Mean-Variance Analysis in Temporary Equilibrium: Chaotic Transients and Crises*

Michael T. Rauh†
Department of Economics, Management School, University of Liverpool, Liverpool L69 7ZH England.

Giulio Seccia
Department of Economics, University of Southampton, Southampton SO17 1BJ England.

Abstract. Mean-variance analysis (MVA) relies crucially on the assumption that returns are normally distributed. Since real-world returns are non-normal, it would seem that MVA leads to potentially large, systematic mistakes. In this paper, we embed MVA in the standard OLG model with one asset, money. In temporary equilibrium, agents mistakenly assume that the distribution for the real return on money is normal but make correct conjectures about its mean and standard deviation. Simulations show that this leads to chaotic transients via a boundary crisis. From the agents' perspective, the deviations from normality of the true distribution of returns appear small and unsystematic.

Keywords: mean-variance analysis, temporary equilibrium, boundary crisis.

JEL Classification Numbers: D81, G11.

---

* We thank participants at the 2002 WEHIA conference in Trieste. We also thank David Brookfield, Robert Hilborn, Brendan McCabe, and seminar participants at the University of Liverpool for comments. The usual disclaimer applies.

† Corresponding author.

1. Introduction

Mean-variance analysis (MVA) is one of the principal models of portfolio choice in finance. We informally describe this model as follows.[1] An agent has wealth W which she can invest in J assets. The return R_j on each asset j is a random variable. Let α_j denote the fraction of W invested in asset j, where negative α_j corresponds to short selling. A *portfolio* is a vector $\alpha = (\alpha_1, \ldots, \alpha_J)$ such that $\sum_{j=1}^{J} \alpha_j = 1$. The return on portfolio α is the random variable $R_\alpha = \sum_{j=1}^{J} \alpha_j R_j$.

We now come to the main assumption of MVA; indeed, the assumption that gives MVA its name: the agent only cares about mean and variance. More specifically, the expected utility of wealth for portfolio α depends only on the mean and variance (or standard deviation) of R_α. In contrast, in state-preference theory expected utility depends on the complete specification of the distribution for R_α, not just a few summary statistics like mean and standard deviation. E.g., the agent may care about the third (skewness) or higher-order moments.

MVA can be reconciled with state-preference theory by assuming that the return R_j on each asset is normally distributed and that the agent knows the mean and standard deviation of R_j, as well as its covariances with all other assets.[2] These assumptions imply that R_α is normally distributed, the agent can compute its mean μ_α and standard deviation σ_α, and the agent's expected utility depends only on $(\mu_\alpha, \sigma_\alpha)$ since the normal distribution is characterized by these parameters.

In this paper, we focus on two seemingly problematic aspects of MVA. First, it is a static model in a dynamic world. Second, MVA assumes that returns are normal when it is well-known that real-world returns are non-normal with "fat tails". This has a theoretical parallel in general equilibrium frameworks where the real return on any nominal asset is endogenous[3] so one cannot impose *a priori* restrictions on real returns such as normality.

[1] E.g., see Copeland and Weston (1988).

[2] MVA can also be reconciled with state-preference theory by assuming that the von Neumann-Morgenstern utility function is quadratic in wealth but it seems perilous to base an entire theory of portfolio choice on a single utility function. Furthermore, quadratic utility only makes sense for low enough wealth levels and even then it exhibits increasing absolute risk aversion, which seems implausible.

[3] The real return on any nominal asset must be endogenous because prices are. But returns may be endogenous for other reasons; e.g., the return on a particular stock depends on the behavior of the firm.

Even if agents believe that returns are normal, in general the true (equilibrium) distribution of returns will be non-normal. Despite its status as one of the principal portfolio choice models in finance, it would seem that MVA leads to potentially large, systematic mistakes.

