A Dynamic Systems Model of Parent-Child Interaction in the Development of Antisocial Behavior

January 2011

Portions of this paper were initially
presented as a poster to the
Society for Research in Child Development
at its Biennial Meeting, April 2001

David A. Boulifard
Psychology Department
Tillett Hall, Livingston Campus
The State University of New Jersey
RUTGERS

This paper presents a computerized model of parent-child interaction based on social learning principles. The model situates these principles within a dynamic systems framework (Thelen, 1989; van Geert, 1994), and serves to examine the coherence of Patterson's (1982) *coercion theory* as an explanation of antisocial behavior development.

In what follows, I shall briefly discuss developmental models and antisocial behavior. Then I shall explain my efforts to cast ideas from coercion theory into the form of a computationally active, dynamic systems model. Finally, I shall present the findings of a simulation study I conducted, and discuss their implications for the model and the theory.

Developmental Models

How does a developmental account of a phenomenon differ from a purely descriptive account? One important feature is its emphasis on temporal change, by small or large steps, from a more basic to a more advanced condition. Another is its explanation of how potentially complex structure and function originate.

Developmental models address two fundamental questions, namely, (1) what mechanisms or processes account for development, and (2) why are certain developmental pathways traversed, or outcomes reached, when others are conceivable? The answers themselves have evolved over time.

An old and familiar developmental model is *preformationism*, which asserts that an organism begins in a fully structured albeit miniscule condition, and simply grows in size. More flexible are *factorial* models, which allow influences from multiple sources to affect developmental outcomes, either independently (as *main effects*) or in ways that moderate each other's influence (via *interaction effects*).

Still more flexible are *transactional* models (Sameroff, 1975, 1995), in which various agents or factors within a developmental process may influence one another. In focusing attention on what happens at each step in a temporal sequence, transactional models adhere more faithfully to the general concept of a developmental progression than their simpler counterparts.

Self-Organization

Both preformationist and factorial models assume that development occurs when various agents or factors transmit preexisting structure to a developing organism. The transactional model deviates from this scheme, insofar as the organism may change what affects it.

A novel approach to the explanation of structure is to view all the agents and factors in a developmental process as parts of a system, and inquire how order *emerges* within the system (Thelen, 1989). This approach is associated with *dynamic systems theory*, a discipline that studies the time-dependent behavior of systems with potentially many interacting components.

The principle of *self-organization* posits that the features of system components, and the conditions under which they interact, may conspire to amplify some activities and to suppress others. This phenomenon, called *synergy*, can induce regularities in otherwise haphazard component behavior. These regularities *diminish* the system's overall complexity, because they reduce the number of ways its components can move. That is, they compress its *degrees of freedom*, so that its motion may be described by fewer variables than would otherwise be needed.

In a dynamic systems approach, the investigator attempts to identify a system's modes of regularity (organization), and the ways in which component features and environmental conditions may promote the formation and dissolution of its organized states.

Antisocial Behavior

Few social problems have inspired as much consternation, or prompted as much remedial effort, as behavioral tendencies labeled *antisocial* or *conduct disordered*. They are among the most costly to endure in psychosocial and economic terms, and the most difficult to modify once firmly established (Caspi & Moffitt, 1995).

Conceptualizations of antisocial behavior have richly varied (Hinshaw & Anderson, 1996). Individuals may differ from one another systematically, not only in their core symptom presentations but also in the origins, correlates, and courses of their disturbance. Explanations of antisocial behavior must account for temporal patterns of continuity and discontinuity, identifying those factors that promote the emergence, maintenance, remission, or changing expression of disorder.

Many research efforts have been guided by main effects models. For example, irritable temperament and dysfunctional parenting have been identified as significant risk factors for child antisocial behavior (Maughan & Rutter, 1998).

Fewer research efforts have been guided by transactional models. Noteworthy examples include findings by Olweus (1980) of developmental pathways that led from "active, hot-headed" early childhood temperament to adolescent aggressiveness through maternal permissiveness for aggression; and findings by Maccoby and Jacklin (1983) of reciprocal, inverse influence between "difficult" infant temperament and maternal "socialization pressure" across a six-month period.

<u>Coercion Theory</u>

Some of the strongest empirical evidence of transactional processes in the etiology of antisocial behavior has emerged from research by Gerald R. Patterson (1982) and his colleagues. Viewing overt social interactions among family members from a social learning perspective, he argued that children develop conduct disorders when they learn to escape from unpleasant situations (such as parental discipline) through aversive behavior.

Patterson analyzed three-step event sequences, each comprising an *antecedent (A)*, a *response (R)*, and a *consequence (C)*. Key hypotheses concerned the ways in which pleasing or displeasing consequences could alter *functional relations*; that is, conditional probabilities $p(R|A)$ that particular responses would follow particular antecedents.

His most remarkable findings flowed from studies of bi-directional influence in sequential interactions. He showed, for example, that family members of children referred for treatment of conduct problems *negatively reinforced* each other's aversive behaviors more frequently than those of children without conduct problems.

Figure 1 presents a sequence of alternating parent and child behaviors as a series of overlapping A-R-C segments, revealing the multiple roles individual events can play in altering functional relations. This example illustrates a hypothetical scenario Patterson discussed (pp. 143-145) to explain the [negative] *reinforcement trap*, in which a person escapes temporarily from a relationship partner's aversive behavior (e.g., whining) through a palliative response (e.g., soft talking), but ironically encourages the partner's future, instrumental repetition of the aversive behavior (i.e., coerciveness).

Figure 1. A-R-C Analysis Units for Hypothetical Event Sequence

A-R-C Unit	Messy Room (aversive)	Mother Scolds (aversive)	Child Whines (aversive)	Mother Talks (neutral)	Child Quiet (neutral)	Direct Effect of Consequence
1	A_1	R_1	C_1			Punish R_1 [Decrease $p(R_1|A_1)$]
2		A_2	R_2	C_2		Negatively Reinforce R_2 [Increase $p(R_2|A_2)$]
3			A_3	R_3	C_3	Negatively Reinforce R_3 [Increase $p(R_3|A_3)$]

A Dynamic Systems Model

Implicit in coercion theory is the notion that distinctive patterns of behavioral interaction can arise within families through self-organizing processes. The project described here attempts to realize this interpretation explicitly, by constructing a computer model that dynamically represents social learning processes within the context of parent-child interaction.

