The Evolution of Institutions: the medium, the long, and the ultra-long, run

DESIREE A. DESIERTO
University of the Philippines
School of Economics, Diliman, Q.C. 1101
E-mail: dadesierto@econ.upd.edu.ph

and

JOHN V. C. NYE
Economics Department, George Mason University
Carow Hall MSN 1D3, 4400 University Drive
Fairfax VA 22030
E-mail: jnye@gmu.edu

Version: 30 April 2009

How do institutions form and how do they change in the presence of random shocks and uncertainty? We formulate stochastic dynamics with drift and time-varying mutation to provide a general theory that explains how institutions are created and how they survive in the medium, long, and ultra-long run.

1. INTRODUCTION

This paper has a two-fold objective. One is to provide a formal model of institutional change that is rooted in the New Institutional Economics (NIE) literature. The other is to formulate (evolutionary-game) dynamics capable of accommodating the effect of stochastic 'drift' from, and time-varying 'mutation' into and out of, strategies. Such aims actually serve to motivate each other.

Scholars have proposed a number of game-theoretic formulations to study institutional change (i.e. Greif, Aoki, Young) without coming to a clear consensus. We hold the view that institutional change is best modeled in an evolutionary framework for several reasons. One is that we are able to relax assumptions of altruism and hyperrationality implicit in other dynamic (i.e. repeated) games by allowing the game to be played by 'shorter-sighted' individuals who only care about their own expected payoffs and who can also make mistakes. Another obvious reason is that we can analyze strategy selection in the longer run and refine Nash equilibria (NE) in terms of asymptotic stability (AS), and stochastic stability (SS). Lastly, because of the infinite-time horizon in which the evolutionary game is played, we are readily able to capture the idea that institutions keep on building from previous ones, and avoid the problem of infinite institutional regress by assuming just one initial (exogenous) starting point of the game.

We first set out by recognizing that there are institutions whose emergence and disappearance are occurrences over which the particular players of the game have no control. Such exogenous institutions are herein defined as the 'Meta Set'. Faced with an evolving Meta Set, agents then play an evolutionary game of choosing whether to keep upholding and/or rejecting it. That is, as new institutions are formed and existing ones are destroyed or replaced, the players choose whether or not to use such institutions, thereby strengthening or weakening the 'de facto existence' of institutions. Note, then, that de facto existence is endogenously determined, while the evolution of the Meta Set is not. While there may be some degree of path dependence, we let both types of evolution be influenced by stochastic processes.

Path dependence in the evolution of the Meta Set is captured by letting new institutions build up on previous ones. However, the appearance of new institutions and disappearance of old institutions are random shocks in that the timing of the creation of new, and destruction of old, institutions is unpredictable. The timing may be linked to (random) changes in the environment, thereby allowing changes in the the Meta Set to be interpreted as a way of adapting to the new environment. Such adaptation, however, may be 'correct' or 'incorrect', depending on whether or not agents uphold the Meta institutions via the game determining their de facto existence.

The evolution of de facto existence is path dependent to the extent that agents consider the history of play in calculating for expected payoffs from their strategies, but it is partly stochastic since these agents can also deviate from expected behavior because of two things. One reason is that they sometimes 'experiment' or switch to a strategy which they have not been pre-programmed or pre-conditioned to play - this is interpreted as 'mutation'. Another is that the evolution of the Meta Set itself makes them uncertain as to whether it is best to uphold or reject the Meta Set at all times, or to switch strategies at some point, since the Meta Set can randomly change at any time. This causes 'drift' of players into and out of strategies, and is a larger source of uncertainty than 'mutation' since all players share the same Meta Set are are thus affected by changes therein.

Analyzing the stochastic elements in the dynamics of institutional change poses some technical difficulties. Essentially, one needs to be able to model two sources of randomness: large 'macro' shocks (i.e. changes in the Meta Set) faced by all players and can thus cause 'drift'; and mutation by individual players to capture experimentation or mistakes at the 'micro' level. In addition, one needs to allow players' mutation behavior to be systematically related to the source of drift. Intuitively, the evolving pattern of creation/destruction of institutions may continuously influence the extent of uncertainty of players and, hence, affect how mutation evolves over time. While there exist stochastic dynamics that can be used to model drift and mutation, these models typically employ Wiener processes and/or assume fixed mutation rates. What we develop is a version that accommodates more general stochastic processes for drift. More importantly, we let mutation rates be related to the drift and be time-varying. Such a model then allows us to examine equilibria over three periods:

In the medium run, we assume the Meta Set to be approximately stable/fixed. That is, there are no macro shocks and, hence, there is no drift and mutation rates are fixed. The result obtained is consistent with most evolutionary-game models in that the equilibrium corresponds to the risk-dominant strategy – agents uphold the Meta Set when it is a risk-dominant strategy to do so, even when there is constant mutation. This implies that risk dominance in the payoffs overcomes all other 'small' uncertainty from mutation/experimentation.

Over the long run, we let the Meta Set evolve via random shocks that create new, or destroy old, institutions. We find that the equilibrium can still be the risk-dominant strategy even in the presence of both mutation and drift. However, this result is not generalizable but may only be a special case obtained when the evolution of the Meta Set happens to follow a stochastic process with zero mean, which makes the drift equal to zero on 'average'. To verify the general case, we look at the ultra-long run which we particularly define as the period during which the Meta Set will have evolved sufficiently many times so as to reveal the underlying distribution governing such evolution.

We propose a new technique for analysing the ultra-long run. (It is new because most evolutionary-game models have heretofore considered only fixed mutation rates even in the very long run.) In our version of the ultra-long run, that is, after the Meta Set has undergone many changes, players are now assumed to be able to form expectations as to the nature of this evolution and choose their strategy accordingly. It is then as though the uncertainty from the evolution of the Meta Set disappears and, hence, the source of the drift is fixed at its expectation. Any remaining mutation behavior can then follow a fixed rate. We show that our ultra-long run equilibrium coincides with the risk-dominant strategy only when (a) the Meta Set is stable on expectation or it eventually becomes less complex; or (b) the Meta Set becomes more and more complex, but at a slow enough rate. In both cases (a and b), the uncertainty from the evolution of the Meta Set, or the drift, is not too large so as to overcome risk-dominance in payoffs. When the Meta Set becomes more and more complex at a fast rate, the drift overcomes risk-dominance in payoffs, and the equilibrium becomes the non-risk dominant strategy. That is, players become too uncertain, they mistrust the risk-dominance of a strategy, which causes much drift out of it and into the non-risk dominant strategy. Intuitively, then, assuming that the risk-dominant strategy is to uphold the Meta Set, de facto existence of the Meta Set can only be maintained in the ultra-long run when the Meta set is either stable, eventually becomes less complex, or does not become too complex too fast.

For ease of exposition, the next sections 2 and 3 first propose new stochastic dynamics with drift and mutation and derive some general results, which are then applied in section 4 to the particular process of institutional change. Section 5 concludes.

2. STOCHASTIC DYNAMICS WITH DRIFT AND MUTATION

Let a large population of players be randomly and continuously drawn to play a symmetric 2×2 (evolutionary) coordination game $G_I = \{N, S, A\}$, where N defines the set of players, $S = \{1, 2\}$ is the strategy set common to all players, and $A = \{a_{ij}\}$, $i, j = \{1, 2\}$ is the corresponding (doubly-symmetric) payoff matrix where element a_{ij} is the payoff of adopting strategy i against an individual adopting strategy j[1]:

$$A = \begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix},$$

and where strategy 1 is risk-dominant if $a_1 > a_2$.

Let the outcome of play at each time t be described by the proportions of players adopting each strategy. That is, (dropping time subscripts), let the state of the population at t be described by $V = (v_1, v_2)$, where $v_1 = \frac{n_1}{N}$, and n_1 is the number of players that have adopted strategy 1 as of t. Let the payoff be denoted as $u(e^1, e^1)$ for pure strategy 1 (and analogously for 2), and expected payoff to pure strategy 1 when the population is in state V as $u\left(e^1, V\right)$. The population average payoff is $u\left(V, V\right) = \sum v_i u\left(e^i, V\right), i = \{1, 2\}$.

