Endogenous Transition Dynamics in Corruption: An Agent-Based Computer Model

Ross Hammond*

Center on Social and Economic Dynamics
Working Paper No. 19
December 2000

Abstract

Corruption is modeled as a game-theoretic micro-level interaction, using an agent-based computer model with heterogeneous agents. Emergent macro-level behaviors differ from traditional literature on the subject, and suggest an endogenous social transition from a high-corruption state to a low-corruption state is possible. The paper explores the conditions necessary for such a transition, as well as related dynamics in the model.

---

* For helpful conversations and comments throughout the research process, the author thanks Jonathan Conning and Douglas Gollin at Williams College and Robert Axtell and Joshua Epstein at Brookings. For technical feedback and support, special thanks to Miles Parker.

The author is currently with the Brookings Institution, Washington, DC 20036.
E-mail: rhammond@brook.edu or rhammond@wso.williams.edu

Corruption in any society is the result of both individual choices and social norms. Together, these give rise to an equilibrium social level of corruption. This equilibrium level differs in various regions of the world-- from very high to very low. In some specific cases, a *transition* from very high levels of corruption to relatively low levels can be observed over time.

The majority of existing models of corruption in economics and political economy, however, do not focus explicitly on such a transition case, but rather assume static or stable corruption and focus instead on its effects and internal processes[1]. Existing models that do address transition almost exclusively explain a shift in corruption levels as the result of a one-time, *exogenous* "shock" to the original system in which corruption thrived[2]. In such models, changes in legal systems and enforcement structures—the result of deliberate social policy-- drive transitions.

The present study offers an alternative explanation for shifts in corruption levels, based on a new, agent-based model of corruption as a simple, game-theoretic repeated interaction on the micro level over time. The emergent social behaviors described here differ markedly from the results of previous models, in that transitions occur *endogenously*, without any change in system or structure. Instead, this endogenous transition occurs as the result of a cascade of micro-level events, set in motion by a chance coincidence of interactions. This result suggests that apparently stable corruption can, in fact, be *un*stable and break down over time. Indeed, it can be postulated that a transition to lower levels of corruption will inevitably occur, under certain necessary conditions.

[1] Acemoglu and Verdier 1997, Dey 1989, Gray and Kaufmann 1998, Rose-Ackerman 1978, Goudie and Stasavage 1997, Klitgaard 1988 and 1995, Leff 1964, Beenstock 1979, Andvig 1991, Bardhan 1997, et al.

[2] Acemoglu and Verdier 1997, et al.

Model structure

The model is based around a simple game-theoretic basic framework similar to that of previous researchers[3], but adds a substantial degree of heterogeneity and opportunities for dynamic behavior.

Two similar but distinct agent populations are created—citizen agents (Citizens) and government agents (Bureaucrats). Every round, each citizen agent will "play" a randomly chosen government agent in a simple simultaneous game with fixed (pure) strategies—"Corrupt" (C) or "Non-Corrupt" (NC). The interaction is modeled on a real world situation resembling tax collection, with potential private gains to be realized through successful collusion[4]. Each agent decides his or her strategy immediately before each interaction, according to a decision rule and payoff matrix discussed in detail below. Play is simultaneous once strategies have been chosen. Corrupt agents risk going to "jail" for a fixed sentence if apprehended, during which time they are removed from play and cannot interact or gain payoffs. The enforcement mechanism is also detailed below.

Agent Characteristics

Both groups are populated using a basic agent "prototype" with the same basic sets of characteristics and information, although the populations are heterogeneous (specific contents of these characteristics differ). Each agent has a **social network** of "friend" agents of his/her same

---

[3] Sah 1988, Cadot 1987

[4] In the case of a tax collection, a strategy of "Non-Corrupt" would represent simple compliance (payment) for the citizen and honest reporting (and transfer of tax revenue) for the bureaucrat; a "Corrupt" strategy, if chosen by both parties, would involve collusion in which the citizen pays the bureaucrat a bribe *less* than the taxes due, which the bureaucrat keeps for himself or herself and does not pass along to society.

type (citizen or bureaucrat). These networks are of fixed (standard) size, but the specific contents of each agent's network is randomly assigned during initialization. The agent uses the social network to dynamically gather information about his/her surroundings (particularly with regard to enforcement)—this information is a very important part of the motivation for agent's behavior. Each agent also has a **memory** of past interactions (again of a fixed size). The memory contains a record of the strategy chosen by the opponent in each of the last *n* interactions, and is filled with (uniform) random values during initialization before the first round of play. Finally, an agent has an individual **inherent predisposition** to be either more corrupt or more honest (implementation explained in more detail in the decision rule). This predisposition is randomly (uniformly) distributed through the population of agents.

Payoff Matrix

The micro corruption game has a 2x2 payoff matrix as follows:

	C	NC
C	x	y
NC	y	y

Only when *both* players choose "Corrupt" can they realize the gains from successful collusion. In this case, both players are assumed to be personally better off (receiving $\boldsymbol{x} > \boldsymbol{y}$). This basic game offers solutions under traditional game theory[5].

However, the game is more complicated. Although the payoff x in the above matrix is of fixed value for the duration of the game, it is *perceived* differently by each individual player. Agents in this model are moral beings, and each individual has a characteristic **Inherent Propensity for Honesty** (***i***)—a decimal between 0 and 1 (randomly assigned at the agent's creation during model initialization) which is intended to represent that agent's "moral stance". This decimal is used to measure the "moral cost" a more honest agent incurs by choosing a corrupt strategy. For an ***i*** of 1 (perfect Honesty), this cost is assumed to be very high, while for an ***i*** of 0 (perfect Corruption) it is assumed to be zero. Each agent, then, perceives their own ***x*** to be:

$$\boldsymbol{x_i = (1\text{-}i)\, x}$$

In other words, a perfectly Honest agent gains zero benefit from a corrupt action, a perfectly corrupt agent receives the true ***x*,** and everyone else receives a proportionally diminished ***x***. In all cases except that of successful corrupt collusion, the payoffs are the same (***y***,***y***).

"Mismatches" and the Enforcement Technology

In the case of a mismatch {C,NC} or {NC,C}, there is an additional consequence, built into the model but *uncertain* from the agent's point of view. The honest agent who chose "NC" is

[5] Under traditional game theory, {C,C} is a strict Nash equilibrium and {NC,NC} is a simple Nash; under evolutionary game theory, {C,C} is evolutionary stable, while {NC,NC} is not. An analytical solution of this basic game under evolutionary game theory, then, would predict an "invasion" by C. (π_1 (C) is > π_1 (H) ; π_2 (C) is > π_2 (H))

assumed to "turn in" the corrupt agent who offered to take or receive a bribe-- either through a direct action (e.g., spreading rumors or a report to a superior) or through a willingness to reveal or confirm this agent's corrupt activities when questioned.

In the model, an enforcement entity tracks each agent's actions, noting each time the agent is turned in for corruption by an honest opponent. At the end of each round, this "global policeman" examines all citizens and bureaucrats. If an agent has exceeded some specified number of reported instances of corruption, this agent is sent to jail for a finite number of periods.

The agents themselves, however, lack information on exactly how the system of enforcement works. They are aware that it is possible to go to jail for corrupt activities-- indeed they may have observed their "friends" being sent to jail for such activities. However, they have no real certainty about how likely this outcome might be. They do *not* know the exact number of offenses that will lead to incarceration; indeed, they do not even know that the enforcement mechanism works in this way. Such knowledge would give them a much too precise and scientific measurement of something as (realistically) uncertain in their lives as enforcement effectiveness.

Instead (and perhaps more realistically), the agents make an educated guess of their chances of being sent to jail for a corrupt act based on *their observations* of the world around them. Agents use their social network (of "friend" agents), and their memory of previous interactions-- observations that may change as time progresses in the model—to calculate probabilities and construct a weighted payoff, as detailed below.

The Decision Rule

Each agent first checks on the current set of payoffs ($\boldsymbol{x}$ and $\boldsymbol{y}$), and on the current length of the jail term (set during the initialization of the model by the user, and known to all agents).

Next, the agent computes a subjective probability assessment, **A**, of the likelihood of encountering a corrupt agent. This involves an analysis of their memory (of fixed size **N**). The probability assessment is based on the number of corrupt agents, **n**, with whom the agent has been matched in **N** previous interactions; i.e. **A = n/N**.

Agents also analyze their social network-- a list of "friend" agents (always within their own population... fellow Citizens or fellow Bureaucrats) whose actions and status they monitor. While no agent has access to global (population) data about the total number of corrupt agents or agents in jail, each agent *does* have access to this information about all "friend" agents.

Agents count how many of their "friends" are in jail, and how many were corrupt in the last round. They use these two values to construct a subjective estimate, **B** of their *perceived chances* of being caught for a corrupt action this round. If **m** denotes the number of the agent's friends in jail and **M** is the number who were corrupt in the last round, then **B = m/M**.

The agent uses all of the above pieces of information (**A**, **B**, $\boldsymbol{x}$, $\boldsymbol{y}$) to make a quasi-rational decision by constructing weighted payoffs for each of the two strategies available. An agent of type $\boldsymbol{i}$ faces an expected payoff from corruption of:

$$\mathbf{(1\text{-}B)\ [A}\boldsymbol{x_i}\mathbf{ + (1\text{-}A)}\boldsymbol{y}\mathbf{]\ \ +\ \ B\ [}\boldsymbol{y}\mathbf{ - k}\boldsymbol{y}\mathbf{]}$$

where **B** is the subjective probability that the agent will be caught if s/he is corrupt; **A** is the subjective probability of being matched with a corrupt agent; and **k** is the length of the jail term. Agents, if apprehended, are removed from play for the duration of the jail term—receiving the payoff $\boldsymbol{y}$ before beginning their sentence. The cost of a jail sentence is taken to be **k**$\boldsymbol{y}$.[6]

[6] This is a lower bound on the agent's opportunity cost of going to jail (an agent loses *at least* $\boldsymbol{y}$ for each of **k** rounds spent in jail). An upper bound would be **k**$\boldsymbol{x}$, yielding an expected payoff to corruption of **(1-B)**

The "weighted" payoff for "**H**onesty" is simply ***y***.

The agent, then, makes a rational welfare-maximizing decision by comparing the weighted payoffs and choosing the higher of the two.

Assumptions of limited information

Implicit in the model are several important knowledge restrictions placed on both global and local information.

Each agent knows the payoff matrix (as influenced by his/her own individual inherent propensity), and each agent knows the length of the jail term imposed on those apprehended for corruption.