In this paper, we investigate the consequences of using MVA in a dynamic general equilibrium environment where the true cumulative distribution function (CDF) for returns is actually non-normal. We consider the standard pure exchange overlapping generations model with one consumption good and one asset, money.[4] At each date t, the young agent decides how much money to hold assuming that the real return on money r_{t+1} at $t+1$ is normally distributed on $[0, \infty)$. Furthermore, in accordance with MVA, the agent makes correct conjectures about the mean μ_t and standard deviation σ_t of r_{t+1}. Note that the agent *assumes* that r_{t+1} is normally distributed rather than contemplating market-clearing conditions for future dates and states as she would under rational expectations. The relevant equilibrium concept is therefore that of Hicks-Grandmont *temporary equilibrium* (TE): at each date t, the markets for consumption and money clear and the young agent makes correct conjectures about μ_t and σ_t but markets need not clear at future dates and states from the perspective of date t.[5] Although r_{t+1} does indeed have mean μ_t and standard deviation σ_t, it is an endogenous random variable (a function of μ_t, σ_t, and the stochastic endowment) whose induced CDF is non-normal.

The model leads to a deterministic system of nonlinear difference equations in (μ_t, σ_t). We use simulations to study the dynamics of the system as a certain control parameter a (a parameter for the CDF for endowments) is varied. As in Grandmont (1985), we are only able to compute explicit backward orbits. For low values of a, we observe convergence to a stable node in the backward dynamics (divergence from an unstable node when the orbits are read forwards). For higher values of a, we observe *chaotic transients* which wander about chaotically in a neighborhood of the unstable node for finitely many periods before hitting an "escape region" after which they rapidly diverge.[6] The duration of these

[4] One might object to a one-asset framework since MVA is a theory of portfolio choice. But in this paper, we are more concerned about mean-variance decision-making than about portfolio composition *per se*.

[5] For more on TE theory, see Grandmont (1977, 1985, 1998) and the references therein.

[6] The model is bounded so the divergence cannot continue forever.

chaotic transients seems to increase exponentially with a. This "route to chaos" matches the *boundary crisis* scenario developed by Grebogi, Ott, and Yorke (1982, 1983) which has been observed in the logistic and Hénon maps as well as physical laboratory experiments [e.g., see Hilborn (1985)]. In this scenario, the average duration of the chaotic transients depends on the control parameter(s) and can be made *arbitrarily large.* Furthermore, when the scenario is played in reverse, these chaotic transients foreshadow the sudden appearance of a chaotic attractor.[7]

There are two main economic implications. First, chaotic dynamics in (μ_t, σ_t) mean that the deviations from normality of the true CDF for returns will appear *unsystematic* from the agents' perspective. Second, we find that these deviations are small in the L^1 norm along the TE dynamics. Hence, agents' observations provide no reason for them to be dissatisfied with the MVA decision rule, which is simpler and easier to implement than state-preference or rational expectations decision rules. We also point out that many of the techniques which are being used to look for chaotic behavior in financial time series (e.g., attractor reconstruction methods, dimension calculations) can only detect chaotic *attractors* and not chaotic *transients.*

The paper should also be of interest to macroeconomists and economic theorists interested in the literature on rational expectations and learning dynamics briefly surveyed in Grandmont (1998). We emphasize that in this paper we do not consider learning because MVA itself does not (it is a static model) and we wish to evaluate MVA on its own terms. But from the perspective of that literature, the TE dynamics in this paper can be viewed as recasting MVA as the sort of "learning equilibrium" which is the focus of much of that literature: a non-rational expectations equilibrium in which agents make small, apparently unsystematic mistakes.

We should emphasize what the paper is *not* about. One can choose the CDF for endowments to make the true CDF for returns normal but the dynamics are trivial in this case and it would be factually incorrect. It might be possible to finesse the beliefs to match the true, non-normal CDF for returns but that would not be MVA. Our objective is not

[7] For textbook discussions of the boundary crisis scenario, see Hilborn (2000, Chapter 7), Medio (1992, Chapter 9), and Peitgen, Jurgens, and Saupe (1992, p. 646-650).

to "improve" MVA by suggesting more accurate beliefs or introducing learning. Instead, our objective is to study the consequences of using MVA in a dynamic world where the true distribution of returns is non-normal.

The plan for the rest of the paper is as follows. In section 2 we present the model and results. Section 3 concludes.

