existing values, past experiences, and needs of potential adopters; *complexity* is the degree to which a technology is perceived as relatively difficult to understand and use; *trialability* is the degree to which a technology may be experimented with on a limited basis; and *observability* is the degree to which a technology's results are visible to others. Technologies perceived "as having greater relative advantage, compatibility, trialability, and observability and less complexity will be adopted more rapidly than other [technologies]" (Rogers 2003, p. 16). We include these well-established factors to assess whether the slopes of our curves are at least partly predictable.[2] If, as stated, collecting cumulative timed intentions encourages greater cognitive effort, one would expect to see that reflected in the relationship between the slope of adoption curves and common antecedents of adoption like these. Therefore, finding strong associations would lend credence to our approach.

ESTIMATING INDIVIDUAL-LEVEL, CUMULATIVE ADOPTION LIKELIHOOD CURVES

Data considerations

Cumulative timed intent measurement entails asking respondents to inform the researcher about the expected likelihood that they will have purchased a technology at several points in the future, for example, within one month, 6 months from now, 12 months from now, and so forth, conditional on their not having already adopted (that is, at "zero months from now"). These measures are collected at one point in time. As part of a cross-sectional survey, each respondent is presented with the total time horizon and all time intervals at once (in a cumulative format) and asked to provide an expected likelihood score for each time interval. An example using a three-year horizon, 10% probability intervals, and (mainly) six-month time slices appears in Table 1. There is no requirement that the time points be equally spaced, or indeed that each respondent be presented with the same grid, although we have not attempted to optimize these across respondents or make them adaptive.

[include Table 1 about here]

To assess the performance of the cumulative timed intent measure, we first need to estimate the individual-level cumulative adoption likelihood curves. As discussed in the previous section, our goal is to provide a heterogeneous account of individual-level probabilistic growth curves as

[2] It is critical to realize that including or excluding even potentially relevant (Level 2, in hierarchical language) variables has little or no effect on the predictive accuracy of our method. That is, the estimated slope of an individual's adoption curve does not change just because another variable may be used to help explain it.

a function of time and other relevant covariates and examine their predictive accuracy. Because it is critical to allow the degree of time-dependence to vary across people, we adopt a hierarchical Bayes formulation. While these data may seem fairly straightforward from a data collection point-of-view, they pose several challenges with regards to estimating adoption likelihood curves based on them. We briefly summarize these modeling challenges as a prelude to a discussion of how the proposed model resolves them.

The first and main challenge is that the "dependent variable" is a cumulative probability curve. This implies that any model should not even in principle predict outside the unit interval. Furthermore, because each individual's curve is *cumulative*, it is monotonic: stated probabilities cannot decrease with time. The model must therefore allow for flexibility across respondents in shape—how quickly a curve can approach "certainty" (stated probability = 100%)—but not allow curves that peak and then fall. A second challenge is that respondents actually state probabilities of both zero and one. A typical approach in the multilevel modeling literature (e.g., Goldstein 2003; Raudenbush and Bryk 2002) is to transform the dependent variable so that it spans the entire real line, so that a standard heterogeneous hierarchical linear model can be applied. However, common (inverse) transforms like logit and clog map the observed endpoint values of zero (0%) and one (100%) to negative and positive infinity, invalidating typical multilevel formulations. Although arcsine transformations have occasionally been used in heterogeneous formulations in marketing (see, for example, Batra et al. 2008), these do not solve the out-of-range prediction problem, and introduce other known scaling artifacts. Third, because the data are individual-level, a multilevel model is needed to account for heterogeneity, and moreover to *explain* that heterogeneity in terms of covariates (e.g., the technology adoption antecedents discussed previously): how do respondents differ in the degree of 'rise' in their adoption curves? Fourth, due to boundedness of the dependent variable, there is intrinsic nonlinearity: a change from a stated probability of 90% to 100% is likely to be a stronger statement than that from 40% to 50%; some form of nonlinearity must be encoded into the response model. Fifth, there are scaling challenges: responses may appear to be interval-scaled, but in fact lie on a discrete scale whose granularity requires that participants round presumably interval-scaled 'underlying' probability estimates to the provided scale points. For example, participants will often appear to plateau at a specific value (e.g., 50%) for several time periods, only to begin rising again, but it is highly doubtful that their true adoption probability stalled at that precise value. We therefore need a model that relies on some *latent* variable, which is then translated into what we observe on the discrete, ordered scale.

In sum, modeling cumulative timed intent data is apparently novel so far as the multilevel literature goes[3], although throughout we call on established specifications and model estimation techniques. As discussed further below, although it is not a strict requirement, the proposed model should be suitably parsimonious, so that key tests can be mapped unambiguously to specific model parameters, not subtle patterns of variation across sets of them.

We build the model in three stages, as follows, in accordance with general linear model and hierarchical Bayes logic (Rossi and Allenby 2003). The main idea is to have a standard heterogeneous linear regression formulation underlie the response scale represented in the grid of Table 1. Toward this end, the first stage maps covariates to (latent, linear-additive) adoption *propensities* (Kamakura and Wedel 1995) that account for each subject's time trend. The second stage maps these (latent) propensities onto a (latent) *probability* scale, using a probit transform. The third stage maps these (latent) probabilities to *stated* intentions, using a "multiplicative" or "rank order" binomial specification. As such, it is the first of these stages that completely

[3] Goldstein, Harvey; personal communication.

accounts for individual-specific covariates' effects on adoption propensity; the two later stages simply transform these propensities to probabilities and then to intentions scale points.

We discuss each of these stages in turn to help clarify and motivate the specific hierarchical Bayes formulation for the intentions scale underlying our empirical work. Before doing so, we first point out that we do not *impose* monotonicity on any particular respondents' cumulative response pattern, because the cumulative timed intent grid of Table 1 does not force subjects to respond 'correctly' with a non-decreasing response curve; only the directions supplied with the scale help do that (we later examine compliance with monotonicity for all three empirical studies). Our goal, rather, is to both account for and *explain* the heterogeneity in these response curves, while examining compliance with the scale usage directions empirically, along with the individual (posterior) curves stemming from the hierarchical Bayes (HB) formulation, below.

Stage 1: Mapping Covariates to (Latent) Propensities and "Slopes". A key modeling issue is how to account for differing time trends in stated intentions across the respondent pool. As described previously, we adapt a number of formulations from latent growth curve modeling to the standard format for HB models in marketing applications. Specifically, we presume that each respondent has a latent time 'slope'. The "Level-II" model specifies that these slopes are each a linear function of covariates ("explained heterogeneity") plus a residual ("unexplained heterogeneity"), which we take to be (univariate) normal, as is also standard. Although the slopes are heterogeneous, we used a homogeneous intercept for two reasons: first, it is known with certainty that, at time "zero", none of the respondents had adopted, so we have a common degree of knowledge about them; second, estimating a heterogeneous intercept produced inferior values of the integrated likelihood for the focal models (presented later).[4] The general formulation is easily

[4] We examined several alternative formulations – including a hierarchical set up for both the slope and intercept – but did not find them to provide enhanced predictive accuracy, although they did substantially complicate interpretation.

adapted to heterogeneous intercepts, even though we found them to be unnecessary for our specific applications; model estimation and inference would be unchanged. The general specification is therefore consistent with models in the abundant literature on growth curves (e.g., Duncan et al. 2006; Singer and Willett 2003), with the added benefit of including unobserved heterogeneity. Specifically, we propose that the latent adoption propensity for subject *i* at time *t* is given by

$$\text{Propensity}_{it} = \beta_0 + \beta_{1i} t \, . \tag{1}$$

This tells us how the propensity changes over time for an individual, but not how it varies *across* individuals. For that, we must specify the "heterogeneity model",

$$\begin{aligned} \beta_{1i} &= \Delta z_i + u_i \\ u_i &\sim \mathrm{N}\left(0, \Omega_u\right) \end{aligned}, \tag{2}$$

which relates slopes (β_{1i}) to individual-level covariates (z_i) and associated coefficients (Δ). We always include a constant in $\{z_i\}$. Note that Δ being indistinguishable from zero does not mean that respondents failed to differ in their adoption curves (which would manifest in a small variance parameter, Ω_u); rather, it means we cannot systematically *explain* this variation in terms of included covariates. So, the elements of Δ will be important in theory-testing. Finally, we stress that we do not assume that the distribution of latent slope coefficients, β_{1i}, is normal across the population; its distribution will in fact depend on the distribution of covariates, z_i, as well as how they are amplified by the elements of Δ. Rather, it is only the *unobserved* heterogeneity, represented by u_i, that is presumed normal, and Ω_u will account for the degree of this unexplained variation.

