Draft 9/11/08

Personal versus Situational Dynamics: Implications of Barry Richmond's Models of Classic Experiments in Social Psychology[*]

James K. Doyle, Khalid Saeed, Jeanine Skorinko
Department of Social Science and Policy Studies
Worcester Polytechnic Institute

ABSTRACT

There is a long-standing debate in the field of social psychology as to which is the primary determinant of behavior, the situation or system in which people act or the personalities of the role players. Psychologists have long studied this problem with controlled experiments on human subjects, and have now come to a general resolution of the debate. However, the field of psychology still lacks an efficient method for teasing apart the relative contributions of personal and situational variables in applied domains. An alternative to human subjects experiments is to employ system dynamics models of role systems, as was demonstrated by Barry Richmond when he attempted to model two classic experiments in social psychology: the Milgram and Stanford Prison experiments. In this paper, we replicate and discuss Barry Richmond's models to present them to a new audience. In addition, we use the models as a springboard to explore the relationship between social psychology and system dynamics and the potential for useful collaboration between the two fields.

1. Introduction

The disciplines of system dynamics and social psychology would appear to have a natural affinity for each other. Despite their widely varying methodologies (computer simulation versus experimentation on human subjects), both disciplines offer similar explanations for human decision making and behavior. In system dynamics modelers constantly make the claim that "structure drives behavior," meaning that the physical structure of the system, how the stocks

[*] Based on an essay written by the authors for a volume in memory of Barry Richmond to be published by isee systems, inc.

and flows are interconnected, is often a stronger determinant of system behavior over time than the specific values of parameters. Similarly, in social psychology experiments often reveal that the situation or role structure in which behavior occurs is a stronger predictor of cognition and behavior than the precise makeup of an individual's character, personality, or habits of thought. Thus, both fields conceive of the individual decision maker as being embedded in an external system whose properties can have a strong influence over behavioral outcomes.

But despite the similarity in theoretical approach, there has historically never been much interaction between the two fields. Few psychologists have made the effort to learn the system dynamics approach and apply it to the understanding of social behavior. Social psychologists, in fact, only rarely use any form of computer simulation to develop and test their theories. At the same time, few system dynamicists have attempted to apply the system dynamics method to such socially relevant topics as the nature of obedience, group dynamics, persuasion, aggression, or prejudice, that is, to the standard topics of social psychological inquiry. When such topics are included in system dynamics models, they are often done so without reference to the large experimental literature in social psychology.

There has, however, been one particularly notable exception to this lack of synergy between the two fields: Barry Richmond's attempt to replicate classic experiments in social psychology using system dynamics. Barry joined the system dynamics doctoral program at MIT in 1975 and, along with his fellow doctoral students in the Sloan School of Management (including one of the present authors), took a mandatory seminar in behavioral science. A particularly fascinating part of this seminar was the presentations by Professor Edgar Schein of two highly influential and somewhat disturbing social psychological studies: Stanley Milgram's obedience experiment (Milgram, 1963, 1974) and Philip Zimbardo's prison simulation study, known as the Stanford Prison Experiment (Haney et al, 1973; Haney and Zimbardo, 1976). [Readers interested in learning more about the Milgram obedience studies would do well to begin with the work of Thomas Blass, a Milgram scholar who maintains a website on Milgram and his work at http://www.stanleymilgram.com and who has published the definitive biography of Milgram, *The Man Who Shocked the World* (Blass, 2004). Philip Zimbardo maintains a website on the Stanford Prison Experiment at http://www.prisonexp.org/ and has recently for the first time published his personal recollections of the experiment in his book, *The Lucifer Effect* (Zimbardo, 2007a).] These two studies are often cited as the prototypical examples of a widespread phenomenon in social behavior that psychologists refer to as "the power of the situation," that is, the ability of circumstantial or situational factors external to the individual to overwhelm the influence of individual differences and personality factors, resulting in identical behavior from very different people. Both experiments placed their human subjects in a difficult situation and set up a conflict between the values of the individual subjects and the demands of the situation. In the Milgram obedience studies, subjects were asked to obey the instructions of the experimenter to deliver painful electric shocks to a fellow subject. In the Zimbardo experiment, subjects were asked to play the role of prison guard in a simulated prison environment. Which would matter more? The personal traits of the individuals or the situational pressures created by the roles they were asked to assume? Would the subjects follow their individual values or adapt their behavior to fit the demands of the experimenter or the role they were asked to play? The surprising and disturbing result in both cases was that a majority of subjects took actions to harm or denigrate others that prior to the experiment they would never have dreamed they were capable of doing.

Barry, who also held a Master's degree in psychology, walked away from these presentations with what, for the time, was a very radical idea, namely, that the experiments could conceivably be conducted on a computer using system dynamics modeling rather than employing real human subjects. Despite being new to the field, Barry took on the challenge of building a system dynamics model of the Milgram experiment, and produced a fascinating technical report describing the model and his experimentation with it (Richmond, 1977). Later in his career, with more experience under his belt, Barry also built a more elegant model of the Zimbardo experiment. This model was never fully written up but was made available to the authors with the kind assistance of *isee systems*. It is briefly described in a paper published by Barry in the *System Dynamics Review* (Richmond, 1993).

This paper attempts to bring these two previously unpublished social psychological models to light and evaluate their potential for guiding future research and education in socialpsychology.Toward that end the models have been translated from the original Dynamo to iThink, updated according to current system dynamic practice, and modified to allow easier experimentation with parameters. For both the Milgram experiment and the Stanford Prison

Experiment, an introduction, focusing on the experiences of the subjects involved, is provided for readers who are unfamiliar with the details of the studies. Each introduction is followed by a presentation of Barry Richmond's system dynamic model of the experiment, along with sample computer simulation runs. Finally, in the concluding sections the models are used as a springboard to address several important questions at the interface of system dynamics and social psychology of current relevance, including: Why has there historically been such little interaction between the two fields? What would be the benefits of more interaction in the future? How suitable is the system dynamics methodology for exploring social psychological topics, compared to other modeling paradigms?

2. The Milgram Experiment

After World War II, the world, shocked by the atrocities committed by Nazi soldiers onto the German Jews, wondered what could make the Nazis such an evil group of people. In the midst of this outrage against Nazi Germans, Stanley Milgram, interested in the influence of the situation on individuals, wondered if it was that the Nazi's were naturally evil people or whether anyone, given the right situation, might commit similar acts. It was from this question that he designed his now-famous obedience experiments.

In the first experiment, Milgram advertised in the local New Haven, Connecticut newspaper, inviting individuals to be paid to participate in an experiment on learning. Interested participants arrived at the Yale University lab, and were greeted by the Experimenter (dressed in a laboratory coat), and met another "participant", Mr. Wallace. Unbeknowst to the participant, Mr. Wallace actually worked for the experimenter and was only pretending to be another participant (called a confederate).

After the introductions, the Experimenter informed both the participant and Mr. Wallace that the study is about learning and punishment. The experimenter also informed them that one of them would be randomly assigned to be the Teacher and the other the Learner. While the participant and Mr. Wallace drew sticks to see who would play each role, the experiment was set up such that the participant always ended up being the Teacher and Mr. Wallace ended up as the Learner.

The Experimenter then explained that the Teacher would be teaching the Learner word-pairs, and the Learner would need to accurately recall the word pairs or they would be punished via electric shocks. With each incorrect response, the intensity of the shocks increased by 15 volts.

After explaining the tasks, the Learner took a seat, and the Experimenter strapped electrodes onto the Learner. While being strapped in, the Learner informed both the Experimenter and the Teacher that he had a heart condition, and the Experimenter assured the Learner that while he would feel pain from the shocks that his heart would experience no adverse effects. To ensure that the Teacher understood the intensity of the shocks (and to provide the Teacher with a mental cue of the pain the Learner would be withstanding), the Experimenter gave the Teacher a sample low-intensity shock (that was still quite painful).

After getting the Learner set up in their room, the Experimenter showed the Teacher into another room to see the "shock generator" (Kassin, Fein, & Markus, 2008). The generator had switches that increased in 15 volt increments from 15-450 volts, and the highest voltage (450 volts) was labeled "XXX". The Teacher was reminded to increase the shocks by 15 volts each time the Learner incorrectly answers a word-pair. Then the experiment began.

At first, the experiment went smoothly, as the Learner recalled the word-pairs correctly, and the Teacher probably felt little if any anxiety. But as time went on, the Learner's performance declined, and the shock intensity needed to be increased, and the anxiety felt by the participant (the Teacher) most likely increased as well. By observing the behavior of the Teacher and measuring the maximum voltage that the Teacher would give the Learner before stopping the experiment, Milgram hoped the experiment would enlighten him on how far would be too far to go when administering the shocks, or when might the anxiety of doing something harmful to another person override obeying an authority figure.

To aid in the reality of the situation (and make the harm on the other person more apparent), the Teacher began to hear reactions from the Learner based on the shock's intensity. When the voltage ranged from 75-135, the Teacher heard the Learner yell, "Ugh". When the intensity increased to 150 volts, the Learner screamed "Ugh! Experimenter! That's all. Get me out of here. I told you I had heart trouble. My heart's starting to bother me now. Get me out of here, please. My heart's starting to bother me. I refuse to go on. Let me out" (Milgram, 1963).