2. The Model and Results

We consider the standard pure exchange overlapping generations model with one consumption good and one asset, money. At each date $t \geq 1$ a single agent is born and lives for two periods. At $t = 1$ there is one old agent endowed with one unit of fiat money. We focus on the young agent at t and let c_t denote her consumption at t and c_{t+1} her consumption at $t+1$. All agents are *ex ante* identical with von Neumann-Morgenstern utility function

$$U(c_t, c_{t+1}) = \sqrt{c_t} + \beta\sqrt{c_{t+1}} \tag{1}$$

where $0 < \beta \leq 1$.

At each date t, the young agent is endowed with ω_t units of the good, where ω_t is an i.i.d. random variable with beta probability density function (PDF)

$$g(\omega_t|a,b) = \frac{1}{B(a,b)}\omega_t^{a-1}(1-\omega_t)^{b-1} \tag{2}$$

on the support $I = [1,5]$.[8] The old agent is endowed with none. We chose the beta PDF because it is computationally convenient and because it is very flexible and can assume a wide variety of shapes. To keep notation at a minimum, throughout the paper we only consider the case $a = b$.[9] When $a = b$, the beta PDF is symmetric about its mean and increasing a pinches the PDF in towards the mean, with less mass in the tails.

Since ω_{t+1} is random, the price p_{t+1} of the good at $t+1$ is also random and so is the real return on money $r_{t+1} = 1/p_{t+1}$. To decide on her money holdings M_t at t, the young agent must form a belief about the CDF for r_{t+1}. In accordance with MVA, and the fact

[8] The choice $I = [1,5]$ is arbitrary. A longer I, such as $I = [1,10]$, leads to larger values for μ_t and σ_t and more prominent fluctuations in them but also increases the running time of simulations.

[9] Simulations involving $a \neq b$ produced nothing new.

that r_{t+1} is nonnegative, agents assume that its CDF is given by the truncated normal CDF on $[0, \infty)$ with PDF

$$\psi(r_{t+1}|\mu_t, \sigma_t) = \frac{1}{\sigma_t\sqrt{2\pi}K(\mu_t, \sigma_t)} \exp\left[-\frac{(r_{t+1} - \mu_t)^2}{2\sigma_t^2}\right] \tag{3}$$

where

$$K(\mu_t, \sigma_t) = \int_0^\infty \frac{1}{\sigma_t\sqrt{2\pi}} \exp\left[-\frac{(r_{t+1} - \mu_t)^2}{2\sigma_t^2}\right] dr_{t+1} \tag{4}$$

is the normalization factor. Let $m(\mu_t, \sigma_t)$ denote the mean and $s(\mu_t, \sigma_t)$ the standard deviation of (3), which are in general *not* equal to μ_t and σ_t, respectively, but are instead somewhat involved functions of those parameters.[10] Also in accordance with MVA, we assume that agents know μ_t and σ_t which implies that agents know $m(\mu_t, \sigma_t)$ and $s(\mu_t, \sigma_t)$.

The maximization problem of the young agent at t is therefore

$$\max \sqrt{c_t} + \beta \int_0^\infty \sqrt{c_{t+1}}\, d\psi(r_{t+1}|\mu_t, \sigma_t) \tag{5}$$

subject to

$$p_t c_t + M_t \leq p_t \omega_t \tag{6}$$

$$p_{t+1} c_{t+1} \leq M_t \tag{7}$$

$$c_t, c_{t+1}, M_t \geq 0. \tag{8}$$

This simple problem has a unique interior solution where the constraints (6) and (7) hold as equalities. We can therefore re-write the problem as

$$\max \sqrt{\omega_t - r_t M_t} + \beta\sqrt{M_t}\Psi(\mu_t, \sigma_t) \tag{9}$$

where

$$\Psi(\mu_t, \sigma_t) = \int_0^\infty \sqrt{r_{t+1}}\, d\psi(r_{t+1}|\mu_t, \sigma_t). \tag{10}$$

The demand for money is

$$M_t = M(r_t, \omega_t, \mu_t, \sigma_t) = \frac{\omega_t(\beta\Psi(\mu_t, \sigma_t))^2}{r_t(r_t + (\beta\Psi(\mu_t, \sigma_t))^2)}. \tag{11}$$

[10] Nevertheless, for ease of reference we will often call μ_t and σ_t the "mean" and "standard deviation", respectively. Note that the map $(\mu_t, \sigma_t) \mapsto (m(\mu_t, \sigma_t), s(\mu_t, \sigma_t))$ is one-to-one because it is impossible for two distinct normal PDFs to become the same truncated normal PDF after truncation at zero.