Agents are *not*, however, aware of the exact nature or *structure* of enforcement efforts. Agents do not know how many times they have been "turned in" previously, nor do they know the exact enforcement limit beyond which they will go to jail. Moreover, agents do not know that such a limit exists at all—they are unaware that each mismatch counts as one additional "strike" against them, building toward a fixed limit at which they will be apprehended. Instead, agents make estimations of their chances of being caught in any given round by observing the environment around them (through their social network).[7]

Agents also lack data on aggregate facts. Agents do not know the size of the population at any given point, or any sort of statistic (number of agents choosing corrupt, number in jail,

---

$\mathbf{[Ax_i + (1\text{-}A)}\boldsymbol{y}\mathbf{] + Bk}\boldsymbol{x}$. The exact value of this cost depends on future (random) interactions, and is thus complicated to calculate—using the lower bound makes no qualitative difference to my results.

[7] This restriction on agents' knowledge seems a realistic one. Real-world agents rarely, if ever, have complete information about the exact mechanism of enforcement or have a mathematically rigorous way to determine their exact chances of getting caught. Instead, they might base decisions about corrupt behavior on the basis of observation of their society and surroundings, and resulting crude estimates of risk levels.

distribution of inherent propensities, etc.) for the population as a whole.[8] Instead, they possess some such information about their own circle of friends, which they use to subjectively estimate the state of the larger population.

Agents in this model are boundedly rational—they do not have unlimited capacity for computation and information storage. Instead the tools they use to construct the weighted probabilities in their decision rule (memory and social network) are of limited (bounded) size[9].

Model Results and Discussion

Many sets of parameterizations for the above model (those with relatively "low" payoff ratios and relatively shorter jail terms), simply result in an equilibrium of rather immediate and stable honesty. However, changes in both model parameters and in basic model assumptions yield quite different results.

Endogenous Transition

One particular class of parameterizations yields especially interesting results. With a fairly substantial payoff ratio and with a fairly short jail term, a "transition to honesty" emerges:

After an initial period of fluctuation in the first few rounds, Corruption (with a few holdouts with very high inherent propensity to be Honest) dominates both populations. The actual number of corrupt agents fluctuates slightly, but remains the vast majority (circa 90%) of each

[8] This, again, seems a realistic assumption—most real-world people acknowledge that there is some continuum of more or less "honest" people in society, but few (if any) could tell you the exact distribution of population within this continuum. Accurate aggregate statistical information (especially where corruption is concerned) is often hard to come by.

[9] The argument for such "bounded rationality" has been made convincingly elsewhere (Simon et al.).

population. The number of agents in jail from each population also varies at a fairly low level, with occasional peaks at slightly higher levels.

This seemingly stable state persists for some period of time (varying in each run of the model) until, **inevitably**, a "fault point" is reached. At this point, a spike in jailed agents from one of the populations triggers a chain reaction as follows:

1. the jailed agents in the peak of "arrests" happen to constitute a large portion of the social network of a noticeable number of agents
2. these agents change their strategy to honesty as a result
3. this change results in more mismatches, and thus more jailed agents
4. this accentuates the "fear" transmitted through social networks: more and more agents fear apprehension as larger percentages of their social networks end up in jail; they shift their behavior to honesty accordingly
5. the cycle feeds back and intensifies

The net result is that, within 5-10 rounds, all agents of both populations have become Honest and will continue to choose Honesty indefinitely. The above transition can be seen in Figure 1. On the x-axis are iterations ("rounds"), essentially representing time. On the y-axis is the number of agents in a given population (here out of a population of 300) meeting each of the four color coded criteria in the key at the bottom—the purple line represents the number of corrupt citizens at any point in time, the pink line the number of citizens in jail, and so on.

This result represents an **endogenous** "transition" between two equilibria (Corruption and Honesty) with a rather abrupt shift, *in the absence of any structural change or exogenous shock*. The transition takes more or less time (rounds), as summarized in Table 1[10]-- but it *always*

[10] This table represents 35 separate runs of the model (with a different random seed each time). Thirty-five runs is a sufficient number to smooth out statistically irregularities and provide an accurate sample of the model's behavior. The first column in Table 1 represents the number of the run (1 to 35). For each run, column #2 represents the iteration (round) at which the transition to honesty was completed (in both

arrives, given enough time. Corruption, then, is internally unstable as an equilibrium in this result—a result which contradicts the implicit conclusion of existing political economy and economics literature that transitions from corruption to honesty are always the result of exogenous changes in government policy or structures. My model suggests that dramatic transitions to honesty can be endogenous and spontaneously occurring, and do not necessitate any external policy or structural change. This has fundamentally optimistic implications for countries currently experiencing high levels of corruption, and suggests that such corruption may not be as internally stable as first appears.

The exact parameters generating this result are detailed in Table 2. The basic transition result responds to marginal changes in the payoff ratio, the jail term, and the number of mismatches resulting in jail—the transition point is, on average, longer or shorter depending on these changes, but the basic behavior is qualitatively unchanged.

Of additional interest is the stability of Honesty as an equilibrium, once reached. Figure 1 illustrates only the transition itself, but Honesty is a stable equilibrium and continues indefinitely (out to thousands of rounds). This is in keeping with a traditional "development" story in which a norm of honesty, once achieved, never reverts to corruption.[11] In the language of evolutionary game theory, Honesty appears to be an "evolutionarily stable" equilibrium, able to withstand an "invasion" by corruption. This seems intuitively correct: once all agents are honest, agent's memory and social network are quickly filled with honest agents (in the language of the model, this means **A** and **B** are both equal to zero). If a few corrupt agents "invaded", they would very quickly reach the limit of mismatches and end up in jail—any agents who happened to have the newcomers as "friends" would be deterred from following their strategy choice.

---

populations). At the bottom, an analysis of these 35 transition points is presented with mean, variance, and standard deviation. Columns #3 and #4 are for calculation purposes.

[11] It would be hard to imagine, for example, Norway (or even the United States) reverting to the situation of Zaire, without major structural changes.

Also of interest are the large “dips” in corruption that can be observed on Figure 1 between, for example, iterations 51 and 61 (in the government population), or between 66 and 76 (in the citizens). These dips are the result of circumstances similar to those leading to transition: once again, a small peak of arrests spreads “fear” to some agents through their social networks, leading these agents to switch to honesty. In the case of the dips, however, the circumstances involved aren’t quite enough to tip the system into the transition behavior—either too few agents change strategy or those that do share relatively self-contained social networks and the fear doesn’t spread to the rest of society.

The dips suggest spontaneous and endogenously occurring “waves” of honesty and/or corruption. This has an interesting implication for the interpretation of the history of corruption in any given country. Often “waves” of corruption or honesty are assumed to be the result of an exogenous force (a new election, government policy, or economic shock). The results of the model presented here suggest that such waves might instead be endogenous and a normal part of a pattern of corruption.

Model sensitivity to assumptions

The base case above makes use of a *limited local information* restriction—as described in the model structure above, agents have a limited memory (here 5 rounds) and social network (here 10 agents). The transition to Honesty is sensitive to this restriction, as follows:

When agents have a Memory of 500 and a Social network of 299 (i.e. “everyone” in their own population), the “fault point” described appears not to arrive in an observable time—near-universal corruption seems to be a sustainable equilibrium. This can be seen in Figure 2, which resembles Figure 1, but lacks both the “dips” and the transition to honesty. Further experimentation reveals that “infinite” memory is by itself sufficient for sustainable corruption;

"infinite" social network size requires a certain size of memory to be sufficient (Figure 3a,b and Table 3a,b).[12]

Clearly, limited local information is important to the base case transition result. The "fear" of enforcement which spreads through the population in the base case plays a crucial role in the eventual transition. A long memory or a large social network "dilutes" the impact of a "friend" agent going to jail or an opponent switching strategy—beliefs are slow to change and the "fear" cannot spread very quickly.

An important policy implication of this result is that an "inscrutable" enforcement system may be more effective than a transparent one. There exists some historical evidence for such an assertion[13], although it contradicts the majority of existing literature which tends to exhort enforcement agencies to increase transparency (make the "rules" clear) and disseminate information about enforcement.[14] A lack of adequate communication systems (infrastructure or communication technology) in the real world would also closely parallel the condition of limited local information; the model would suggest such a lack would make corruption easier to sustain. The sensitivity of the base case result to social network size also implies that breaking up existing social networks would be an effective enforcement tool (as evidenced in many totalitarian political systems).

[12] Axes of Figure 3a represent increasing Memory Size (social network size held constant) and the percentage (of each 35 run sample) exhibiting a transition to honesty, respectively. Figure 3b is similar, with social network size varying along the x-axis and memory size held constant (at 5). Tables 3a and 3b show the data from which these charts were derived.

[13] The case of Jusice Plana in the Philippines is instructive. An important part of Plana's strategy was, in essence, to induce greater "fear" of enforcement without necessarily dramatically increasing enforcement. Instead, Plana arrested a few very high profile officials, and "interviewed" government officials at random to spread the sense that enforcement was being taken more seriously. (Klitgaard 1988, Chapter 3). Arguably, the U.S. IRS employs similar tactics.

Model sensitivity to population parameters

In the base case result, "honesty" (measured by inherent propensity) is *uniformly* distributed over the interval [0,1] amongst agents. Returning the model to the base case conditions, but fixing the inherent propensity at $i = .5$ for all agents (same mean, 0 variance) has a profound impact on the model outcome: all agents are corrupt indefinitely (once again, corruption becomes sustainable). This is illustrated in Figure 4. Not only are the "dips" and the eventual transition to honesty missing, but the number of corrupt agents is fully 100% in both populations, and is immune from the slight noise evident in the previous two results. Note that this outcome is an analytically soluble *Nash equilibrium,*[15] and note that the predicted Nash ***breaks down*** in deviations from this outcome (such as the base case). The result can be generalized for any level of *homogenous* inherent propensity in the population.[16]

Clearly, a heterogeneous population (in terms of inherent propensity) is also extremely important to the base case result. This has optimistic implications, given an assumption of heterogeneous (or at least less than perfectly homogenous and "indifferent") populations in the real world.[17]

---

[14] See Tanzi 1998 for an extensive discussion of this line of argument.

[15] Once all agents are corrupt, both the memory and the social network of each individual agent is filled with other corrupt agents. Thus, in the language of the model, **A** = 1 and **B** = 0 (the agent feels certain to encounter another corrupt agent, and has no perceived chance of being apprehended). Given $x_i > y$, the agent has no motivation to deviate from a corrupt strategy. As long as no agent deviates, there is no actual chance of any agent playing a mismatch and ending up in jail—the system is stable.

[16] Given an appropriately high payoff to corruption (such that $x_i > y$) any value of inherent propensity less than 1 will yield this result.

[17] If corruption is, indeed, so entrenched that society is indifferent, there may be an important role for politics to play. A certain line of argument in the literature suggests that politicians and leaders may be able to change *attitudes* toward corruption over time (see Klitgaard 1988, p. 58 for example). In the terms of my model, such a change would move the population from the homogeneous (indifferent) case to the "base case" (and thus from the stable corruption result to the transition result).

Model sensitivity to enforcement structure

Agents in the base case model have very restricted knowledge of enforcement mechanism. The transition result appears to be marginally sensitive to this assumption.