And, at 330 volts, the Learner yelled, “Let me out of here” repeatedly. After 330 volts, the Learner went silent. Thus, the Teacher most likely began to feel two types of increasing anxiety: one over the fact they had to administer increasingly powerful shocks to the Learner, and two over the reactions and screams given by the Learner to the increased intensity of the shocks. It is important to point out that after the Teacher and the Experimenter left the Learner in the room, the Learner (Mr. Wallace), removed himself from the shocking machine and played a tape recording of his reactions given the voltage being administered.

Often after the first few reactions by the Learner and most likely due to increased felt anxiety due to both the reactions and the increasing intensity of shocks needing to be administered, the Teacher would look to the Experimenter for guidance or advice on what to do, especially after the Learner asked to be let out of the experiment (at 150 volts). If the Teacher looked to the Experimenter for guidance, the Experimenter responded in one of four scripted responses: “Please continue”, “The experiment requires you continue”, “It is absolutely essential that you continue”, or “You have no other choice, you must go on”. The Experimenter did nothing more than make one of those four statements. Yet, the insistence by the Experimenter to continue the experiment provided the participant (the Teacher) with one additional avenue for feeling increased anxiety. The Teacher needed to weigh the anxiety they felt for simply administering painful shocks, the anxiety they felt for the pain they were clearly inducing on the Learner, the anxiety they felt for obeying an authority figure (the Experimenter) who stated that no real harm was being done, and the anxiety they felt for not letting the Experimenter down and helping them conduct their experiment.

What happened? Twenty-six out of 40 men (or 65 percent of the participants) “shocked” the Learner to the maximum 450 volts. Were there any limits to these findings? Milgram (1974) ran several replications to this original study. One key factor that did not play a role was gender. In one subsequent study he found that sixty-five percent of the female participants shocked to the maximum voltage—replicating the results when using male participants.

But, several other factors did result in lowered obedience by participants. First, the proximity to the victim slightly reduced compliance. When the Learner (or victim of the shocks) was simply in the same room as the Teacher, 40% of the participants gave the maximum shock. And, if the

Teacher had to physically touch the Learner, and place the Learner's hand on the shock plate, then 30% of the participants continued to the maximum 450 volts. Felt anxiety may explain these results. The more the Teacher had to face the victim (the Learner) and the pain they were administering, then the more anxiety they felt and the more likely they were to want to witness the harm they were inflicting—regardless of the presence of the Experimenter and anxiety felt for trusting and keeping the Experimenter happy.

In addition, factors relating to the authority figure also contributed to reductions in obedience. First, the presence of an authority figure played a key role. If the Experimenter left the room to leave an ordinary member of their staff to run the experiment, then only 20% of the participants shocked all the way. Likewise, the proximity of the Experimenter played a key role as well. If the Experimenter administered the experiment from a different room, then again only 20% of the participants obeyed and shocked to 450 volts. And, if the Experimenter never gave any of the four prods (or commands described above), than very few (less than5%) of the participants administered the maximum shock.   Again, felt anxiety may help explain these results, as less anxiety would be felt if the participant did not have to face the disappointed Experimenter.

And, the presence of other dissenting/rebelling participants led to increased rebellion on behalf of the participant. If the experiment was run such that confederates pretending to be other participants refused to continue shocking, then less than 10% of the participants would continue shocking (especially to 450 volts) as well. In this case, the social influence of seeing someone else refuse to continue provided the participant with an excuse to stop for both normative and informational reasons. First, they may not have wanted to be the only ones not to stop running the study. And, second, they may have believed that the other participants who were refusing to continue had a reason to do so, or knew something about the situation that they did not know.

In conclusion, Milgram's obedience studies have had important implications for a variety of social science disciplines and remain "the most famous series of experiments in social psychology" (Brown, 1985). First, the findings raised awareness of the importance of the situation in possibly overriding personality traits. Second, it also raised the issue of ethics in conducting research and raised questions about the delicate balance between the risks to participants (e. g, the emotional distress they felt at harming another individual) and the benefits gained from the empirical exploration. In reaction to the findings, the Milgram study has been replicated investigating many different factors (many explained above) in many different countries, and at many different times, and each time the results are replicated. For instance, ABC primetime (2007) replicated the original Milgram study (but stopped the study at 150 volts), and found similar results to Milgram. Sixty-five percent of men and 73 percent of women administered painful shocks to the Learner. And, surprisingly, even when they ran a version of the study where a confederate refused to continue the experiment, 63 percent of the participants continued to administer painful shocks. In addition, Meeus & Raaijmakers (1995) replicated the original Milgram study, but instead of relying on administering physical pain on an ostensible other participant, they were instructed to administer psychological pain. In this study, the participants were told that the researchers were interested in the effects of stress on test taking. Therefore, the participants were instructed to distract a test-taker with harassing remarks. Ninety-two percent of the participants made all 15 harassing statements--even with protests and pleas from the applicant to stop. Thus, the take home message is that given the right situation, any individual might do things that even they could not imagine they would do.

3. The Milgram Model

Barry originally formulated the dynamic hypothesis underlying the decision process in Milgram's experiment as shown in Figure 1. It consisted of two major negative feedback loops that assured delivery of the voltage up to the implicit goal of the system as determined by the personality-related parameters of the subject and the specification of the experimenter.

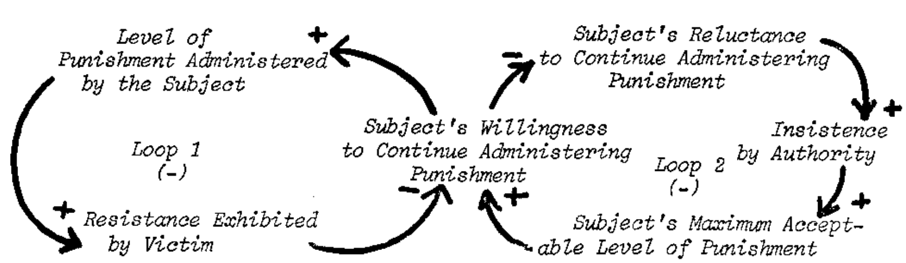

Fig. 1. Barry Richmond's original causal loop diagram describing behavior in the Milgram experiment. From Richmond (1977).

We have reformulated the model in four sectors using the *iThink* hierarchy: Punishment delivered, Decision to Continue, Anxiety, and Termination of the Experiment. The structure and information relationships among the Punishment Delivered and Decision to Continue sectors are shown in Fig. 2. The experiment is terminated when the total punishment delivered to the victim exceeds a value

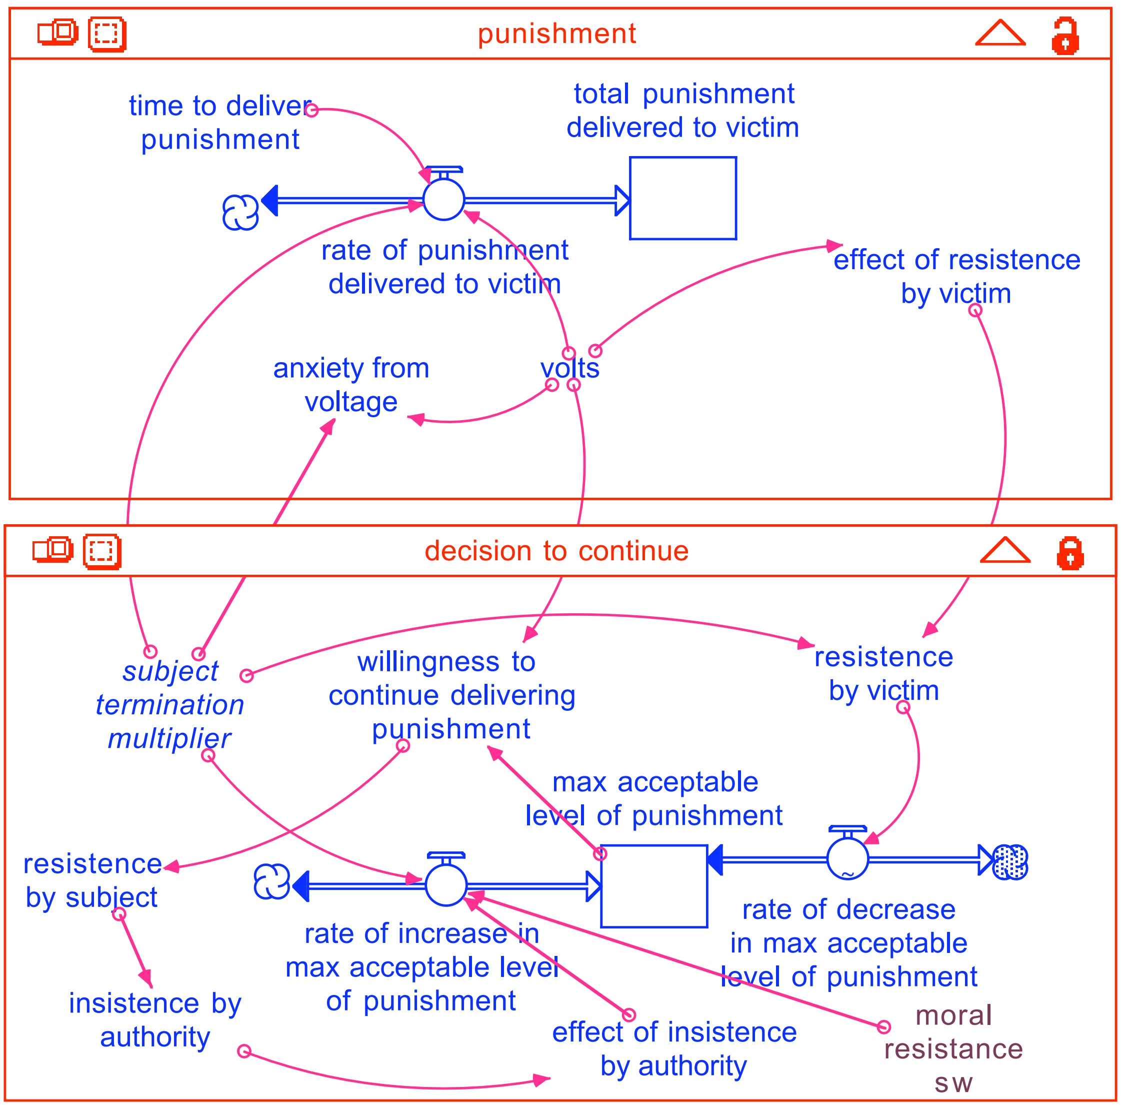

Fig. 2. Punishment Delivered and Decision to Continue sectors of the Milgram model.