The model is implemented using spreadsheet software (van Geert, 1994). Figures 2 and 3 respectively present a diagram of the program flow and a portion of the spreadsheet layout.

Figure 2. Model Program Flow

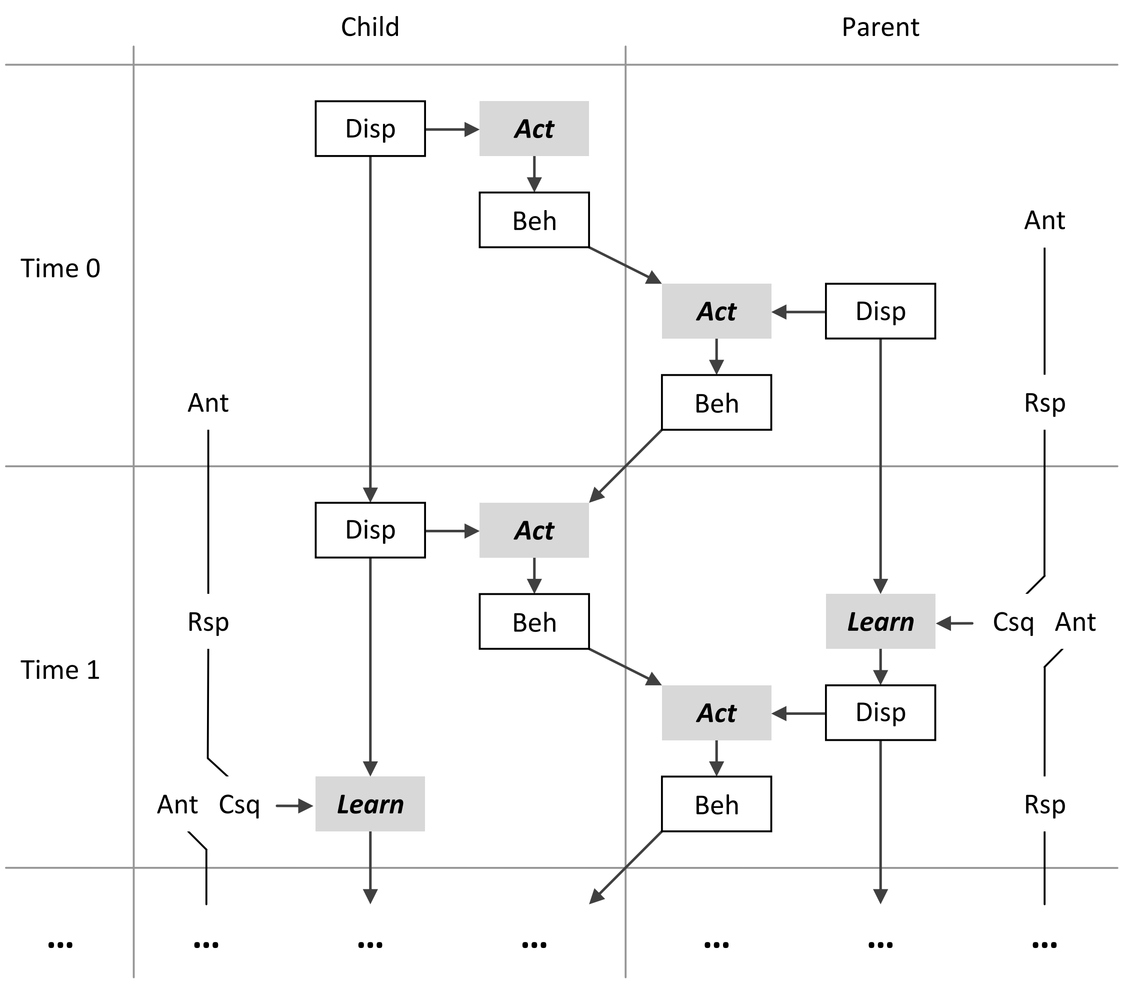

Note. Items labeled in plain font are data values; those in italic font are computational processes.

Figure 3. Model Spreadsheet Layout

Sample Columns for One Participant (Child or Parent)

Antecedent		Aversive (A)			Neutral (N)			Prosocial (P)			
Response		A	N	P	A	N	P	A	N	P	
		Learning Parameters (Consequence Valuations)									
Consequence	A	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	
	N	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
	P	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	
Time		Response Dispositions (Behavior Probabilities)									Behavior
0		0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33	A
1		0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33	N
2		0.33	0.33	0.33	0.34	0.31	0.34	0.33	0.33	0.33	P
...		...	...	...	...	...	...	...	...	...	...
N		0.08	0.60	0.32	0.02	0.55	0.44	0.15	0.37	0.48	N

The model produces a cascade of alternating parent and child behaviors, each defined as an *aversive (A)*, *neutral (N)*, or *prosocial (P)* event. Rows in the lower portion of the spreadsheet store these data. Each one contains the *response dispositions* (conditional behavior probabilities) and behaviors for an alternation cycle, in column groups allocated respectively to the parent and child. The computer regenerates the entire time series whenever the user instructs it to recalculate the spreadsheet.

The model subdivides the time series into sequences of a user-specified length. It randomly chooses the beginning behavior in each sequence using response dispositions averaged across antecedent types, and randomly chooses each subsequent behavior using response dispositions applicable to the relationship partner's preceding behavior. For example, the average of the child's dispositions $p(R_N|A_A)$, $p(R_N|A_N)$, and $p(R_N|A_P)$ specifies the probability that the child will behave neutrally at the beginning of a sequence, whereas the parent's disposition $p(R_P|A_N)$ specifies the probability that the parent will behave prosocially after the child behaves neutrally.

