

Emergent Consequences: Unexpected Behaviors in a Simple Model to Support Innovation Adoption, Planning, and Evaluation

H. Van Dyke Parunak¹, Jonathan A. Morell²

¹Soar Technology, Inc.
3600 Green Court, Suite 600
Ann Arbor, MI 48105
van.parunak@soartech.com

²Fulcrum Corporation
4075 Wilson Blvd, Suite 550
Arlington, VA 22203
jmorell@fulcrum-corp.com
jmorell@jmorell.com

Abstract. Many proven clinical interventions that have been tested in carefully controlled field settings have not been widely adopted. We study an agent-based model of innovation adoption. Traditional statistical models average out individual variation in a population. In contrast, agent-based models focus on individual behavior. Because of this difference in perspective, an agent based model can yield insight into emergent system behavior that would not otherwise be visible. We begin with a traditional logic of innovation, and cast it in an agent-based form. The model shows behavior that is relevant to successful implementation, but that is not predictable using the traditional perspective. In particular, users move continuously in a space defined by degree of adoption and confidence. High adopters bifurcate between high and low confidence in the innovation, and move between these groups over time without converging. Based on these observations, we suggest a research agenda to integrate this approach into traditional evaluation methods.

Keywords: Agent-based models, emergent behavior, innovation adoption, clinical evaluation

1 Introduction

There is an extensive literature on factors that facilitate the adoption of innovations [8]. Two research traditions underlie this work. (1) Traditional statistical analysis applies familiar statistical methods to test hypotheses concerning the implementation and continued use of innovations in social and organizational settings. (2) Case studies perform in-depth observation and analysis on innovation cases.

Collectively, this research has successfully revealed which innovations will succeed. Despite this research, however, funders of research in health, mental health,

substance abuse, and social betterment face a dilemma. Even with decades of funding, many interventions that have proved to be successful in carefully controlled field settings have not been widely adopted. Why is this so, and what can be done about it? Traditional models of innovation adoption have not helped.

Our research applies systems thinking to innovation adoption in clinical settings, and is part of a trend in applying systems methodologies to study health and public health interventions [2]. In particular, we apply agent-based modeling to address two questions. First, does an agent-based approach reveal anything about innovation adoption that traditional research does not? Second, can agent-based models add value to traditional efforts to evaluate exercises in innovation adoption?

Agent-based models might add value because they can be based on knowledge from traditional research, but they look at phenomena in a way that is fundamentally different from statistical methods. [4,5]. Traditional statistics obviates the contribution of individual variation in a population. In contrast, the agent-based view is based on the interaction of individual behaviors. We seek to use this interaction to explore emergent system-level behavior not accessible to a mean-field model. (We do not claim to predict the detailed behavior of an individual human therapist.)

This paper presents a laboratory-based proof of concept that agent-based modeling can reveal unanticipated behavior that may affect adoption. Section 2 introduces a simple model of innovation adoption. Section 3 documents our agent-based implementation of this model, and Section 4 describes its behavior. Section 5 discusses the research implications of the observed behavior from two perspectives: the adequacy of our underlying logic model, and the implications for clinical management of innovation, including introduction, motivation, and evaluation.

2 A Logical Model of the Adoption of Innovations

Fig. 1 illustrates the type of program theory that might characterize evaluation of a program designed to facilitate the adoption of a new, proven best practice in a clinical setting. This model depicts a logic in which therapists' use of a new treatment is based directly on his or her confidence that the treatment will work, combined with organizational support for use. Confidence in turn is a function of the influence of a therapists' colleagues, quality of training in the new treatment, and the therapist's conclusion about success or failure in clinical outcome. Colors reflect the psychological (salmon), organizational (blue), and social psychological (yellow) factors that combine to facilitate use of the innovation. As with all evaluation "logic models," Fig. 1 is based on input from the stakeholders involved in planning the program, and on relevant research findings in health care and in many other domains. That research makes it clear that