set by the experimenter. It can also be terminated due to the anxiety level of the subject or his/her willingness to continue to inflict punishment. Termination, in turn, ends anxiety, punishment, and the decision to continue. Both punishment and the decision to continue fuel anxiety. The instantaneous rate of punishment delivered to the victim is measured in volts/unit time. This punishment is accumulated into a stock for driving the experimenter's termination process. Volts increase over time and are an exogenous input to the system as the experiment continues. Volts create anxiety on the part of the subject, while also invoking resistance by the victim and together with the criterion " max acceptable level of punishment" determine the willingness of the subject to continue punishment. Anxiety from voltage is eliminated when the subject decides to terminate the process. The maximum acceptable level of punishment is a stock whose increase is determined by insistence by authority, and can also change when the subject resists or his willingness to continue wanes. The maximum acceptable level of punishment also drains as the victim pretends to show resistance. We have added to Barry's original model a "moral resistance switch" that slows down the rate of increase of the maximum acceptable level of punishment. This allowed us to change the intrinsic value of the moral resistance a subject brings into the experiment. Barry probably would have approved of this alteration since his model of the Stanford Prison Experiment, which he developed much later, has built-in switches that allow a user to change personality factors in the role system the model portrays.

Fig. 3 shows the anxiety sector of the model. The growth in the level of anxiety is driven by the existing level of anxiety, the rising voltage, the resistance shown by the victim, and the insistence by the authority. Anxiety decays as a first order function of the anxiety stock, but is drained rapidly when the subject terminates the experiment. We have added to the original model an "empathy switch" that allows us to specify another personality factor -- the empathetic nature of a subject. We assume that anxiety will grow faster in case of a more empathetic subject. Additional personality factors could be added to the model in terms of parameters that can be exogenously changed, although we restricted our current experiments to changes only in empathy and moral resistance levels of the subjects.

Figure 4 shows how the experiment might end -- through anxiety and unwillingness to continue by the subject or when the experimenter has determined that the maximum level of punishment has been reached.

The model equations are placed at Annex A.

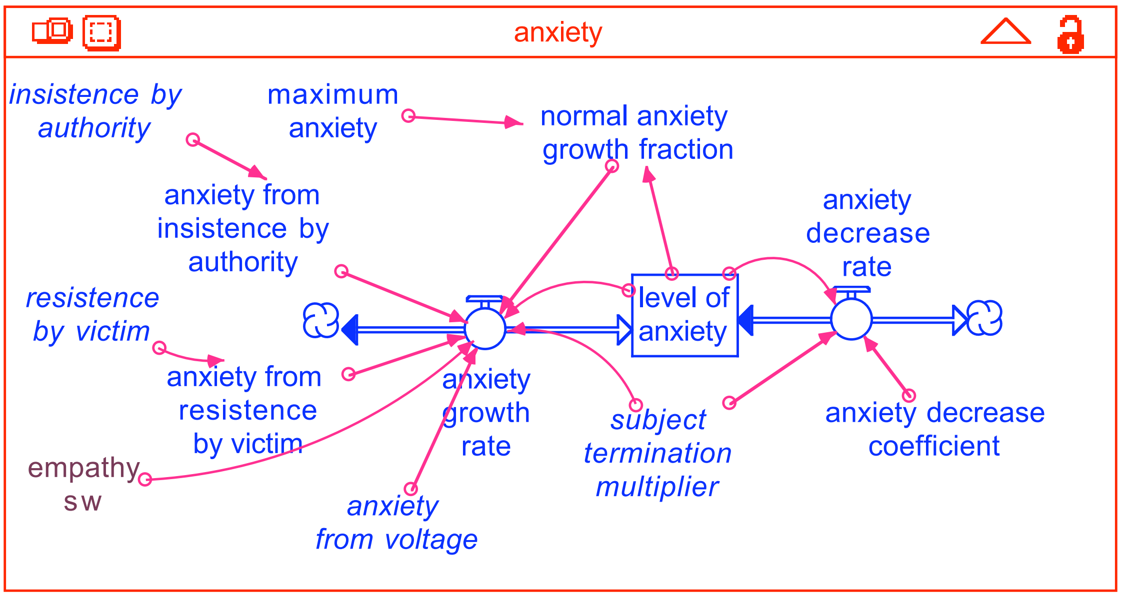

Fig. 3. Anxiety sector of the Milgram model.

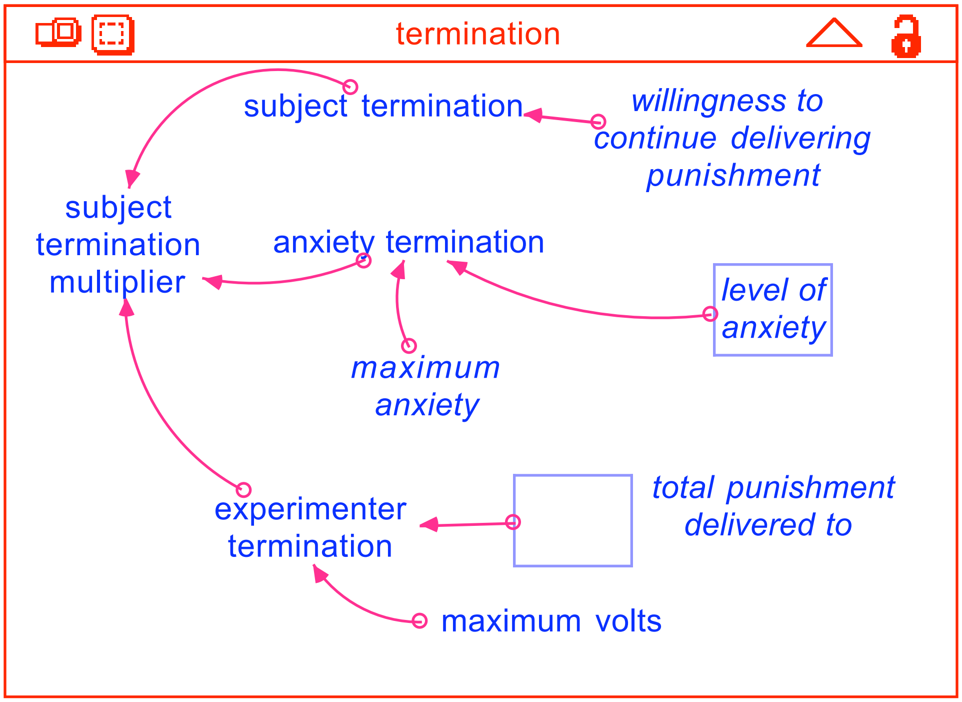

Fig. 4. Termination of the Experiment sector of the Milgram model.

Figure 5 shows a computer simulation run of the experiment. An exogenous growth in volts delivered drives the system. The level of anxiety, the maximum acceptable level of punishment, the total punishment delivered to the victim, the resistance by the victim and the insistence by the authority grow concomitantly until the experiment is terminated. The graph representing insistence by the authority has a peculiar pattern that results from three nonlinear graphical functions operating in tandem that create a complex net shape. Barry as a mature practitioner of system dynamics would probably frown upon such a structure, but remember that this model was developed during his first year at MIT!