In TE, the date t markets for consumption and money clear. By Walras' Law, we need only consider equilibrium in the money market. Setting supply equal to demand (the stock of money is constant at one unit), we obtain the TE real return on money

$$r_t = r(\omega_t|\mu_t,\sigma_t) = (1/2)\beta\Psi(\mu_t,\sigma_t)\left[\sqrt{4\omega_t + (\beta\Psi(\mu_t,\sigma_t))^2} - \beta\Psi(\mu_t,\sigma_t)\right]. \tag{12}$$

Equation (12) is invertible

$$\omega_t = r^{-1}(r_t|\mu_t,\sigma_t) = \left(\frac{r_t}{\beta\Psi(\mu_t,\sigma_t)}\right)^2 + r_t. \tag{13}$$

Hence, the true CDF of returns is endogenous and is given by

$$F_T(r_t|\mu_t,\sigma_t,a) = G(r^{-1}(r_t|\mu_t,\sigma_t)|a) \tag{14}$$

where G is the CDF for the beta distribution.[11] The mean of F_T at $t+1$ is given by

$$\Phi_\mu(\mu_{t+1},\sigma_{t+1}|a) = \int_I r(\omega_{t+1}|\mu_{t+1},\sigma_{t+1})\, g(\omega_{t+1}|a)\, d\omega_{t+1} \tag{15}$$

with standard deviation

$$\Phi_\sigma(\mu_{t+1},\sigma_{t+1}|a) \tag{16}$$

defined similarly.

We now formally define the natural equilibrium concept in this context.

Definition. *A* **TE** *at date $t \geq 1$ is a price p_t for the consumption good and conjectures (μ_t,σ_t) such that the date t markets for consumption and money clear and the conjectures (μ_t,σ_t) are correct*

$$\begin{aligned} m(\mu_t,\sigma_t) &= \Phi_\mu(\mu_{t+1},\sigma_{t+1}|a) \\ s(\mu_t,\sigma_t) &= \Phi_\sigma(\mu_{t+1},\sigma_{t+1}|a). \end{aligned} \tag{17}$$

In accordance with MVA, agents assume that the CDF of returns is normal and equations (17) require that their conjectures about (μ_t,σ_t) be correct. The expressions on

[11] To make F_T normal, set $G = F_N \circ r$, where F_N is the truncated normal CDF and r is defined in (12). The resulting model is trivial and factually incorrect.

the left-hand side of (17) are the mean and standard deviation of the truncated normal CDF while those on the right-hand side are the mean and standard deviation of the true CDF of returns F_T in (14).[12] Although the correct conjectures condition (17) ensures that these two CDFs have the same mean and standard deviation, in general F_T is non-normal.

Equations (17) govern the TE dynamics of the model. Note that this system is deterministic and bound to be highly nonlinear. Furthermore, the system only determines the dynamics for (μ_t, σ_t) *implicitly.* If we substitute an initial condition into the functions on the left-hand side of (17) and we are then able to solve for $(\mu_{t+1}, \sigma_{t+1})$ *uniquely* on the right-hand side, then this determines a *forward dynamic* with which we can compute *forward orbits* for various initial conditions. Due to the complexity of the functions Φ_μ, Φ_σ we have not been able to compute forward orbits nor can we ascertain whether or not a forward dynamic exists; i.e., whether the map $(\mu_t, \sigma_t) \mapsto (\mu_{t+1}, \sigma_{t+1})$ implicitly defined by (17) is actually a function.[13] On the other hand, if we substitute $(\mu_{t+1}, \sigma_{t+1})$ into the functions on the right-hand side and we are then able to solve for (μ_t, σ_t) uniquely on the left-hand side, then this determines a *backward dynamic* Σ

$$(\mu_{t+1}, \sigma_{t+1}) \mapsto_\Sigma (\mu_t, \sigma_t) \tag{18}$$

with which we can compute *backward orbits.*[14] This we were able to do. A system such as the Hénon map which has both forward and backward dynamics is called *invertible.* This is a useful property because one can use the forward dynamic to map out attractors and the backward dynamic to detect repellors, which can have an important influence on the dynamics. Obviously, any backward orbit read forward will satisfy (17) and in that sense represents the forward evolution of the system. For lack of alternatives, we will often