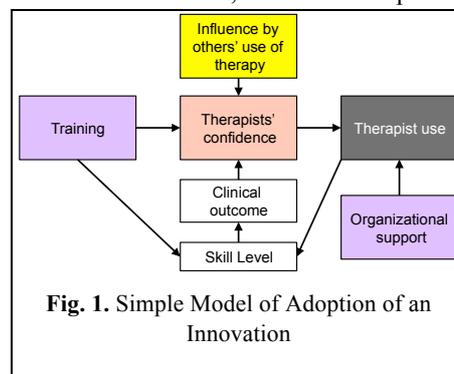


Fig. 1. Simple Model of Adoption of an Innovation

successful innovation is a function of the characteristics of: 1) adopters, 2) the innovation, and 3) the setting in which adoption takes place [1,6]. Building models such as Fig. 1 is standard practice in the field of evaluation, a practice with a long history of success in leading to insightful understanding of program behavior and outcome. Such models help planners articulate their (often implicit) assumptions about what they are doing and why; provide a framework for drawing on the results of prior research; identifying the constructs that need to be operationalized and measured; and assuring that all parties involved have a common understanding of what the program will be and how it will be evaluated.¹

This approach has limitations. The first is epistemological, i.e., innovation efforts based on theories such as these have provided disappointing results. Clearly, they do not provide the knowledge needed to achieve the desired level of implementation of best practices. Second, they are limited in their capacity to help stakeholders consider the likely behavior of the programs they are implementing. The work reported here probes the value of improving the application of traditional logic models by integrating an agent-based simulation into the development of those models.

3 An Agent-Based Implementation

We implement this model in NetLogo [7], a framework for multi-agent modeling that is widely used in the social sciences. Each therapist t_i is a software agent with six characteristics. We represent them as functions over the therapist index. Where there is no ambiguity about the therapist in question, we refer simply to the function name.

Three agent characteristics are assigned when the agent is initialized.

- **Group membership $G(t_i)$ = set of therapists:** Therapists exchange their experiences only within their group.
- **Adaptability $a(t_i) \in [0, 1]$:** Indicates how susceptible a therapist is to the opinions of her peers.
- **Training $t(t_i) \in [0, 1]$:** Models the quality of the training received by the therapists in how to apply the innovative best practice. Each therapist is assigned a level of training, which can modulate his initial confidence and his skill in applying the innovation. Training is a number in a specified range, and can be assigned uniformly randomly over all therapists, randomly chosen between the maximum and minimum level, or held constant within agent groups.

Three characteristics vary during the therapist's career:

- **Skill $s(t_i) \in [0, 1]$:** This variable is initialized to the therapist's training level, then increases as the therapist exercises the innovation and decays while she does not

¹ In a real evaluation this model would be more detailed. For instance "confidence" could be modulated by a random variable representing clinical judgment about the value of the new treatment with respect to a specific client. The level of detail presented here is more than adequate for our objective of showing how modeling can yield insights not available through the static model alone.

exercise it. The higher the therapist's skill, the more successful an exercise of the innovation is likely to be.

- **Confidence $c(t_i) \in [0, 1]$** : The central variable in Fig 1 is confidence, driven by the therapist's training, clinical experience, and input from colleagues in her group.
- **Adoption $d(t_i) \in [0, 1]$** : This variable records the percentage of opportunities on which the therapist decides to use the innovation.

In addition, 11 variables govern the entire model:

- the total number of therapists
- the number of groups
- **Organizational Support $o \in [0, 1]$** : the degree to which the therapists' employer encourages the innovation
- **Skill Success $\tau \in [0, 10]$** : how skill influences success
- **Maximum Success $m \in [0, 1]$** : the maximum success rate for this innovation
- **Success impact $i_s \in [0, 1]$** : how much a successful outcome increases c
- **Failure impact $i_f \in [0, 1]$** : how much an unsuccessful outcome decreases c
- **Skill increment $k_i \in [0, 1]$** : how much exercising an innovation increases s
- **Skill decrement $k_d \in [0, 1]$** : how much s decreases an innovation is not exercised
- **History length $h_t \in [1, 20]$** : how many rounds contribute to adoption level
- **History decay $h_d \in [0, 1]$** : the rate at which history decays

In the real world, adoption of innovation depends on more variables than these. However, these are adequate to capture our logic model (Fig. 1), and even this simplified model generates behavior that we might not have anticipated from the logic model.