In Table 4, the enforcement mechanism is restructured to conform more closely to agent's expectations. A fixed level of corrupt agents (here 0.5%) go to jail each round, irrespective of their opponent's choice. The result of this change is to mix stable corruption (out past thousands of iterations) with a transition resembling the base-case behavior. In a 35-run analysis,[18] corruption was stable 24 times (68.5%), while a transition occurred 11 times (31.5%). The average transition time was 379 iterations. Raising the enforcement rate (from 0.5% upwards) very quickly leads to a rapid and predictable transition to honesty.

This suggests that a level of uncertainty with regard to the enforcement mechanism is fairly important for a reliable transition. With the new fixed enforcement rate, agents are more often (68.5% of the time) able to "guess" the enforcement strategy after some time, and are less prone to the fear of apprehension that leads to the transition in the base case.

Another means of increasing agent's awareness of the enforcement mechanism is to include in their decision rule an understanding that mismatches increase their chances of being caught. The new decision rule is as follows:

A (the subjective estimate of being matched with a corrupt agent) now becomes the central variable. The first half of the equation for the weighted payoff to corruption, representing the contingency of the agent both *successfully colluding* and *getting away with it* is now just $\mathbf{A}\boldsymbol{x}_{\mathbf{i}}$.

[18] Table 4. The first column gives the run number (1 to 35). The second column, for each run, notes the iteration number at which a transition took place, or notes that such a transition did not take place.

B (an analysis of the social network) is no longer used by the agent as a subjective estimate of the probability of being caught. This estimate is now taken to be **(1-A)**—an agent can only be caught if their opponent plays Honest. Agents now know this, and have incorporated it into the decision rule. Agents still do *not* know, however, exactly how many mismatches are required to land them in jail. To estimate *this* probability, they use **B** = **m/M**. So the weighted value of the payoff to corruption becomes:

$$\mathbf{A}x_{\mathbf{i}} \;+\; \mathbf{(1\text{-}A)\ [B(}y\mathbf{\text{-}k}y\mathbf{)]}$$

Interestingly, such a change has little or no effect on the model: the results look very similar to those of the base case. Again, as in the base case, Corruption persists for some period of time, and then reaches a “fault point” and switches abruptly to Honesty. Note in the attached data summary (Table 5) that a 35-run compilation yields a very similar mean value for the transition point (291) to that of the base case. In this configuration, however, one “population” (either the citizens or the bureaucrats) can take much longer to react to a shift to honesty in the other population. Attached examples show a run with an almost simultaneous transition (Figure 5), a transition in which the citizens shift first and the bureaucrats lag (Figure 6), and one in which the citizens lag considerably behind the bureaucrats (Figure 7).

Removing limited information sets by increasing memory and social network size in this configuration of the decision rule yields the same result as in the base case—stable corruption. Interestingly, populations of agents with the new decision rule can “withstand” much higher levels of jailed members without (or before) reverting to honesty. This is to be expected—agents with this new decision rule are much *less* susceptible to the “fear” that could be transmitted through social networks in the base case (although such transmission still plays a role).

This result qualifies a conclusion presented earlier (in the “base case”)—an enforcement strategy need not be *structurally* inscrutable in an absolute fashion to be effective, but simply be relatively opaque on a local level.

Conclusion

This paper has outlined a micro-level agent-based model of corruption as a simple game-theoretic repeated interaction over time. The fully heterogeneous agent population and random pairing of interactions used in the model would make it very difficult to solve, if not intractable, in a traditional analytical framework. Instead, the model makes use of the unique dynamics of the agent-based technique to allow a transition behavior that differs substantially from existing literature to emerge, and to qualify the conditions for its emergence. The results suggest several important points:

1. A transition from corruption to honesty can happen *endogenously*. The base case transition in the model occurs *inevitably* and without structural or exogenous changes. Instead, the transition is the result of a cascade of micro-level events, set in motion by a chance coincidence of interactions.

2. Limited local information is extremely important for the rapid transmission of "fear of enforcement" which makes the transition behavior possible. In fact, limited local information seems to be a necessary condition for the transition result. This suggests that breaking up social networks and limiting local information are effective enforcement strategies. The model results also suggest that opacity with regards to the nature or structure of enforcement is effective in reducing corruption, by promoting fear and apprehension among agents.

3. Agents in the model are heterogeneous with respect to their inherent moral attitudes, and this diversity is important to the transition result. This suggests that social norms and/or socialization in the real world may be extremely important in the dynamics of corruption, and that diversity may be important in eliminating corruption.

These points suggest an unproved theorem: An *endogenous* transition to honesty can occur in a corrupt system, given the sufficient (and necessary) conditions of a heterogeneous population (with respect to attitudes toward corruption) and limited local information and social networks. No structural change or exogenous influence is necessary for such a transition, which arises spontaneously given enough time.

Selected Bibiliography

Acemoglu, D. and T. Verdier. "The Choice Between Market Failures and Corruption". *Centre National de la Recherche Scientifique, Ecole Normale Superieure and Ecole Nationale des Ponts et Chaussees*. Working Paper WP D1 97-06 (Document No. 97-06). 1997.

Andvig, Jens Christopher. "The Economics of Corruption". *Studi Economici*. Vol. 46, No. 43 (1991), pp. 57-94.

Bardhan, Pranab. "Corruption and Development: A Review of Issues". *Journal of Economic Literature*. Vol. XXXV (Sept. 1997), pp. 1320-1346.

Beenstock, Michael. "Corruption and Development". *World Development*. Vol. 7 (Jan. 1979), pp. 15-24.

Cadot, Olivier. "Corruption as a Gamble". *Journal of Public Economics*. Vol. 33 (1987), pp. 223-244.

Dey, Harendra Kanti. "The Genesis and Spread of Economic Corruption: A Microtheoretic Interpretation". *World Development*. Vol. 17, No. 1 (January 1989), pp. 503-511.

Goudie, Andrew W. and David Stasavage. "Corruption: The Issues". *OECD Technical Paper* No. 122 (January 1997).

Gray, Cheryl W. and Daniel Kaufmann. "Corruption and Development". *Finance and Development*. March, 1998, pp. 7-14.

Klitgaard, Robert. *Controlling Corruption*. Berkeley: University of California Press, 1988.

Klitgaard, R. "Institutional Adjustment and Adjusting to Institutions". *World Bank Discussion Paper* No. 303, September 1995.

Leff, Nathaniel H. "Economic Development Through Bureaucratic Corruption". *The American Behavioral Scientist*. Vol. VIII, No. 3 (November 1964), pp. 8-14.

Rose-Ackerman, Susan. *Corruption: A Study in Political Economy*. New York: Academic Press, 1978.

Sah, Raaj Kumar. "Persistence and Pervasiveness of Corruption: New Perspectives". *Economic Growth Center, Yale University*: Center Discussion Paper No. 560. August, 1988.

Tanzi, Vito "Corruption Around the World: Causes, Consequences, Scope and Cures". *International Monetary Fund Staff Papers*, v. 54 i. 4 (December 1998), p. 558.

Fig. 1

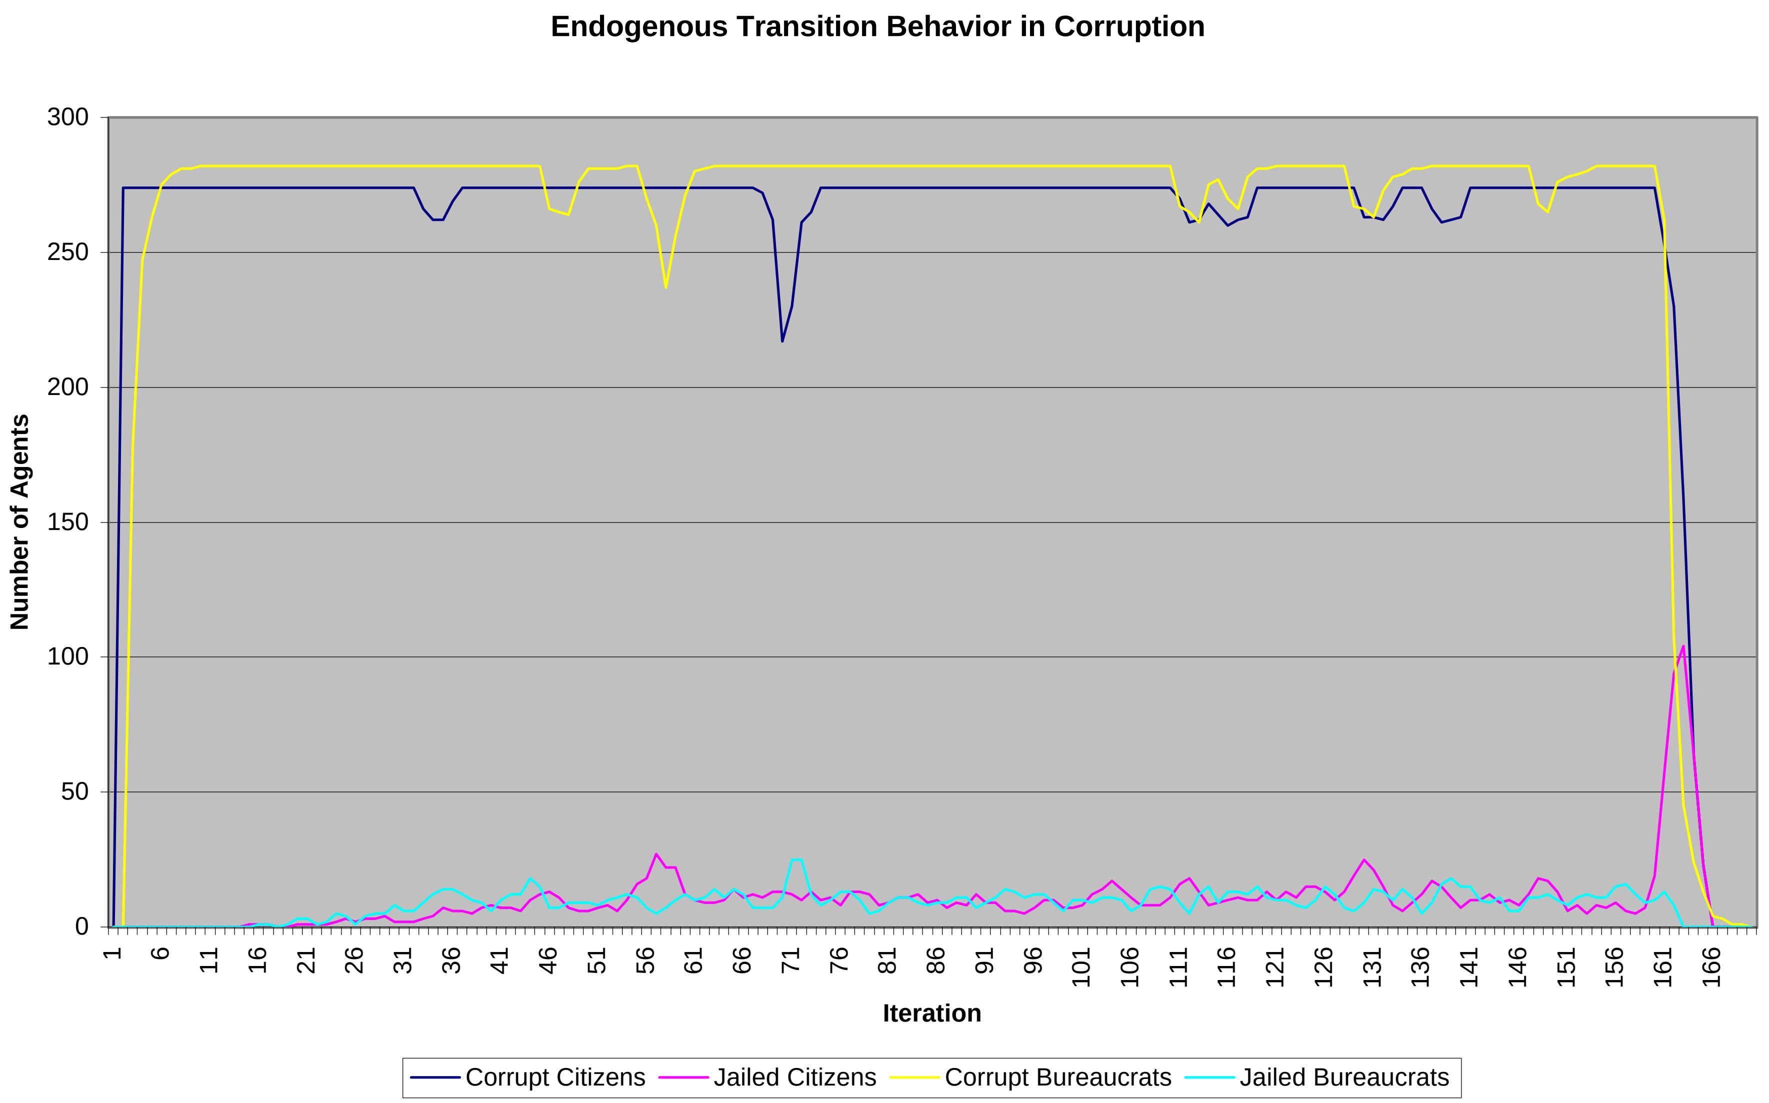
Endogenous Transition Behavior in Corruption
Number of Agents
300
250
200
150
100
50
0
1
6
11
16
21
26
31
36
41
46
51
56
61
66
71
76
81
86
91
96
101
106
111
116
121
126
131
136
141
146
151
156
161
166
Iteration
Corrupt Citizens
Jailed Citizens
Corrupt Bureaucrats
Jailed Bureaucrats