[12] The functions $m(\mu_t, \sigma_t)$ and $s(\mu_t, \sigma_t)$ have closed-form expressions. We do not have closed-form expressions for Φ_μ and Φ_σ but we can compute their values precisely given numerical values for their arguments since the Integrate routine in Mathematica is able to calculate the relevant indefinite integrals in (15) and (16) symbolically. Hence, working with system (17) does not involve any extraordinary round-off or other machine errors. In contrast, the NIntegrate routine introduces an artificial smoothing of the integrands and produces erroneous results.

[13] We tried several different combinations of utility functions and CDFs for endowments in an attempt to get a forward dynamic but were unsuccessful because the $t + 1$ variables are deeply embedded in the complicated integrals in (15) and (16).

[14] Recall that in Grandmont (1985) the dynamics are explicitly backward dynamics but he uses these to construct what he calls the "true" dynamics. In this paper, we only analyze the explicit backward dynamics.

use the terms "forward orbit" and "forward dynamics" in this sense, but formally this is incorrect since it is an open question as to whether a forward dynamic actually exists. For a discussion of these issues in the context of Grandmont (1985), see Medio (1992, p. 223).

We now come to a trivial but important point. From (17), it is clear that if the true CDF for returns were truncated normal then the dynamics would be trivial since in that case $\Phi_\mu = m$ and $\Phi_\sigma = s$. So the *positive* observation that MVA assumes normality when in fact real-world returns are non-normal is a crucial one for this model. In particular, any deviation from normality, no matter how small, produces a qualitative change in the dynamics.

We now turn to the simulations since we cannot study the system (17) analytically.[15] We set $\beta = 1$ for all simulations. Figure 1 below depicts a typical backward orbit for the case $a = 8$.

Figure 1 Goes Here

The initial conditions are $\mu_0 = 0.8$ and $\sigma_0 = 0.15$. In the top panel, we plot μ_t and in the bottom panel σ_t, both against time t. Although the orbits were computed backward, we depict them as moving forward. Hence, the fixed point in Figure 1 is unstable.

In fact, it is an unstable *node*. This can be seen by applying the backward dynamic Σ to the set of initial conditions enclosed by the rectangle in Figure 2 below.

Figure 2 Goes Here

The image of the rectangle under Σ is the one-dimensional curve within the rectangle. I.e., for any initial condition in the rectangle, the first iteration in the backward dynamic involves a large "jump" to this curve. The fact that the curve lies within the rectangle suggests convergence in the backward dynamic and, indeed, further iterations compress the curve lengthwise. Hence, in a neighborhood of the fixed point, all orbits approach it from a certain direction [c.f. Figure 2.12a in Medio (1992, p. 52)] which indicates that

[15] The simulations were performed using Mathematica Version 4 on an Apple Cube G4 500 MHz computer. In our case, Mathematica is a powerful tool since we have relied on it to calculate several complicated closed-form expressions symbolically.

the fixed point is indeed a node. Since these are mean-variance dynamics, backward orbits must approach the node from the southwest and northeast since an increase (decrease) in μ_t must be accompanied by an increase (decrease) in σ_t if the agent is to be content holding the one unit of money.

In Figure 3, $a = 9$ and the initial conditions are $\mu_0 = 1.5$ and $\sigma_0 = 0.14$.[16]

Figure 3 Goes Here

We now include the phase (or state) space in the top panel, where the node is indicated by the larger dot singled out by the arrow. Once again, we observe divergence from the unstable node but now the orbit fluctuates irregularly and briefly (for about 20 periods) in a neighborhood of the node before diverging off.