At each cycle of the model, each therapist takes all four of the following actions. Then the next therapist takes them, and so forth. In incrementing or decrementing bounded variables, we modulate the impact by the amount of range left to move in the direction of the change, consistent with commonly observed ceiling and floor effects.

1. Update confidence level based on peer interaction. The therapist's confidence is moved toward that of her peers by a fraction of how far it is from theirs:

$$\gamma = \alpha * \left(\frac{\sum_{t_i \in G} c(t_i)}{|G|} - c \right) / 2 \quad (\text{Equation 1})$$

If $\gamma > 0$ (indicating that others in the group are, on average, more confident than ego), c is replaced by $c + \gamma(1 - c)$. If $\gamma < 0$, c is replaced by $c(1 + \gamma) < c$.

2. The therapist decides whether to adopt the innovation on her next engagement, by computing a probability of adoption p_d and then comparing it to a uniformly distributed random number in $[0, 1]$. The model provides two ways to compute p_d : either as the average of c and o , or as a logistic function of their sum:

$$p_d = 1 / (1 + e^{-s(c+o-1)}) \quad (\text{Equation 2})$$

s is steepness. For $c + o = 1$, $p_d = 0.5$, and ranges from 0 to 1 as the sum of these values moves from 0 to 2. The logistic models widely-observed saturation effects.

3. If the therapist decides to adopt the innovation:
 - (a) Her adoption level increases. We provide two ways to monitor adoption. An array remembers whether (1) or not (0) the therapist used the innovation on the last h_i trial, and we compute h as the average of this array. Alternatively, we increment d by $h_d(1 - d)$, which realizes an exponential weighted adoption rate.
 - (b) Her skill is incremented by $k_i(1 - s)$.
 - (c) The intervention succeeds with probability $x(1 - e^{-\tau*s})$.
 - (d) Optionally, her confidence is incremented or decremented depending on whether the intervention was successful or not. On success, confidence is incremented by $i_s(1 - c)$. On failure, confidence is decremented by $i_f c$.
4. If the therapist decides not to adopt the innovation:
 - (a) Her adoption level decreases, by recording a 0 in the history or decrementing d by $h_d d$.
 - (b) Her skill is decremented by multiplying it by k_d .

Our dependent variables are adoption and confidence of the therapists, and evolve as the model executes. We do not track outcome, since it is not an emergent effect of therapist interactions, but is determined directly by individual skill and training level.

Our implementation is much more detailed and specific than the logical model of Fig 1, and another realization of that model might behave very differently than what we report in the next section. We have posted the full model online,² and encourage other researchers to explore variations of our implementation.

4 Model Behavior

At a gross level, confidence and adoption behave as one might expect. For example,

- Increased training raises the highest level of confidence attained;
- Increased organizational support raises both the upper and the lower limits of the adoption observed;
- High adaptability (susceptibility to colleagues' opinions) narrows the distribution of confidence across therapists.

But some other behaviors of the model are counterintuitive.