Figure 6 shows the relationship between anxiety and maximum acceptable level of punishment, which grows concomitantly because insistence by authority mitigates the resistance created by anxiety. Three

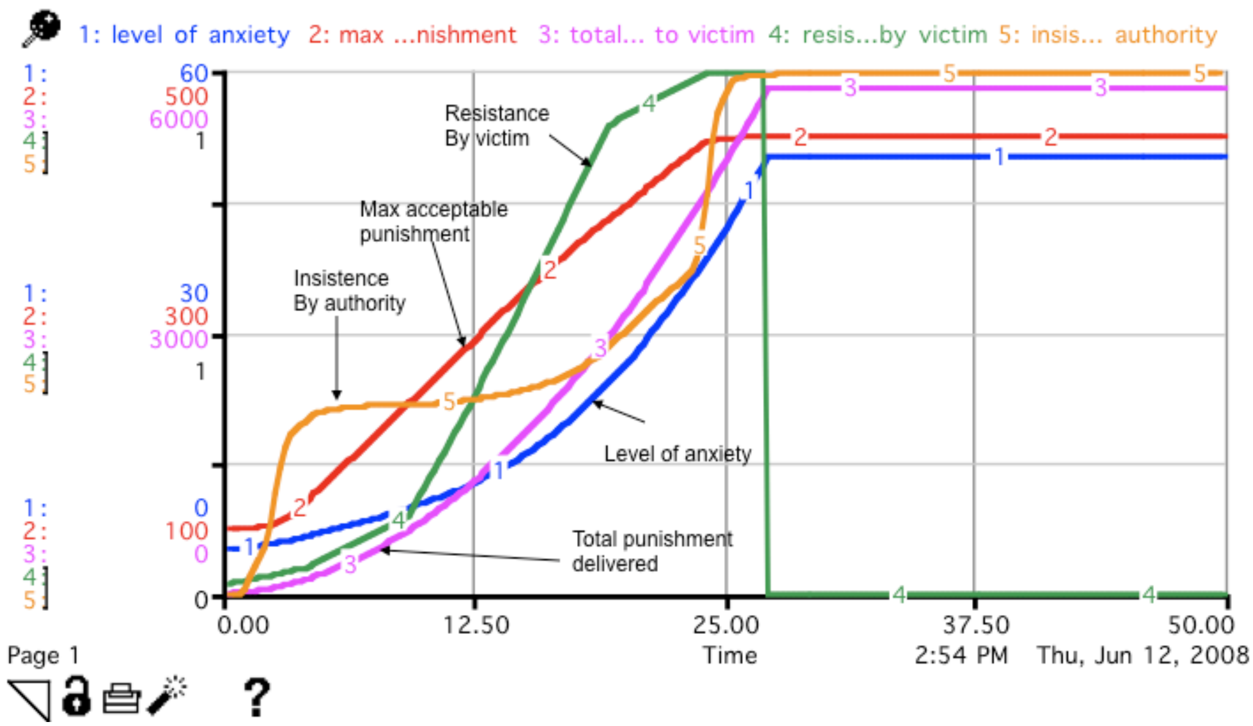

Figure 5: Simulation run of the Milgram model.

cases are shown: a subject with an "average" personality, a subject that is more empathetic than average, and a subject with an above average level of moral resistance. As would be expected, a moral subject terminates the experiment at a low level of punishment; however, an empathetic subject does not terminate the experiment, but only takes longer to get to the design level of maximum punishment. In the latter case, the insistence by the authority mitigates empathetic behavior while in the former case, moral resistance effectively works against the rise in the acceptable level of punishment.

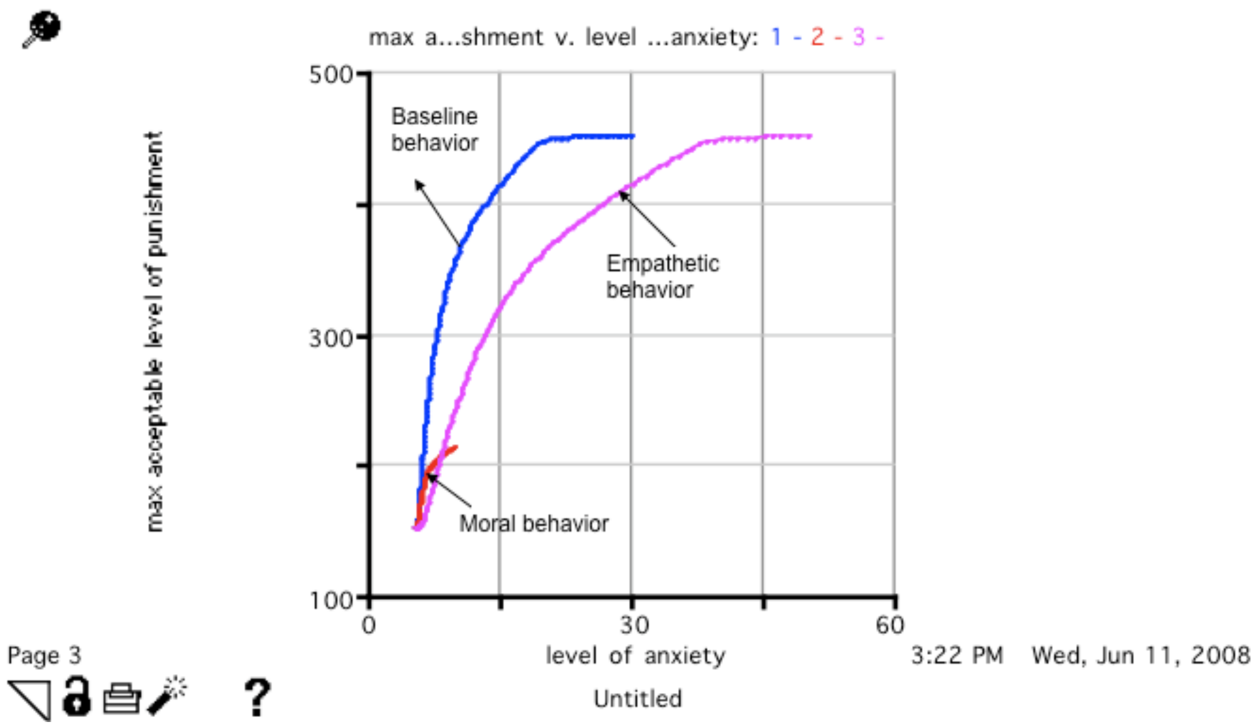

Fig. 6. Relationship between anxiety and maximum acceptable level of punishment in the Milgram model.

4. The Stanford Prison Experiment

Also interested in the effect the social situation has on individuals, Philip Zimbardo and colleagues (1973) decided to try to understand prison behavior. Unable to put participants inside a real prison, Zimbardo created his own prison in the basement of Jordan Hall (the psychology building) at Stanford University. Like Milgram, Zimbardo advertised in newspapers to solicit participants. Interested participants were screened for mental illness and proclivity towards violence, to ensure neither of these variables would influence the results of the study. Those participants who passed these tests were then randomly assigned (by flipping a coin) to be either a prisoner or a guard.

On the first day of the study, the guards reported to Jordan Hall and were briefed. They were informed that there could be no violent acts towards the prisoners, they must call them by their Prisoner ID number, they had to conduct lineups and count the prisoners to ensure no one had escaped, and they were allowed three bland meals and three toilet visits per day. Meanwhile, the prisoners were unexpectedly arrested at their residences by the local police department. On coming into Jordan Hall, the prisoners were then officially booked, strip searched, given a Prisoner ID number (the only way they would be referred to throughout the experiment), given prisoner apparel (looking similar to a smock worn during doctors visits), and had a chain placed around their ankle.

At first, the scene was mundane as the prisoners and guards got situated to the new surroundings and the experiment. But, as time went on, even on that first day, both the prisoners and guards got used to their roles and acted accordingly. The guards started seeing the power they had in the situation and some started taking advantage of their power. They started being cruel and unjust with the prisoners, they would wake the prisoners in the middle of the night to do a lineup, they started making the prisoners do manual labor (e.g., scrubbing the floors). On seeing other guards taking advantage of their power, more guards felt compelled to follow suit, whether to fit in with their fellow guards (normative influence, Kassin, Fein, & Markus, 2008) or because they felt that the other guards must know something more about how to act in the situation then they did (informational influence, Kassin, Fein, & Markus, 2008). One guard, in particular, used the power he had by being a guard to punish and abuse the prisoners much more than the other guards, and for his cruel behavior he earned the nickname “John Wayne”.

In response to the treatment by the guards, the prisoners started at first started to rebel. They refused do the mandated manual labor, and ignored requests to wake up in the middle of the night. As long as one prisoner rebelled, then the other prisoners felt as if they too could rebel. But, as the rebellions increased and the number of prisoners rebelling increased, the guards started to retaliate, especially “John Wayne”. They increased the cruelty in their language and treatment of the prisoners, they woke the prisoners more frequently in the middle of the night, increased the manual labor, and started placing the rebels in solitary confinement. The rebellions and retaliations went back and forth. If the initial retaliation by the guards did not work (e.g., it was not cruel, repressive, or “John Wayne” enough), then they increased the cruelty to squash the rebellion. For example, the leader of the rebellion was Prisoner #416. He objected to the treatment by the guards onto the prisoners and his complaints were received with more abuse by the guards. In rebellion against this treatment, Prisoner #416 announced a hunger strike and refused to eat his meals. In retaliation, the guards locked him in solitary confinement, and attempted to turn the other prisoners against him by announcing that all the prisoners would need to give up their blankets and sleep on a cold floor in order for Prisoner #416 to be released from the solitary confinement.

With each failed rebellion, the prisoners started to feel a sense of learned helplessness (Seligman & Maier, 1967), or that no matter what they did in reaction to the guards' treatment there was nothing they could do to remedy or escape the situation. Eventually, after enough failed rebellions and with excessive increases in the cruelty from the guards, the prisoners stopped rebelling and started to become more reticent, more passive, and more and more unhappy. And, after 36 hours as a prisoner, one prisoner participant was released from the study because he was exhibiting signs of severe depression and anxiety. Several more prisoners left for similar reasons, and after 6 days Zimbardo ended the study. Ironically, the prisoners forgot, most likely due to the situational demands and the effects of repeated failure in the rebellions, that they could have escaped the treatment by asking to leave the study.