After each behavior that neither begins nor ends a sequence, the model adjusts response dispositions to reflect learning from experience. First it changes the

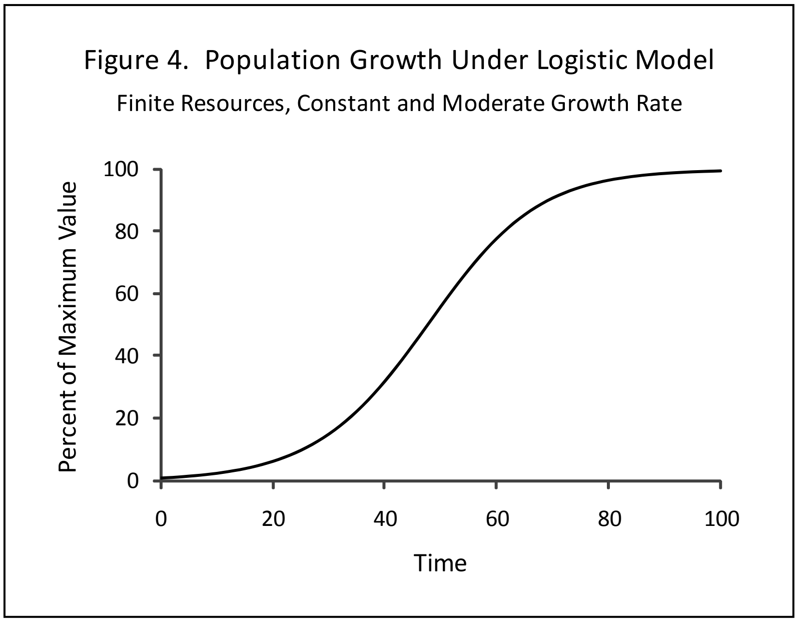

disposition for the occurring response, using a valuation of the consequence. (Rows in the upper portion of the spreadsheet store these valuations as *learning parameters*.) Then it rescales the other two dispositions associated with the antecedent, to restore a sum of unity for the group.

The model calculates the occurring response's disposition change in accordance with a *logistic* model of population growth under finite resources (van Geert, 1994). The disposition corresponds to a population size, and the consequence valuation to a growth rate, which increases or decreases the population size when it is respectively positive or negative.[1] Figure 4 illustrates this type of growth for a situation in which the population is initially small and the growth rate is a positive constant of moderate magnitude.

The model output of primary interest is the time evolution of its response dispositions (i.e., its *learning curves*). Two user-specified model inputs strongly affect this output, namely, the learning parameters and the initial (Time 0) response dispositions. One can alter these features to examine the model's behavior under various circumstances.

[1] For a consequence valuation *cval*, and an occurring response's disposition *disp*, the new disposition *ndsp* is calculated as follows, subject to a minimum of 0 and a maximum of 1: ndsp = disp + cval * disp * (1 - disp).

Example: Therapeutic Parenting

An important test of the model's plausibility is whether it yields developmental trajectories expected under conditions assumed by simpler models. The inputs for this example establish a scenario in which a stable environment transmits structure to a changing organism. I have termed it *therapeutic parenting* because it reflects an established conception of behavioral treatment for a child with conduct problems.

Figure 5 presents the inputs and outputs for an illustrative trial. The first panel shows the consequence valuations and the remaining panels show time series graphs for the response dispositions.

The child's consequence valuations are uniform across antecedent-response combinations; they produce an encouraging effect of prosocial consequence events, and a discouraging effect of aversive ones, on conditional responding. The parent's consequence valuations are uniformly zero; they suppress any effects of experience on conditional responding.

The child's initial response dispositions for all antecedents are biased heavily towards aversive behavior. The parent's initial response dispositions are biased heavily towards matching their antecedents, providing rewards and punishments respectively for the child's prosocial and aversive behaviors.

The series length is 2000 alternation cycles, which generates 4000 behaviors. The sequence length is five behaviors, which matches the length of Patterson's hypothetical example and (as an odd number) alternates sequence beginnings between child and parent. (These choices also apply to configurations introduced below.)

The trial presented evinces output features typical for this configuration. For each antecedent, the child's aversive response disposition sharply declines. Both neutral and prosocial dispositions initially rise, but eventually the prosocial displaces the neutral and locks itself into the maximum possible value of unity. The parent's response dispositions remain constant.

The upper portion of each time series graph shows the *entropy* of the response dispositions. This measure (Krippendorff, 1986) quantifies uncertainty in the random behavior generation. It increases as the response probabilities converge (making behavior more variable) and decreases as they diverge (making it less variable). Increasing or decreasing values respectively indicate an expansion or compression of systemic degrees of freedom; that is, dissolution or formation of an organized state.

For each antecedent, the child's entropy rises to a maximum as the child's initial behavior pattern dissolves, and falls to zero as a new (and completely rigid) pattern forms. The parent's entropy remains constant in accordance with unchanging response dispositions.

Figure 5. Sample Time Series for Therapeutic Parenting

Child's Learning Parameters

Antecedent		Aversive (A)			Neutral (N)			Prosocial (P)		
Response		A	N	P	A	N	P	A	N	P
	A	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10
Consequence	N	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	P	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10

Parent's Learning Parameters

Antecedent		Aversive (A)			Neutral (N)			Prosocial (P)		
Response		A	N	P	A	N	P	A	N	P
	A	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Consequence	N	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	P	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Response Dispositions for Neutral Antecedent