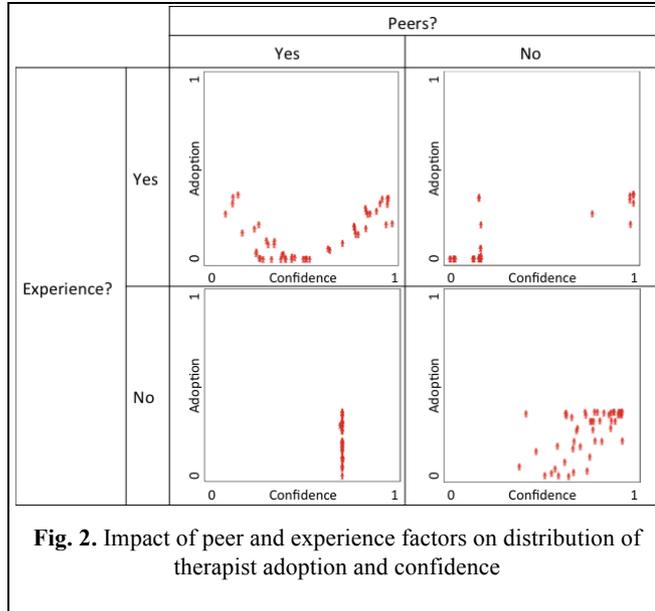
A therapist's confidence is initialized by training level, and then changes in only two ways: by attraction to the opinion of peers, and by experience with the innovation. Fig 2 shows how these effects interact. In these runs, 50 therapists are all in a single group, with training assigned uniformly randomly in $[0, 1]$. While these shapes reflect convergence of the model, individual therapists do not settle down to a single location in the state space of Adoption vs. Confidence, but continue to move.

In the bottom row, therapists' experience with the innovation does not change their confidence. At the right, with no peer influence ($\alpha = 0$ for all therapists), therapists are distributed continuously over both adoption and confidence, with a reasonable in-

² <http://www.abcresearch.org/abcresearch.org/models/InnovationAdoption.nlogo>

crease in adoption with confidence. Therapists change adoption over time, but remain at their initial confidence.

At the lower left, therapists still do not change confidence with experience, but adopt the confidence of their peers (to varying degrees). Eventually all therapists converge to the same confidence, within which they continue to change in adoption. Again, this result is not surprising.



The top row is less intuitive. At the right, confidence changes by experience but not by peer interaction. Confidence and adoption are still correlated, but the distribution is no longer continuous. Instead, the therapists bifurcate in confidence, moving to extremes with a wide unpopulated area in between. While individual therapists continue to move in both dimensions, the bifurcation in confidence remains. Experimentation with the model shows that the bifurcation is the result of fairly high levels of i_f and i_s (0.5 or higher). When the impact of success or failure on confidence is high, a few successive successes or failures quickly drive the therapist to an extreme in confidence. This observation is relevant to training therapists in understanding the dynamics of best practices that demonstrate overall positive impact, but varying consequences for individual patients. We can understand the dynamic, but we would be unlikely to recognize the potential for this behavior from Fig 1 alone.

At the top left, both peers and experience modulate confidence. High levels of confidence show the same bifurcation as in the previous condition. Again, the bifurcation is driven by higher levels of i_f and i_s . However, low adopters tend to have middling levels of confidence. Because they seldom try the innovation, the main influence on their confidence is peer pressure, which pushes them toward intermediate confidence.

Interestingly, while confidence and adoption are positively correlated for therapists on the right-hand arm of the distribution, the relative shift in impact of outcome and peer confidence leads to a *negative* correlation on the left-hand arm. Detailed observation of the agents as the model executes shows that they tend to move counter-clockwise around the space. Consider an agent who starts on the right-hand arm, where adoption and confidence are correlated.

- The agent may move back and forth along this arm, as the peer effect pulls confidence lower, decreasing adoption, and as success pushes confidence higher.
- At some point, a failed adoption cuts the agent's confidence dramatically (because of the high value of i_j), moving the agent to the left-hand arm.
- Due to lower confidence, adoption drops. At the same time, peer pressure pulls confidence higher, moving the agent down the right-hand arm toward the center.
- At some point the agent's confidence grows enough that it resumes the therapy, moving higher. If successful, it climbs the right-hand arm; if not, it moves back to the left-hand arm.

This agent-based model shows that the logic model of Fig. 1 implies dynamics that would not be anticipated by simple examination, and suggests research hypotheses and practical innovations inaccessible from the static model alone.