Thus, Philip Zimbardo, like Stanley Milgram, conducted research that forever shaped the field of social psychology. Like the Milgram experiments, Zimbardo's prison study raised awareness of the importance of the situation. This research demonstrated that taking innocent, mentally healthy, and non-violent individuals, and putting them into even an artificial setting where social roles and expectations prevail can lead to serious ramifications in how people treat each other—regardless of one's personality traits. While not as many replication studies have been conducted on the Zimbardo study, most likely due the complex nature of creating a mock-prison or finding a less complex proxy, Zimbardo (2004) recently argued that the results from his Prison Study predicted, or could have foretold, the atrocities that occurred at Abu Ghraib towards Iraqi prisoners (Stannard, 2004). In conclusion, the Zimbardo prison study, like the Milgram study, was pivotal in the history of social psychology as being some of the first research to demonstrate the power of the situation over personality traits. In addition, the results of both the Milgram and Zimbardo studies have stood the test of time and replication.

5. The Stanford Prison Experiment Model

Barry's original formulation of the Stanford Prison Experiment model has three main stocks representing Prisoner resistance, Prisoner Fear, and Repression by Guards. A fourth stock named Perceived Resistance by Guards is a complex exponential average of the Prisoner Resistance, with perception time also varying with the level of Prisoner Resistance. This means that while it would take the guards some time to perceive the actual prisoner resistance level, low levels of resistance will go undetected and high levels will be quickly detected.

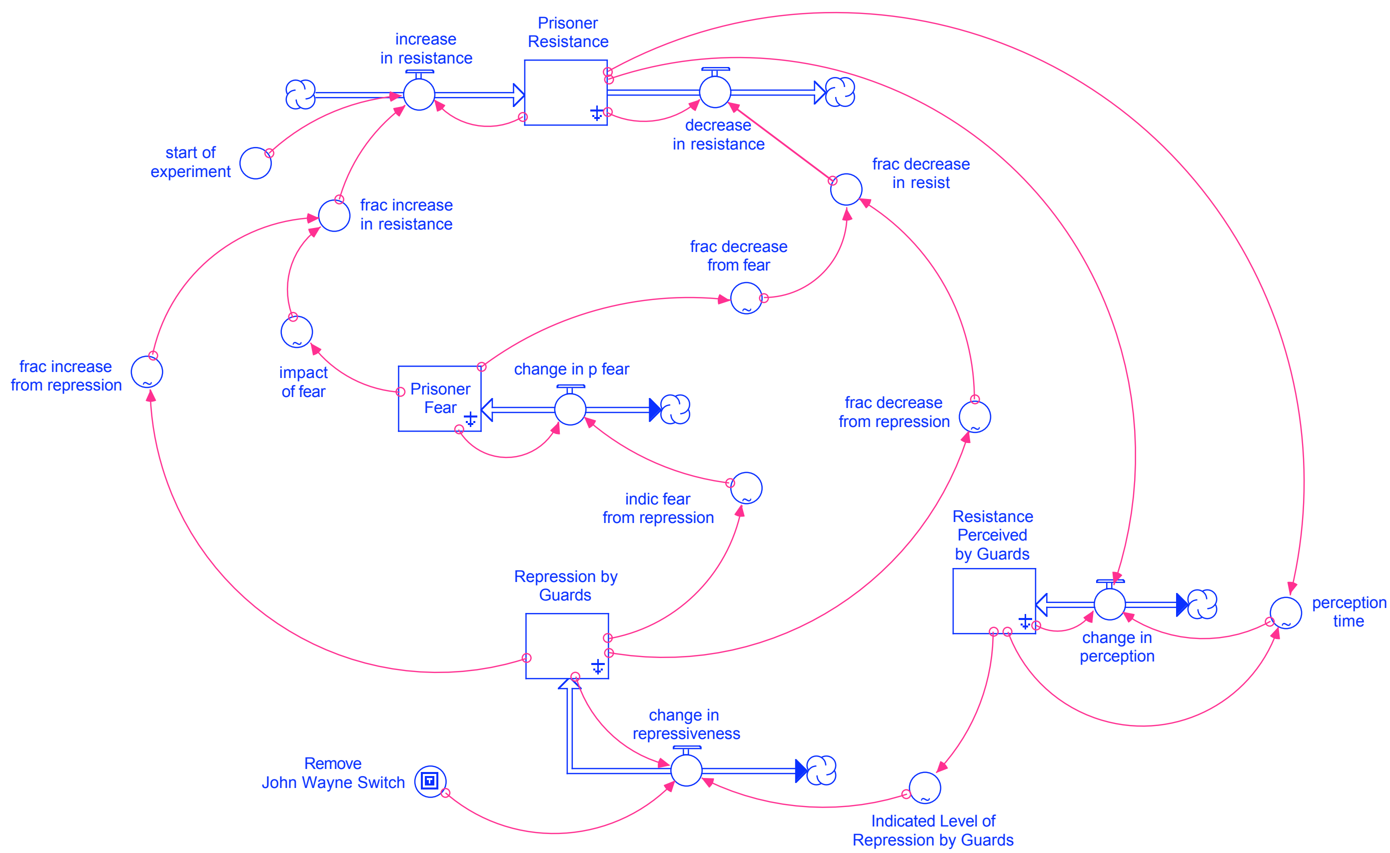

Fig. 7. Structure of Barry Richmond's model of the Stanford Prison Experiment.

Perceived Resistance leads to Repression, which invokes more resistance, creating an escalating process. At the same time, Repression also fuels Prisoner Fear. Both Repression and Fear decay Resistance, thus creating major balancing feedback loops. While a major balancing feedback loop alone can create an overshoot and instability, a positive feedback loop coupled with this process will exacerbate such behavior. Hence, the structure of this system is bound to create escalation of hostilities between the prisoners and guards, irrespective of the personalities of the two parties.

The model includes a “Remove John Wayne Switch” which, when thrown, constrains the growth in repression by the guards. With this switch, the model can be tested with different guard personalities. In fact, the graphical functions representing Indicated Level of Repression by Guards, Perception time for guards to detect Prisoner Resistance, Impact of fear on fractional increase and decrease in resistance, and indicated fear from repression, all represent personality related behavioral responses. Changing the slopes of these functions implies bringing in actors with different personalities. The model equations are placed at Annex B.

Figure 8 compares changes in prisoner resistance with normal and John Wayne behavior. Figure 9 compares repression by guards with the two guard personality profiles.

Figure 8: Changes in prisoner resistance with normal and John Wayne behavior by guards.

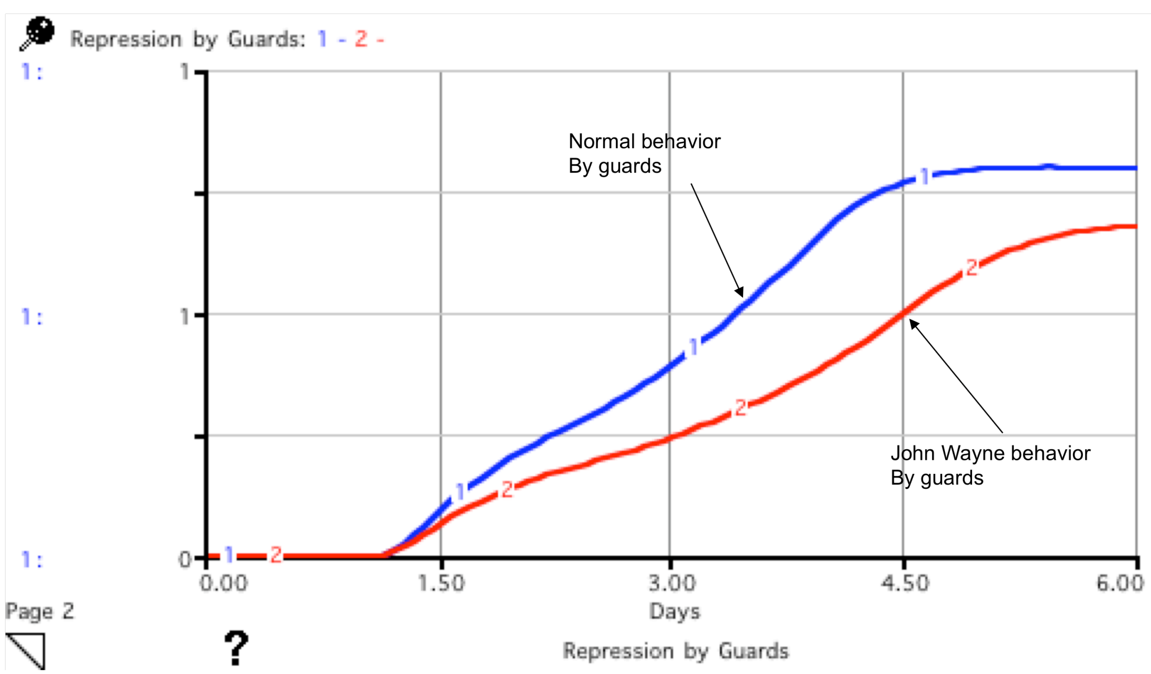

Figure 9: Changes in repression by guards with normal and John Wayne attitude

John Wayne behavior amplifies the repressive response on the part of guards. Normal behavior is John Wayne behavior, with indicated repression of 100%. When the Remove John Wayne switch is turned on, indicated repression is reduced to 70% of the normal value. When guards act like John Wayne, prisoner resistance is broken down more quickly, hence repression by the guards does not rise as much. When guards are more benign, prisoner resistance takes longer to break and rises to a higher peak, and hence, repression rises too.