In Figure 4, $a = 10$ and the initial conditions are $\mu_0 = 1.5$ and $\sigma_0 = 0.15$.

Figure 4 Goes Here

The orbit fluctuates in "chaotic" fashion for about 450 periods before diverging off. Note that the unstable node is now close to the southeastern edge of the irregular activity.

In Figure 5, $a = 11$ and the initial conditions are again $\mu_0 = 1.5$ and $\sigma_0 = 0.15$.

Figure 5 Goes Here

Now the irregular fluctuations persist in the backward dynamics for close to 1,000 periods without reaching the node. The pictures for $a = 12$ with the same initial conditions are similar.[17] For $a = 13$, the backward dynamics become explosive and we could not find any backward orbits of significant length. The problem does not appear to be computational, but rather the non-existence of such orbits.

To determine if the orbits in Figures 4 and 5 are indeed chaotic, we compute the Lyapunov exponents.[18] Note that since these orbits are not on an attractor, the appropriate

[16] We cannot consider values $8 < a < 9$ since Mathematica can only compute orbits for integer values of a.

[17] For $a = 11, 12$ we did not compute orbits much longer than 1,000 periods because the computations were so slow due to the complexity of the system (17) (e.g., it took about two days to generate Figure 5). We were thus unable to locate the node when $a = 11, 12$.

[18] For discussions on Lyapunov exponents, see Hilborn (2000), Medio (1992), and Peitgen, Jurgens, and Saupe (1992).

definition of "chaos" reduces to sensitive dependence on initial conditions. There are several procedures available for computing Lyapunov exponents, but many of these require a closed-form expression for the Jacobian of the system. In our case, this is unavailable and we must treat the computed orbits as two-dimensional experimental time series. The Lyapunov exponents were computed using a program based on the procedure in Eckmann, Kamphorst, Ruelle, and Ciliberto (1986) which uses the available time series to essentially estimate the Jacobian of the system using ordinary least-squares.[19] This procedure seems to be the "state of the art" and is believed to compute both exponents (not just the dominant one) reliably. The results are reported in Table 1 below.

Table 1 Goes Here

The dominant exponent λ_1 is positive in all cases, indicating expansion in at least one direction along the orbit and therefore chaos. For $a = 10$, λ_2 is negative, indicating contraction in one direction, which is picking up the fact that the backward orbit in Figure 4 is convergent. When $a = 11, 12$, λ_2 is also positive, indicating expansion in both directions and no sign of convergence for these orbits.

As a check, in the top two panels of Figure 6 we depict the power spectra (actually the periodograms) for μ_t and σ_t for the orbit with initial conditions $\mu_0 = 1.5$ and $\sigma_0 = 0.15$ when $a = 12$.[20]

Figure 6 Goes Here

For purposes of comparison, in the bottom panel we depict the spectrum for an orbit generated by the random number generator on our computer which was calibrated to produce "random" values similar to those of σ_t. Chaotic data typically have spectra with a small number of sharp peaks with a broad-band "noise floor" in between. This is the case for all of the spectra in Figure 6. Indeed, the orbit for μ_t appears "noisier" than the orbit generated by the random number generator.

---

[19] The Mathematica program is available from the authors upon request.

[20] To remove the spike at zero and improve the view at higher frequencies, both orbits have been normalized by subtracting their means. For a nice discussion on power spectra and their uses, see Medio (1992, Chapter 5).