Stage 2: Mapping Propensities to (Latent) Probabilities. Because the latent propensity is unbounded, we need to map it to a probability representing (latent) adoption probability. There are several transformations commonly used for this task, among them probit, logit, cloglog, and

---

Note that the intercept (β_0) does not represent where the latent line intersects the first observation (i.e., at "T = 1"), but rather where it *would* cross at T = 0, which we do not observe, and where everyone has told us they had not already purchased.

cauchit, with the first two overwhelmingly represented in marketing applications. We choose a probit transform because of its conjugacy properties vis-à-vis Bayesian algorithm design, and because it offered superior fit to the logit in all our applications. That is, the latent adoption propensity is translated to the [0, 1] probability scale via an inverse-normal transform:

(3) $$\pi_{it} = \Phi^{-1}\left[\text{Propensity}_{it}\right]$$

Because this resulting probability (π_{it}) is continuous, but our observable stated intent lies on a discrete scale, one more model stage is required.

Stage 3: Mapping Latent Probabilities to the Observed Adoption Intent Scale. There are many ways to map a continuous latent variable to a given discrete scale. When the scale has ordinal properties, a typical set-up involves estimating cutoffs. For adoption intent scales with a reasonable degree of granularity, like in Table 1, this would require estimating a large number of cutoffs, which would render the resulting model highly non-parsimonious, particularly so compared with all previous intention models. Because Stage 2 provides us with a (latent) probability (π_{it}), we can employ an especially parsimonious transformation, the *rank-ordered binomial*. This mapping (from continuous latent, to discrete observed, probabilities) has a long history in psychometrics and marketing, including a classic paper (Adrich 1978) cataloging its properties in a Rasch model framework, and an extended statistical treatment (Rost 1985) upon which several marketing applications (e.g., Kamakura and Wedel 1995; Kamakura et al. 2003; Ying, Feinberg, and Wedel 2006) have been based. Specifically, to make the discussion concrete, let us assume that probabilities are elicited on a K-point scale (in our applications K is kept constant across data sets, but could be technology-specific). Denote the observed intent scale category as Y_{it}, where $Y_{it} = k$ if, in time period t, respondent i chooses probability category k ($k = 1,\ldots, K$). For simplicity of notation, we presume that each respondent uses the same grid of time and probability points, as per Table 1, resulting in a balanced data set; this can easily be relaxed if the application somehow requires it. The rank-order binomial mapping is as follows:

(4) $$p\left(Y_{it}=k\right)=\binom{K-1}{k-1}\pi_{it}^{k-1}\left(1-\pi_{it}\right)^{K-k},k=1,\ldots,K \ .$$

The above representation not only ensures logical consistency (Naert and Bultez 1973), but is especially parsimonious: it requires many fewer parameters than an unconstrained or rank-order multinomial specification (e.g., Chien and George 1999). It also has a special, and desirable, property: the expected value of the observed adoption intent response is equal to π_{it}. Thus, the rank-order binomial *itself* does not 'shift' the probability curve given by the first two model stages, only maps it onto the given adoption intent scale.

Final Model. Conjoining all three model stages and Equations 1 to 4 yields the hierarchical Bayes formulation estimated in the empirical studies to follow:

(5) $$\begin{aligned} &\text{Level I:} && p\left(Y_{it}=k\right)=\binom{K-1}{k-1}\pi_{it}^{k-1}\left(1-\pi_{it}\right)^{K-k},k=1,\ldots,K \\ & && \text{Probit}\left(\pi_{it}\right)=\beta_0+\beta_{1i}t \\ &\text{Level II:} && \beta_{1i}=\Delta z_i+u_i \\ & && u_i \sim \mathrm{N}\left(0,\Omega_u\right) \end{aligned} \ .$$

Inferences about all model parameters are made based on posterior distributions obtained via MCMC, using 100,000 thinned draws for inference. Additional details on data preparation, as well as on model estimation, inference, and checking appear in Web Appendix A., together with a link that provides the code and (partly synthetic, as required by our sponsoring firm) data in three 'independent' systems (MLWin, WinBUGS, and SAS).

STUDIES 1 – 3: EXAMINING CUMULATIVE TIMED INTENT AMONG MANAGERIAL AND CONSUMER DECISION MAKERS

To examine cumulative adoption likelihood curves in more detail, three studies were conducted. The data were collected among real managerial and consumer decision makers as part of a longitudinal research collaboration with a global Fortune 100 Company operating in the heavy equipment industry. The objective of Studies 1 and 2 was to provide a heterogeneous account of individual-level adoption likelihood curves as a function of time and the antecedents of technology adoption, as identified by Innovation Diffusion Theory, in two unique and highly distinct contexts; the first of these involved 212 superintendents of U.S. golf courses deciding on the adoption of an advanced mower, the second 266 U.S. farm operators deciding on the adoption of an autoguidance system. The objectives of Study 3 were to extend our investigation of cumulative timed intent from managerial to consumer decision makers, who decided whether to adopt a GPS cell phone, using a shorter total time horizon and smaller time-intervals; to examine the accuracy of the estimated cumulative adoption likelihood curves for predicting whether and when people adopt; and compare its predictive performance to that of two single-intent measures (Juster 1966; Morwitz 1994) and two well-known intent models (Bemmaor 1995; Morrison 1979).

Study 1: Golf Course Superintendents Deciding on the Adoption of an Advanced Mower

Study 1 involved 212 superintendents of U.S. golf courses, who were surveyed about their intention to adopt an advanced mower using the cumulative timed intent measure. Conventional mowers typically have a gas or diesel engine that powers the riding unit as well as the alternator. This alternator powers the cutting units independently of propulsion speed. The alternator of the advanced mower is powered by electricity, allowing all the hydraulics in the cutting units to be removed, effectively eliminating the risk of detrimental oil leakage on the golf course.

The questionnaire was sent to a random sample of superintendents of 3000 golf courses in the U.S.; contact information was retrieved from the National Golf Foundation Database. All people in the sample were offered an opportunity to enter a sweepstakes for a $20 gift certificate, to be handed out to a total of fifty respondents. In total, 212 superintendents responded to the survey (response rate = 7.6%). A follow-up with non-respondents and a comparison of early and late respondents suggested that non-response biases appear marginal at best (see Web Appendix B for details). The majority of these managerial decision makers, 95.7%, indicated that they would be making or have a significant influence on the final decision of adopting this technology.

Study 2: Farm Operators Deciding on the Adoption of an Autoguidance System

Study 2 involved 266 U.S. farm operators, who were surveyed about their intention to adopt an autoguidance system for tractors using the cumulative timed intent measure. An autoguidance system automatically steers farm machinery using Global Positioning System (GPS) satellite signals. It improves accuracy (reducing overlaps and skips when working the land) and reduces operator fatigue. The questionnaire was sent to a random sample of 3000 U.S. farm operators; contact information was retrieved from a publicly available database. In total, 266 farm operators responded (response rate = 9.1%). As in Study 1, non-response biases appear marginal at best (see Web Appendix B for details). Close to 100% of the participants indicated that they would be responsible for or have a significant influence on the final decision to adopt this technology.

Study 3: Students Deciding on the Adoption of a GPS Cell Phone

Study 3 surveyed a sample of 354 student participants at a large U.S. university on their intentions to adopt cell phones with GPS technology. We selected this particular context because of its relevance among college students. After the participants entered the lab, they were informed that a growing number of companies are introducing cell phones with GPS technology and that we were interested in their intent to buy one. Half the participants were asked to express their intent to adopt using two traditional single-intent measures (Juster 1966; Morwitz 1994), while the other half was asked to express their intent to adopt using cumulative timed intent measures. At the end of the study, all participants were asked to participate in a longitudinal follow-up study on the same topic. Participants were informed that they would be re-contacted via email every six months for at least two years and that those responding to all future surveys would receive \$20. 40.4% of the participants agreed and participated in all four follow-up surveys; 49.0% of those who agreed and participated were in the single-intent measure condition, 51.0% were in the cumulative timed intent measure condition. A variety of checks were conducted to examine if participants in the follow-up study differed from those who decided not to participate in the first place. The results suggest that non-response biases are at most modest (see Web Appendix B for details).

Measures Used in Studies 1, 2, and 3

Dependent Variables. After participants read the cover story describing the relevant technology in objective terms, intention measures were collected. For the golf course superintendents in Study 1 and the farm operators in Study 2, a 5-year time span was identified as a realistic investment horizon for the equipment under consideration. Cumulative timed intent was measured for both technologies using 6-month increments (with the exception being the first, which was labeled “end of this month,” and coded as $t = 1/12$ year). The response scale ranged from 0% = “I will not have bought one” to 100% = “I will have bought one,” with 10% intervals. The 10% interval width was selected as a trade-off between scale granularity and the potential for being taxing to respondents (Juster 1966).