Many other experiments can be performed with this model by changing the slopes of the graphical functions that represent personal behavioral characteristics of the subjects. Such experimentation can help to further clarify the contribution of the personal and situational factors to generating the dynamics observed. While we have explored some of these experiments, we have not documented them in this article.

6. System Dynamics and Social Psychology: Unrealized Potential

Barry Richmond's attempt to model the underlying dynamics of subjects' behavior in the Milgram obedience experiments and the Stanford Prison experiment demonstrates the feasibility of combining the study of system dynamics and social psychology and the valuable insights that can be gained by doing so. Yet 30 plus years later, it remains a rare example of such interdisciplinary work. To our knowledge, in the intervening years only one other effort has been made to build a dynamic simulation model of behavior in either of these two classic experiments (see Rochat *et al*., 2000).

In the field of system dynamics, the psychological literature on judgment and decision making has made its way into modeling practice, but other subdisciplines of psychology have generally been overlooked. In the system dynamics literature the social psychology of group behavior is widely used in the design of group model building interventions (see Vennix, 1996), but models designed to explain experimental results in social psychology or models that incorporate those results in addressing social or business problems are few and far between [however, see

Repenning and Sterman (2000) for an example of a model that incorporates a classic social psychological result called the "fundamental attribution error"].

In social psychology, the dynamic perspective has been an important part of the field since at least the appearance of Kurt Lewin's book *The Dynamic Theory of Personality* (1935), and can even be traced back to William James (1890). And, many current social psychologists realize that the outcomes of their experiments have dynamic implications, and have started to rely on Structural Equation Modeling (SEM) to better understand how and to what extent the independent variables influence the dependent variables (see Kenny, 2008; Maydeu-Olivares, McArdle, 2005; and Ridgon, 2008). In other words, social psychologists use SEM to determine if a model is valid based on the results and also to determine which model best fits the data. However, the use of computer simulation modeling to build and test theories has not become part of mainstream social psychology — even though there have been arguments for more widespread use of dynamic simulation modeling in social psychology, particularly in recent years (see, e.g., Palys, 1979; Vallacher et al., 2002; Mason et al., 2007; and Smith and Conrey, 2007). And surprisingly, the most influential book on the topic, *Dynamical Social Psychology* (Nowak and Vallacher, 1998), does not discuss or demonstrate any awareness at all of Forrester-style system dynamics modeling.

The lack of communication and interaction between social psychology and system dynamics is striking. On the one hand is a field, system dynamics, which has for decades developed and improved a technique for modeling economic, policy, and business problems that are inherently social but has never fully taken advantage of the enormous empirical literature on human social behavior. On the other is the field of social psychology, which has a long tradition of constructing theories of behavior that are inherently dynamic, but generally resists the use of computer simulation as an aid to theory building and seems almost completely unaware of the system dynamics modeling tradition. How did this happen?

The answer, we believe, lies mainly in the fact that researchers and practitioners in the two fields go about their business very differently. Social psychologists focus on conducting controlled experiments that change one variable at a time and employ random assignment of subjects in order to allow causality to be unambiguously determined. For example, in Milgram's replications of his original obedience study, he systematically varied such factors as proximity to the victim while keeping all other aspects of the experiment (e.g., the instructions, the sequence of events, the experimenter's prompts, the location, the personnel) the same. In Zimbardo's prison experiment, subjects were randomly assigned to the play the role of guard or prisoner. To a social psychologist, verifying the nature and magnitude of the relationships in a system dynamics model would represent a lifetime of work. In fact, many social psychologists have devoted their entire careers to trying to definitively pin down the relationship between two variables (e.g., does watching televised violence cause aggressive behavior?, or does contact with outgroup members reduce prejudice?). Furthermore, the amount of time, effort, and expensive involved in conducting experiments on human subjects means that the kind of time series data that would lend itself to modeling is not often collected; most social psychology experiments are one-time snapshots of behavior or, at best, two snapshots (one before the experimental task, and one after).

System dynamicists, by contrast, set aside the question of precise values for parameters and skip right to the big picture task of identifying the structure of the system in which behavior takes place. Since they theorize that it is primarily structure that drives system behavior, they believe it is unnecessary to identify the precise magnitude of relationships between variables up front; in fact, sensitivity analyses may ultimately show that system behavior is unchanged even with widely varying assumptions about parameter values.

Thus, system dynamicists would suggest that social psychologists are probably collecting a lot of unnecessary data. To which social psychologists would respond with the common argument that system dynamics is "measurement without data." With a long history within the field of combating "common sense" assumptions about human behavior that controlled experimentation proves to be false, the system dynamics attitude of "model first, measure later" would be viewed as counterproductive and potentially harmful by many social psychologists.

Another important difference between the two fields is in the relative importance assigned to different types of data for input to theories and/or models. Controlled experiments are held in high esteem by social psychologists for their ability to determine causality. However, external validity (whether results obtained in the lab also apply in real life) is an important concern as well, and thus social psychologists employ a variety of additional data gathering techniques that vary in the degree to which they offer experimental control versus external validity, including survey research, naturalistic observation, field and case studies, and interviewing techniques.

However, no matter the technique there is always an overriding concern for systematicity in data collection and representativeness in the human subjects under study. System dynamicists, by contrast, often display little interest in formal or systematic methodologies or the empirical psychological literature and simply go straight to the managers or decision makers who have the information they need and ask them for it. The modeling process determines the need for information, and the pressure to solve a practical policy problem often does not leave time to consult the empirical psychological literature. This approach suggests a high level of confidence in the ability of people to accurately remember and report on their own thinking and behavior; however, psychologists, based on decades of uncovering counterexamples in their experiments, are not as confident in this type of approach. To a psychologist, the chance that information collected so informally will turn out to be inaccurate or misleading is simply too high.

A third difference between the two fields is the relative emphasis placed on theoretical versus applied work. Social psychologists conduct applied research on a wide variety of socially important topics, but the emphasis is on developing and testing theories of human social behavior. The applied work that is done is not primarily conducted in corporate or management settings. In comparison, system dynamics modelers are concerned almost exclusively with social problem solving as opposed to theory, and most of this work is done in collaboration with, or in consultation for, businesses and other large organizations. Hence, another reason for lack of collaboration between social psychologists and system dynamicists is simply that their paths do not cross very often. In addition, neither field is a typical (or even atypical) component of academic training in the other: there is a profound, mutual lack of awareness between the two fields.

To fully realize the potential for productive interaction between system dynamics and social psychology, social psychologists will need to a) become more open to the advantages of the more wholistic approach that system dynamicists can provide, b) start applying the research they have already conducted and associated findings to many different models of human behavior, and c) begin to be more willing to leap into the modeling enterprise without waiting for data that may take years, if ever, to materialize. Even to date, social psychologists may use bootstrapping simulations (see Efron &Tibshirani, 1993) in order to determine whether increasing the sample size may increase statistical power and the likelihood of significant effects, but they rarely use computer simulations based on their own data for any other purpose.

System dynamicists, following Barry Richmond's example, will need to learn to find value in exploring and applying the empirical literature in social psychology and to discover the advantages of more systematic and representative data collection, and consequently more complex (and possibly accurate) models. And of course researchers in both fields will need to find ways to educate themselves about the methods and of the other domain.

The benefits of such interaction would in many ways be similar to the benefits offered by dynamic modeling in any domain. The discipline imposed by modeling would doubtless uncover errors, omissions, and inconsistencies in mental models of social psychological theories, as it has in many other fields. It could also help strengthen the existing theories on human behavior (e.g., Barry's model on Milgram strengthened Milgram's claim that the situation played more of a role than personality). It would also help to make the theories less conceptual and more operational. In addition, with the cognitive assistance provided by computer models, social psychological theories could become more dynamic and complex. In addition to these general benefits, there are some advantages to dynamic modeling that are unique to social psychology (and other experimental disciplines within psychology). Given the time-consuming, labor-intensive nature of each empirical test of a relationship in a psychology research program, use of simulation runs to test alternate hypotheses before running human subjects would greatly aid in preserving limited resources for use in the most promising live experiments. It could also enhance understanding of human behavior theories by being able to combine many theories and variables at one time that could not be included in one parsimonious experimental design. In addition, as Barry Richmond (1977) suggested in his Milgram model manuscript, the simulation model could be used to test experimental conditions that would be difficult or impossible to reproduce in the laboratory or to put simulated subjects through procedures that would be unethical to duplicate in real life.