The orbits in Figures 4 and 5 therefore appear to be chaotic. There are several known "routes to chaos" including the well-known period-doubling scenario. One of these, the *boundary crisis* scenario (also known as a "blue sky catastrophe"), was developed by Grebogi, Ott, and Yorke (1982, 1983) and has been observed in the logistic and Hénon maps as well as in physical laboratory experiments [e.g., Hilborn (1985)]. This scenario involves an unstable fixed point or periodic cycle and a chaotic attractor. As a parameter a of the system is reduced, the two move closer together and eventually collide when $a = a_c$.[21] This collision destroys the attractor and it suddenly (discontinuously) disappears. For values of a below a_c, the attractor is gone but nearby orbits nevertheless rapidly move towards where it used to be and remain there, mapping out the old attractor, for finitely many periods before hitting an "escape region" and diverging off. These are called *chaotic transients* because the behavior of the orbits on the "ghost" of the old attractor appear chaotic. If we reverse the process, increasing a towards a_c from below, the duration of these chaotic transients increases without bound; i.e., their duration can be made arbitrarily long by choosing a close enough to a_c. Furthermore, the increase in the duration of the chaotic transients foreshadows the sudden appearance of the attractor. This is precisely what we seem to be observing in Figures 1 and 3-5. In particular, for $a = 11$ both Lyapunov exponents are positive and there is no sign of convergence after 1,000 periods.

Unlike Grebogi, Ott, and Yorke (1983), who were working with the invertible Hénon map, we are unable to present pictures of the unstable node colliding or about to collide with a chaotic attractor. First of all, since we only have the backward dynamic, we cannot observe any attractors (even presuming that a forward dynamic exists). Even with just the backward dynamic, we might still hope to obtain a picture of the node just after the collision, burrowing its way into the ghost of the old attractor. E.g., in the case of the Hénon map, for values of the control parameter just after the boundary crisis value (and holding the value of the other parameter fixed), the chaotic transients continue to map out the old attractor with much of the fractal structure still evident. In Figure 7, we track the movement of the unstable node as a is reduced from 10 to 5.

[21] Compared with the usual textbook discussion, we are reversing the order of the movement in a to conform with our system.

Figure 7 Goes Here

The node moves in the northwesterly direction as indicated by the arrow. Returning to Figure 4, we observe that for $a = 10$ the node appears to have made substantial progress inside the southeastern edge of what may be the ghost of the old attractor, although no fractal features are present. In Figure 3, for $a = 9$ the node appears to have moved to a more central position with respect to the irregular fluctuations but it is difficult to be certain since very little of the old attractor remains. Recall that we were unable to locate the node for $a = 11, 12$ so we are unable to observe the scene closer to the supposed collision. Note that for $a = 11, 12$ the ghost of the old attractor is still a formless mass, which may indicate that $a = 12$ is not yet close to the boundary crisis value. Although we have not been able to present conclusive evidence for a collision, Figures 3 and 4 do show that the chaotic transients occur in a neighborhood of the unstable node.

For the Hénon map, Grebogi, Ott, and Yorke (1983) show that the boundary crisis is accompanied by the appearance of horseshoe-like objects. If these were present for our system, repeated iteration of the backward dynamic Σ starting from a rectangle of initial conditions as in Figure 2 should produce an increasing number of ever-thinning horizontal strips.[22] The fact that we have not observed this should not be too surprising since the boundary crisis for the Hénon map involves a saddle fixed point rather than an unstable node and the horseshoe-like objects are caused by the development of complex relationships between the stable and unstable manifolds of the saddle.

We conclude that the above evidence is consistent with the boundary crisis scenario and none other. We observe chaotic transients always in the immediate vicinity of the unstable node. These chaotic transients are always followed by divergence and not a return to regular behavior (as would be the case for *intermittency*). Finally, the duration of the chaotic transients seems to increase exponentially with a.

We now turn to the economic implications of these results. First, from (12), we see that $r(\omega_t|\mu_t, \sigma_t)$ has a deterministic component (μ_t, σ_t) as well as a purely stochastic component ω_t. If the deterministic component followed some simple orbit then agents

[22] E.g., see the discussion in Hilborn (2000, p. 199).

would make systematic mistakes. E.g., for a fixed point the true CDF for returns would be constant and non-normal all along the orbit. In contrast, when the orbit is chaotic these mistakes would appear *unsystematic* from the agents' perspective.