The growth of proportions can be captured by the canonical Replicator Dynamic (RD):[2]

$$\dot{v}_i = \left[u\left(e^i, V\right) - u\left(V, V\right)\right] v_i,$$

by which a strategy grows when it does better than the population average. Such a dynamic approximates 'deterministic' behavior, i.e. as if players choose strategies only according to expected payoffs.[3] A stochastic version of this, however, is easily obtained by explicitly adding a 'mutation' term. That is, denoting λ_{ij} as the probability of switching from strategy j to i, the net flow to v_i is captured by $\dot{v}_i = \left[u\left(e^i, V\right) - u\left(V, V\right)\right] v_i +$

[1] A is constructed from subtracting a_{21} from column 1 and a_{12} from column 2. Note that $(a_{11} - a_{21})(a_{22} - a_{12}) > 0$.

[2] See, for instance, Taylor and Jonker (1978), Weibull (1995), Vega-Redondo (1996), Sandholm (2007) on how the RD is explicitly derived.

[3] By the law of large numbers, players behave as expected. See Boylan (1992, 1995) for an analysis.

$\Sigma_{j=1}^{h}[\lambda_{ij}v_j - \lambda_{ji}v_i]$ as in Cabrales (2000). Letting $\lambda_{ij} = \lambda_{ji} = \lambda$, and using the fact that $v_2 = (1 - v_1)$, we obtain:

$$\dot{v}_1 = \left[u\left(e^1, V\right) - u\left(V, V\right)\right] v_{21} + \lambda(1 - 2v_1) \tag{1}$$

Alternatively, we can add stochastic noise in the growth of strategies by introducing a 'drift' term as in Foster and Young (1990), Fudenberg and Harris (1992) and Cabrales:

$$\dot{v}_1 = \left[u\left(e^1, V\right) - u\left(V, V\right)\right] v_1 + \sigma\dot{\mu} \tag{2}$$

where σ is a positive constant. The last term ($\dot{\mu}$ weighted by σ) describes the random flow into v_1 caused by an underlying Weiner process W_t which is distributed as normal $(0, t)$. Thus $\dot{\mu} = \dot{W}$ captures Brownian motion.[4] Allowing for both drift and mutation, the total flow into v_1 is[5]:

$$\dot{v}_1 = \left[u\left(e^1, V\right) - u\left(V, V\right)\right] v_1 + \sigma\dot{\mu} + \lambda(1 - 2v_1) \tag{3}$$

Equation (3) then describes the growth of v_1 in terms of three things. The first term captures the deterministic flow, while the second and third are the noise components in the form of drift and mutation, respectively.

Our model modifies the Foster and Young - Cabrales dynamic in several ways. First, we eschew the use of Brownian motion and instead allow for any stochastic process to capture the drift term. Second, instead of treating the two sources of randomness as independent of each other, we let mutation and drift be defined by some variable q and its time derivative $\dot{q}$, respectively, thereby positing, as it were, a common source of 'uncertainty' that affects the dynamic. The material implication is that the probability of mutation is not time constant, and nor can we let it arbitrarily vanish to obtain a stationary limiting distribution and readily analyze stability in terms of stochastically stable (SS) strategies. Lastly, while we treat the third term as a net mutation flow into strategy 1, we assume the drift to be a random flow *out of* it. The idea is that while the current state of q positively affects players' propensity to 'experiment' regardless of player-type - q increases first-order uncertainty, greater second-order uncertainty in the form of $\dot{q}$ induces flow out of one strategy (strategy 1 in this case) and into the other (strategy 2). That is, increasing complexity of q, captured by positive $\dot{q}$, tends to select one strategy over the other (regardless of strategies' expected payoffs).

Consider then the following specifications.

---

[4] Generally, for strategy i, $\dot{v}_i = v_i\{\left[u\left(e^i, V\right) - u\left(V, V\right)\right] + \sigma\Gamma(V)\dot{\mu}\}$, where $\Gamma(V)$ is continuous in V and $V^T\Gamma(V) = [0, 0, ..0]^T$. The particular 2×2 case is analyzed in Foster and Young, while Cabrales generalizes.

[5] We focus on the RD for strategy 1 since the RD for 2 is easily obtained from the fact that growth rates sum to zero.

DEFINITION 1. *Drift.* Define the flow of 'drifting' players out of strategy 1 to be $\dot{\mu} \equiv g(\dot{q})$, where $\dot{q}$ is a realization at t of a random variable following some exogenous stochastic process.[6] Denoting $\dot{Q} \subset \Re$ as the set of all values for $\dot{q}$, let function $g : \dot{Q} \to \Re_-$ be a (linear) mapping into the set of non-positive real numbers, with $g(0) = 0$, $g' < 0$, $g'' = 0$. Denoting $\Sigma \subset \Re_+$ as the set of all values for σ, we can then define the set $X \subset \Sigma \times \dot{Q}$ of all values of $x = -\left|\sigma g(\dot{q})\right| = (\sigma\dot{\mu})$ as the set of non-positive values $X = \{x \in \Re_- : 0 < \frac{-\dot{v}_1 + x + \lambda}{-u(e^1,V)+2\lambda+u(V,V)} < 1\}$, for $\lambda \neq \frac{[u(e^1,V)-u(V,V)]}{2}$. This condition guarantees that outward and inward drifts are small enough to keep v_1 between 0 and 1 at all times.[7]

DEFINITION 2. *Mutation.* Let the rate of mutation for both player-types be defined by $\lambda \equiv e(q)$ and, denoting $Q \subset \Re_+$ as the set of all values for q, let $e : Q \to [0, 1]$, $e(0) = 0$, $e' > 0$ and $e'' = 0$.[8] The mutation flow, in terms of the proportion of players into strategy 1, is given by $\lambda(1 - 2v_1)$, but the mutation/switching *rate* is given by $\lambda \equiv e(q)$.[9]

Note that the drift represents a larger outflow than (net) mutation:

LEMMA 1. *The absolute value* $\mid \sigma g(\dot{q}) \mid > \lambda(1 - 2v_1)$.

Proof. Re-writing the condition $0 < \frac{-\dot{v}_1+x+\lambda}{-u(e^1,V)+2\lambda+u(V,V)} < 1$ from the definition of X, where $x = -\left|\sigma g(\dot{q})\right|$, we have: $(\dot{v}_1 - \lambda) < -\left|\sigma g(\dot{q})\right| < [-u(e^1, V) + \lambda + u(V, V) + \dot{v}_1]$. From the first inequality, we deduce that $\lambda > \dot{v}_1$ (since $\dot{v}_1$ can be positive and $(\dot{v}_1 - \lambda)$ has to be negative). Thus, since λ cannot be negative, we know that in the second inequality $(\lambda+\dot{v}_1) > 0$. Note that $[-u(e^1, V) + \lambda + u(V, V) + \dot{v}_1] > 0$ if $(\lambda + \dot{v}_1) > [u(e^1, V) - u(V, V)]$. If this is indeed the case, the second inequality is non-binding since $-\left|\sigma g(\dot{q})\right| < 0$, and the absolute value $\mid \sigma g(\dot{q}) \mid$ can take on any real value greater than $(\dot{v}_1 - \lambda)$. On the other hand, if the second inequality is binding, i.e. $(\lambda + \dot{v}_1) < [u(e^1, V) - u(V, V)]$, the magnitude of $\mid \sigma g(\dot{q}) \mid$ is bounded from above by the absolute value $\mid [-u(e^1, V)+\lambda+u(V, V)+\dot{v}_1] \mid$

---

[6] While $\dot{q}$ need not be an underlying Wiener/Brownian motion, we assume that $\dot{q}$ can be approximated as a continuous (random) flow.

[7] That is, expressing $\dot{v}_1 = \left[u\left(e^1, V\right) - u\left(V, V\right)\right] v_1 - \left|\sigma g(\dot{q})\right| + \lambda(1 - 2v_1)$ in terms of v_1 yields $\frac{-\dot{v}_1 - \left|\sigma g(\dot{q})\right| + \lambda}{-u(e^1,V)+2\lambda+u(V,V)}$, which we contain between 0 and 1.

[8] While we do not explicitly define the functional relationship between $\dot{q}$ and q, it is enough to specify that while the change $\dot{q}$ can be any real number, q itself is always contained within $\Re_+$.

[9] Such may be called 'state-independent' mutation in the sense of Bergin and Lipman (1993).