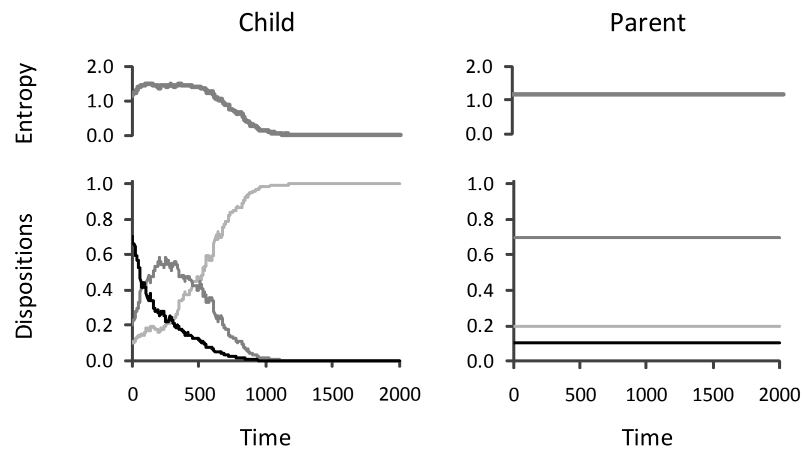

Response Dispositions for Aversive Antecedent

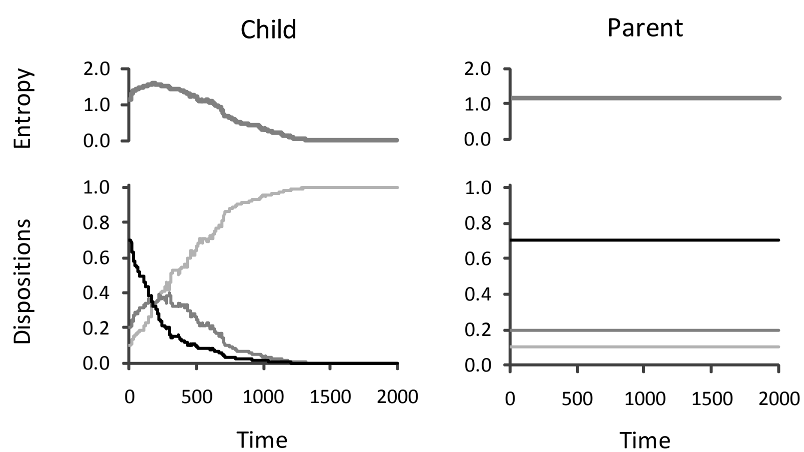

Response Dispositions for Prosocial Antecedent

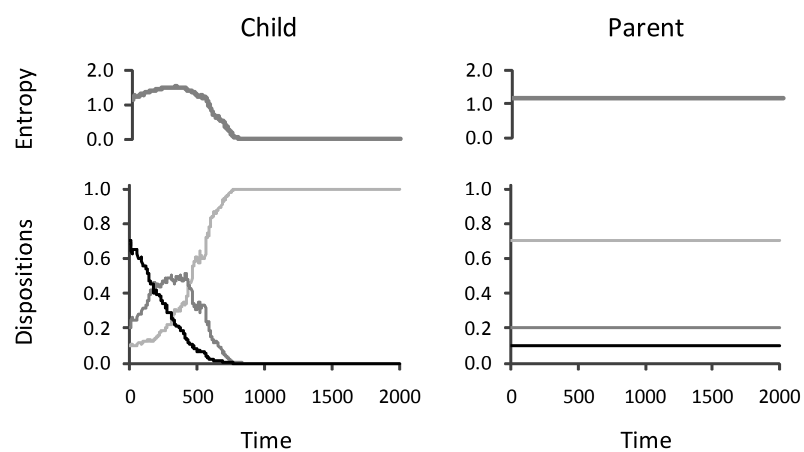

Dispositions: A N P

Example: Matched Learning Parameters

The model's performance attribute of primary interest is its ability, under appropriate circumstances, to exhibit self-organizing behavior. The inputs for this example establish a scenario in which order may emerge through the interaction of system components. I have termed it *matched learning parameters* because it frames the child and parent as active relationship partners with identical learning abilities.

Figure 6 presents the inputs and outputs for an illustrative trial. The panel arrangement replicates that of Figure 5. Consequence valuations for both partners replicate those for the child in the preceding example. Initial response dispositions have all been made identical, to establish thoroughly undifferentiated behavioral orientations.

The outputs for this model configuration exhibit substantial variety across trials, precluding the characterization of any one as typical. Those for this trial, however, evince several commonplace features. These include learning curves that grow more or less monotonically, reverse direction, meander, spurt, or freeze.

A crucial question regarding these curves is whether they are more than a disjointed assortment of *random walks* (van Geert, 1994). How might we discern the presence of genuine organization?

One indication may be linkages among learning curves belonging to different antecedents or partners. For example, the parent's neutral response dispositions peak distinctively in this trial, near the 750th alternation cycle, for prosocial and neutral antecedents. Because regularities of this type are not incorporated into the model's design, they appear to reflect actual synergies.[2]

The child's aversive response disposition, and the parent's neutral one, rise more or less steadily for aversive antecedents in this trial. This outcome pattern is consistent with the process of child antisocial behavior development that coercion theory posits. Trials likewise occur, however, in which the parent's aversive response disposition and the child's neutral one grow large for aversive antecedents.

Efforts to understand and summarize the model's performance by examining individual trials are susceptible to perceptual biases. One must investigate the model's performance systematically to identify robust outcome patterns and ascertain their implications for coercion theory.

---

[2] By contrast, relationships among the response dispositions for any particular antecedent do reflect deliberate design; their values necessarily sum to unity.

Figure 6. Sample Time Series for Matched Learning Parameters

Child's Learning Parameters

Antecedent		Aversive (A)			Neutral (N)			Prosocial (P)		
Response		A	N	P	A	N	P	A	N	P
Consequence	A	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10
	N	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	P	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10

Parent's Learning Parameters

Antecedent		Aversive (A)			Neutral (N)			Prosocial (P)		
Response		A	N	P	A	N	P	A	N	P
Consequence	A	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10
	N	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	P	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10