Data Analysis of Transition Result (35 individual runs)

Run #	Transition Completed (Iteratior	std. Deviatio	std. deviation^
1	73	156.2702703	24420.39738
2	169	60.2702703	3632.505482
3	11	218.2702703	47641.9109
4	915	-685.7297297	470225.2622
5	245	-15.7297297	247.4243964
6	169	60.2702703	3632.505482
7	73	156.2702703	24420.39738
8	12	217.2702703	47206.37036
9	226	3.2702703	10.69466784
10	118	111.2702703	12381.07305
11	45	184.2702703	33955.53252
12	14	215.2702703	46341.28928
13	16	213.2702703	45484.20819
14	374	-144.7297297	20946.69466
15	1522	-1292.72973	1671150.154
16	115	114.2702703	13057.69467
17	158	71.2702703	5079.451429
18	261	-31.7297297	1006.775747
19	434	-204.7297297	41914.26222
20	349	-119.7297297	14335.20817
21	64	165.2702703	27314.26225
22	103	126.2702703	15944.18116
23	187	42.2702703	1786.775751
24	14	215.2702703	46341.28928
25	50	179.2702703	32137.82981
26	384	-154.7297297	23941.28925
27	261	-31.7297297	1006.775747
28	312	-82.7297297	6844.208176
29	456	-226.7297297	51406.37033
30	120	109.2702703	11939.99197
31	104	125.2702703	15692.64062
32	259	-29.7297297	883.856828
33	363	-133.7297297	17883.64061
34	200	29.2702703	856.7487234
35	307	-77.7297297	6041.910879

MEAN	229.2702703
VARIANC	79631.75953
STD DEV	282.1909983

Table 2

Model Parameterization for the "Base Case"

Default Values:

Parameter	Value
x (base payoff to corruption)	20
y (payoff to honesty)	1
i (inherent propensity)	uniformly distributed through population over [0,1]
N (size of memory)	5 rounds
M (size of social network)	10 agents
k (length of jail term)	2 rounds
# of mismatches to get caught	

In all runs of model, each agent's memory is filled with random values during initialization before the first round.

The specific members of any given agent's social network are also randomly assigned during initialization.

Direction of Model Response to Marginal Shifts:

Parameter	Marginal shift up	Marginal shift down
x*:*y	transition takes longer	transition disappears; all honest very quickly
k	transition more quickly	transition takes longer
# mismatches	transition takes longer	transition more quickly

Fig. 2

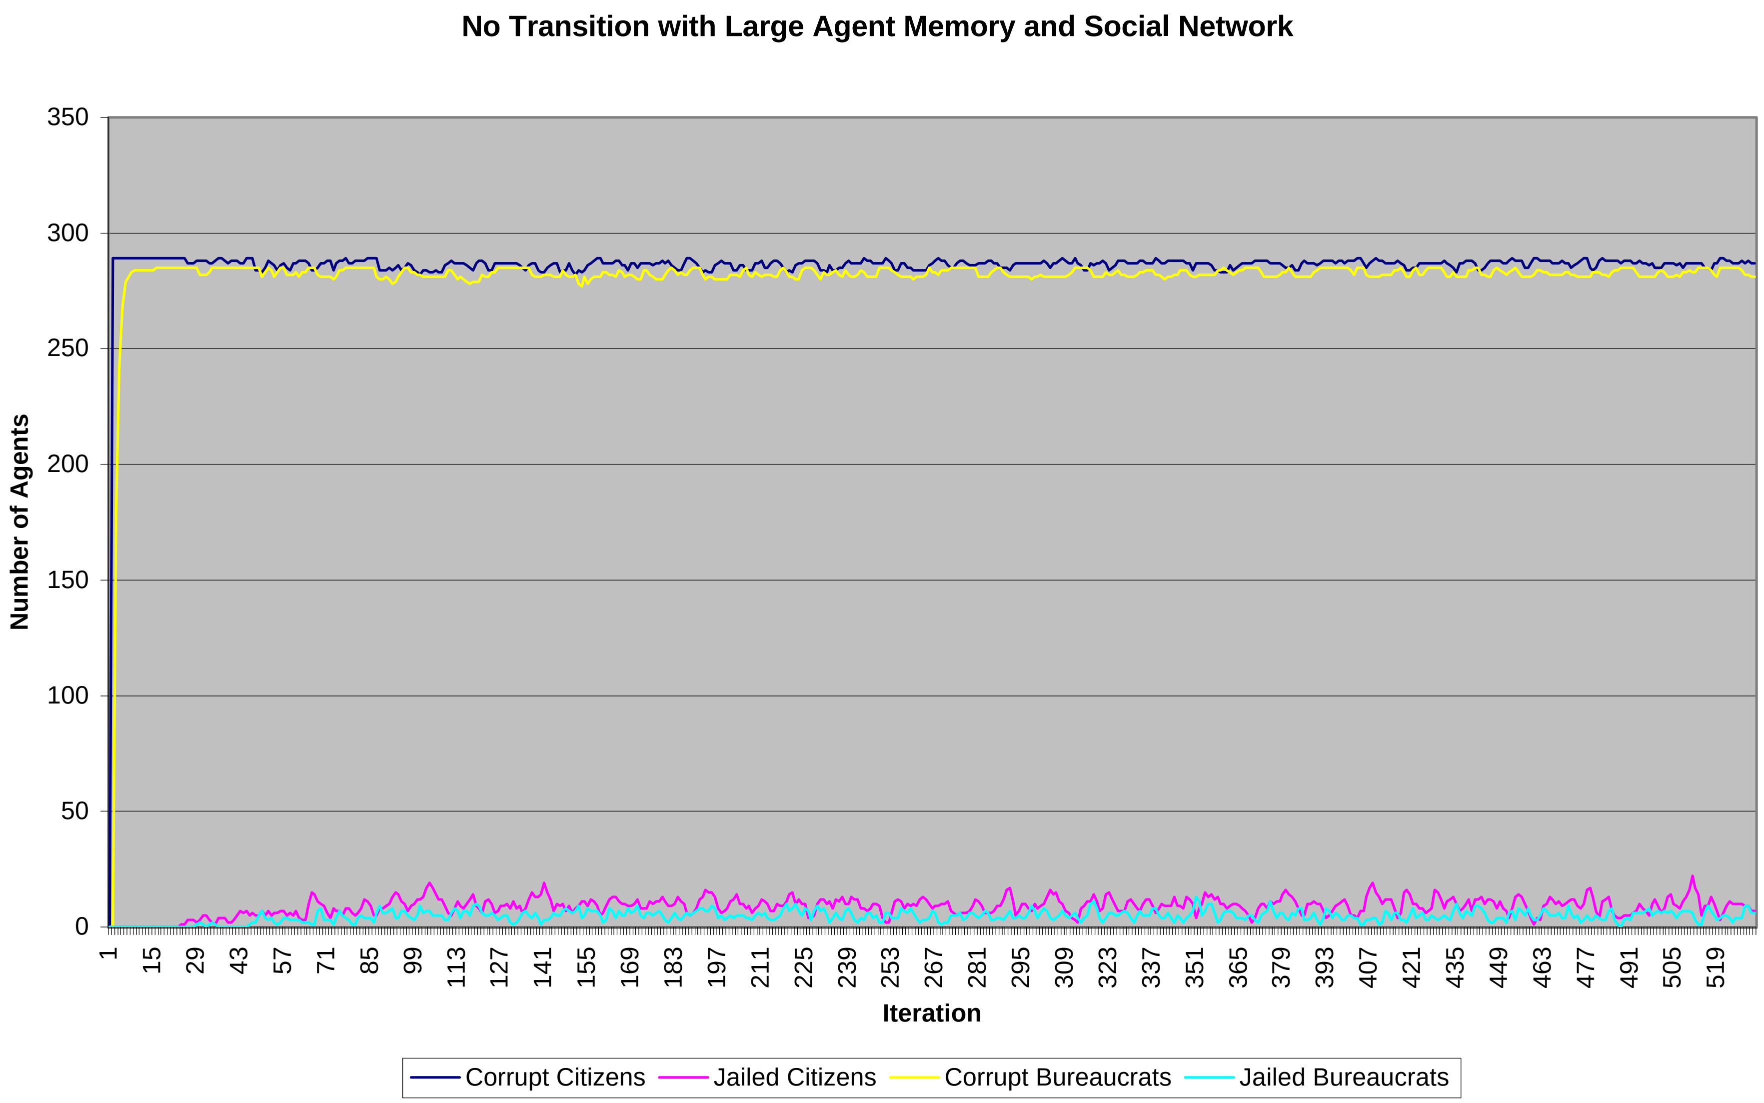
No Transition with Large Agent Memory and Social Network
Number of Agents
350
300
250
200
150
100
50
0
1
15
29
43
57
71
85
99
113
127
141
155
169
183
197
211
225
239
253
267
281
295
309
323
337
351
365
379
393
407
421
435
449
463
477
491
505
519
Iteration
Corrupt Citizens
Jailed Citizens
Corrupt Bureaucrats
Jailed Bureaucrats