With respect to the students of Study 3 adopting GPS cell phones, a 36-month time horizon was selected, based on an average cell-phone replacement rate of 24 months. Cumulative timed intent was measured among half the sample using 3-month increments. The other half expressed their intent to adopt a GPS cell phone in the next 36 months using two single-intent measures. Participants were first asked for the likelihood that they will have purchased a GPS cell phone in the next 36 months using a response scale that ranged from 0% = "I will not have bought one" to 100% = "I will have bought one," with 10% intervals (Juster 1966). Next, they were asked in an open-ended format how many months from now they expect they will have purchased the GPS cell phone (cf., Morwitz 1994). Participants were able to respond "I will never purchase a GPS cell phone."

[include Tables 2 and 3 about here]

Feedback received during informal debriefings, as well as high compliance rates, consistently suggested no difficulties in responding to the cumulative timed intent question. All participants reporting any intent to adopt exhibited monotonic, non-negative cumulative timed intent, and only 1.7% provided the same probability for all time intervals. Participants who did plan to adopt the technology beyond the provided time horizon shared, during debriefing, that they reported an expected likelihood on the final time interval well below 100%. Finally, we found no evidence to suggest that not having a "do not know" category adversely influenced responses.

Independent Variables. To measure the independent variables proposed by Innovation Diffusion Theory, we relied on published, validated scales (see Table 2). All scales were pre-tested among the respective participants for clarity prior to being included in the survey. All items were randomized, and participants responded using a 5-point response scale (1 = total disagree; 5 = totally agree). The internal reliability of the scales—Cronbach's α, calculated based on the average inter-item correlations—ranges from reasonable to very good across technologies and decision makers. Orthogonalized weighted averages were used in model estimation.

Results – Antecedents of the Slope of Cumulative Timed Adoption Curves (Studies 1, 2, and 3)

To examine whether the steepness of slopes of the estimated cumulative timed adoption curves behaves according to expectations, we first examine the cumulative timed intent measures for all three studies; results appear in Table 3.

The 'baseline' model includes only a constant among the z_i in Equation (2); the Time coefficient has no predictors and is therefore just a pure normal random effect across respondents. As the significant effect of Time in Table 3 demonstrates, the underlying (i.e., latent) likelihood of adoption increases over time, over and above any other model covariates, as one might expect.

Innovation Diffusion Theory suggests that technological adoption speed is positively influenced by perceived relative advantage, compatibility, trialability, and observability, and negatively by perceived complexity. Consistent with this suggestion, we find that *relative advantage* positively relates to the adoption speed for all three technologies. *Compatibility* also positively relates to the speed of adoption among most decision makers, although only marginally among golf course superintendents. This makes sense when one considers that the advanced mower can be used in and by itself, without greatly affecting existing operations, while an autoguidance system is added onto a tractor. Next, we find that perceived *complexity* negatively relates to rate of technology adoption among managerial decision makers. Because these technologies are critical to their daily operations, it is important for them to believe that they are able to make such technology work for them. Perceived *trialability*, in turn, also positively relates to the adoption rate of technologies among both managerial decision makers, but not among the consumer decision makers. This may be due, in part, to the technology itself. Most consumer decision makers were familiar with GPS technology for location and direction purposes. Finally, perceived *observability* relates to the adoption rate of the advanced mower only. This finding lines up with the strong desire and common practice of "demo-ing" new technologies before deciding to adopt them among golf course superintendents.

The results presented here suggest that the slopes of our curves are at least partly predictable, and thereby lend credence to the general approach. The confirmation of the importance of the five technology characteristics identified by Rogers *across* these studies provides a degree of construct validation for our cumulative timed intent measure.

Next, the predictive accuracy of the proposed cumulative timed intent measure is examined using the behavioral data collected during the two-year longitudinal follow up of Study 3. We assess the accuracy of the proposed measure in predicting individual-level adoption (whether and when), as well as cumulative adoption, and compare it to the predictive performance of eight alternative measures and methods, as discussed below.

Measures and Models to Examine the Predictive Accuracy (Study 3)

Model M1. In predicting adoption, we will assess the extent to which the proposed measure predicts (1) *whether* the consumer decision makers of Study 3 actually adopted a GPS cell phone during the first two years following the original survey, and (2) if they adopted during that time frame, how accurate our measure is in predicting *when* exactly they adopted. We will use a respondent's *individual-level* estimates—based on the HB model of Equation (2)—to predict his or her actual time to adoption. This requires that we include both the individual's self-reported probability curve *and* the covariates (for that individual) in each of the models compared. However, we found empirically that, conditional on that individual's self-reported adoption curve and the HB model estimates, the (individual's) covariates added little, if anything, to predictive accuracy. Thus, we present results only for the "baseline" HB model, with no covariates (results for the other models are available from the authors). In predicting expected time-to-adoption, we use a 50% probability cut-off for each individual-level curve; because this corresponds to a probit transform of zero, Equation (2) makes this simple to calculate: $E[T_i] = -b_{0i} / b_{1i}$ for each respondent, *i*. Decision makers for whom the predicted timing of adoption falls outside the 2-year window will be classified as "predicted not to adopt during the time frame."

Model M2. A first benchmark is to compare the predictive accuracy of the estimated cumulative adoption likelihood curves with the predictive accuracy of the *raw cumulative timed intent scores* based on which these curves are estimated. In predicting expected time-to-adoption, we use a 50% probability cut-off, which is inferred assuming a linear relationship between time and adoption likelihood – i.e., if an individual reports a likelihood of 30% at t = 3 and a 60% likelihood at t = 4, the expected time-to-adoption is inferred as t = 3.67. More sophisticated, non-linear methods of interpolation made virtually no difference in reported figures. We expect M1 to outperform M2, for two reasons: first, M1 predicts adoption based on the estimated "latent" intention, as opposed to 'raw' and possibly noisy reported intention used by M2; second, M1's hierarchical nature draws stability from the full set of respondents. That is, M1 'smoothes' individual respondents' likelihood estimates over time (via the Level-I "utility" model), and 'shrinks' them across respondents (via the Level-II "heterogeneity" model), unlike M2, which simply takes each respondent at his or her word.

Model M3. M3 is a *covariate-based model* (i.e., those allowing for observed heterogeneity via estimated values for Δ). This approach is equivalent to assuming that each of the respondents is "new;" that is, we observe only that respondent's covariates, *not* his or her self-reported adoption curves, and use the hyperparameters of the HB model to make predictions (we leave out the details of Bayesian computation here, as they are standard, involving summing over thinned draws from the MCMC sampler for the individual random-effects and model hyperparameters). In sum, then, for the "existing" customer analysis, we look only at the baseline HB model, while for the "new" customer comparison, we consider both models, the baseline and the IDT-based model we have previously detailed. In reality, of course, the two sets of customers are identical, and we have merely excised the self-reported adoption curves in the "new" customer case to allow meaningful comparisons and as a nod to the relatively small number of respondents who agreed to provide subsequent data. M1 is anticipated to outperform M3, as M1 (in addition to the benefits relative to M2 discussed previously) predicts adoption based on individual-level intent data, while the latter uses aggregate-level estimates to predict behavior using individual-level data on *determinants of* intentions instead.

Model M4. M4 uses respondents' intentions to adopt expressed using a traditional *single-intent measure* – the *expected likelihood of adoption* (0%-100%). To infer the 50% probability cut-off for the single-intent likelihood measure, a linear relationship between time and adoption likelihood is assumed – i.e., if an individual reports a likelihood of 80% at t = 36 months, the inferred likelihood at 24 months is 53.3% and the expected time-to-adoption is inferred as 22.5 months. [This approach yielded more favorable results than using common non-linear curves for interpolation purposes.] Since asking decision makers for their expected likelihood that they will have purchased a technology at several points in the future should encourage greater cognitive effort (than to asking this for a single point in time), one might expect M1 to provide more accurate predictions than M4.

Model M5. M5 uses an alternative *single-intent measure* that asks respondents to indicate the *expected timing of adoption* (*how many months from now they will adopt the technology*). The open-ended responses are used to assess whether and when respondents intent to adopt the technology during the first two years following the survey. Anticipated differences in cognitive effort suggest that M1 should, as was predicted for M4, outperform M5.