Response Dispositions for Neutral Antecedent

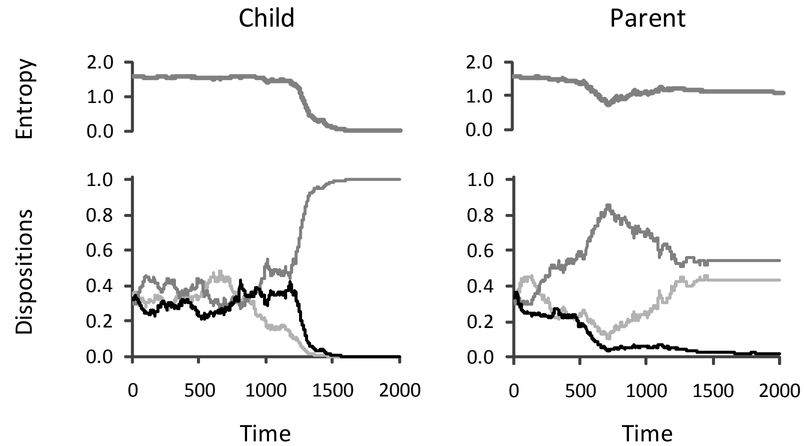

Response Dispositions for Aversive Antecedent

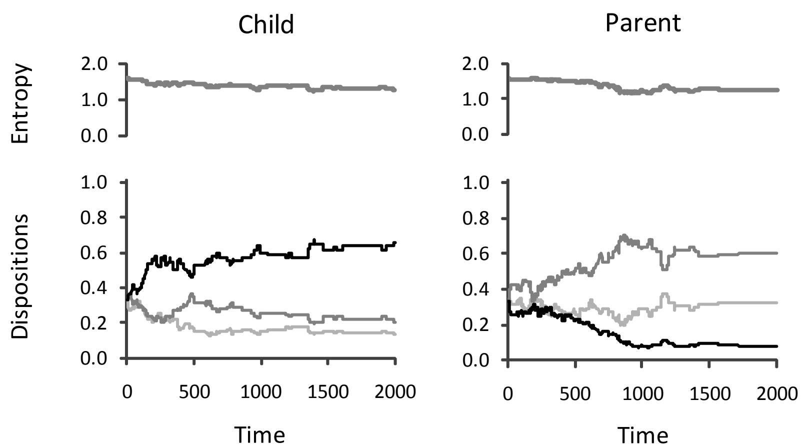

Response Dispositions for Prosocial Antecedent

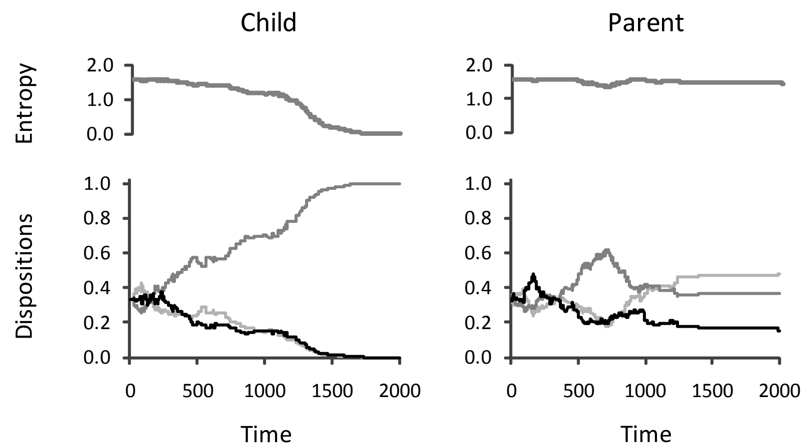

Dispositions: A N P

Simulation Study

Each model trial is a random experiment. In the matched learning parameters configuration, this experiment's outcome patterns are difficult (if not impossible) to deduce. The study described here seeks alternatively to discover outcome patterns by observing a large number of independent trials.

Uniform Learning

Of special interest in regard to coercion theory are response dispositions for aversive antecedents. Are particular dispositions likely to grow large? Are they likely to grow concomitantly, for other antecedents in the same individual or aversive antecedents in the relationship partner?

To answer these questions, I ran 2500 trials of the model configuration shown in Figure 6, and gathered final (Time 2000) variable values for analysis. I classified a response disposition as *large* if it was more than or equal to 0.67, and thus more than twice the size of either remaining disposition for the same antecedent.

Table 1 presents frequencies (rounded to the nearest integer) of large response disposition occurrences per 100 trials. Entries on the main diagonal count overall occurrences by the model's eighteen distinct dispositions; those off the main diagonal count joint occurrences by pairwise combinations of these dispositions. For example, the third and sixth columns of the third row respectively indicate that the child developed large prosocial response dispositions for aversive antecedents on 16 of every 100 trials, and for neutral as well as aversive antecedents on 11 of these 16 trials.

Entries in Table 1 are grouped by antecedent into 3 x 3 subtables. For each antecedent, the three possible types of large response disposition and the possible absence of any large disposition define a complete set of mutually exclusive outcome categories. One may thus consider each subtable to be the first three rows and columns of a 4 x 4 contingency table that cross-classifies all trial outcomes; the fourth row and column of this implicit table count outcomes lacking large dispositions. (Figure 7 presents the completion of one such contingency table.)

A strong similarity in the outcome profiles of the child and parent is discernible in Table 1. This similarity is evident in the numerical proximity of entries indexed by interchanging the roles of child and parent. It reflects both the model's configuration symmetry (each partner is equally likely to evolve in a particular way on a particular trial) and the statistical precision achieved by running numerous trials (matching entries estimate the same population value, and the standard error of this estimate is less than 1).

Examination of the main diagonal frequencies reveals only modest differences among those for aversive antecedents ($2 \leq |d| \leq 6$). More interesting is a progressive shift in the balance among large dispositions to match the social tone of antecedents. This pattern is discernible among the lightly shaded bars of Figure 8, which chart frequencies from the main diagonal of Table 1. Correlation coefficients, computed between disposition type and antecedent type by coding aversive, neutral, and prosocial categories respectively as -1, 0, and 1, reflected a moderate but statistically robust association (for child and parent each, $r = .15$, $p < .0001$).

Examination of the off-diagonal disposition frequencies in Table 1 reveals that large dispositions of each type tended to evolve concomitantly across antecedents within an individual. This pattern is evident in the elevation of frequencies on the main diagonals of intra-individual subtables. Unweighted kappa statistics (Fleiss, 1981), computed for the corresponding 4 x 4 contingency tables, confirmed that these frequencies were greater than expected by chance ($.44 \leq \kappa \leq .56$, $p < .0001$).