Fig. 3a

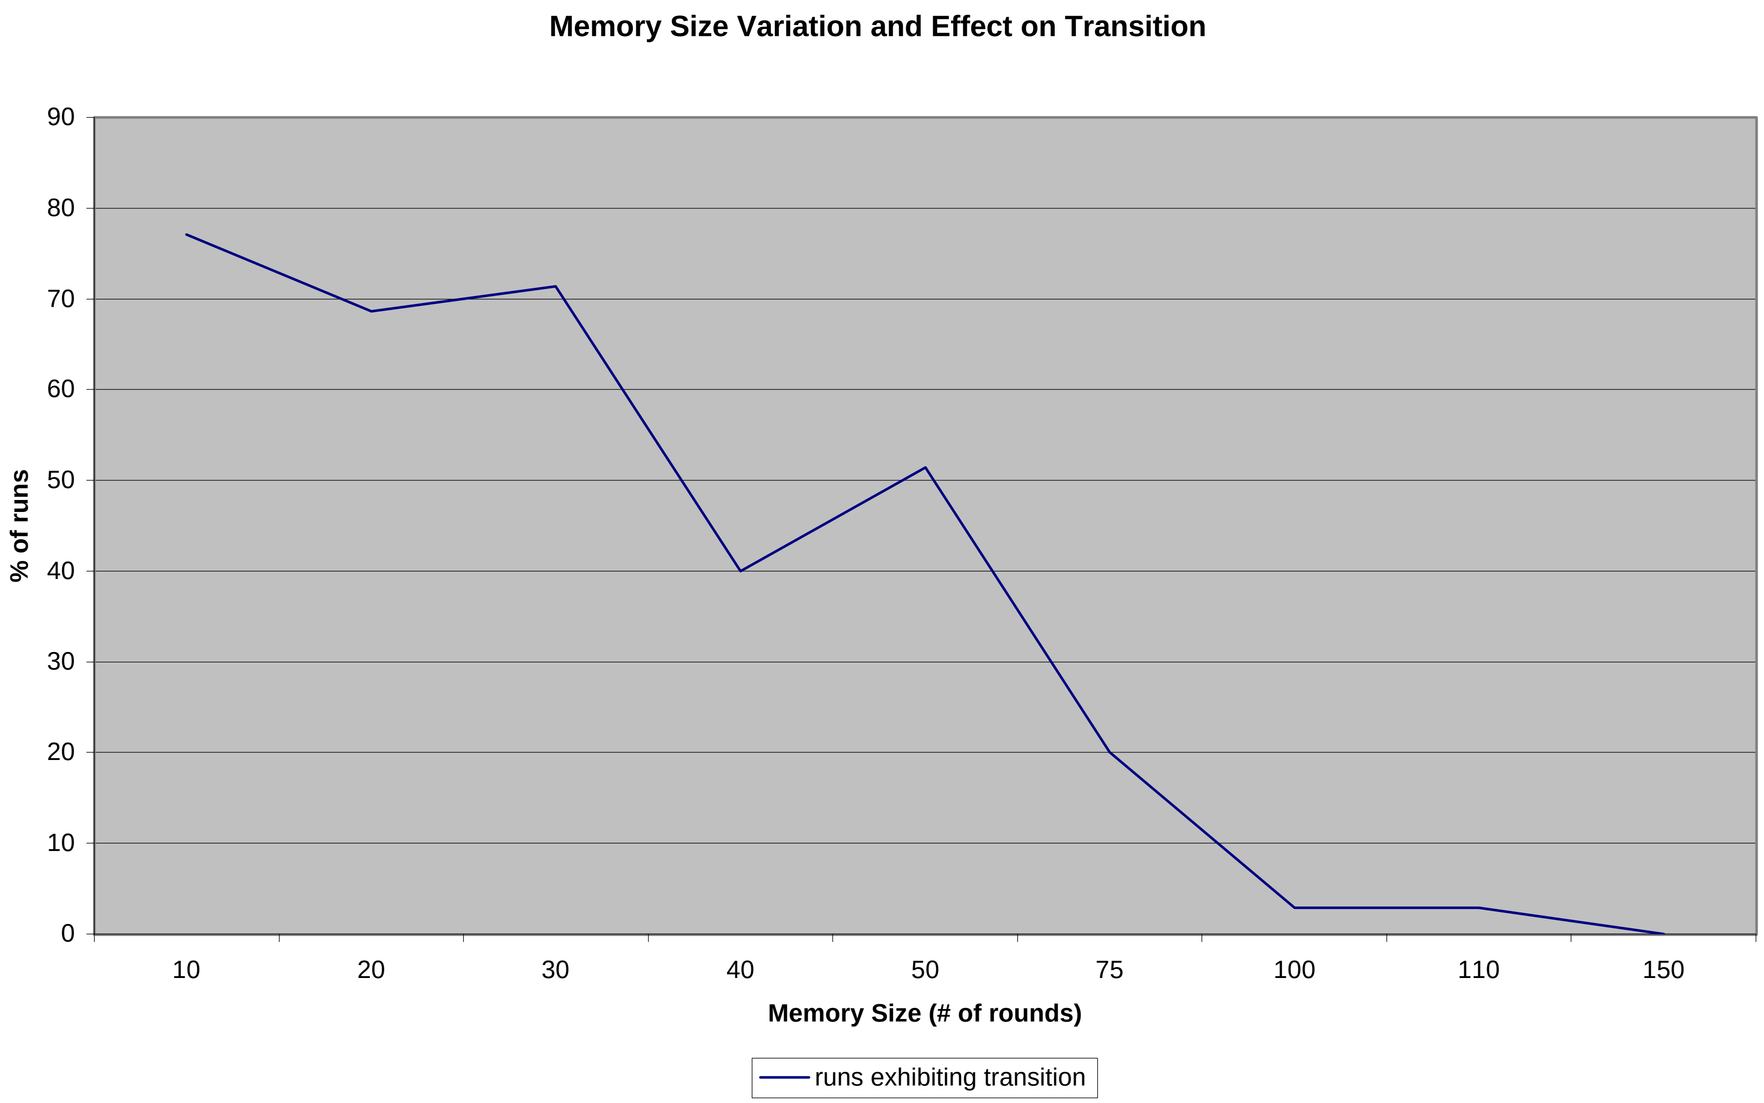
Memory Size Variation and Effect on Transition
90
80
70
60
50
40
30
20
10
0
% of runs
10
20
30
40
50
75
100
110
150
Memory Size (# of rounds)
runs exhibiting transition

Memory Size Variation and Effect on Transition Behavior

Memory Size	Observed Transitions (out of 35 ru	Non-Transitio	% Transitions
10	27	8	77.1
20	24	11	68.6
30	25	10	71.4
40	14	21	40
50	18	17	51.4
75	7	28	20
100	1	34	2.9
110	1	34	2.9
150	0	35	0