Models M6 and M7. M6 and M7 use the beta-binomial model of Morrison (1979) (Kalwani and Silk 1982) to predict cumulative adoption for the first two years. M6 uses the reported likelihood at t = 24 months that was collected as part of all cumulative timed intent measures. Note that, because M6 thus benefits from the posited cognitive-effort advantages of the proposed method (though without availing of the full set of individual-level cumulative timed likelihoods), whereas M7 estimates the Morrison model based on the reported likelihood using the single-intent measure, M6 is expected to offer superior performance. M1 should handily outperform M6 (and thus M7): as discussed previously, M1 predicts adoption based on "smoothed and shrunk" individual-level cumulative adoption data, whereas M6 cannot.

Models M8 and M9. Finally, M8 and M9 use Bemmaor's (1995) intent model to predict cumulative adoption for the first two years using the same two likelihood scores as M6 and M7, respectively. Therefore, using the same logic as above (i.e., M1 vs. M6 vs. M7), we anticipate that M8 will outperform M9, and M1 will outperform them both.

[include Table 4 about here]

Results – Predictive Accuracy of Cumulative Timed Adoption Curves (Study 3)

Individual Hit Rates. With regards to predicting whether decision makers will adopt during the first two years of technology availability, Table 4 shows that the proposed measure (M1) is highly accurate: M1 correctly classified 87.2% of the consumer decision makers as non-adopters and 80.6% as adopters, for an overall "hit rate" of 84.3%. M1 in fact substantially outperforms the other measures and models estimated (M2-M9). As shown in Table 4, M1's relative strength lies in correctly classifying adopters, which, generally speaking, is more elusive than predicting status quo (cf., non-adoption).

Let us examine the proposed model's relative performance in more detail. First, the cumulative timed intent measure, coupled with the proposed HB framework, offers greater accuracy than the raw data on which it was estimated (84.3% for M1 vs. 78.6% for M2; $z = 1.26$, $p = .10$).[5] The difference in performance is mainly driven by M1's relative performance in correctly classifying adopters (80.6% vs. 71.0%; $z = 1.73$, $p < .05$). The statistical strength of these results, of course, is compromised by modest sample size and the nature of binary, across-group comparisons.

M1 significantly outperforms the covariate-based model (M3), which achieves only a 61.4% hit rate ($z = 2.85$, $p < .01$). The estimated cumulative adoption curves (M1) more accurately classify non-adopters (87.2% vs. 61.5%; $z = 2.36$, $p < .01$) and adopters (87.2% vs. 61.3%; $z = 1.65$, $p < .05$) than the covariate-based model (M2). Taken together, the M1-M2 and M1-M3 comparisons suggest that collecting self-reported cumulative timed intent data and using estimated adoption likelihood curves to predict future behavior is a particularly expedient way to extract predictive accuracy from cumulative timed intent data.

The results further show that the cumulative timed intent measure outperforms both of the single-intent measures (M4, M5). The predictive accuracy of M1 significantly outperforms that of the adoption likelihood question (84.3% vs. 60.2% for M4; $z = 3.01$, $p < .001$) as well as the timing-based question (57.5% for M5; $z = 3.33$, $p < .001$). As Table 4 shows, M1 is more than twice as accurate in classifying adopters as both single-intent measures (80.6% vs. 35.3%; $z = 3.41$, $p < .001$).

Finally, we find that the difference in predictive accuracy between M1 and the four intent models (M6–M9) largely depends on which intent data are used to estimate the intent models.

[5] Recall that M1, M2, M3, M6, and M8 are based on the cumulative timed intent measure, collected for one respondent sample. Comparing them thus requires a test of "correlated proportions" (a McNemar test). M4, M5, M7, and M9 are based on single intent measures collected for a different respondent sample, so comparisons to M1 require tests of two "independent proportions."

Recall that M6 and M8 are estimated using the reported likelihood (at $t = 24$ months) collected as part of all cumulative timed intent measures. As a result, M6 and M8 'benefit' from the advantages of cumulative timed intent data, apparent in their respective columns of Table 4. Nevertheless, M1 outperforms M6 and M8 (84.3% for M1 vs. 78.6% for M6 and M8; $z = 1.26$, $p = .10$), albeit marginally. The differences in performance are much more pronounced when comparing the predictive accuracy of the proposed measure to M7 and M9, which represent the Morrison and Bemmaor models estimated in a traditional way, using a single-intent measure, in which case M1 substantially outperforms both models (84.3% for M1 vs. 67.1% for M7 and M9; $z = 2.19$, $p < .01$. This is also confirmed by the reported log-likelihoods.) That the predictive performance of both intent models improves by replacing the traditionally used *single* intent measure with data collected as part of the *cumulative* timed intent measures (from 67.1% for M7, M9 to 78.6% for M6, M8; $z = 1.42$, $p < .10$) helps corroborate the differential value of the newly proposed measure.

Accuracy of Predicted Timing. A comparison of the absolute differences in the predicted and actual *timing* of adoption, as well as the RMSEs, confirms the superior performance of cumulative timed intent measures. M1 most accurately predicts the timing of adoption: for a 24-month time period, the Mean Absolute Deviation of the cumulative timed intent measure is a relatively small 3.6 months. While M1 outperformed M2 in correctly classifying adopters (80.6% vs. 71.0%; $z = 1.73$, $p < .05$), M1 is only directionally more accurate in predicting when these correctly classified adopters adopt than M2 (3.6 vs. 4.2 mos; $t = 1.24$, $p > .10$). M1 is, however, significantly more accurate than M3 (3.6 vs. 8.0 mos.; $t = 3.59$, $p < .001$), M4 (3.6 vs. 5.5 mos.; $t = 2.8$, $p < .01$), and M5 (3.6 vs. 8.8 mos.; $t = 3.80$; $p < .001$) in predicting when correctly classified adopters actually adopt the technology.

Cumulative Adoption. Although the objectives of collecting cumulative timed intent measures are to predict whether and when *individuals* adopt a new technology and to study the *heterogeneity* in whether and when people adopt, it is useful to be able to predict cumulative adoption during a certain time frame (even though a cumulative adoption prediction does not reveal the distribution during that time frame).

Table 4 demonstrates that the proposed measure M1 performs only marginally better than M2 (-1.4% vs. -4.3%; $z = .17$, $p > .10$) and M3 (-1.4% vs. +4.3%; $z = .51$, $p > .10$), but significantly better than M4 (-1.4% vs. -20.6%; $z = 1.94$, $p < .05$) and M5 (-1.4% vs. -17.8%; $z = 1.85$, $p < .05$) in predicting cumulative adoption during the first two years after the survey. Figure 1 shows the actual cumulative adoption and the predicted cumulative adoption for M1 to M5. The results demonstrate the superiority of M1 and helps highlight why M2-M5 may not perform so well, by comparison. For instance, M2 tends to overpredict cumulative adoption throughout the two-year time frame, while M3, M4, and M5 underpredict (with only M3 surging at the very end to overprediction). The 'staircase' appearance of the cumulative adoption line for M5 suggests that, when participants are asked in an open-ended format how many months from now they intend to adopt, they favor responding in 6-month intervals (e.g., 6 months from now, 12 months from now). This may contribute to the poor performance of this measure.

[include Figure 1 about here]

As the results in Table 4 demonstrate, M1's superiority to M6–M9 in predicting cumulative adoption at the end of the two-year time horizon is marginal at best ($p > .10$). This suggests that the proposed approach, which estimates individual-level cumulative adoption likelihood curves calibrated only on self-reported likelihood scores (and not on actual behavioral data), performs at least as well in predicting cumulative adoption as the Morrison and Bemmaor models. Combined with the superior results of the proposed method for predicting individual-level adoptions, as well as its ability to predict whether *and* when people adopt, the results suggest the proposed approach 'outbenefits' both prior ones.

We summarize that the predictive accuracy of the proposed approach (1) outperforms the raw data model M2 and the covariate-based model M3 in predicting individual-level adoption; (2) outperforms two single-intent measures (M4, M5) for predicting individual and cumulative-level adoption; and, finally, (3) outperforms two widely-validated intent models in predicting individual-level adoption. In drawing the final conclusions, it is important to note that correctly predicting behavior at the individual level is more complex than making aggregate predictions, as the latter can benefit from averaging out under- and over-predictions. Furthermore, an important reason for collecting individual-level data is that it allows for studying *sources* of heterogeneity. Finally, the individual-level cumulative timed intent measure allows for jointly predicting both *whether* and *when* an individual will adopt; separate, aggregate models of these phenomena miss their intrinsic interrelatedness. With this in mind, we conclude that these results, taken together, suggest great promise for future research on using and optimizing cumulative timed intent measures across a variety of application domains.