Patterns linking response dispositions between partners are more complex, but two are highly distinctive. Frequencies of 0 in the tenth column of the first row, and of 26 in the eighteenth column of the ninth row, respectively indicate that large aversive dispositions for aversive antecedents did not evolve concomitantly across relationship partners, whereas large prosocial dispositions for prosocial antecedents did.

Enhanced Escape-Avoidance Learning

An important element of coercion theory is the robustness of the [negative] reinforcement trap. Patterson asserted (pp. 144-148) that acquisition efficiency and extinction resistance of escape avoidant responding contributed significantly to the development of antisocial behavior. The model's explicit control of learning rates through parameter settings permits an exploration of its performance in this respect.

Accordingly, I repeated the procedure described in the preceding section, after changing the valuations of aversive, neutral, and prosocial consequences for aversive antecedents respectively from -0.10, 0.00, and 0.10 to -0.05, 0.05, and 0.15. To distinguish the original and altered configurations, I termed the former *Uniform Learning (U)* and the latter *Enhanced Escape-Avoidance Learning (E)*.

Table 2 presents outcome data from these trials in the same format as Table 1. The heavily shaded bars of Figure 8 chart frequencies from its main diagonal.

Large dispositions occurred more frequently in the enhanced than the uniform learning condition for aversive (but not other) antecedents (for the child, $f_E = 75$, $f_U = 58$, $d = 17$, $p < .0001$; for the parent, $f_E = 78$, $f_U = 58$, $d = 20$, $p < .0001$). Their distribution also shifted towards more socially negative types for each antecedent, as demonstrated by correlations between disposition type (coded as above) and learning

Table 1. Simulation Study for Matched Learning Parameters: Uniform Learning
Large Response Dispositions per 100 Trials (Grouped by Antecedent)

	Child									Parent								
	Aversive (A)			Neutral (N)			Prosocial (P)			Aversive (A)			Neutral (N)			Prosocial (P)		
	A	N	P	A	N	P	A	N	P	A	N	P	A	N	P	A	N	P
Child																		
Aversive																		
A	20	0	0	12	3	1	10	4	3	0	3	6	1	3	5	0	2	3
N		22	0	1	16	1	0	16	3	3	1	2	0	3	11	3	2	3
P			16	1	1	11	0	0	13	6	1	3	6	2	4	0	3	13
Neutral																		
A				15	0	0	7	3	3	0	0	6	1	0	4	0	0	3
N					25	0	1	17	3	3	3	2	0	5	11	2	3	3
P						24	0	1	13	5	9	3	4	10	4	0	11	12
Prosocial																		
A							12	0	0	0	3	0	0	2	0	0	2	0
N								24	0	2	2	4	0	3	12	2	3	1
P									27	4	2	13	3	3	14	0	1	26
Parent																		
Aversive																		
A										21	0	0	12	3	1	10	4	4
N											20	0	1	16	1	0	16	2
P												17	0	1	13	0	0	13
Neutral																		
A													16	0	0	8	3	3
N														24	0	1	17	2
P															26	0	1	14
Prosocial																		
A																13	0	0
N																	23	0
P																		27

Note. Counts are based on 2500 trials. The standard error for each entry in this table is less than 1.

Table 2. Simulation Study for Matched Learning Parameters: Enhanced Escape-Avoidance Learning
Large Response Dispositions per 100 Trials (Grouped by Antecedent)

	Child									Parent								
	Aversive (A)			Neutral (N)			Prosocial (P)			Aversive (A)			Neutral (N)			Prosocial (P)		
	A	N	P	A	N	P	A	N	P	A	N	P	A	N	P	A	N	P
Child																		
Aversive																		
A	36	0	0	18	7	4	21	5	4	0	16	4	0	8	5	1	7	4
N		31	0	2	15	5	1	14	5	17	7	2	3	10	9	7	6	5
P			9	0	0	6	0	0	6	5	1	1	3	1	2	0	2	6
Neutral																		
A				22	0	0	17	1	2	1	3	3	3	0	2	2	0	2
N					24	0	1	16	2	9	9	1	0	7	9	4	3	2
P						18	0	1	10	6	9	1	2	9	3	0	8	9
Prosocial																		
A							22	0	0	1	6	0	1	3	0	1	3	0
N								20	0	8	6	2	0	3	9	3	3	0
P									18	5	5	6	2	2	10	0	1	18
Parent																		
Aversive																		
A										39	0	0	21	8	4	23	5	5
N											31	0	3	15	5	1	13	5
P												8	0	0	6	0	0	6
Neutral																		
A													24	0	0	19	1	2
N														23	0	1	15	2
P															19	1	1	10
Prosocial																		
A																25	0	0
N																	19	0
P																		18

Note. Counts are based on 2500 trials. The standard error for each entry in this table is less than 1.

Figure 7. Contingency Table
Completion for Excerpt from Table 1

Child: Aversive Antecedent	Child: Neutral Antecedent: A	N	P	None	Total
A	12	3	1	4	20
N	1	16	1	4	22
P	1	1	11	4	16
None	2	5	11	25	42
Total	15	25	24	36	100

Note. Arithmetic discrepancies reflect rounding error.

Figure 8. Simulation Study for Matched Learning Parameters:
Large Response Dispositions per 100 Trials (Grouped by Antecedent)