We are also interested in the size of the mistakes that agents are making. Let $F_N(r|\mu_t, \sigma_t)$ denote the normal CDF on $[0, \infty)$. The average error on $[0, R]$ is

$$(1/R) \int_0^R |F_N(r|\mu_t, \sigma_t) - F_T(r|\mu_t, \sigma_t, a)| \, dr \tag{19}$$

which is essentially the L^1 norm. In Table 2 below we give the average error for selected points along the orbit for $a = 12$ with initial conditions $\mu_0 = 1.5$ and $\sigma_0 = 0.15$.[23]

Table 2 Goes Here

The largest average error in the table is approximately 0.004. To interpret this, randomly select a point x in the support of both CDFs. Then on average, the difference in the probability assigned to the event $[0, x]$ by the two CDFs will be less than half a percent. So agents in this model make very small mistakes in the TE dynamics.

Finally, when $a = 13$ the backward dynamics become explosive and we are unable to compute backward orbits of any significant length. This raises the tantalizing possibility that the chaotic attractor (assuming the existence of the forward dynamic) has reappeared at $a = 13$. If so, then nearby initial conditions would move towards the attractor and if convergence were reasonably strong, lengthy backward orbits would not exist.

3. Conclusions

Despite its status as one of the principal portfolio choice models in finance, MVA seems to suffer from two major theoretical problems. First, it is a static model in a dynamic world. Second, it crucially assumes that asset returns are normally distributed

[23] We did not compute the average error all along the orbit for the following technical reasons. The procedure was as follows. For a selected period t, we first plotted the integrand in (19). Since the CDFs typically cross two or more times, this introduces singularities into the integrand which need to be taken into account in order to accurately compute the integral in (19). In each case, R was chosen as the point where the normal CDF became 1. Although the normal PDF is positive on the entire real line, it asymptotes to the horizontal axis exponentially so at some point R the normal CDF becomes 1 in machine terms.

when, in fact, real-world returns are non-normal. It would therefore seem that MVA leads to potentially large, systematic mistakes.

In this paper, we embedded MVA in the standard OLG model with one consumption good, stochastic i.i.d. endowments, and one asset, money. In accordance with MVA, agents assume that the distribution for the real return on money is (truncated) normal and make conjectures about its mean and standard deviation. In temporary equilibrium, the current markets for consumption and money clear and these conjectures are correct. Although agents make correct conjectures about its mean and standard deviation, the true distribution of returns is non-normal. It is precisely this deviation from normality which opens the door to potentially interesting dynamics. The model leads to a deterministic system of nonlinear difference equations in the mean and standard deviation. Simulations revealed the presence of chaotic transients seemingly produced by a boundary crisis. As a result, deviations from normality of the true distribution of returns will appear to be unsystematic. We also found that they are typically small. Hence, agents' observations provide no reason for them to be dissatisfied with MVA decision rules, which are simpler and easier to implement than state-preference or rational expectations decision rules.

References

Copeland, T.E. and J.F. Weston (1988): *Financial Theory and Corporate Policy* (3rd Edition). New York: Addison-Wesley.

Eckmann, J.P., S.O. Kamphorst, D. Ruelle, and S. Ciliberto (1986): "Liapunov Exponents from Time Series," *Physical Review A,* 34, 4971-4979.

Grandmont, J.M. (1977): "Temporary general equilibrium theory," *Econometrica,* 45, 535-572.

Grandmont, J.M. (1985): "On Endogenous Competitive Business Cycles," *Econometrica,* 53, 995-1045.

Grandmont, J.M. (1998): "Expectations Formation and Stability of Large Socioeconomic Systems," *Econometrica,* 66, 741-781.

Grebogi, C, E. Ott, and J.A. Yorke (1982): "Chaotic Attractors in Crisis," *Physical Review Letters,* 48, 1507-1510.

Grebogi, C, E. Ott, and J.A. Yorke (1983): "Crises, Sudden Changes in Chaotic Attractors, and Transient Chaos," *Physica,* 7D, 181-200.

Hilborn, R.C. (2000): *Chaos and Nonlinear Dynamics: An Introduction for Scientists and Engineers* (2nd Edition). Oxford: Oxford University Press.

Medio, A. (1992): *Chaotic Dynamics: Theory and Applications to Economics.* Cambridge, UK: Cambridge University Press.

Peitgen, H.O., H. Jurgens, and D. Saupe (1992): *Chaos and Fractals: New Frontiers of Science* New York: Springer-Verlag.