(and from below by $\mid (\dot{v}_1 - \lambda) \mid$). Note, however, that in either case, $\mid \sigma g(\dot{q}) \mid > \lambda$, which is greater than $\lambda(1-2v_1)$ for any value $v_1 \in (0,1)$. ∎

Using game G_I, and plugging in our particular definitions for $(\sigma\dot{\mu})$ and λ, we obtain the following stochastic dynamic for strategy 1:

$$\dot{v}_1 = \left[a_1 v_1 - a_1 v_1^2 - a_2(1-v_1)^2\right] v_1 - \sigma g(\dot{q}) + e(q)(1-2v_1) \qquad (4)$$

3. MAIN RESULTS

To analyze equilibrium properties of equation (4), we look at three specific cases:

3.1. Zero Drift and Time-Constant Mutation Rates

Here we approximate the dynamic when mutation rates are constant. Setting $\dot{q} = 0$ fixes q, which reduces equation (4) to equation (1). Assuming that strategy 1 is risk-dominant, the dynamic will select 1 as the equilibrium strategy in the following sense:

Proposition 1. *An ergodic distribution exists, whose mass is concentrated over the risk-dominant strategy.*

Proof. Suppose strategy 1 is risk-dominant. We give a rough proof for the above proposition by showing that, regardless of initial values, v_1 always stays in the basin of attraction $(\frac{a_2}{a_1+a_2}, 1)$, than in the basin $(0, \frac{a_2}{a_1+a_2})$ of $v_1 = 0$, where $\frac{a_2}{a_1+a_2}$ is the mixed Nash equilibrium. (The mixed NE is less than $1/2$ since $a_1 > a_2$, that is, the basin for 1 is larger than for 2.) If this is the case, then the dynamic is always pulling towards $v_1 = 1$ and spends most of its time near it than away from it. Note, then, from equation (1) that for any value $\lambda \in [0,1]$ v_1 grows, i.e. $\dot{v}_1 > 0$, as long as $\lambda > \frac{-\left[u\left(e^1,V\right)-u(V,V)\right]v_1}{(1-2v_1)}$, decreases when the inequality is reversed, and is stationary at the exact equality defined by $\xi \equiv \frac{-\left[u\left(e^1,V\right)-u(V,V)\right]v_1}{(1-2v_1)}$. Since 1 is risk-dominant, the numerator of ξ is always negative when $\frac{a_2}{a_1+a_2} < v_1 \leq 1$. The denominator is positive when $v_1 \leq 1/2$ and, hence, $\xi < 0$, and v_1 keeps increasing since any value of λ will always be greater than ξ. When $v_1 > 1/2$, $\xi > 0$, making the inequality $\lambda > \frac{-\left[u\left(e^1,V\right)-u(V,V)\right]v_1}{(1-2v_1)}$ binding, but the point at which v_1 can start decreasing is always above $1/2$ and, hence, above $\frac{a_2}{a_1+a_2}$. (At $v_1 = 1$, $\lambda < 0$, so that v_1 starts decreasing at $v_1 = 1$.) Thus, the stable point is always in $(1/2, 1]$, which is in the basin $(\frac{a_2}{a_1+a_2}, 1)$. When $v_1^0 < \frac{a_2}{a_1+a_2}$, the numerator of ξ is positive while the denominator is positive, so the inequality is binding and v_1 keeps increasing. (At $v_1 = 0$,

$\lambda > 0$, preventing v_1 to decrease).[10] Thus, in all cases, i.e. regardless of initial values, v_1 is always within the basin $(\frac{a_2}{a_1+a_2}, 1)$. ∎

(Insert figure 1 here.)

A similar result is derived in most evolutionary-game models, e.g. Foster and Young, Fudenberg and Harris, Cabrales, Kandori, Mailath and Robb (1993) (hereafter KMR) and Young (1993, 1998). There, however, the stationary distribution is clearly obtained by letting mutation rates vanish to $\lambda \to 0$, and the distribution collapses to a point mass over the risk-dominant strategy.[11] In contrast, our mutation rate is fixed all throughout, since $\dot{q} = 0$.

3.2. Drift and Time-Varying Mutation Rates

Suppose $\dot{q} \neq 0$ and that the relationship between q and $\dot{q}$ is given by $q_t = q_{t-1} + \dot{q}_t$. Let $\dot{q}_t$be a particular realization at t of a random variable that is assumed to follow a uniform distribution over some interval. Thus, the (time-varying) mutation rate is $\lambda_t = e(q_t) = e(q_{t-1} + \dot{q}_t)$. To simplify, let $g(\dot{q}_t) = \dot{q}_t$. Then equation (4) takes the particular form:

$$\dot{v}_{1t} = \left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1 - v_{1t})^2\right] v_{1t} - \sigma \dot{q}_t + e(q_{t-1} + \dot{q}_t)(1 - 2v_{1t}) \quad (5)$$

PROPOSITION 2. *An ergodic distribution does not exist.*

Proof. We prove by providing the following example in which either strategy 1 or 2 is selected depending on the initial value v_1^0. ∎

EXAMPLE 1. We simulate the path of proportion $v_{1t} = v_{1t-1} + \dot{v}_{1t-1}$ when strategy 1 is assumed to be risk-dominant, function e takes the simple form $\lambda_t = \epsilon q_t = \epsilon(q_{t-1} + \dot{q}_t)$, and $\dot{q}$ is drawn at each t from a uniform distribution over interval $[-1, 1]$.[12] Figure 2 simulates the evolution of v_1 assuming different initial values, and it is seen that only for large enough initial values does the dynamic approach $v_1 = 1$.

(Insert figure 2 here.)

The above proposition is inconsistent with KMR, etc. in which the risk-dominant strategy is always deemed to be stochastically stable. Of course,

---

[10] In this sense, $v_1 = 1$ is 'absorbing' in that no escape is possible.

[11] Bergin and Lipman show that this result is generated because mutation rates are the same. When mutation is state-independent, the stationary distribution (obtained by taking the limit $\lambda \to 0$) puts probability 1 on the state in which everyone adopts the risk-dominant equilibrium since this has the larger basin of attraction. In contrast, state-dependent mutation can always obtain either strategy as the long-run equilibrium depending on how the relative mutation rates are modeled.

[12] We also assume initial $q_0 = 3$, $a_1 = 2$, $a_2 = 1$, $\sigma = 0.03$ and $\epsilon = 0.0001$.

while KMR, etc. establish stochastic stability by using constant mutation rates and obtaining the stationary distribution as $\lambda \to 0$, equation (5) features time-varying mutation rates. However, even with time-varying mutation, Robles (1998) shows that if $\lambda_t \to \widehat{\lambda} > 0$ (where $\widehat{\lambda}$ is some fixed rate to which time-varying rate will converge), and a stationary distribution exists for the process defined by $\widehat{\lambda}$, then a (modified) long-run equilibrium (MLRE) exists. This MLRE just coincides with the equilibrium of the process defined by fixed $\widehat{\lambda}$ (and can thus be made consistent with KMR, etc.). By a different method, i.e. computing for compound probability distribution, Houba and Tieman (2001) also show that the equilibrium obtained by KMR coincide with the equilibrium obtained from re-casting KMR using time-varying mutation rates. Neither Robles nor Houba and Tieman, however, consider drift as an additional source of noise and nor do they relate mutation to drift.

3.3. Expected Drift and Mutation

Consider some 'average' dynamic in which, with very many draws for $\dot{q}$ and corresponding values for q, drift and mutation rates behave as expected.[13] That is, the pattern of drift and the probability of mutation center around their expected value. Growth $\dot{v}_{1t}$ (conditional on current state v_{1t}) then follows as expected, where the expectation is taken over the distribution of $\dot{q}$ and q. Making the simplifying assumption $\lambda_t = \epsilon q_t = \epsilon(q_{t-1} + \dot{q}_t)$, we let $\dot{v}_{1t} = E(\dot{v}_{1t} \mid v_{1t})$:

$$\dot{v}_{1t} = E\{\left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1 - v_{1t})^2\right] v_{1t} \mid v_{1_t}\} - E[(\sigma \dot{q}_t) \mid v_{1t}] + E\{[\epsilon(q_{t-1} + \dot{q}_t)(1 - 2v_{1t})] \mid v_{1t}\}. \tag{6a}$$

That is, $\dot{v}_{1t}$ behaves as expected - not only in terms of expected payoffs, but also of expected drift and mutation. Assuming independence of q (and hence, $\dot{q}$) from past and current v_1[14], equation (6a) reduces to:

$$\dot{v}_{1t} = \left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1 - v_{1t})^2\right] v_{1t} - \sigma E(\dot{q}_t) + E[\epsilon(q_{t-1} + \dot{q}_t)(1 - 2v_{1t})]. \tag{6b}$$

Now, using our assumption of a uniform distribution for $\dot{q}$, we know that $E(\dot{q}_t) = \frac{a+b}{2}$, for any interval $[a, b]$. If we assume that $\dot{q}$ takes on values from $[-1, 1]$ (as in Example 1), then $E(\dot{q}_t) = 0$. Generally, the range can be larger, but as long as $a = -b$, then $E(\dot{q}_t) = 0$. Of course, $\dot{q}$ can follow other distributions with mean zero and/or non-zero. However, when the drift has mean zero, equation (6b) simplifies to $\dot{v}_{1t} =$

[13] The idea is perhaps similar to Houba and Tieman, although the latter do not consider drift and use the KMR model instead of the Replicator Dynamic.