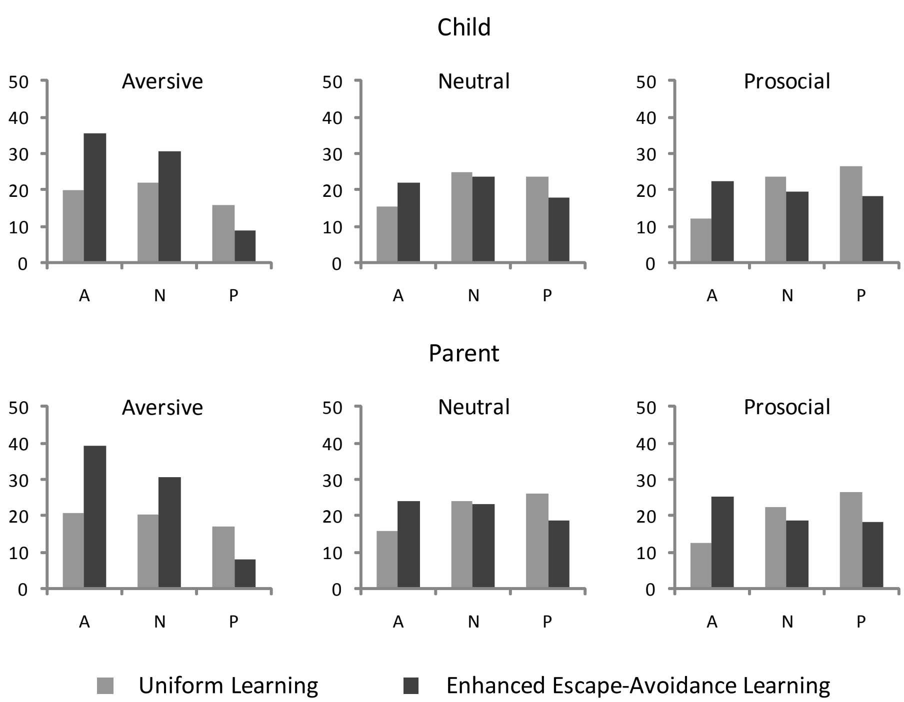

condition (for the child, $r_A = .19$, $r_N = .12$, $r_P = .19$, $p < .0001$; for the parent, $r_A = .22$, $r_N = .15$, $r_P = .21$, $p < .0001$). Correlations between disposition type and antecedent type in the enhanced condition virtually equaled those in the uniform condition (for child and parent each, $r = .16$, $p < .0001$). These patterns converged in markedly elevated frequencies of aversive and (less prominently) neutral dispositions for aversive antecedents.

Three patterns observed in the uniform learning condition continued in the enhanced condition. Large dispositions of each type tended to evolve concomitantly across antecedents within an individual ($.44 \leq \kappa \leq .57$, $p < .0001$); and large aversive dispositions for aversive antecedents did not evolve concomitantly across relationship partners, whereas large prosocial dispositions for prosocial antecedents did.

A novel pattern linking response dispositions between partners emerged in the enhanced condition. Frequencies of 16 in the eleventh column of the first row, and of 17 in the tenth column of the second row, indicate that large neutral dispositions for aversive antecedents were likely to evolve in an individual when large aversive dispositions for aversive antecedents evolved in the relationship partner.

Discussion

The computerized model of parent-child interaction presented in this paper situates social learning principles within a dynamic systems framework. Early applications of these principles treated development as a unilateral process, but Patterson's antisocial behavior studies demonstrated their compatibility with the concept of reciprocal influence. Dynamic systems modeling of his theory constitutes a natural extension of his effort.

Model Performance

This paper's model needed to perform two complementary tasks. The first was to generate outcomes expected from simpler developmental theories in a configuration reflecting their assumptions. The *therapeutic parenting* example achieved this by treating unilateral influence as a special case of reciprocal influence. Organized behavior in a parent whose dispositions were stabilized reliably changed organized behavior in a child who learned from experience.

The model's second task was to generate outcomes expected from more advanced developmental theories in configurations reflecting their greater generality. The *matched learning parameters* example partly achieved this by orchestrating truly reciprocal influence. Organized behavior seemed to evolve from undifferentiated behavior in partners equally suited to learn from experience, and it seemed to exhibit synergy because linkages occurred among model components that were free to vary independently. Diverse trial outcomes, however, marked a need for systematic investigation.

Simulation Study

A study conducted to meet this need tabulated large response dispositions occurring over multiple trials, in each of two model configurations. One of them, termed *uniform learning*, stringently tested self-organizing behavior; it equalized learning efficiencies across responses to all antecedents. The other, termed *enhanced escape-avoidance learning*, specifically tested the role of negative reinforcement in promoting coercive behavior; it biased learning towards the acquisition and maintenance of responses yielding escapes from aversive antecedents.

Statistically robust patterns emerged in the outcomes for both configurations. One pattern was a moderate but reliable tendency of response dispositions to match antecedents in social tone. More striking were tendencies of aversive, neutral, or prosocial dispositions to generalize across antecedents within an individual; and of prosocial dispositions for prosocial antecedents, but not aversive dispositions for aversive antecedents, to generalize across relationship partners.

Robust differences also emerged between outcomes in the uniform condition and those in the enhanced condition. Greater frequencies of aversive and neutral dispositions for aversive antecedents occurred in the latter, along with a socially negative shift in disposition types for all antecedents; and neutral dispositions for aversive antecedents tended to accompany aversive ones for aversive antecedents across relationship partners.

These outcome patterns strongly imply that the model exhibited self-organizing behavior in both configurations: their mere existence discounts the view of its learning curves as unconnected random walks. Questions that remain are how this organization occurred, and what the patterns it yielded imply for coercion theory.

Self-Organization

One might answer the first of these questions by reviewing the model's design. Of particular importance is the relationship between response generation and learning. The former process requires (at most) a single preceding behavior. The latter process requires a sequence of three behaviors formed by chaining successive responses. What begins as fortuitous learning may therefore become systematic, as predictability in response generation becomes predictability in sequence formation.

Complementary aspects of this predictability may foster the differentiation of dispositions. Reliability in an individual's response to a particular antecedent promotes learning opportunities for the antecedent-response combination, whereas reliability in the relationship partner's ensuing reaction yields consistent consequences for it. If these consequences provide reinforcement, the antecedent's recurrence favors the combination's disposition growth.

Of special interest in this regard are circumstances in which the relationship partner's own developing dispositions promote the antecedent's recurrence. Positive feedback resulting from this type of interaction seems especially likely to foster the growth of large response dispositions.

Consequence reliability may furthermore promote the generalization of large dispositions across antecedents within an individual. A response that elicits a given consequence may be preceded by different types of antecedent, consigning the disposition for each antecedent-response combination to a similar fate.