GENERAL DISCUSSION AND CONCLUSIONS

A common approach to reducing the uncertainty about whether and when people will adopt new technologies is to collect self-reported adoption intentions. Self-reported intentions are among the most widely applied proxy measures for any type of behavior, including the adoption of new technologies. Although multiple meta-analyses have demonstrated that intention measures do correlate with actual behavior, the predictive accuracy of these measures tends to remain limited.

Instead of developing a formal statistical model that attempts to correct for biases that limit the predictive accuracy of intention measures, we proposed a response scale that helps reduce these biases at the point of collection. More specifically, responding to suggestions that the integration of time into intention measurement may help improve the predictive accuracy of these measures, we proposed a novel measure, called *cumulative timed intent*, for assessing whether and when people intend to adopt a new technology. To enable estimation of individual-level cumulative adoption likelihood curves, we asked decision makers to inform us about their expected likelihood that they will have purchased a technology at several points in the future, for example, within one month, 6 months from now, 12 months from now, and so forth, conditional on their not having already adopted. We stress again that individual-level cumulative adoption likelihood curves can be estimated without actual sales data, so thus can be determined prior to launch, and allow for accurately predicting whether *and* when people intend to adopt a relevant technology.

We first demonstrated the method in two unique and highly distinct contexts, the first involving 212 superintendents of U.S. golf courses deciding on the adoption of an advanced mower, the second 266 U.S. farm operators deciding on the adoption of an autoguidance system. We estimated the individual-level cumulative adoption likelihood curves—curves calibrated on self-reported adoption likelihoods for cumulative time intervals across a fixed horizon—of these 478 managerial decision makers, and, using a hierarchical Bayes formulation, provided a heterogeneous account of the curves as a function of time and common antecedents of technology adoption as proposed by Rogers' Innovation Diffusion Theory. For example, consistent with expectations, we found that the rate of adoption likelihood increases with the perceived relative advantage of the technology. The perceived complexity of a technology, in turn, has the opposite effect, reducing the rate of adoption likelihood. These results suggest that the slopes of our curves are at least partly predictable and lend credence to the general approach.

A third study generalized these results among 354 *consumer* decision makers and, using behavioral data collected during a two-year longitudinal study involving a subsample of 143 consumer decision makers, provided empirical evidence for the accuracy of cumulative adoption likelihood curves for predicting whether and when a technology is adopted as well as for predicting cumulative adoption. Cumulative adoption likelihood curves are shown to outperform two single-intent measures (Juster 1966; Morwitz 1994) as well as two widely-validated intent models (Bemmaor 1995; Morrison 1979) for predicting individual-level adoption for a fixed time period of two years. Overall, we conclude that using individual-level adoption curves can provide superior accuracy in predicting whether and when people will adopt a new technology.

FUTURE RESEARCH

Because our approach is new, our empirical findings, while clear and unequivocal taken on their own, should be considered suggestive of a future research program addressing cumulative, timed intent measures in a variety of substantive settings. Intentions are measured across a variety of academic disciplines to examine and predict a rich array of behaviors, and future research may examine the applicability and performance of the proposed cumulative timed intent measure in those contexts. As the cumulative timed intent measure not only more accurately predicts *whether* people behave in a certain way in the future, but also *when* this behavior is most likely to occur, future research may also address the relative utility of existing behavioral theories for predicting the timing of various actions. For instance, to what extent do the antecedents specified by the Theory of Reasoned Action (Fishbein and Ajzen 1975) predict when people will display the behavior under study? In the specific context of technology adoption, it would be intriguing to examine how well the popular Technology Acceptance Model (Davis 1989) predicts when people actually do adopt a specified technology.

To optimize new product and brand introductions, intentions are also collected to determine the importance of product attributes among consumers deciding whether or not to purchase the focal product or brand (Harte and Koele 1995). The cumulative timed intent measure allows for examining the importance of product attributes in deciding exactly when to purchase the product or brand in question. Besides collecting intentions to predict first-time purchase behavior, intentions are collected to predict *repeat* purchases. For instance, using satisfaction ratings, firms often try to predict customer repurchase intentions. The proposed cumulative timed intent measure may mitigate some of the challenges faced when collecting satisfaction and intent data in the same survey (see Mittal and Kamakura 2001), allowing researchers to examine the relationship between satisfaction and future behavior more closely, and accurately. Furthermore, the proposed approach may be used in the large body of experimental research that collects longitudinal (i.e., within-subject, over time) rating or other such scaled data; the hierarchical nature of the model allows researchers to assess the effects of specific subject-specific covariates, such as demographics, on the (latent) steepness of the individual-level curves (cf., Dubé and Morgan 1996).

In benchmarking cumulative timed intent measures, different aspects can be examined: Does the response scale type affect predictive accuracy (e.g., a probability vs. a binomial intent scale)? To what extent would *non*-cumulative time intervals offer up results similar to those reported here? Furthermore, must respondents really respond to each time interval, or would it suffice for them to identify “the most likely time period” and provide an intention measure for that period only? Finally, more research on optimizing critical scale elements is desirable. For example, as we have presented it, the scale—the total time horizon, time intervals used, and size of the probability blocks—needs to be constructed by the researcher through deep knowledge of the product or technology whose acceptance is being analyzed. Adaptive methods, like those used in conjoint, might allow each respondent to quickly zero in on the ‘sweet spot’, the temporal region in which the adoption curve shows the most pronounced S-shape, and thereby offer the most information through the fewest requested data points.

Besides conducting new studies to further benchmark cumulative timed intentions, simulations may be used to examine the methodological and managerial implications of potential response biases and the impact of various modeling assumptions. For example, despite the inherent flexibility of the hierarchical Bayes approach, a host of assumptions need to be invoked, such as for priors, the specific form of the heterogeneity model (here, a single normal component, as opposed to latent classes or a general mixture), the covariates entering the Level-II model, the 'links' from the latent linear specification to the latent probability and then to the observed scale value, etc. There are also alternative forms of 'bias', as catalogued in much prior research, including those induced by the scale itself and those intrinsic to having people express uncertain intentions. It is possible that the much-touted flexibility of our approach interferes with its robustness, for example. On the other hand, we deliberately sought a parsimonious specification, avoiding 'nuisance parameters' in the form of estimated cutoffs for the ordinal intentions scale, for example. Dedicated simulations could examine how well our specific formulation holds up to these known sources of difficulty in intentions measurement, compared with extant methods, as well as with alternative, less parsimonious, specifications.

While it is notoriously difficult to assess the size of non-response biases, they did not appear pronounced when using the proposed method. As actual adoption is predicted based on individual-level cumulative timed intentions, non-response biases should not grossly affect predictive accuracy. However, non-response biases may influence the relationship between correlates of intentions and adoption. More research on the effects of such non-responses biases is clearly warranted.

Future research would do well to examine potential estimation biases in the method. For example, we did notice a mild tendency towards overprediction of adoption timing: on average, the proposed model predicted adoption take place 3.0 months before it actually did. Whether or not this is a function of the type of products used, an artifact of some aspect of the scale, or a perceptual or cognitive effect remains to be determined.

It would also be fruitful to examine in more detail if and how predictive accuracy can itself be improved. Prior to responding to cumulative timed intent questions, decision makers could be asked to list reasons why they may want to adopt a technology, as well as reasons why they may not be able to actually purchase it (Castano et al. 2008; Hoch 1984). By rendering abstract benefits and concrete constraints more salient, the accuracy of cumulative timed intent measured may be further enhanced. The proposed model could also be extended to account for *uncertainty* in self-reported adoption likelihood, by explicitly requesting this information from respondents. Incorporating this type of individual-level uncertainty may help improve the predictive accuracy of the proposed method. For example, some respondents may be certain they will purchase a new product at a precise time in the future, due to a specific need or event, while others are only sure they will need it *within* a certain time frame, and yet others are only guesstimating. It would not be difficult to incorporate this information into the model's likelihood function, but it may place a higher cognitive burden on respondents, who would need to specify *two* probabilities for each time point: likelihood of purchase, and certainty of estimate. Of course, estimating such a model would require dedicated data, and is not possible, to our knowledge, with data collected in any prior intentions study, ours included.

In sum, there is ample opportunity to examine cumulative timed measures in a variety of substantive settings, particularly so for purchase intent. We hope that the approach presented here will not only stimulate academic research on cumulative timed intent measures, but will also encourage practitioners to start utilizing them. As cumulative timed intentions can be collected prior to launch at relatively low costs, they can be used to reduce the uncertainty of whether and when the target market will adopt a new technology, allowing marketing managers to better plan a technology's production, pricing, distribution, and promotion.