[14] That is, changes in q affect the next/future values of v_1 through drift and mutation, but not its current nor past values.

$\left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} + E[\epsilon(q_{t-1} + \dot{q}_t)(1-2v_{1t})$, which we can keep iterating until we get:

$$\dot{v}_{1t} = \left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} + \epsilon(q_0)(1-2v_{1t})], \qquad \text{(6c)}$$

which is just equation (4) evaluated at initial measure q_0. Thus, by Proposition 1,

PROPOSITION 3. *If $E(\dot{q}_t) = 0$, an ergodic distribution exists, whose mass is concentrated over the risk-dominant strategy.*

Proof. See proof of Proposition 1. ∎

If the drift has non-zero mean, two cases are possible: $E(\dot{q}_t) < 0$ or $E(\dot{q}_t) > 0$:

PROPOSITION 4. *If $E(\dot{q}_t) < 0$, an ergodic distribution exists, whose mass is concentrated over the risk-dominant strategy.*

Proof. We prove roughly as in Proposition 1 by showing that v_1 always evolves towards $v_1 = 1$ regardless of initial state v_1^0. The dynamic when $E(\dot{q}_t) < 0$ is $\dot{v}_{1t} = \left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} + \sigma E(\dot{q}_t) + [\epsilon(q_0)(1 - 2v_{1t})]$ which will always be positive as long as $\left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} + \sigma E(\dot{q}_t) + [\epsilon(q_0)(1-2v_{1t})] > 0$, or $\frac{-\left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} - \sigma E(\dot{q}_t)}{(q_0)(1-2v_{1t})} < \epsilon$. Note that the inequality is non-binding when initial $\frac{a_2}{a_1+a2} < v_1^0 \leq 1/2$, since for all t the denominator is positive while the numerator is negative, and we know that ϵ cannot be negative. (The first term in the numerator is negative for v_1 greater than the mixed Nash Equilibrium $\frac{a_2}{a_1+a2}$.) Thus, when $\frac{a_2}{a_1+a2} < v_1^0 \leq 1/2$, v_1 increases. When $v_1 > 1/2$, the denominator becomes negative and the inequality becomes binding. We know, however, from Lemma 1 that $[-\sigma E(\dot{q}_t)] < [\ -\epsilon(q_0)(1-2v_{1t})]$. Thus, (rearranging the inequality)$\{ -\left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} - \sigma E(\dot{q}_t)\} < [\ -\epsilon(q_0)(1-2v_{1t})]$, and v_1 still increases beyond 1/2. (At $v_1 = 1$, we still have $0 - \sigma E(\dot{q}_t)] < -\epsilon(q_0)$, which prevents v_1 from decreasing.)[15] When $v_1^0 < \frac{a_2}{a_1+a2}$, $\{-\left[a_1 v_1^0 - (a_1 v_1^0)^2 - a_2(1-v_1^0)^2\right] v_1^0\} < 0$, while $(q_0)(1 - 2v_{1t}) > 0$. Thus, the inequality is non-binding, and v_1 always increases. Thus, whatever the value of v_1^0, v_1 always grows towards $v_1 = 1$. ∎

(Insert Figure 3 here.)

PROPOSITION 5. *If $E(\dot{q}_t) > 0$, an ergodic distribution exists. However, its mass concentrates over the risk-dominant strategy only for small enough $\sigma E(\dot{q}_t)$, otherwise, it concentrates over the other strategy.*

[15] In this sense, $v_1 = 1$ is 'absorbing' when $v_1^0 > 1/2$ in that no escape is possible.

Proof. Note that v_1 grows when $\frac{-\left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1-v_{1t})^2\right] v_{1t} + \sigma E(\dot{q_t})}{(q_0)(1-2v_{1t})} < \epsilon$. We know from Lemma 1 that $\sigma E(\dot{q_t}) > \epsilon(q_0)(1-2v_{1t})$ for all v_1. Now there exists a set Ω_1 of some small values of $\sigma E(\dot{q_t})$ such that at initial $t = 0$, $[\sigma E(\dot{q_t}) - \epsilon(q_0)(1 - 2v_1^0)] < \left[a_1 v_1^0 - a_1(v_1^0)^2 - a_2(1 - v_1^0)^2\right] v_1^0$, and even when $v_1 > 1/2$, we still have $[\sigma E(\dot{q_t}) + \epsilon(q_0)(1-2v_{1t})] < \left[a_1 v_{1t} - a_1 {v_{1t}}^2 - a_2(1 - v_{1t})^2\right] v_{1t}$ at all t. For this set Ω_1, v_1 always increases, regardless of its initial value. (When $v_1 = 1$, $\sigma E(\dot{q_t}) - \epsilon(q_0) > 0$ since by Proposition 1, $\sigma E(\dot{q_t}) > \epsilon(q_0)$. Thus, v_1 decreases at $v_1 = 1$, making $v_1 = 1$ stable.) On the other hand, there exists a set Ω_0 of large values of $\sigma E(\dot{q_t})$ such that $[\sigma E(\dot{q_t}) - \epsilon(q_0)(1 - 2v_{1t})] > \left[a_1 v_{1t} - a_1 v_{1t}^2 - a_2(1 - v_{1t})^2\right] v_{1t}$ at all t and v_1 always decreases, regardless of initial values. (At $v_1 = 0$, $\sigma E(\dot{q_t}) - \epsilon(q_0) > 0$. Thus v_1 keeps decreasing.)[16] ∎

(Insert Figure 4 here.)

Note that Propositions 3 and 4 are consistent with KMR, etc. only because the drift eventually disappears or is expected to disappear - either the underlying distribution is centered around zero[17], or $E(\dot{q_t}) < 0$. In this case, any remaining noise captured by mutation need not overcome the risk-dominance of a strategy.[18] (Of course, the stationary distribution cannot put probability 1 on the risk-dominant strategy as there will always be some (fixed) mutation rate.) Proposition 5 further highlights the importance of drift - too much of it overrides the risk-dominance of a strategy.

4. APPLICATION TO NIE

To apply the dynamics in section 2, and the results in section 3, to the process of institutional change, we first offer formal definitions as to what institutions generally are. In 4.1, we make a broad distinction between the 'Meta Set' of institutions that are available for use/adoption by a given population of agents, and the 'de facto institutions' that such agents actually adopt. The former type might then be thought of as formal rules while the latter might correspond to informal norms. These two types might converge or diverge, depending on the outcome of an evolutionary game which we pattern to the model of sections 2 and 3. That is, in 4.2 we analyze the extent of de facto survival/existence of the (evolving) Meta Set through the medium, long, and 'ultra-long', run, using the dynamic of 3.1, 3.2 and 3.3, respectively. Subsection 4.3 then relates our approach to

[16] In this sense, $v_1 = 0$ is absorbing. In contrast, the system is deflected from $v_1 = 1$.

[17] Recall that Example 1 uses a uniform distribution over $[-1, 1]$.

[18] It still remains to formally establish the extent to which our result might be the same as the MLRE in Robles and/or the notion of stochastic stability as defined in the evolutionary game literature. To do this, one might need to explicitly propose an underlying Markov process that could give rise to the (continuous) dynamics proposed herein.

some leading theories of institutional change.

4.1. Institutions

Recall that in our model of stochastic drift and mutation, the crucial elements that drive the uncertainty are the variable q and its time derivative $\dot{q}$. In the following exposition, we propose how q and $\dot{q}$ might be interpreted in the context of institutional change.