Fig. 3b

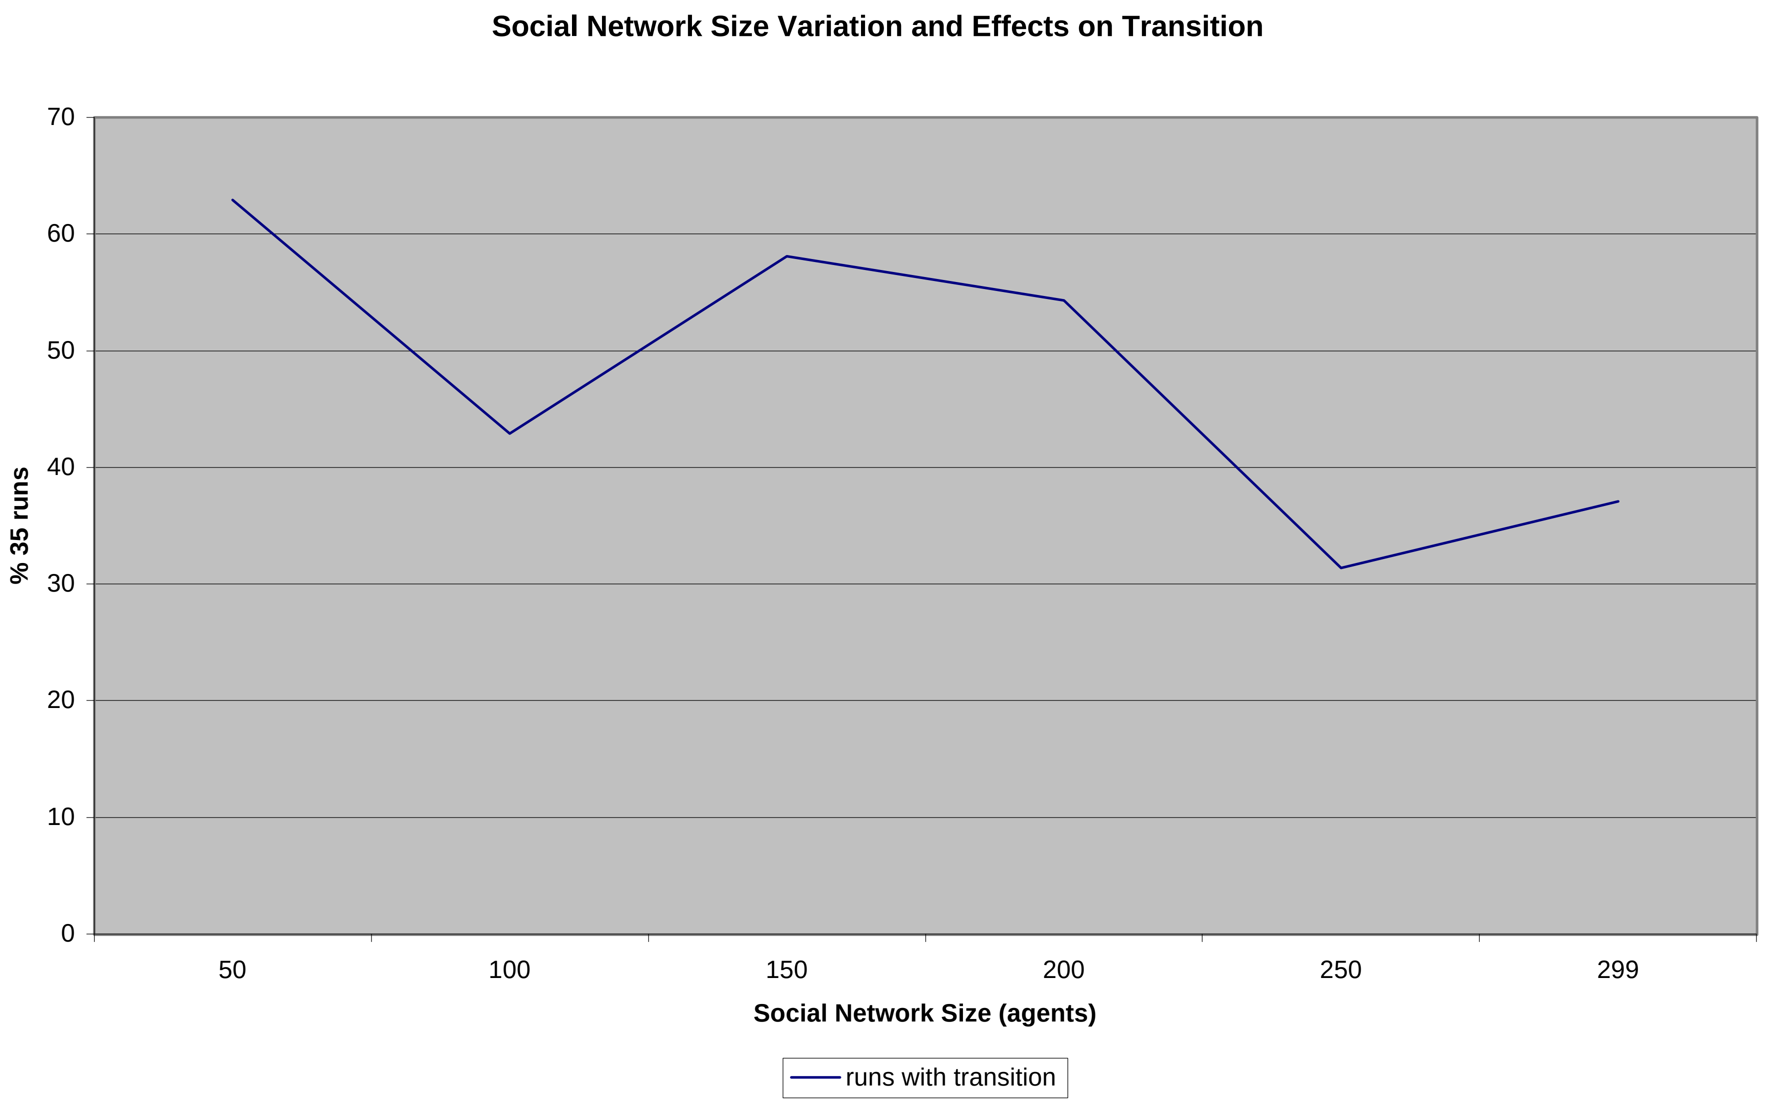
Social Network Size Variation and Effects on Transition
70
60
50
40
30
20
10
0
% 35 runs
50
100
150
200
250
299
Social Network Size (agents)
runs with transition

Table 3b

Social Group Size Variation and Effect on Transition Behavior

Social Group Siz	Observed Transitions (out of 35 ru	Non-Transition	% Transitions
50	22	13	62.9
100	15	20	42.9
150	20	15	58.1
200	19	16	54.3
250	11	24	31.4
299	13	22	37.1

Fig. 4

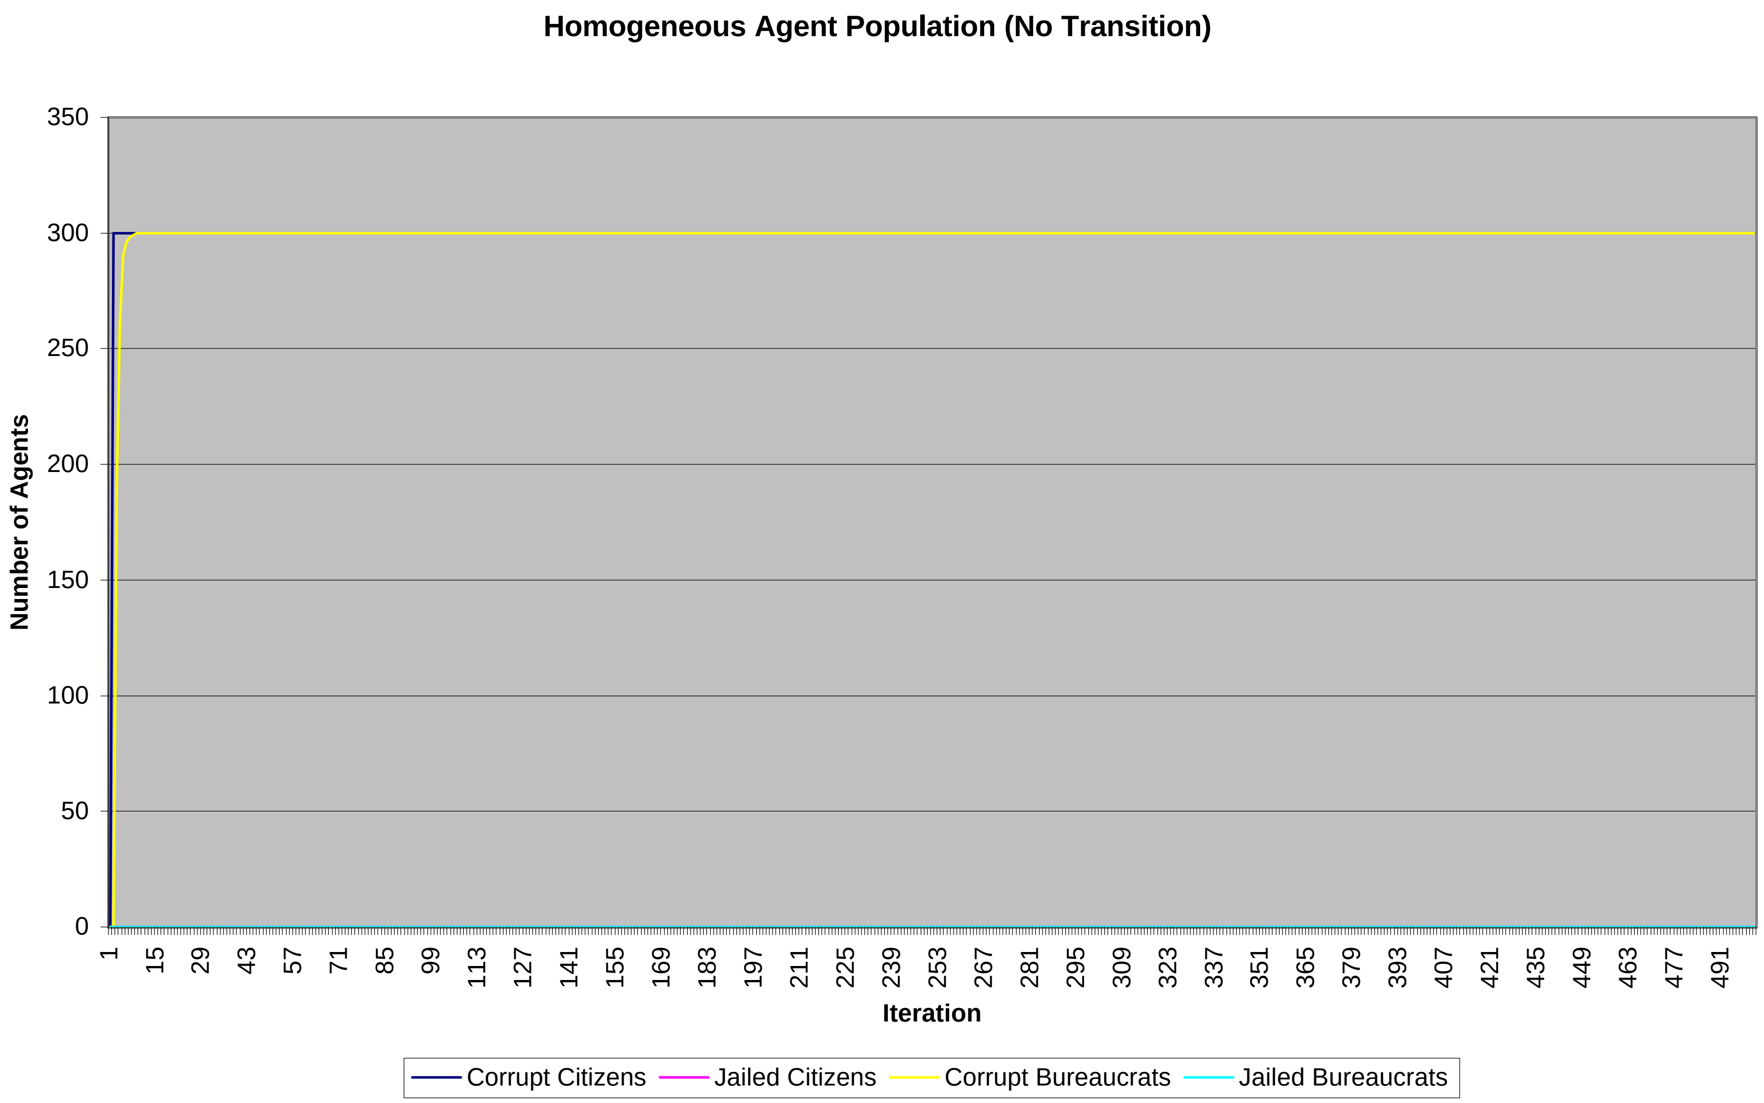
Homogeneous Agent Population (No Transition)
Number of Agents
Iteration
Corrupt Citizens
Jailed Citizens
Corrupt Bureaucrats
Jailed Bureaucrats

Table 4

Data Analysis of Behavior Under Fixed Enforcement Level

Run #	Transition Observed (Iteration #):	Population Exhibiting Transition:
1	none[1]	
2	403	Citizens only
3	none	
4	447	Bureaucrats only
5	none	
6	none	
7	none	
8	none	
9	none	
10	none	
11	none	
12	497	Bureaucrats only
13	907	Bureaucrats only
14	309	Bureaucrats only
15	412	Bureaucrats only
16	none	
17	none	
18	none	
19	none	
20	none	
21	384	Bureaucrats only
22	none	
23	96	Bureaucrats only
24	none	
25	none	
26	260	Bureaucrats only
27	none	
28	180	Bureaucrats only
29	none	
30	none	
31	none	
32	none	
33	274	Bureaucrats only
34	none	
35	none	