REFERENCES

Abraham, Charles, Paschal Sheeran, Paul Norman, Mark Conner, Nanne de Vries, and Wilma Otten (1999), "When Good Intentions Are Not Enough: Modeling Postdecisional Cognitive Correlates of Condom Use," *Journal of Applied Social Psychology*, 29 (December), 2591-612.

Ajzen, Icek and Martin Fishbein (2005), "The Influence of Attitudes on Behavior," in *The Handbook of Attitudes*, Dolores Albarracín, Blair T. Johnson, and Mark Zanna, eds. Erlbaum, Mahwah, NJ, 173-221.

Albarracin, Dolores, Blair T. Johnson, Martin Fishbein, and Paige A. Muellerleile (2001), "Theories of Reasoned Action and Planned Behavior as Models of Condom Use: A Meta-Analysis," *Psychological Bulletin*, 127, 142-61.

Alexander, David L., John G. Lynch Jr., and Qing Wang (2008), "As Time Goes By: Do Cold Feet Follow Warm Intentions for Really New Versus Incrementally New Products?" *Journal of Marketing Research*, 45 (June), 307-19.

Andrich, David (1978), "A Rating Formulation for Ordered Response Categories," *Psychometrika*, 43, 561-73.

Armor, David A. and Aaron Sackett (2006), "Accuracy, Error, and Bias in Predictions For Real versus Hypothetical Events," *Journal of Personality and Social Psychology*, 91 (4), 583-600.

Armstrong, J. Scott, Vicki G. Morwitz, and V. Kumar (2000), "Sales Forecasts for Existing Consumer Produces and Services: Do Purchase Intentions Contribute To Accuracy?" *International Journal of Forecasting*, 16 (3), 383-97.

Batra, Rajeev, Nilufer Z. Aydinoglu, Fred M. Feinberg and Charles Y. Z. Zhang (2008), "Multi-Country Brands: Does More Consistent Global Imagery Increase Consumer Preference?" *Working Paper*, University of Michigan.

Bemmaor, Albert C. (1995), "Predicting Behavior from Intention-To-Buy Measures - The Parametric Case," *Journal of Marketing Research*, 32 (May), 176-91.

Castano, Raquel, Mita Sujan, Manish Kacker, and Harish Sujan (2008), "Managing Consumer Uncertainty in the Adoption of New Products: Temporal Distance and Mental Simulation," *Journal of Marketing Research*, 45 (June), 320-36.

Chib, Siddhartha (1995), "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association*, 90 (432), 1313–21.

Chien, Yung-Hsin and Edward I. George (1999), "A Bayesian Model for Collaborative Filtering," In Proceedings of the 7th international workshop on artificial intelligence and statistics, San Francisco, CA.

Davis, Fred D. (1989), "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly*, 13, 319-39.

--- and Paul R. Warshaw (1992), "What Do Intention Scales Measure," *Journal of General Psychology*, 119 (4), 391-407.

Dubé, Laurette and Michael S. Morgan (1996), "Trend Effects and Gender Differences in Retrospective Judgments of Consumption Emotions," *Journal of Consumer Research*, 23 (September), 156-62.

Duncan, Terry E., Susan C. Duncan, and Lisa A. Strycker (2006), *An Introduction to Latent Variable Growth Curve Modeling: Concepts, Issues, and Applications*, 2nd ed., Lawrence Erlbaum Associates.

Fishbein, Martin and Icek Ajzen (1975). *Beliefs, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley Publishing.

Goldstein, Harvey (2003), *Multilevel Statistical Models*, 3rd ed. London: Oxford University Press.

Gollwitzer, Peter M. (1999), "Implementation Intentions - Strong Effects of Simple Plans," *American Psychologist*, 54 (7), 493-503.

Harte, Johanna M. and Pieter A. Koele (1995), "A Comparison of Different Methods for the Elicitation of Attribute Weights: Structural Modeling, Process Tracing, and Self-reports," *Organizational Behavior and Human Decision Processes*, 64, 49-64.

Hoch, Stephen J. (1984), "Availability and Interference in Predictive Judgment," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10 (4), 649-62

Hsiao, Cheng and Baohong Sun (1999), "Modeling Survey Response Bias - With an Analysis of the Demand for an Advanced Electronic Device," *Journal of Econometrics*, 89 (1-2), 15-39.

---, Baohong Sun, and Vicki G. Morwitz (2002), "The Role of Stated Intentions in New Product Purchase Forecasting," *Econometric Models in Marketing*, 16, 11-28.

Juster, Thomas (1966), "Consumer Buying Intentions and Purchase Probability - An Experiment in Survey Design," *Journal of the American Statistical Association*, 61 (315), 658-96.

Kalwani, Manohar U. and Alvin J. Silk (1982), "On the Reliability and Predictive Validity of Purchase Intention Measures," *Marketing Science*, 1 (3), 243-86.

Kamakura, Wagner A. and Michel Wedel (1995), "Life Style Segmentation with Tailored Interviewing," *Journal of Marketing Research*, 32 (August): 308-17.

---, ---, Fernando de Rosa, and Jose A. Mazzon (2003), "Cross-selling Through Database Marketing: A Mixed Data Factor Analyzer for Data Augmentation and Prediction," *International Journal of Research in Marketing*, 20 (1), 45-65.

Klein, Gary and Beth W. Crandall (1995), "The Role of Mental Simulation in Problem Solving and Decision Making," in *Local Applications of the Ecological Approach To Human-Machine Systems*, Peter Hancock, ed. Volume 2: Resources For Ecological Psychology. Mahwah, NJ: Lawrence Erlbaum Associates, 324-58.

Koehler, Derek J. and Connie S. K. Poon (2006), "Self-Predictions Overweight Strength of Current Intentions," *Journal of Experimental Social Psychology*, 42, 517-24.

Lipshitz, Raanan and Orna Strauss (1997), "Coping With Uncertainty: A Naturalistic Decision Making Analysis," *Organizational Behavior and Human Decision Processes*, 69, 149-63.

Mittal, Vikas and Wagner Kamakura (2001), "Satisfaction, Repurchase Intent, and Repurchase Behavior: Investigating the Moderating Effect of Customer Characteristics," *Journal of Marketing Research*, 38 (February), 131-42.

Morrison, Donald G. (1979), "Purchase Intentions and Purchase Behavior," *Journal of Marketing*, 43 (2), 65-74.

Morwitz, Vicki G. (1994), "Predicting Whether and When Consumer Will Purchase From Intentions Data," *Working Paper*, Stern School of Business, New York University.

--- and David Schmittlein (1992), "Using Segmentation to Improve Sales Forecasts based on Purchase Intent: Which Intenders Actually Buy?" *Journal of Marketing Research*, 29 (November), 391-405.

---, Joel H. Steckel, and Alok Gupta (2007) "When Do Purchase Intentions Predict Sales?" *International Journal of Forecasting*, 23 (3), 347-64.

Naert, Philippe A. and Alain Bultez (1973), "Logically Consistent Market Share Models," *Journal of Marketing Research*, 10 (August), 334-40.

Osberg, Timothy M. and J. Sidney Shrauger (1986), "Self-Prediction: Exploring the Parameters of Accuracy," *Journal of Personality and Social Psychology*, 51 (5), 1044-57.

Raftery, Adrian E., Michael A. Newton, Jaya M. Satagopan, and Pavel N. Krivitsky (2007), "Estimating the Integrated Likelihood via Posterior Simulation Using the Harmonic Mean Identity," in *Bayesian Statistics 8*, Jose Bernardo, Susie Bayarri, James Berger, A. Philip Dawid, David Heckerman, Adrian Smith and Mike West, eds. Oxford University Press, 1-45.

Raudenbush, Stephen W. and Anthony S. Bryk (2002), *Hierarchical Linear Models: Applications and Data Analysis Methods*, 2nd ed. Newbury Park, CA: Sage.

Rogers, Everett M. (2003). *Diffusion of Innovations*, 5th ed. New York: Free Press.

Rossi, Peter E. and Greg M. Allenby (2003), "Bayesian Statistics and Marketing," *Marketing Science*, 22 (3), 304–28.

Rost, Jürgen (1985), "A Latent Class Model for Rating Data," *Psychometrika*, 50 (1), 37-49.

Salisbury, Linda Court and Fred M. Feinberg (2008), "Future Preference Uncertainty and Diversification: The Role of Temporal Stochastic Inflation," *Journal of Consumer Research*, 35 (2), 349-59.

Sheppard, Blair H., Jon Hartwick, and Paul R. Warshaw (1988), "The Theory of Reasoned Action: A Meta-analysis of Past Research with Recommendations for Modifications and Future Research." *Journal of Consumer Research*, 15, 325-43.