Over time, let new institutions be formed and existing ones destroyed (or replaced), but at the same time, people choose whether or not to use such institutions, thereby strengthening or weakening the latter's 'de facto' existence. Consider the set of all institutions available to a given population of agents at some point in time:

Definition 3. Let vector $\mathbb{I}_t \equiv (\mathbf{I}_t, \mathbf{I}_{t-1}, \mathbf{I}_{t-2}, ...\mathbf{I}_0)$ be called the Meta Set of institutions as of time t, whose elements may be vectors themselves with elements corresponding to single institutions. Thus, at time t, what is available is not only the current set, but all past sets of institutions still in existence. Vectors $\mathbf{I}_t, \mathbf{I}_{t-1}, \mathbf{I}_{t-2}, ...\mathbf{I}_0$ thus have common elements. There are also new ones that can appear (and/or old ones that are replaced) at any time. Suppose that the appearance is approximately continuous. This is possible if we count as a new institution any existing one that has undergone the slightest modification.[19] Counting only the non-redundant elements in $\mathbb{I}_t$, let this measure of 'net additions' as of t to the existing Meta Set be denoted by q, and its time derivative $\dot{q}$. Let the evolution of the Meta Set then be captured by $\dot{q}$, and let the latter be governed by a stochastic process. That is, we are concerned with the rate at which new institutions appear (and/or old ones replaced), and where the timing of appearance is unpredictable.

Given the Meta Set, each agent chooses to either uphold them or not. As more (less) people uphold it, the Meta Set becomes stronger (weaker). That is, institutions are also defined in terms of its 'de facto' existence.

Definition 4. Let de facto existence/survival of the Meta Set be described by the outcome of game G_I (see section 2), where strategy 1 corresponds to the strategy of upholding/using the Meta Set, while strategy 2 is rejecting the Meta Set.[20] Furthermore, let the evolution of de facto survival be captured by equation (4). The Meta Set and de facto institutions converge as proportion $v_1 \longrightarrow 1$, and diverge as $v_1 \longrightarrow 0$.

Remark 1. On the universal applicability of game G_I. The payoffs of the game G_I can then capture the extent of 'fitness' of players, but not

[19] Thus, we capture the notion that new institutions build up on previous ones.

[20] For an agent that adopts strategy 1, the Meta Set is also the de facto set of institutions. If an agent adopts strategy 2, s/he in effect establishes de facto institutions that are different from the Meta Set.

solely in the biological sense. One can think of them as utilities derived from the use of, or having access to, existing institutions that affect such utility. Such definition is general enough to accommodate pure biological survival, social and/or economic 'fitness', or any combination thereof, by choosing the appropriate combination of institutions. For instance, we can consider a subset $\mathbb{I}_t^N \subset \mathbb{I}_t$ as the particular Meta Set that generates the utilities described by payoff matrix A. This also defines membership of the particular population in question, i.e. $\mathbb{I}_t^N$ is the set of institutions that the N players initially face and are choosing to uphold when choosing strategy C and to reject when choosing A. Of course, we can also study the entire set $\mathbb{I}_t$, thereby combining all populations of agents that play the 'institution game' into a single population. The point is that for any set of institutions $\mathbb{I}_t^N \subset \mathbb{I}_t$, the latter's de-facto existence can be determined from the outcome of the game generally described by G_I. Recall from definition 3, however, that we also allow such initial set to evolve. The measure of the set whose members are the non-redundant elements of $\mathbb{I}_t^N \in \mathbb{I}_t$ can then still be denoted by q_t[21], and the change in this measure as $\dot{q}_t$.

Assume that for the N players, the strategy to uphold the Meta Set is risk-dominant. That is, $a_1 > a_2$.[22] This risk-dominance is preserved even while $\mathbb{I}_t^N$ evolves, since A does not change throughout the (evolutionary) game. If/when the measure q is/becomes fixed, the 'de-facto' game still goes on (except in the unlikely case where $\mathbb{I}_t^N$evolves to a null set). To put in another way, the evolution $\dot{q}_t$ of the relevant Meta Set $\mathbb{I}_t^N$ only pertains to creation of new, and/or destruction of old, institutions that preserve the specific values of a_1 and a_2 in the payoff matrix.[23] Other values would describe payoffs obtained from upholding/rejecting another set $\mathbb{I}_t^{H \neq N} \subset \mathbb{I}_t$ and any subsequent additions to, and/or deductions from, that set.[24]

By the use of game G_I, we propose that the coordination game can accommodate de facto evolution of any and/or all sets of institutions. The coordination game is the most appropriate since it is flexible in the sense that it allows multiple equilibria and, in a way, can capture the development of all kinds of institutions, i.e. 'good' or 'bad'. In other words, the coordination game does not make any ex ante judgments as to which types of institutions could and should prevail. The only requirement for institutions to persist is that society coordinates on their behavior to uphold them.[25]

[21] Assuming, say, that the number of non-redundant elements is always a finite set, or countably infinite, measure q_t is then the cardinality of this set at t.

[22] For more on risk-dominance, see Harsanyi and Selten (1988).

[23] Strictly, then, we may use notation q_t^N for q_t and $\dot{q}_t^N$ for $\dot{q}_t$.

[24] We still maintain, however, that the game of upholding/rejecting *any* set of institutions has the coordination-game structure, albeit the risk-dominant strategy may be to uphold for some sets, and to reject for others.

[25] Peyton Young (1998) also uses the coordination game to illustrate the evolution of conventions.

At the very least, we only need to capture two general components to help define institutional change: (a) all institutions represent (collective) reaction of individuals to some stimulus, which get (b) affirmed or become 'stronger' as more and more individuals adopt the same reaction. The idea is that individuals commit themselves to some kind of collective way of doing things - a social technology. Component (b) might then be easily captured by a 2×2 coordination game with strategies 'Uphold' and 'Reject'. Of course, while de-facto evolution can be captured by game G_I (and equation 4), we still have to define the appropriate 'stimuli' that give rise to institutions or induce evolution in the Meta Set. Nevertheless, given these stimuli, a coordination game can generalize reactions to two kinds of strategies that rule individuals' decision-making with regard to the formation of institutions.

One obvious stimulus is the 'environment'. However, the notion of the environment is complex in that it can be 'context specific'. For one, the environment can change as specific institutions are formed inasmuch as the latter help define the new environment that individuals face for future plays of the game. Precisely to abstract from such endogeneity, we let the Meta Set already include all past institutions from which newer institutions build on. This helps to define our notion of the 'environment' as the truly exogenous component of institutional change.

Perhaps the only component of the environment that can be taken as truly exogenous by any individual (and hence, in the formation of any kind of institution) is that which existed before the first individual was born, i.e. the 'physical' environment. If the first institution was a reaction to this initial environment, then in effect all succeeding institutions will also be reactions to this inasmuch as future institutions build up on previous ones. Thus, solving the problem of infinite institutional regress boils down to assuming a first/initial environment beyond which no individual was in existence, since this initial environment cannot have been a 'created' institution.

Given this initial environment, we can then identify the appropriate kinds of reactions. To be applicable to all institutions, we start from assuming a single basic instinct that all individuals (living during any time period) share - survival or self-preservation. The competing 'reactions' to the environment should then reflect the options available to any individual that will foster self-preservation. We posit only two general options that can do this: use existing conventions/norms in society, or 'deviate' by establishing other ways of living in society. The former option can then be tantamount to 'upholding' current institutions, whatever they may be, while the latter corresponds to rejection of the current set. Note, however, that the ultimate aim of both options is to increase the chances for, or extent of, self-preservation. Note also that there is no a priori reason why upholding current norms might be more effective than deviating from them, so we cannot readily assume that the latter strategy is always dominated.

Remark 2. On the applicability of equation (4). In using equation (4) to describe the dynamic of institutional change, we not only model the growth of de-facto institutions, but explicitly tie this with the evolution of the Meta Set through q and $\dot{q}$. While q and $\dot{q}$ are exogenous to the agents, they are 'allowed' to react to these by playing game G_I. Also, inasmuch as q and $\dot{q}$ are random elements, the endogenous reaction of agents are underlined by uncertainty. Thus, evolution of both the Meta Set and de facto institutions are stochastic. Note by definition 1 that positive $\dot{q}$ or an increase in q induces 'drift' *out of* v_1 (while decreases in q induces drift into v_1). The intuition is that agents are more hesitant to 'trust' and uphold an increasing number of institutions (while they find it easier to trust a smaller Meta Set). On the other hand, the specification in definition 2 - that the mutation rate is larger when q is bigger - captures the notion that a larger Meta Set induces more (directionless) mutation. That is, more agents will want to switch from upholding to rejecting, and vice-versa. In this sense, they put less trust on the correctness of upholding and/or rejecting the Meta Set when the latter is large (while it is easier to decide when the Meta Set is small). Lastly, note that by Lemma 1, the *evolution* of the Meta Set generally causes larger random outflows from v_1 in the form of (big) drifts than the *size* of the Meta Set per se which only affects (small) mutation. Hence, de facto survival of the Meta Set is harder to achieve when the Meta Set evolves, i.e. when $\dot{q}$ (and hence $\sigma g(\dot{q})$) is not equal to zero.