Consequence valuations interacting with learning processes seem likely to foster distinctive disposition linkages between partners. For prosocial antecedents in particular, concomitant disposition growth in the partners is compatible with both of them having prosocial dispositions, as each one thereby provides the other with positive reinforcement. For aversive antecedents, however, concomitant disposition growth in the partners is incompatible with both of them having aversive dispositions, as neither one would provide the other with negative reinforcement.

Coercion Theory

The simulation study outcome patterns have several implications for coercion theory. First, the finding that enhanced escape-avoidance learning fostered coercive responding was consistent with the theory's predictions. It thereby demonstrates the possibility (but not the necessity) of coercive behavior emerging through reciprocal interaction governed solely by social learning principles.

Second, the finding that large response dispositions generalized across antecedents within an individual suggests that social learning principles can also explain a broadening of previously established aversive behavior to novel contexts. More generally, it suggests that the modeling process initiated here could, with further development, serve to generate novel predictions from coercion theory.

Third, the finding that coercive behavior *never* evolved concomitantly across relationship partners presents a challenge to the model and (perhaps) the theory. One might ask what model changes would accommodate mutually aversive interactions. A possible approach would be to define multiple classes of aversive, neutral, or prosocial behavior, thereby allowing different interaction patterns to form for different classes. Escalating exchanges of aversive behavior of the kind that Patterson observed in referred families (pp. 155-162), however, are evidently a phenomenon incompatible with the model's present design.

Fourth, the finding that each relationship partner was equally likely to develop coercive behavior highlights the possible role of negative reinforcement in promoting harsh or punitive parenting practices.[3] The unilateral character of this coercion starkly demonstrates how such practices might evolve in the parenting of compliant children (Miller, 1983, 1990). Thus, although the principles of coercion theory were first explored and articulated in the context of antisocial child behavior, they have important applications in the context of child maltreatment.

Conclusion

This paper presented a computerized model of parent-child interaction that situated ideas from Gerald R. Patterson's coercion theory within a dynamic systems framework. In a simulation study, the model generated plausible and comprehensible outcome patterns that were largely consistent with predictions from the theory. The findings implied that essential elements of coercion theory are compatible with a conceptualization of development as a self-organizing process, and with the use of computational modeling as a tool of theoretical exploration.

[3] Notwithstanding my knowledge that identification of the relationship partners as child and parent in the matched learning parameters configuration was arbitrary, I found myself initially reacting to model trials in which coercive behavior developed in the parent, thinking, "It's not supposed to work that way!"

References

Caspi, A., & Moffitt, T. E. (1995). The continuity of maladaptive behavior: From description to understanding in the study of antisocial behavior. In D. Cicchetti & D. J. Cohen (Eds.), *Developmental psychopathology, Vol. 2: Risk, disorder, and adaptation. Wiley series on personality processes* (pp. 472-511). New York, NY, USA: John Wiley & Sons.

Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions* (2nd ed.). New York: John Wiley & Sons.

Hinshaw, S. P., & Anderson, C. A. (1996). Conduct and oppositional defiant disorders. In E. J. Mash & R. A. Barkley (Eds.), *Child psychopathology* (pp. 113-149). New York, NY, USA: Guilford Press.

Krippendorff, K. (1986). *Information Theory: Structural Models for Qualitative Data*. Newbury Park: Sage Publications, Inc.

Maccoby, E. E., & Jacklin, C. N. (1983). The "person" characteristics of children and the family as environment. In D. Magnusson & V. L. Allen (Eds.), *Human development: An interactional perspective* (pp. 75-91). San Diego, CA: Academic Press.

Maughan, B., & Rutter, M. (1998). Continuities and discontinuities in antisocial behavior from childhood to adult life. *Advances in Clinical Child Psychology, 20*, 1-47.

Miller, A. (1983). *For your own good: hidden cruelty in child-rearing and the roots of violence*. New York: Farrar, Straus, Giroux.

Miller, A. (1990). *The drama of the gifted child*. New York: BasicBooks.

Olweus, D. (1980). Familial and temperamental determinants of aggressive behavior in adolescent boys: A causal analysis. *Developmental Psychology, 16*(6), 644-660.

Patterson, G. R. (1982). *Coercive family process*. Eugene, OR: Castalia.

Sameroff, A. J. (1975). Early influences on development: Fact or fancy? *Merrill-Palmer Quarterly, 21*(4), 267-294.

Sameroff, A. J. (1995). General systems theories and developmental psychopathology. In D. Cicchetti & D. J. Cohen (Eds.), *Developmental psychopathology, Vol. 1: Theory and methods. Wiley series on personality processes* (pp. 659-695). New York, NY, USA: John Wiley & Sons.

Thelen, E. (1989). Self-organization in developmental processes: Can systems approaches work? In M. R. Gunnar & E. Thelen (Eds.), *Systems and development. The Minnesota symposia on child psychology, Vol. 22* (pp. 77-117). Hillsdale, NJ, USA: Lawrence Erlbaum Associates, Inc.

van Geert, P. (1994). *Dynamic systems of development: Change between complexity and chaos*. London, England UK: Harvester Wheatsheaf.

Postscript

Although I had intended merely to reformat my poster as a paper, I found myself repeatedly discovering needs for improvement. I resolved this conflict in favor of producing a paper I considered useful to read that was nonetheless faithful to the substance of the poster. Here are the principal changes.

In the poster model, I calculated growth for the occurring response's disposition using the product of the consequence valuation and all three response dispositions for the antecedent. I now think it better to treat the disposition straightforwardly as an independent grower. This approach yields faster growth, so I have reduced the magnitudes of consequence valuations.

In the poster model, I generated a single interaction sequence that spanned 2000 alternation cycles. The current, more plausible approach yields slower growth, presumably because learning occurs on fewer occasions.

In the poster simulation study, I generated 300 trials each for the uniform and enhanced escape-avoidance learning scenarios, and my analysis comprised mainly scatter plots illustrating relationships among response dispositions. The current, more rigorous approach yields findings consistent with but more extensive than the earlier one.

I have lightly edited the introductory sections and thoroughly rewritten the others. I have also added tables and figures for the expanded analysis.