Of course, any of several different modeling paradigms might potentially offer the same list of benefits to the field of social psychology. And for certain social psychological topics, other modeling techniques seem to be obviously better suited than system dynamics (e.g., modeling crowd or large group behavior seems to call for an agent-based modeling approach). Is there anything that makes system dynamics particularly well-suited to the study of social psychology? Certainly the emphasis in the system dynamics paradigm on the inclusion of so-called "soft" variables, and the techniques that have been developed for incorporating them into models, is an advantage for studying social psychology, where concepts such as anxiety, obedience, prejudice, aggression, conformity, perception, and judgment may be the main topics of investigation [an interesting example of modeling soft variables in an organizational context is in Saeed (1998), where he attempts to track changes in role autonomy, citizenship and professional competence in an innovation organization]. It would also allow the merging of theories and concepts from many different domains with social psychology, psychology, and other disciplines (e.g., anthropology, economics, law, sociology). And other central aspects of the system dynamics tradition, including the emphasis on feedback relationships and the operational emphasis, and the ideas that the main purpose of modeling is for learning rather than prediction and that models are most effective when they are designed to help solve problems rather than exhaustively replicate systems, certainly offer an alternative to other modeling paradigms.

However, the best reason to choose a system dynamics approach for the study of social psychology might be its suitability to study a long-standing problem at the heart of the field (and of central importance to an understanding of the Milgram and Stanford Prison studies as well): the relative contribution of the person versus the situation in determining behavior. For many years a lively debate raged between the situationist camp, who emphasized the ability of situational variables to overwhelm individual differences, and the personality camp, who emphasized the important role of enduring character traits in explaining behavior. Today this debate is largely concluded, with the majority of social psychologists supporting the view that the person and situation are roughly equal in importance, and in addition can interact with each other in interesting and sometimes counterintuitive ways (see, e.g., Fleeson and Noftle, 2008). Even the modern interpretations of the Milgram and Stanford Prison experiments have evolved from an almost exclusively situationist perspective to a more balanced view that recognizes the important role of individual differences in providing a comprehensive explanation of the results [see Blass, 1991, for a more modern interpretation of the Milgram experiment; and Donnellan et al. (2007) and Zimbardo (2007) for a discussion of person-situation interactionism in the Stanford Prison Experiment].

However, for any given topic of study or experimental situation, it is still not easy to predict or to tease apart the relative contributions of the person and the situation, how they may interact, and how that interaction may unfold over time. System dynamics modeling could be a very valuable tool to help simplify and systematize this effort. In fact, among the various modeling alternatives, system dynamics seems ideally suited for the task. In system dynamics situational variables can be incorporated into the model structure or constants. Individual difference or personality variables can be represented as adjustable parameters. The test of the relative strength of situational versus personality factors thus reduces to the standard practice of parameter sensitivity analysis. This is in fact precisely the approach that Barry Richmond took in his experiments with the Milgram and Stanford Prison Experiment models. In the case of the models replicated in this chapter, these analyses confirm that the models are relatively insensitive to changes in personality variables, supporting the situationist views espoused by Milgram and Zimbardo.

7. Conclusion

It has been inspiring for the present authors (one system dynamicist and two social psychologists) to have the opportunity to examine and reflect upon Barry Richmond's early effort to combine his knowledge and interest in social psychology and system dynamics. This exercise has greatly aided our thinking about the need for and benefits of productive collaboration between the two fields, and improved our understanding of the obstacles that will need to be overcome to increase such collaboration in the future. We hope other modelers with an interest in psychology will be inspired to follow up or expand upon Barry Richmond's work and, perhaps, explore other theories and applications of social psychology using the powerful tools of system dynamics.

We also hope Barry's models will gain some use in psychological education. The concept of "the power of the situation" is contrary to the mental models people form about human behavior through their life experience. For example, when people hear about the Milgram experiment, everyone thinks they would be one of the ones to resist authority. But the experiment, and its many replications with diverse subjects, prove otherwise. Allowing students to experiment with the simulation model themselves and see the effects of changing personality variables may be an effective way to drive home the point that in certain situations personality just does not matter that much.

We believe there are also larger lessons to be learned from Barry's work. Certainly, it paves the way for system dynamics modelers to make better use of an underutilized information source, namely, the results of controlled experiments on human subjects. In addition, as there are still those in the system dynamics community that are reluctant to model so-called "soft" or psychological variables, Barry's work serves as a reminder that such work, while challenging, can, and must, be done. And, finally, the topics of the Milgram and Stanford Prison Experiment models themselves serve as a useful reminder that the field of system dynamics began by addressing, and still can address, the great problems of society.

Annex A
Equations for Milgram model

anxiety

level_of_anxiety(t) = level_of_anxiety(t - dt) + (anxiety_growth_rate - anxiety_decrease_rate) * dt

INIT level_of_anxiety = 5

INFLOWS:

anxiety_growth_rate = normal_anxiety_growth_fraction_fraction*level_of_anxiety*anxiety_from_insistence_by_authority*anxiety_from__resistence_by_victim*anxiety_from_voltage*subject_termination_multiplier

OUTFLOWS:

anxiety_decrease_rate = level_of_anxiety*subject_termination_multiplier*anxiety_decrease_coefficient

anxiety_decrease_coefficient = .03

maximum_anxiety = 50

anxiety_from_insistence_by_authority = GRAPH(insistence_by_authority)

(0.00, 1.00), (0.2, 1.05), (0.4, 1.20), (0.6, 1.45), (0.8, 1.60), (1.00, 1.70)

anxiety_from__resistence_by_victim = GRAPH(resistence_by_victim)

(0.00, 1.00), (0.2, 1.10), (0.4, 1.30), (0.6, 1.60), (0.8, 2.00), (1.00, 2.50)

normal_anxiety_growth_fraction_fraction = GRAPH(level_of_anxiety/maximum_anxiety)

(0.00, 0.06), (0.25, 0.03), (0.5, 0.015), (0.75, 0.01), (1.00, 0.008)

decision to continue

max_acceptable_level_of_punishment(t) = max_acceptable_level_of_punishment(t - dt) + (rate_of_increase_in_max_acceptable_level_of_punishment - rate_of_decrease_in_max_acceptable_level_of_punishment) * dt

INIT max_acceptable_level_of_punishment = 150

INFLOWS:

rate_of_increase_in_max_acceptable_level_of_punishment = effect_of_insistence_by_authority*subject_termination_multiplier

OUTFLOWS:

rate_of_decrease_in_max_acceptable_level_of_punishment = GRAPH(resistence_by_victim)

(0.00, 0.00), (0.2, 0.25), (0.4, 1.00), (0.6, 2.50), (0.8, 4.50), (1.00, 7.00)

resistence_by_victim = effect_of_resistence_by_victim*subject_termination_multiplier

effect_of_insistence_by_authority = GRAPH(insistence_by_authority)

(0.00, 0.00), (0.2, 7.00), (0.4, 17.0), (0.6, 19.0), (0.8, 12.0), (1.00, 7.00)

insistence_by_authority = GRAPH(resistence_by_subject)

(0.00, 0.00), (0.2, 0.4), (0.4, 0.7), (0.6, 0.9), (0.8, 0.95), (1.00, 1.00)

resistence_by_subject = GRAPH(willingness_to__continue_delivering_punishment)

(0.00, 1.00), (0.1, 0.98), (0.2, 0.97), (0.3, 0.95), (0.4, 0.85), (0.5, 0.7), (0.6, 0.45), (0.7, 0.25), (0.8, 0.05), (0.9, 0.00), (1, 0.00)

willingness_to__continue_delivering_punishment = GRAPH(max_acceptable_level_of_punishment-volts)

(-25.0, -0.1), (0.00, 0.00), (25.0, 0.05), (50.0, 0.2), (75.0, 0.65), (100, 0.75), (125, 0.9), (150, 0.95), (175, 0.98), (200, 1.00)

punishment

total_punishment_delivered_to_victim(t) = total_punishment_delivered_to_victim(t - dt) + (rate_of_punishment_delivered_to_victim) * dt

INIT total_punishment_delivered_to_victim = 0

INFLOWS:

rate_of_punishment_delivered_to_victim = (volts/time_to_deliver_punishment)*subject_termination_multiplier

time_to_deliver_punishment = 1

volts = (TIME+1)*15

anxiety_from_voltage = GRAPH(volts*subject_termination_multiplier)

(45.0, 1.00), (105, 1.10), (165, 1.20), (225, 1.30), (285, 1.40), (345, 1.50), (405, 1.60), (465, 1.70)

effect_of_resistence_by_victim = GRAPH(volts)

(0.00, 0.01), (75.0, 0.05), (150, 0.15), (225, 0.5), (300, 0.9), (375, 1.00)

willingness_to_continue_to_deliver_punishment =
GRAPH(max_acceptable_level_of_punishment-volts)
(-25.0, -0.1), (3.12, 0.00), (31.2, 0.05), (59.4, 0.2), (87.5, 0.65), (116, 0.75), (144, 0.9), (172, 0.95), (200, 0.98)

termination
anxiety_termination = IF(maximum_anxiety-level_of_anxiety) >0 THEN(1) ELSE(0)
experimenter_termination = IF(maximum_volts-total_punishment_delivered_to_victim) > 0 THEN(1) ELSE(0)
maximum_volts = 11475
subject_termination = IF(willingness_to__continue_delivering_punishment) > 0 THEN(1) ELSE(0)
subject_termination_multiplier =
anxiety_termination*experimenter_termination*subject_termination