4.2. Institutional Change

Using dynamic given by equation (4), and its specific cases given in subsections 3.1, 3.2, and 3.3, we analyze the survivability of an evolving Meta Set through three time horizons - the medium, the long, and the ultra-long, run. We define such horizons in the following manner.[26] Generally, let the medium run approximate some (arbitrary) time period in which the system is almost stable if not for some small mutations or 'experimental' behavior undertaken by agents. (The appropriate dynamic is thus given in 3.1.) Over a longer period of time, we allow bigger sources of uncertainty in including drift into and out of strategies. The long-run dynamic thus explicitly accounts for both drift and (time-varying) mutation (which is given in 3.2). In the ultra-long run, after sufficiently many drift and mutation 'experiences', agents can form expectations about the nature of drifts and mutation, and aggregate behavior is as expected. (The dynamic

[26] The manner in which the terms 'medium, long and ultra-long run' are used here are not conventional in the (evolutionary game) literature where 'medium run' typically refers to analyses of a sample path starting from an initial proportion, while long and/or ultra-long run are usually achieved by letting 'the dust clouds settle', as mutation vanishes, i.e. $\lambda \to 0$. Thus, long and/or ultra-long run equilibria correspond to stochastic stability and the stationary distribution as obtained in KMR, etc.

thus takes the forms given in 3.3).

More specifically, we can interpret these horizons according to the evolution of the Meta Set. In the medium run, the Meta Set is approximately stable in that it is not evolving, i.e. $\dot{q} = 0$, and its measure q is thus fixed. There may be some small uncertainty in the form of mutations, but there is no large uncertainty in the form of drift. In the long run, the Meta Set starts to evolve, which now generates drift (and time-varying mutation). In the ultra-long run, the Meta Set has evolved sufficiently many times such that q and $\dot{q}$ become centered on their expected values. In this case, either the evolution of the Meta Set settles down to zero (and its measure becomes fixed), or its growth (and/or destruction) does not taper off, but the point is that agents can already form expectations about the pattern of this evolution and behave accordingly.

4.2.1. The Medium Run

When the Meta Set is not (yet) evolving, we can approximate the dynamic by setting $\dot{q} = 0$ and fixing q. As in 3.1, the general dynamic of institutional change given by equation (4) reduces to equation (1). Proposition 1 implies that:

PROPOSITION 6. *In the medium run, the latter tends to (de facto) survive if it is risk-dominant to uphold it.*

Proof. See proof of Proposition 1. ∎

Risk-dominance of the strategy to uphold the Meta Set implies, roughly, that using available institutions is a 'safer bet' for increasing an agent's utility or for fostering self-preservation. Deviating or establishing new de facto institutions might be the 'pareto dominant' strategy - doing so is rewarding only when encountering other 'deviants' as well, while the punishment is bigger when one encounters 'conformants'. In contrast, conforming to current norms might be better 'on average'. Of course, whether conforming to, or upholding, the Meta Set is a risk-dominant strategy is an empirical question.

It is unlikely, however, that $\dot{q}$ and q remain fixed as $t \to \infty$. The medium run dynamic then becomes a poor approximation as the Meta Set starts to evolve. Nevertheless, we can use equation (1) to describe institutional change that is solely due to the evolution of de facto survival of a *fixed* Meta Set of institutions.

4.2.2. The Long Run

Consider a longer time period during which the Meta Set starts evolving. In this case, $\dot{q} \neq 0$. Suppose, as in 3.2, that $\dot{q}_t$is drawn from a uniform distribution over some interval. We simplify the drift function to be equal to $g(\dot{q_t}) = \dot{q_t}$, and let $q_t = q_{t-1} + \dot{q}_t$,so that the (time-varying) mutation rate is $\lambda_t = e(q_t) = e(q_{t-1} + \dot{q_t})$. The appropriate dynamic to capture institutional change in the long run is then given by equation (5). By Proposition 2:

PROPOSITION 7. *The Meta Set may or may not survive through the long run. One the one hand, even when the Meta Set of institutions is in continuous flux, agents will not necessarily reject it. On the other, the instability of institutions might lead to de-facto rejection. Whether or not conformity to the Meta Set is a risk-dominant strategy is irrelevant.*

Proof. See proof of Proposition 2, particularly noting that for high enough initial value v_1^0, the dynamic evolves towards $v_1 = 1$, while for low enough v_1^0, $v_1 \longrightarrow 0$. ∎

Proposition 7 can explain how agents can still hold on to past institutions and readily accept new ones even when the nature of such institutions keep changing. That is, they can still end up conforming to the Meta Set even as it changes (whether or not conforming is a risk-dominant strategy). Of course, the above is not a general result - for some low enough initial value for v_1^0, the long run dynamic evolves instead towards $v_1 \longrightarrow 0$ and the Meta Set de facto weakens (or alternatively, de facto institutions become increasingly different from the Meta Set). However, this outcome might be less likely the longer the medium run - the time period during which the Meta Set is still fixed. This is because by Proposition 1, it is always the case that $v_1 \longrightarrow 1$ as long as $\dot{q} = 0$ and q is fixed. Thus, the initial value v_1^0 in the long run dynamic, which roughly corresponds to the 'final' value in the medium-run dynamic, might be large enough such that $v_1 \longrightarrow 1$ all throughout the long run. To put it in another way, we could define a set L of all possible minimum values of initial v_1^0 that ensure that the long-run dynamic approaches $v_1 = 1$, and show that whenever the stable value of v_1 obtained in the medium-run dynamic (equation 1) is contained within L, then the Meta Set would survive all throughout the long run. This would require, however, that the stability of the medium-run dynamic first has to be attained before the long-run dynamic begins, that is, before the Meta Set starts evolving. Otherwise, if the Meta Set starts evolving too quickly, the (long-run) dynamic would likely evolve towards $v_1 \longrightarrow 0$ and the Meta Set would de facto weaken (or, equivalently, would be replaced by de facto institutions that are different from the Meta Set). In this sense, rapid complexity in the Meta Set tends to inhibit the latter's (long-run) de facto survival.

At this point, however, we cannot readily define set L since the particular long-run dynamic we obtained as equation (5) assumes a uniform distribution for $\dot{q}$. It is not altogether clear to what extent other types of distributions can allow long-run evolution of v_1 towards one. But what is important to note from Propositions 2 and 7 is that the survival of the

Meta Set through the long run is not guaranteed, even if conformity to the Meta Set is a risk-dominant strategy. To generalize to all types of distributions, we propose instead to look at the 'ultra-long' run horizon, in which the expected value of $\dot{q}$ is revealed.

4.2.3. The Ultra-long Run

The 'ultra-long' run begins when the Meta Set has evolved sufficiently many times, such that the pattern/trend of its evolution becomes 'clear'. By using dynamic 3.3 to model the growth of v_1 in the ultra-long run, we approximate such pattern/trend by its 'expected' value. Thus, the equilibrium outcome of the ultra-long run dynamic is seen as some 'average' result after considering very many experiences of changes in the Meta Set, that is, many draws for $\dot{q}$ and hence values for q. We use equation (6b) to generally describe institutional change in the ultra-long run, which reduces to equation (6c) for the following specific case:

DEFINITION 5. When the expected change in the measure of the Meta Set of institutions is zero, we say that institutional change is *neutral*. In this case, either there are no 'creation' and 'destruction' of institutions in the ultra-long run, or whatever is created just exactly replaces what is destroyed.