5 Implications

Committing to action is a leap of faith that commits material resources and intellectual capital, incurs opportunity costs, and forces planners and evaluators to confront unintended consequences [3]. In theory, innovative programs can be changed if they do not work, and program evaluation provides data so that rational decisions can be made about program improvement. In practice, commitment makes change difficult. Therefore it is important to do the most rigorous possible planning, and to conduct evaluations whose results will be as potent as possible. We have shown that one way to increase rigor and potency is to reformulate the kind of model commonly used in planning and evaluation practice as an agent-based simulation.

This justification for modeling acknowledges that no matter how much research we draw upon, and no matter how well we conduct group processes with experts, we may miss important insights. Models are useful for catching some of what we have missed, but another possibility may be in play. We may have erred not because we failed to understand what we have, but because what we need to know is not yet known.

Consider two dimensions. First, neither previous research nor expert judgment has considered all the variables that matter, those that could account for unexplained variance in observed outcomes, or that could lead to consideration of variables that account for some of the unexplained variance. Our program theory may either be incorrect, or not strong enough to help us make a big enough practical difference. For instance, existing theory about "adoption" does not consider the built ecology in which therapists work, though traffic patterns affect interactions between therapists and their colleagues, their patients, and their patients' families. Second, traditional theorizing about social phenomena does not consider the importance of emergent dynamics. The *only* way to consider relevant dynamics is to employ an agent based model. All of these possibilities are illustrated in our findings about adoption and confidence.

Fig. 1 does not explicitly model peer interaction. Perhaps it should have been considered, given what we know about the importance of boundary spanners, mavens, leaders, and other such roles. But hindsight is seductive. Given the results of the simulation, *of course* planners should have considered role relationships. We do not know

whether paying attention to these roles would affect the particular group behavior observed, but the results of the model might have led planners to consider peer dynamics more than they otherwise would have. In any case, who could have foreseen that the model would lead people to think of peer behavior? For all anyone knew, the model might have yielded surprising results with respect to clinical judgment of success, or critical thresholds of effectiveness between the old and new treatment, or any one many behaviors that the model might have unearthed. Agent-based models are valuable for planning and evaluation precisely because these dynamics are not easily predicted, despite the best expert judgment and extensive literature reviews.

Our results demonstrate that a simple logic model may have unexpected dynamics. The degree to which our results (and the underlying model) reflect the behavior of real therapists is an important question that lies beyond this paper, but should be confirmed by a broader survey of the empirical literature and clinical validation.

We have presented this model in a clinical context. However, discussions with practitioners in other areas (e.g., educational reform) suggest that the same basic pattern of influences will be widespread, and our research agenda includes studying the degree to which this model can in fact be generalized across disciplines.

6 References

- [1]D.M. Berwick. Disseminating Innovations in Health Care. *Journal of the American Medical Association*, 289(15):1969-1975, 2003.
- [2]P.L. Mabry, B. Milstein, et al. Opening a Window on Systems Science Research in Health Promotion and Public Health. *Health Education and Behavior*, 40(1S):5S-8S, 2013.
- [3]J.A. Morell. *Evaluation in the Face of Uncertainty: Anticipating Surprise and Responding to the Inevitable*. New York, NY, Guilford Press, 2010.
- [4]J.A. Morell, R. Hilscher, et al. Integrating Evaluation and Agent-Based Modeling: Rationale and an Example for Adopting Evidence-Based Practices. *Journal of MultiDisciplinary Evaluation*, 6(14):35-37, 2010.
- [5]H.V.D. Parunak, R. Savit, et al. Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide. In N. Gilbert, R. Conte, and J. S. Sichman, Editors, *Multi-Agent Systems and Agent-Based Simulation, First International Workshop, LNCS*, pages 10-25. Springer, Berlin, Germany, 1998.
- [6]E. Rogers, M. *Diffusion of Innovations*. 5th ed. New York, NY, Free Press, 2003.
- [7]U. Wilensky. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL, 1999. <http://ccl.northwestern.edu/netlogo>.
- [8]J.P. Wisdom, K.H.B. Chor, et al. Innovation Adoption: A Review of Theories and Constructs. *Administration and Policy in Mental Health and Mental Health Services Research*:(in press), 2013.