Silk, Alvin J. and Glen L. Urban (1978), "Pre-Test-Market Evaluation of New Packaged Goods - Model and Measurement Methodology," *Journal of Marketing Research*, 15 (2), 171-91.

Singer, Judith D. and John B. Willett (2003), *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence*. New York: Oxford.

Sun, Baohong and Vicki G. Morwitz (2009), "Stated Intentions and Purchase Behavior: A Unified Model," *Working Paper*.

Sutton, Stephen (1998), "Predicting and Explaining Intentions and Behavior: How Well Are We Doing?" *Journal of Applied Social Psychology*, 28 (15), 1317-38.

Trope, Yaacov and Nira Liberman (2003), "Temporal Construal," *Psychological Review*, 110 (3), 403-21.

Venkatesh, Viswanath, Likoebe M. Maruping, and Susan A. Brown (2006), "Role of Time in Self-Prediction of Behavior," *Organizational Behavior and Human Decision Processes*, 100, 160-76.

Wright, Peter and Mary Ann Kriewall (1980), "State-Of-Mind Effects on the Accuracy with Which Utility-Functions Predict Marketplace Choice," *Journal of Marketing Research*, 17 (3), 277-93.

Ying, Yuanping, Fred M. Feinberg, and Michel Wedel (2006), "Leveraging Missing Ratings to Improve Online Recommendation Systems," *Journal of Marketing Research*, 43 (August), 355-65.

Young, Martin R., Wayne S. DeSarbo, and Vicki G. Morwitz (1998), "The Stochastic Modeling of Purchase Intentions and Behavior," *Management Science*, 44 (2), 188-202.

Figure 1.
ACTUAL AND PREDICTED CUMULATIVE ADOPTION (%) FOR M1-M5

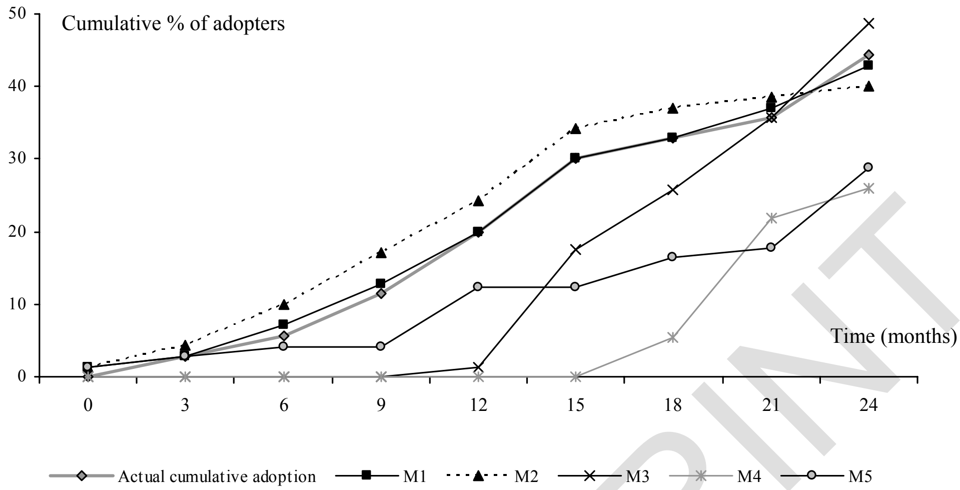

Table 1
EXAMPLE OF CUMULATIVE TIMED INTENT MEASURE

	I will not have bought one										*I will have bought one*
1 month from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
6 months from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1 year from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
1 ½ years from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
2 years from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
2 ½ years from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
3 years from now	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%

Table 2.
ANTECEDENTS OF TECHNOLOGY ADOPTION PROPOSED BY ROGERS' INNOVATION DIFFUSION THEORY

Antecedents	*Definition*	*Scale items*	*Cronbach's α*			*Relationship with steepness of slope*
			Mower	*Autoguidance*	*GPS*	
Perceived relative advantage[a]	the degree to which a technology is perceived to be better than the one it supersedes	Using [technology] in my life/work would increase my productivity If I use [technology] I will increase the quality of output Using [technology] increases my productivity	.81	.88	.93	+
Perceived compatibility	the degree to which a technology is perceived as consistent with the existing values, past experiences, and needs of potential adopters	Using [technology] is compatible with all aspects of my life/work I think that using [technology] fits well with the way I like to live/work Using [technology] fits into my life/work style	.74	.86	.92	+
Perceived complexity	the degree to which a technology is perceived as relatively difficult to understand and use	Using [technology] would take too much time from my normal activities Working with [technology] would be so complicated, it would be difficult to understand what is going on Using [technology] would involve too much time doing mechanical operations	.66	.70	.72	-
Perceived trialability	the degree to which a technology may be experimented with on a small scale	I can use [technology] on a trial basis to see what it can do It is easy to try out [technology] without a big commitment	.60	.65	.72	+
Perceived observability	the degree to which results of a technology are visible to others	I have no difficulty telling others about the results of using [technology] I believe I could communicate to others the consequences of using [technology] The results of using [technology] are apparent to me	.71	.64	.83	+

[a] *While the original operationalization by Rogers was relative, many studies measure it in a more absolute sense.*

Table 3.
ESTIMATION RESULTS FOR STUDIES 1, 2, AND 3

		Study 1: Advanced Mower			*Study 2: Autoguidance System*			*Study 3: GPS cell phone*		
		Estimate	Error	p-value	Estimate	Error	p-value	Estimate	Error	p-value
Baseline Model	Intercept	-1.8225	.0260	.0000	-1.6516	.0275	.0000	-1.9067	.0256	.0000
	Time	.5470	.0541	.0000	.4631	.0815	.0000	.7329	.0564	.0000
	log of integrated likelihood[a]	-5603.1			-4549.9			-5439.5		
Innovation Diffusion Theory Model	Intercept	-1.8216	.0259	.0000	-1.6459	.0274	.0000	-1.9046	.0257	.0000
	Perceived relative advantage	.1843	.0454	.0000	.3661	.0698	.0000	.1932	.0448	.0000
	Perceived compatibility	.0848	.0447	.0578	.5474	.0658	.0000	.3387	.0438	.0000
	Perceived complexity	-.1619	.0423	.0001	-.1243	.0660	.0598	-.0754	.0457	.0988
	Perceived trialability	.1157	.0429	.0069	.1543	.0620	.0129	.0847	.0447	.0582
	Perceived observability	.1718	.0453	.0002	.0686	.0618	.2668	-.0056	.0435	.8975
	Time	.5500	.0442	.0000	.4231	.0623	.0000	.7320	.0465	.0000
	log of integrated likelihood	-5579.1			-4516.6			-5391.8		

[a] *For calculation details, see Raftery et al. (2007). "Small" differences in the integrated likelihood (less than 3 or so) are not considered to indicate differences in model performance.*

Table 4.
THE ACCURACY OF THE CUMULATIVE ADOPTION LIKELIHOOD CURVES FOR PREDICTING WHETHER AND WHEN A TECHNOLOGY IS ACTUALLY ADOPTED DURING THE FIRST TWO YEARS FOLLOWING THE ORIGINAL SURVEY

	Intent Measures					*Intent Models* [a]			
	M1	*M2*	*M3*	*M4*	*M5*	*M6*	*M7*	*M8*	*M9*
	Cumulative adoption likelihood curves (N=70)	*Raw cumulative intent scores* (N=70)	*Observed heterogeneity only models* (N=70)	*Single-intent measure "Likelihood"* (N=73)	*Single-intent measure "Timing"* (N=73)	*Morrison model (cumulative timed intent)* (N=70)	*Morrison model (single intent)* (N=73)	*Bemmaor model (cumulative timed intent)* (N=70)	*Bemmaor model (single intent)* (N=73)
Individual-level adoption									
Whether									
Total Hit Rate	84.3%	78.6%*	61.4%***	60.2%****	57.5%****	78.6%*	67.1%**	78.6%*	67.1%**
Hit Rate for Non-Adopters (N=39/39)[b]	87.2%	84.6%^{ns}	61.5%***	82.1%^{ns}	76.9%^{ns}	79.5%^{ns}	76.9%^{ns}	89.7%^{ns}	76.9%^{ns}
Hit Rate for Adopters (N=31/34)[b]	80.6%	71.0%**	61.3%**	35.3%****	35.3%****	77.4%^{ns}	55.9%**	64.5%^{ns}	55.9%**
When									
MAD |Predicted – actual timing|	3.6 mos	4.2 mos^{ns}	8.0 mos****	5.5 mos***	8.8 mos****	N/A	N/A	N/A	N/A
RMSE	4.9 mos	5.7 mos	9.2 mos	6.9 mos	10.2 mos	N/A	N/A	N/A	N/A
Cumulative adoption									
Actual cumulative adoption	44.3%	44.3%	44.3%	46.6%	46.6%	44.3%	46.6%	44.3%	46.6%
Predicted cumulative adoption	42.9%	40.0%	48.6%	26.0%	28.8%	45.6%	44.9%	43.3%	45.0%
Predicted-actual cumulative adoption	-1.4%	-4.3%^{ns}	+4.3%^{ns}	-20.6%**	-17.8%**	+1.3%^{ns}	-1.7%^{ns}	-1.0%^{ns}	-1.6%^{ns}
Log-Likelihoods						-36.957	-45.077	-35.738	-45.096

* $p \leq .10$, ** $p \leq .05$, *** $p \leq .01$, **** $p \leq .001$, ^{ns} = not significant

Note. The predictive accuracy of M2-M9 is tested against the predictive accuracy of M1, the proposed cumulative timed intent measure.