Note that by Proposition 3, we can derive the following result:

PROPOSITION 8. *When institutional change is neutral, the Meta Set survives through the ultra-long run (or, equivalently, the Meta Set and de facto institutions converge) if it is a risk-dominant strategy to conform to the Meta Set. Otherwise, if it is risk-dominant to deviate from the Meta Set, then the Meta Set does not survive (or, equivalently, the Meta Set and de facto institutions diverge).*

Proof. See proof of Proposition 3. ∎

If institutional change is non-neutral, two cases are possible: when $E(\dot{q_t}) < 0$, the Meta Set tends to become simpler over time; when $E(\dot{q_t}) > 0$, it tends to become more complex. By Proposition 4,

PROPOSITION 9. *If the Meta Set tends to become simpler over time, the Meta set survives through the ultra-long run if it is a risk-dominant strategy to conform to the Meta Set. (Otherwise, the Meta Set does not survive.)*

Proof. See proof of Proposition 4. ∎

Note that the above Propositions 8 and 9 imply that when complexity of the Meta Set is bounded from above, ultra-long run survival of the Meta Set depends on whether or not it is risk-dominant to keep conforming to the Meta Set. The following Proposition, which follows from Proposition 5, makes risk-dominance irrelevant and instead points to the size of the drift, or the rate at which the Meta Set becomes more and more complex, as the crucial factor determining ultra-long run survivability of the Meta Set:

PROPOSITION 10. *If there is no upper bound for the extent of complexity of the Meta Set, the Meta Set survives through the ultra-long run only when it does not become too complex too fast, (that is, only for small enough* $\sigma E(\dot{q}_t)$*).*

Proof. See proof of Proposition 5. ∎

Taken together, the results in Propositions 6 to 10 all seem to substantiate the intuition that too much instability in the Meta Set can lead to its deterioration or, equivalently, to the establishment of de facto institutions that are different from the Meta Set. Note that the medium-run result favoring de facto survival of the Meta Set assumes that the Meta Set is fixed (and hence, stable) while survival through the long run requires enough stability to begin with, i.e. initial v_1^0 is large enough. Lastly, the ultra-long run dynamic allows survival of the Meta Set only when the Meta Set does not become too complex too fast. While the results rely on what happens to the Meta Set (and the amount of drift/uncertainty it causes), we have not fully motivated or explained its evolution. The next subsection proposes an interpretation that could be consistent with the NIE literature.

What Drives the Evolution of the Meta Set? We have assumed that the evolution of the Meta Set is exogenous to the game of de-facto upholding institutions. This, however, is not the same as saying that the same agents are not responsible for creating new, or destroying old, institutions. It is only that whenever they create or destroy, it is assumed to be a 'random' occurrence in that the timing is unpredictable. (An alternative 'endogenous' model would perhaps time an occurrence or disappearance upon reaching a particular 'threshold' state $\overline{v_C}$.)

Nevertheless, we can let $\dot{q}$ be tied with the motion governing an underlying stochastic process $\mathcal{E}_t$ with mean zero, even if in the ultra-long run $E(\dot{q}_t)$ is not zero. Specifically, we can define $\dot{q}$ to be a linear function $f : \mathbf{E} \to Q$, $f' \gtrless 0$, $f'' = 0$, i.e. $\dot{q}_t = a\mathcal{E}_t + b$, where $a \gtrless 0$ and $b \gtreqless 0$. Since f is linear in $\mathcal{E}_t$, we can then obtain the expected value $E[f(\mathcal{E}_t)] = f[E(\mathcal{E}_t)] = f(0) = E(\dot{q}_t) \gtreqless 0$.

We can interpret the process $\mathcal{E}_t$ as the (stochastic) process governing incremental changes to the 'environment'. Thus, *changes* in the environment are random occurrences that are captured by a random variable that is distributed with mean 0.[27] On expectation, the environment stays the same, i.e.$E(\mathcal{E}_t) = 0$, which implies that 'additions' to the original environment are, 'on average', offset by 'destruction'.

The 'environment', in turn, can include not only geography but also other constructs that are exogenous to agents as they play game G_I, and which affect the evolution of the Meta Set (that is, the creation of new and/or destruction of old, institutions by the same agents). Interpreting $\mathcal{E}_t$ as changes in the environment adds two final features that make the model a comprehensive tool for analyzing institutional change: first, we capture the notion that institutions are formed in order to adapt to the environment; and, second, we solve the problem of infinite institutional regress by assuming some given initial environment that is exogenous to agents and from which the first Meta Set was developed.

4.3. A New Approach

The literature on institutional analysis has by now grown quite large. However, there is still some disagreement as to the best way to define and model institutions. The canonical approach derives from North's view of institutions as "the rules of the game." However, as Greif has noted (Greif, 2005) the institutional arrangements that emerge in stable equilibria are themselves an endogenous response to existing conditions. While some institutions can be thought of as formal rules or laws that condition behavior, they are only viable if the enforcement characteristics lead to compliance. In some cases, the institutional rules may not even be directly relevant to the norms and behaviors that emerge. In others, seemingly "inconsistent" rules (when judged by their match with informal norms of behavior) can become established if the state or society pays a high enough price in terms of forcing compliance to the point where people learn to consider the binding rules as part of the background environment. Obviously, a truly satisfactory model of institutions will capture the differences between the initial choices and conditions faced by the economic actors and the equilibrium behaviors that will emerge when we take into consideration how the initial rules are enforced and what choices the actors make when responding to both general socio-economic incentives and the specific enforcement mechanisms used to establish the variety of formal laws and norms. As Greif has noted, "If prescriptive rules of behavior are to have an impact, individuals must be motivated to follow them."

Our particular focus has been on reconciling the views of North and Greif. In particular, we are interested in studying how institutions evolve when we distinguish between the Meta Set of all existing and potential institutional arrangements faced by individuals and the de facto set of institutional arrangements that actually persist when taking into account the

[27] The initial environment is itself may be exogenous and given, but the subsequent changes to this are random.

rules and norms that the agents actually choose to uphold in equilibrium. Individuals decide at each point in time whether to conform to pre-existing rules and institutions, thereby de facto enforcing them, or whether to adopt an individualistic set of behaviors that abandon the conformity to the collective rules and behaviors.

Greif and his followers have worked hard to demonstrate the usefulness of classical game theory in modeling individual responses to institutional rules, even using results from dynamic (i.e. repeated) games to capture the notion of path dependence in institutional change. Path dependent stories typically give greater weight than static analyses to the possibility that chance initial events lead to divergent outcomes. However, Greif's framework is still incomplete because the arrival of chance events is not clearly modeled and their effect on the behavior of agents is not shown in an explicit dynamic context. Furthermore, the very uncertainty induced by chance casts serious doubt as to the aptness of classical game theory which assumes perfect rationality of agents. Since the appearance and disappearance of institutions, and whether the latter end up being effectively enforced or not, can seldom be predicted (ex ante), there is a possibility that experimention/mutation might lead to the 'correct' strategy (ex post).

A persuasive argument can be made that a more limited, bounded rationality model of human behavior might better capture the observed empirical regularities regarding the emergence and persistence of historical institutions. Evolutionary game theory thus seems like a more appropriate framework. Indeed, H. Peyton Young (1998) argues for evolutionary models as more closely following the spirit of classical economics with its notion of individuals and societies adapting to changed circumstances in coherent fashion without imposing a full calculative, all knowing, hyper-rationality on the part of all players involved.

In this sense, Aoki is probably more encompassing than Greif in that he does not eschew the role of mutations or bounded rationality and, hence, uncertainty, in the evolution of the 'rules of the game'. Aoki distinguishes between institutions that are deemed exogenous, and those that evolve within a particular game (in a manner akin to our distinction between the Meta set and de facto institutions). Aoki's proposed methodology is comprehensive - it covers all kinds of institutions; recognises the 'embeddedness' and linkages in the evolution of institutions across domains (i.e. organizational, economic, social and political spheres) and the strategic complementarities across these domains that are necessary for institutions to be enforceable. It allows for possible multi-directional endogeneities in the evolution of such institutions. However, the comprehensivenss of the framework makes it difficult to obtain specific results.

The difficulty stems from the fact that while Aoki outlines the players, payoffs, and forms of games for each kind of domain, the treatment of strategies is somewhat loose and confusing. That is, what, exactly, are the strategies and how are they defined? It seems that they are not specifically institutions, yet adopting them gives rise to institutions.

In other words, it is not clear how a particular action can produce a specific institution - is it that a strategy is literally an 'arrangement' which, as more and more players adopt it, becomes an institution? In this case, then, a strategy is, literally, a 'potential' institution which only "de facto" becomes so as it becomes an equilibrium. (This interpretation would be closer to our model.) Or does creation of an institution follow a more indirect route, i.e. is it a by-product of playing the game? Players choose strategies recursively and the evolving state of play embodies information from which players can deduce the 'rules' governing the game. That is, it is the collective information, or the reinforced belief of how the game ought to be played, that defines an institution, and not the strategies per se. This latter interpretation seems to be a more plausible and consistent reading of Aoki. For if Aoki meant the former 'direct' route where strategies are literally institutions, then another, more problematic issue is apparent:

Are the games, i.e. the number of strategies, finite or infinite? With the 'indirect' route, it is easy to restrict strategies to a finite set and still allow possibly infinite by-products or institutions. However, with the direct route, one might need to allow for infinite strategies, for it seems hard to limit ex ante the set of institutions to a finite set. And even if we did, it seems difficult to define (ex ante) the finite set of strategies. How would anyone know which particular institutions can possibly come out of a game? It is arguably more realistic to allow the emergence/occurrence of at least some institutions to be unpredictable and which, cannot be defined ex ante. Furthermore, it is difficult to justify restricting the timing of occurrences only to the period before the game is played. One can easily envision new institutions emerging as the game is played repeatedly. The point is that if strategies were literally institutions, as in the 'direct' route, one would need to devise a technical framework where new strategies emerge (and possibly stochastically) as the game is repeated.