TOTALS

no transition: **24 (68.5%)**
transition: **11 (31.5%)**
avg. transition time (for observed transitions): **379**

[1] "None" indicates that no transition had occurred by round 1500

Data Analysis of New Decision Rule Behavior

Run #	Transition Completed (Iteration	std. deviatio	deviation^2
1	307	15.6	243.36
2	238	-53.4	2851.56
3	195	-96.4	9292.96
4	168	-123.4	15227.56
5	383	91.6	8390.56
6	1502	1210.6	1465552.36
7	146	-145.4	21141.16
8	147	-144.4	20851.36
9	16	-275.4	75845.16
10	36	-255.4	65229.16
11	341	49.6	2460.16
12	441	149.6	22380.16
13	16	-275.4	75845.16
14	16	-275.4	75845.16
15	35	-256.4	65740.96
16	264	-27.4	750.76
17	129	-162.4	26373.76
18	19	-272.4	74201.76
19	463	171.6	29446.56
20	1218	926.6	858587.56
21	14	-277.4	76950.76
22	16	-275.4	75845.16
23	644	352.6	124326.76
24	40	-251.4	63201.96
25	223	-68.4	4678.56
26	93	-198.4	39362.56
27	819	527.6	278361.76
28	132	-159.4	25408.36
29	732	440.6	194128.36
30	16	-275.4	75845.16
31	86	-205.4	42189.16
32	474	182.6	33342.76
33	269	-22.4	501.76
34	314	22.6	510.76
35	247	-44.4	1971.36

mean	291.4
variance	112939.4971
std. dev.	336.0647217

Fig. 5

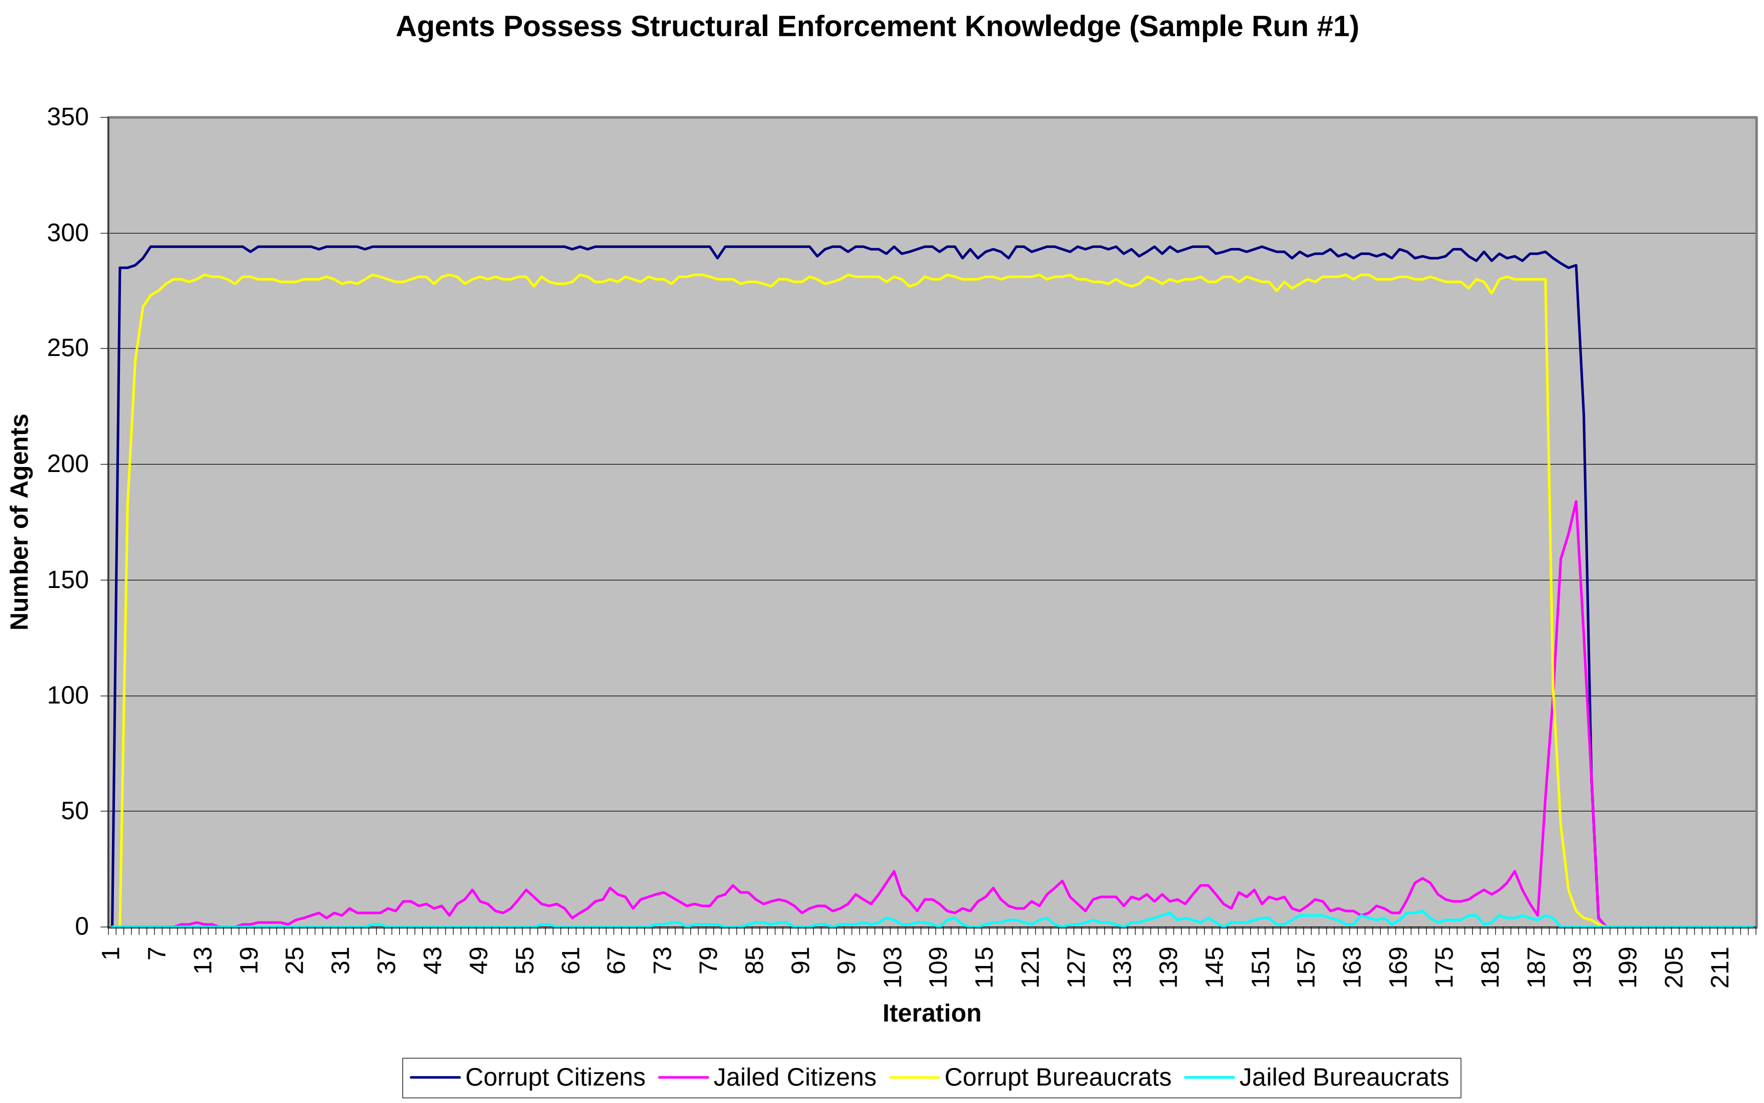
Agents Possess Structural Enforcement Knowledge (Sample Run #1)
Number of Agents
350
300
250
200
150
100
50
0
1
7
13
19
25
31
37
43
49
55
61
67
73
79
85
91
97
103
109
115
121
127
133
139
145
151
157
163
169
175
181
187
193
199
205
211
Iteration
Corrupt Citizens
Jailed Citizens
Corrupt Bureaucrats
Jailed Bureaucrats