Annex B
Equations for Zimbardo Model

Prisoner_Fear(t) = Prisoner_Fear(t - dt) + (change_in_p_fear) * dt
INIT Prisoner_Fear = 0
INFLOWS:
change_in_p_fear = (indic_fear_from_repression-Prisoner_Fear)/.5
Prisoner_Resistance(t) = Prisoner_Resistance(t - dt) + (increase_in_resistance - decrease_in_resistance) * dt
INIT Prisoner_Resistance = 0
INFLOWS:
increase_in_resistance = (Prisoner_Resistance*frac_increase_in_resistance)
+ PULSE(.1,start_of_experiment,100)
OUTFLOWS:
decrease_in_resistance = Prisoner_Resistance*frac_decrease_in_resist
Repression_by_Guards(t) = Repression_by_Guards(t - dt) + (change_in__repressiveness) * dt

INIT Repression_by_Guards = 0
INFLOWS:
change_in__repressiveness = IF Remove_John_Wayne_Switch = 0 THEN
((Indicated_Level_of_Repression_by_Guards-Repression_by_Guards)/.5)
ELSE (((Indicated_Level_of_Repression_by_Guards*.7)-Repression_by_Guards)/.5)
Resistance_Perceived_by_Guards(t) = Resistance_Perceived_by_Guards(t - dt) +
(change_in_perception) * dt
INIT Resistance_Perceived_by_Guards = Prisoner_Resistance
INFLOWS:
change_in_perception = (Prisoner_Resistance-
Resistance_Perceived_by_Guards)/perception_time
frac_decrease_from_fear = GRAPH(Prisoner_Fear)
(0.00, 0.00), (0.1, 0.02), (0.2, 0.065), (0.3, 0.23), (0.4, 0.445), (0.5, 1.00), (0.6, 1.33), (0.7, 1.50),
(0.8, 1.50), (0.9, 1.50), (1.00, 1.50)
frac_decrease_from_repression = GRAPH(Repression_by_Guards)
(0.00, 0.00), (0.1, 0.01), (0.2, 0.01), (0.3, 0.02), (0.4, 0.03), (0.5, 0.08), (0.6, 0.24), (0.7, 2.83),
(0.8, 3.00), (0.9, 3.00), (1.00, 3.00)
frac_decrease_in_resist = frac_decrease_from_fear+frac_decrease_from_repression
frac_increase_from_repression = GRAPH(Repression_by_Guards)
(0.00, 0.00), (0.1, 0.285), (0.2, 0.605), (0.3, 0.82), (0.4, 0.92), (0.5, 0.97), (0.6, 0.985), (0.7,
1.00), (0.8, 1.00), (0.9, 1.00), (1.00, 1.00)
frac_increase_in_resistance = (frac_increase_from_repression*impact_of_fear)
impact_of_fear = GRAPH(Prisoner_Fear)
(0.00, 1.00), (0.1, 1.00), (0.2, 0.99), (0.3, 0.98), (0.4, 0.915), (0.5, 0.765), (0.6, 0.65), (0.7,
0.445), (0.8, 0.275), (0.9, 0.16), (1.00, 0.00)
Indicated_Level_of_Repression_by_Guards = GRAPH(Resistance_Perceived_by_Guards)
(0.00, 0.00), (0.1, 0.27), (0.2, 0.485), (0.3, 0.63), (0.4, 0.74), (0.5, 0.85), (0.6, 0.925), (0.7, 0.95),
(0.8, 0.975), (0.9, 0.985), (1, 1.00)
indic_fear_from_repression = GRAPH(Repression_by_Guards)
(0.00, 0.02), (0.1, 0.03), (0.2, 0.045), (0.3, 0.055), (0.4, 0.08), (0.5, 0.225), (0.6, 0.45), (0.7,
0.645), (0.8, 0.77), (0.9, 0.9), (1, 0.97)

perception_time = GRAPH(Prisoner_Resistance/(Resistance_Perceived_by_Guards+.001))

(0.9, 30.0), (0.92, 30.0), (0.94, 30.0), (0.96, 22.1), (0.98, 10.7), (1.00, 0.25), (1.02, 0.25), (1.04, 0.25), (1.06, 0.25), (1.08, 0.25), (1.10, 0.25)

Remove_John_Wayne_Switch = 0

start_of_experiment = 1

References

ABC Primetime (2007). *The Science of Evil*. Retrieved September 10, 2008. Access: http://abcnews.go.com/Primetime/story?id=2765416&page=1

Blass, T. (1991). Understanding behavior in the Milgram obedience experiment: The role of personality, situations, and their interactions. *Journal of Personality and Social Psychology*, *60*(3), 398-413.

Blass, T. (2004). *The Man Who Shocked the World: The Life and Legacy of Stanley Milgram*. New York: Basic Books.

Brown, R. (1985). *Social Psychology*, 2nd ed. New York: The Free Press.

Donnellan, M. B., Fraley, R. C., and Krueger, R. F. (2007). Letter: Not so situational. *APS Observer*, *20*(6).

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Boca Raton, FL: Chapman & Hall/CRC.

Fleeson, W., and Noftle, E. (2008). The end of the person-situation debate: An emerging synthesis in the answer to the consistency question. *Social and Personality Psychology*, *2*(4), 1667-1684.

Haney, C., Banks, W. C., and Zimbardo, P. G. (1973). Interpersonal dynamics in a simulated prison. *International Journal of Criminology and Penology*, *1*, 69-97.

Haney, C., and Zimbardo, P. G. (1976). Social roles and role-playing: Observations from the Stanford prison study. In E. P. Hollander and R. G. Hunt (Eds.), *Current perspectives in social psychology* (4th ed.; pp. 266-274). New York: Oxford University Press.

James, W. (1890). *Principles of Psychology*. New York: Holt.

Kassin, S., Fein, S., and Markus, H. R. (2008). *Social Psychology, 7th Edition*. New York: Houghton Mifflin Co.

Kenny, D. A. (2008). *David Kenny's Quantitative Methods Home Page*. Retrieved September 10, 2008. Access: http://www.davidakenny.net/kenny.htm.

Lewin, K. (1935). *A Dynamic Theory of Personality*. New York: McGraw-Hill.

Mason, W. A., Conrey, F. R., and Smith, E. R. (2007). Situating social influence processes: Dynamic, multidirectional flows of influence within social networks. *Personality and Social Psychology Review*, *11*(3), 279-300.

Maydeu-Olivares, A., and McArdle, J. J. (2005). *Contemporary Psychometrics*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Meeus, W. H. J., & Raaijmakers, Q. A. W. (1995). Obedience in modern society: The Utrecht studies. *Journal of Social Issues, 51*, 15-175.

Milgram, S. (1963). Behavioral study of obedience. Journal of Abnormal and Social Psychology, 67(4), 371-378.

Milgram, S. (1974). Obedience to Authority. New York: Harper & Row.

Nowak, A., and Vallacher, R. R. (1998). *Dynamical Social Psychology*. London: The Guilford Press.

Palys, T. S. (1979). Simulation methods and social psychology. *Journal for the Theory of Social Behaviour*, *8*(3), 341-368.

Repenning, N. and J. Sterman (2000). Getting quality the old fashioned way: Self-confirming attributions in the dynamics of process improvement. In Cole, R.B. and R. Scott (Eds*.), Improving Theory and Research on Quality Enhancement in Organizations*. Thousand Oaks, CA: Sage.

Richmond, B. (1977). "Generalization with Individual Uniqueness: Modeling the Milgram Experiments." Technical Report D-2508-2, System Dynamics Group, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Richmond, B. (1993). Systems thinking: Critical thinking skills for the 1990s and beyond. *System Dynamics Review*, *9*(2), 113-133.

Rigdon, E. (2008). *Ed Rigdon's Structural Equation Modeling Page*. Retrieved September 10, 2008. Access: http://www2.gsu.edu/~mkteer/index.html.

Rochat, F., Maggioni, O., and Modigliani, A. (2000). The dynamics of obeying and opposing authority: A mathematical model. In Blass, T. (Ed.), *Obedience to Authority: Current Perspectives on the Milgram Paradigm*. New York: Psychology Press.

Saeed, K. 1998. Maintaining Professional Competence in Innovation Organizations. *Human Systems Management.* 17(1): 69-87.

Smith, E. R., and Conrey, F. R. (2007). Agent-based modeling: A new approach for theory building in social psychology. *Personality and Social Psychology Review*, *11*(1), 87-104.

Stannard, M. B (2004, May 8). Stanford experiment foretold Iraq scandal. *San Francisco Chronicle*. Retrieved: September 10, 2008. Access: http://www.sfgate.com/cgi-bin/article.cgi?file=/chronicle/archive/2004/05/08/MNGN76IG761.DTL

Vennix, J. A. M. (1996). *Group Model Building: Facilitating Team Learning Using System Dynamics*. New York: Wiley.

Zimbardo, P. (2007a). *The Lucifer Effect: Understanding How Good People Turn Evil*. New York: Random House.

Zimbardo, P. (2007b). Letter: Person x situation x system dynamics. *APS Observer*, *20*(8).

Zimbardo, P. G. (2004, May 9). Power turns good soldiers into 'bad apples'. *The Boston Globe.* Retrieved: September 10, 2008. Access: http://www.boston.com/news/globe/editorial_opinion/oped/articles/2004/05/09/power_turns_good_soldiers_into_bad_apples/