Thus, within Aoki, it is difficult to formulate any testable hypotheses/predictions on either the particular equilibria attained in a game or the path and speed by which they are attained.

In a way, Young offers a good compromise between Greif and Aoki. He models uncertainty, and proposes exact (evolutionary) games and (stochastic) dynamics by which institutions (conventions, norms, contracts) are formed. While our model is perhaps closest to Young, there are two major limitations in Young that we overcome. First, as he himself admits, there are higher level institutions in which lower level institutions are 'embedded', but whose evolution can all be combined in one game. "Bargaining over contract forms takes place within the shadow of the law, and the law operates within the penumbra of morality, morality is colored by religious belief....Doubtless all of these interactions could be written down as one large game." (Young 1998). Our model addresses this simultaneous notion of hierarchy and comprehensiveness through our Meta Set, which is defined as a combination of all available institutions which could have been built from previous ones, and our singular 'de facto' game of players choosing to uphold these Meta institutions or not, whatever these institutions are, and as they change over time.

The other limitation is that the bounded rationality in Young is arbitrary. In our model, bounded rationality is manifested not only in the 'small' mutation/experimental behavior of agents but also in the 'large' 'drift' that ensues from having the Meta Set evolving as well. This specification is far from innocuous. In fact, it can generate equilibrium predictions that are different from those implied by the standard notion of stochastic stability in current evolutionary-game models including Young. By defining the evolution of the Meta Set as an explicit source of uncertainty, the extent of bounded rationality in our model becomes intuitively time-varying. Thus, unlike most evolutionary-game models, we do not just assume some fixed mutation rate and obtain the (stochastic) equilibrium by examining what happens when we let this mutation rate disappear (at an arbitrary rate). We offer a simple yet intuitive technique for solving for (ultra-long run) equilibria when bounded rationality is time-varying and when they may or may not disappear – we take the expectation of the dynamic over the distribution of the source of the uncertainty. A general result is that depending on this distribution, equilibrium in a 2×2 coordination game may not necessarily correspond to the risk-dominant strategy, unlike in Young and most other evolutionary-game models.

5. CONCLUSIONS

Our proposed stochastic dynamics are based on Foster and Young and Cabrales, but we do not specify a Wiener process to model drift, and we relate drift and mutation through a common variable driving 'uncertainty'. While there exist many other stochastic dynamics which capture mutation (and/or drift)[28], our models of mutation and drift are adaptive and time-varying. To our best knowledge, apart from Robles and Houba and Tieman, results from evolution with time-varying mutation rates (and/or time-varying drift) are still not clearly generalizable.

Using a (evolutionary) coordination game and our proposed stochastic dynamics to derive equilibria relating to institutional change, our bounded rationality framework and methodology are closer to Young than Aoki or Greif. Greif captures the notion of self-reinforcing equilibria, but still in the classic game-theoretic framework. Aoki does allow for bounded rationality, but he models institutional change as a 'by-product', rules governing *any* game, that is, any type of strategic interaction, which also evolve as the

[28] See KMR, Young, Samuelson, Binmore, Samuelson and Vaughan, Blume (2003), Amir and Berninghaus (1996), Fudenberg and Harris (1992), Robson and Vega-Redondo, and Vega-Redondo (1996, 2003) for a survey.

game is played. In contrast, our 'institution game' is *literally* the game of choosing to uphold or reject institutions.

Intuitively, however, we draw more heavily from the NIE literature (North and Greif) than Young in order to formalize concepts. Thus, our model is more exact and specific to institutional change. At the same time, results pertaining to institutional change are obtained as applications of the more general results we derive for 2×2 evolutionary coordination games whose outcome can be described by the stochastic dynamics we have proposed.

The paper's contributions are thus both methodological and conceptual. While we provide generic results for evolutionary games, our model is also specific enough to the phenemenon of institutional change. We believe that the framework serves as an effective formalization of the critical intuitive observation in the NIE literature that the problems inherent in reforming instutions emerges when there is a serious disjunction between formal rules and informal norms of behavior. Our model allows us to examine cases when there is a disjunction between formal rules and informal norms by considering the problem of how agents choose to comply or ignore existing rules. Even while the two might initially match in equilibrium, random shocks (in the form of drift and mutation) can drive a wedge between the two, which might or might not lead back to convergence of formal and informal norms. This not only serves as a model of how unanticipated shocks change institutions but also of how important purposeful changes (such as "shock therapy" or new enforcement mechanisms) have to be to create sustainable reform. Above all, it provides an intuitively plausible formal model of the various notions of path dependence and non-ergodicity that are an important but unmodeled subset of the ideas of the New Institutional Economics.

REFERENCES

Aoki, Masahiko

Amir, Madjid and Berninghaus, Siegfried K., 1996. Another Approach to Mutation and Learning in Games. Games and Economic Behaviour 14, 19-43

Bergin, James and Lipman, Barton L., 1996. Evolution with State-Dependent Mutations. Econometrica 64:4, 943-956.

Binmore, Kenneth G., Samuelson, Larry, and Vaughan, Richard, 1995. Musical Chairs: Modeling Noisy Evolution. Games and Economic Behaviour 11, 1-35.

Blume, Lawrence E., 2003. How Noise Matters. Games and Economic Behaviour 44, 251-271.

Boylan, Richard T., 1995. Continuous Approximation of Dynamical Systems with Randomly Matched Individuals. Journal of Economic Theory 66, 615-625.

Boylan, Richard T., 1992. Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals. Journal of Economic Theory 57, 473-504.

Cabrales, Antonio, 2000. Stochastic Replicator Dynamics. International Economic Review 41:2, 451-481.

Foster, Dean, and Young, Peyton, 1990. Stochastic Evolutionary Game Dynamics. Journal of Theoretical Population Biology 38:2, 219-232.

Friedman, Daniel, 1991. Evolutionary Games in Economics. Econometrica, 59:3, 637-666.

Fudenberg, Drew, and Harris, Christopher, 1992. Evolutionary Dynamics with Aggregate Shocks. Journal of Economic Theory 57, 420-441.

Gihman, I. and Skorohod, A.V., 1972. Stochastic Differential Equations. Springer-Verlag, Berlin.

Greif, Avner, 2006. Institutions and the Path to the Modern Economy: Lessons from Medieval Trade. Cambridge University Press.

Harsanyi, John C., and Selten, Reinhard, 1988. A General Theory of Equilibrium Selection in Games. The MIT Press, Cambridge, Mass.

Houba, Harold and Tieman, Alexander F., 2001. Idiosyncratic and Aggregate Time-Varying Mutation Rates in Coordination Games. Tinbergen Institute Discussion Paper TI 2001-041/1.

Kandori, Michihiro, Mailath, George J., and Rob, Rafael, 1993. Learning, Mutation, and Long Run Equilibria in Games. Econometrica 61:1, 29-56.

North, Douglass C.

Robles, Jack, 1998. Evolution with Changing Mutation Rates. Journal of Economic Theory 79, 207-223.

Robson, Arthur J., and Vega-Redondo, Fernando, 1996. Efficient Equilibrium Selection in Evolutionary Games with Random Matching. Journal of Economic Theory 70, 65-92.

Skorohod, A.V., 1989. Asymptotic Methods in the Theory of Stochastic Differential Equations. American Mathematical Society, Providence, RI.

Taylor, P. and Jonker, L., 1978. Evolutionary Stable Strategies and Game Dynamics. Mathematical Biosciences 40, 145-56.

Vega-Redondo, Fernando, 2005. Economics and the Theory of Games. Cambridge University Press, UK

Vega-Redondo, Fernando, 1996. Evolution, Games, and Economic Behaviour. Oxford University Press, UK.

Weibull, Jorgen W., 1995. Evolutionary Game Theory. The MIT Press, Cambridge, Mass.

Young, H. Peyton, 1993. The Evolution of Conventions. Econometrica 61:1, 57-84.

Young, H. Peyton, 1998. Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press, UK.