[a] The models proposed by Morrison and Bemmaor do not predict *when* adoption will occur.

[b] Of the 70 participants for whom cumulative timed intent data were collected, 31 (39) participants did (not) adopt the technology. Of the 73 participants for whom single intent data were collected, 34 (39) participants did (not) adopt the technology.

Cumulative Timed Intent: A New Predictive Tool for Technology Adoption

KOERT VAN ITTERSUM and FRED M. FEINBERG

Web Appendix A: Model Estimation, Inference, and Convergence Details

As discussed extensively by Raudenbush and Bryk (2002), both the estimation and interpretation of multilevel models are sensitive to how input data are scaled. In accordance with their prescriptions, all covariates were mean-centered except for Time, which was converted to years (e.g., "one month" = 1/12, "six months" = ½, etc.), but with its zero point retained.

Inferences about all model parameters are made based on posterior distributions obtained via MCMC, using diffuse conjugate priors (flat for coefficients and inverted gamma for variance: $\beta_i \propto 1,\ \sigma^2 \propto \mathrm{IG}\left(10^{-3},\ 10^{-3}\right)$. Twenty thousand iterations were used for burn-in and 100,000 for inference, thinned by 10 to reduce autocorrelation. Convergence was monitored via trace plots, Gelman-Rubin and Geweke diagnostics, and by computing autocorrelation functions and effective sample sizes. Stationary distributions were apparent for all parameters; smoothed kernel densities were plotted for reported quantities and were used to compute posterior credible intervals and stated "standard errors" and "*p*-values" (which, though not strictly Bayesian, are nonetheless useful benchmarks). Model comparison and selection is carried out via the log of integrated (marginal) likelihood (Chib 1995), which suitably rewards parsimonious models and is well-suited to random coefficients frameworks.

The code and (partly synthetic, as required by our sponsoring firm) data are available in three ‘independent’ systems (MLWin, WinBUGS, SAS) at www.CumulativeTimedIntent.com.

Web Appendix B: Information on How Non-response Biases Were Examined

STUDY 1:

GOLF COURSE SUPERINTENDENTS DECIDING ON THE ADOPTION OF AN ADVANCED MOWER

To examine potential non-response biases among the approached golf course superintendents, a telephonic follow-up with 26 non-respondents was conducted. Of these, 69.2% indicated that they did not receive a questionnaire and 26.9% mentioned they were too busy at the time. The size of the golf course may be to blame. First, larger companies collect, filter, and sort mail centrally, oftentimes not forwarding promotional materials and survey requests. Second, superintendents of larger courses tend to have more responsibilities and therefore less time than superintendents of smaller golf courses. Comparing the size of the golf courses for non-respondents and respondents revealed no differences (17.3 vs. 19.8 holes; $t = 1.59$, $p > .10$).

We also compared early and late respondents and found that both work at equally sized courses (19.6 vs. 20.1 holes; $t = .35$, $p > .10$), have the same amount of industry experience (20.7 vs. 21.6 years; $t = .64$, $p > .10$), and have similar perceptions of the technology: relative advantage (-.07 vs. .07; $t = -.92$, $p > .10$), compatibility (.04 vs. -.04; $t = .52$, $p > .10$), complexity (.06 vs. -.06; $t = .85$, $p > .10$), trialability (-.05 vs. .05; $t = -.78$, $p > .10$), and observability (.06 vs. -.06; $t = .78$, $p > .10$). Non-response biases therefore appear marginal at best.

STUDY 2:

FARM OPERATORS DECIDING ON THE ADOPTION OF AN AUTOGUIDANCE SYSTEM

To examine potential non-response biases among the approached farm operators, a telephonic follow-up with 24 non-respondents was conducted: 54.2% mentioned that they did not receive a questionnaire; 41.6% indicated that they were too busy. A comparison in farm size, using the number of acres of corn grown as a proxy (the only data available for both groups), between non-respondents and respondents revealed no significant differences (913 vs. 882 acres; $t = .20$, $p > .10$).

A comparison of early and late respondents revealed no differences in total farm size (2,296 vs. 2,530 acres; $t = .47$, $p > .10$). Furthermore, they share the same amount of industry experience (34.9 vs. 35.0 years; $t = .04$, $p > .10$). Finally, no differences in the perceptions of the technology were found: relative advantage (.08 vs. -.08; $t = 1.12$, $p > .10$), compatibility (.02 vs. -.02; $t = .34$, $p > .10$), complexity (-.04 vs. .04; $t = -.46$, $p > .10$), trialability (-.05 vs. .05; $t = -.14$, $p > .10$), and observability (-.05 vs. .05; $t = -.67$, $p > .10$). Non-response biases again appear marginal at best.

STUDY 3:

STUDENTS DECIDING ON THE ADOPTION OF A GPS CELL PHONE

A variety of checks were conducted to examine if participants in the follow-up study differed from those who decided not to participate in the first place. The results in Table B-1 suggest that non-response biases are at most modest.

Table B-1.
COMPARISON OF RESEARCH CONDITIONS TO CHECK FOR POTENTIAL NON-RESPONSE BIASES IN STUDY 3

	Single-intent measure (N=178)		Cumulative timed intent measure (N=176)		F-value (1, 350)		
	Intent only[a] (N=105)	Intent & Behavior (N=73)	Intent only[a] (N=106)	Intent & Behavior (N=70)	Single vs. CTI measures	Intent only vs. Intent & Behavior	Interaction
Gender (% females)[b]	41.0%	38.6%	41.5%	42.9%	.07 ^{ns}	.03 ^{ns}	.40 ^{ns}
Age	20.3	20.5	20.4	20.5	.01 ^{ns}	.77 ^{ns}	.12 ^{ns}
	(1.50)	(1.36)	(1.14)	(1.16)			
Knowledge	-.02	.06	-.09	.12	.05 ^{ns}	1.99 ^{ns}	.40 ^{ns}
	(.99)	(1.00)	(1.07)	(.91)			
Experience	.04	.08	-.05	-.06	1.10 ^{ns}	.03 ^{ns}	.06 ^{ns}
	(1.03)	(1.02)	(1.02)	(.92)			
Relative Advantage	-.09	.08	-.08	.06	.25 ^{ns}	2.61 ^{ns}	.38 ^{ns}
	(.95)	(.92)	(1.08)	(1.01)			
Compatibility	.02	.08	-.08	.07	.24 ^{ns}	1.70 ^{ns}	1.87 ^{ns}
	(1.03)	(1.01)	(.95)	(.99)			
Complexity	.05	-.04	.08	-.10	.14 ^{ns}	2.36 ^{ns}	.40 ^{ns}
	(1.03)	(.88)	(1.03)	(1.02)			
Trialability	.03	.03	-.08	-.04	.26 ^{ns}	.27 ^{ns}	.28 ^{ns}
	(1.02)	(.98)	(1.01)	(1.00)			
Observability	.04	.01	-.05	.01	.20 ^{ns}	.02 ^{ns}	.19 ^{ns}
	(1.08)	(1.12)	(.87)	(.95)			
% actual adopters [c]		46.6%		44.3%	.08 ^{ns}		
Timing of actual adoption (months after survey) [d]		15.7		13.3	2.59 ^{ns}		

[a] This column represents the participants who only participated in the first part of the study during which the intent data were collected, as opposed to those who are presented under the "Intent & Behavior" column who additionally participated in the follow-up during which the behavioral data were collected.
[b] Wald statistics reported. [b] Chi-square analyses. [c] F(1, 63).