Fig. 6

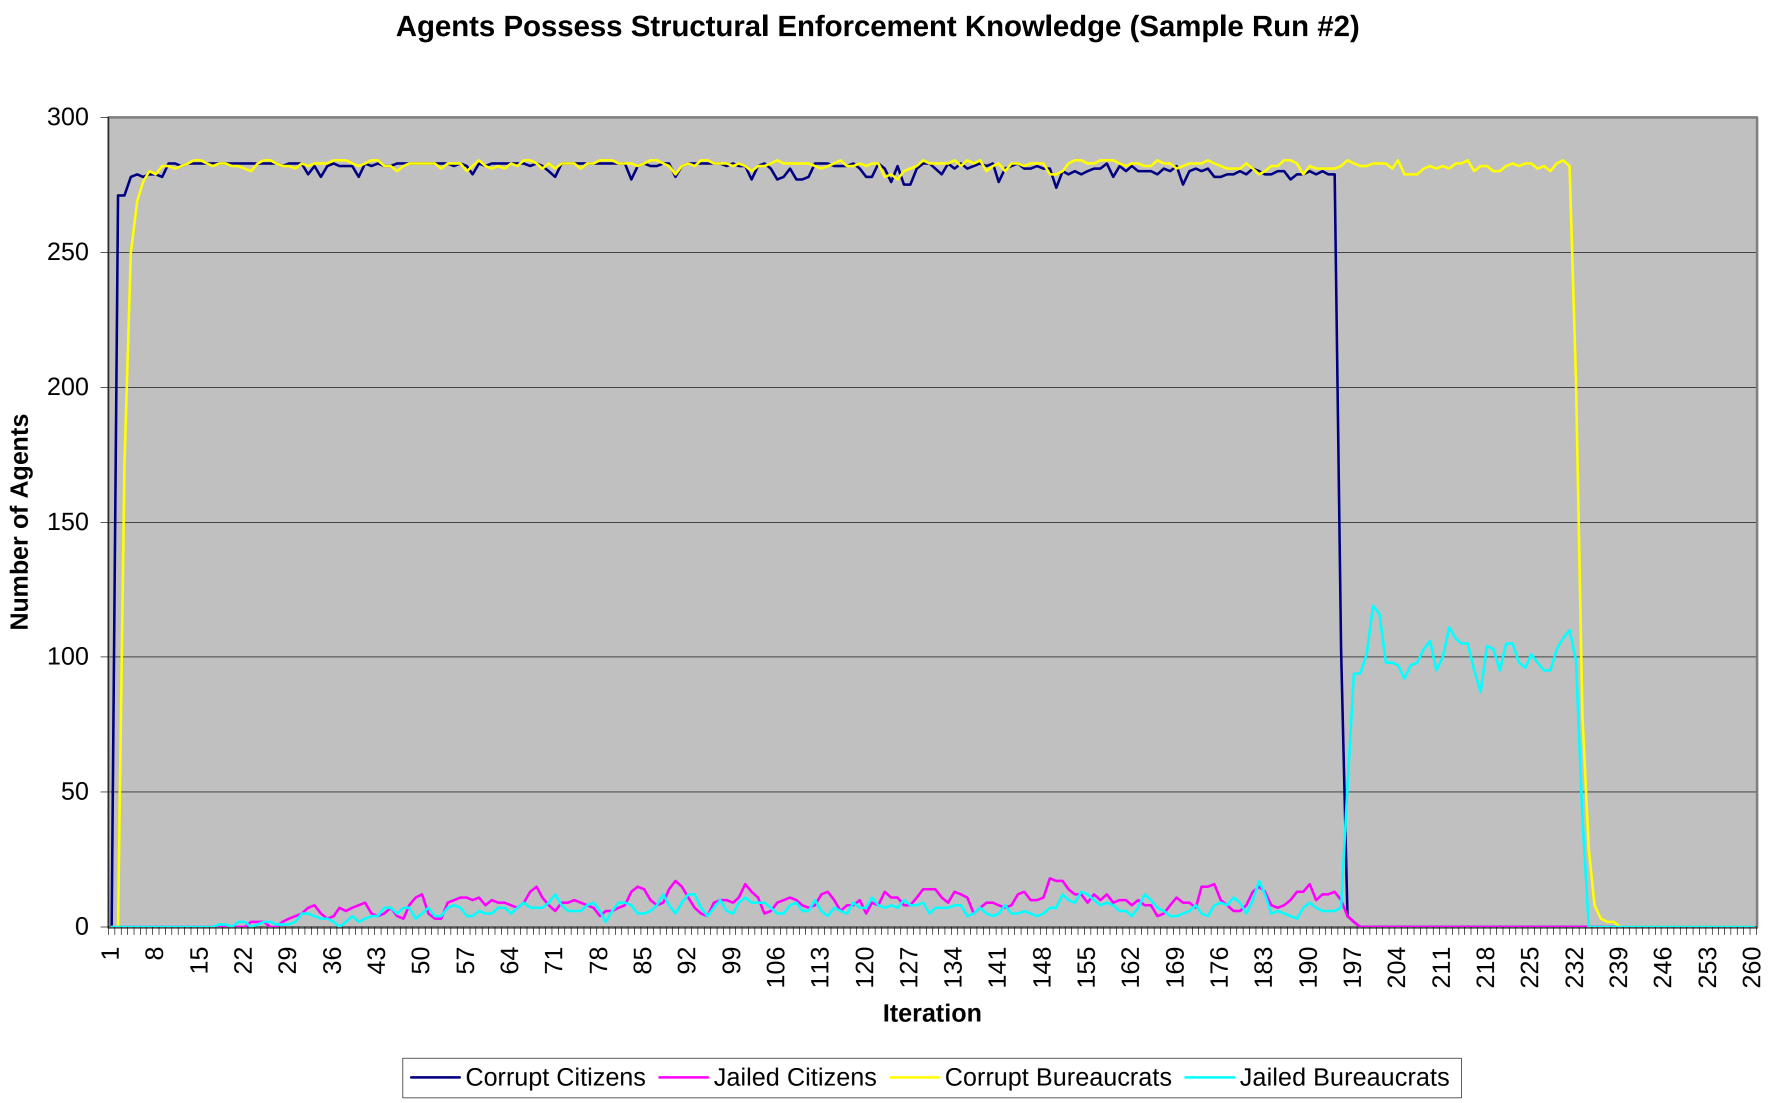
Agents Possess Structural Enforcement Knowledge (Sample Run #2)
Number of Agents
300
250
200
150
100
50
0
1
8
15
22
29
36
43
50
57
64
71
78
85
92
99
106
113
120
127
134
141
148
155
162
169
176
183
190
197
204
211
218
225
232
239
246
253
260
Iteration
Corrupt Citizens
Jailed Citizens
Corrupt Bureaucrats
Jailed Bureaucrats

Fig. 7

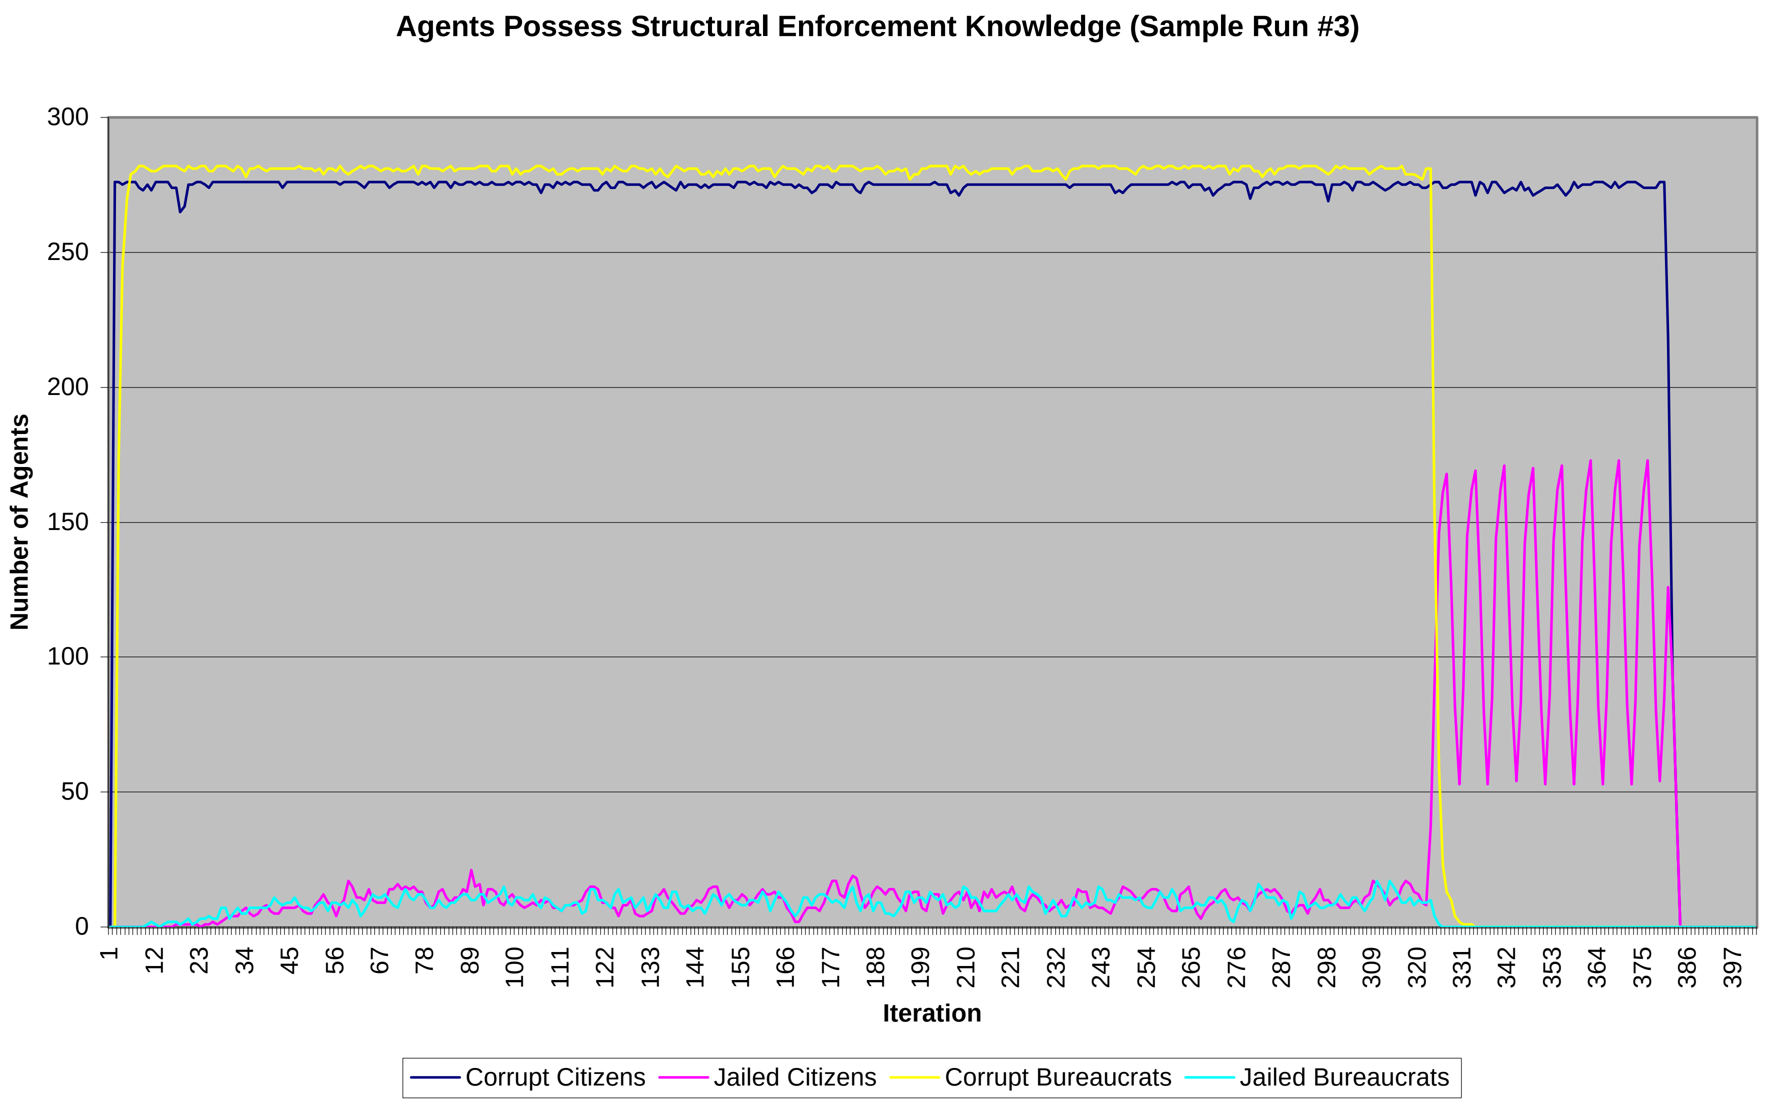
Agents Possess Structural Enforcement Knowledge (Sample Run #3)
300
250
200
150
100
50
0
Number of Agents
1
12
23
34
45
56
67
78
89
100
111
122
133
144
155
166
177
188
199
210
221
232
243
254
265
276
287
298
309
320
331
342
353
364
375
386
397
Iteration
Corrupt Citizens
Jailed Citizens
Corrupt Bureaucrats
Jailed Bureaucrats