CCHE CCES

Canadian Centre for Health Economics
Centre canadien en économie de la santé

Working Paper Series
Document de travail de la série

IS THE RATIONAL ADDICTION MODEL INHERENTLY IMPOSSIBLE TO ESTIMATE?

Audrey Laporte, Adrian Rohit Dass, Brian Ferguson

Working Paper No: 150011

www.canadiancentreforhealtheconomics.ca

June, 2016

Canadian Centre for Health Economics
Centre canadien en économie de la santé
155 College Street
Toronto, Ontario

CCHE/CCES Working Paper No.150011
June, 2016

Is The Rational Addiction Model Inherently Impossible To Estimate?

Audrey Laporte[a,b,*], Adrian Rohit Dass[a,b], Brian Ferguson[b,c]

[a] *Institute of Health Policy, Management and Evaluation, University of Toronto, Toronto, Canada*
[b] *Canadian Centre for Health Economics, Toronto, Canada*
[c] *Department of Economics and Finance, University of Guelph, Guelph, Canada*

JEL Classification: I12; C22; C23

Keywords: Rational Addiction Model; Dynamic Time Series; Dynamic Panel

*Correspondence to: Institute of Health Policy, Management and Evaluation, University of Toronto, Health Sciences Building, 155 College Street Suite 425, Toronto, Ontario, M5T 3M6, Canada. E-mail: audrey.laporte@mail.utoronto.ca

I Introduction

It is probably fair to say that the Becker-Murphy model of Rational Addiction (RA) is one of the most influential, and one of the most commonly empirically implemented, frameworks in health economics. It has been used to estimate demand functions for alcohol, tobacco and addictive drugs, but also for the demand for coffee (Olekalns and Bardsley (1996)), movies (Sisto and Zanola (2010)), credit card debt (Shen and Giles (2006)) and attendance at National Football League games (Spenner, Fenn and Crooker (2010)). Despite its popularity, though, it is probably also fair to say that there is no consensus about its empirical validity. Melberg (2008) summarizes the results of a survey of academics who have written about RA by saying:

> A majority of the respondents believe the literature is a success story that demonstrates the power of economic reasoning. At the same time they also believe the empirical evidence is weak, and they disagree both on the type of evidence that would validate the theory and the policy implications. Taken together, this points to an interesting gap. On the one hand most of the respondents claim that the theory has valuable real world implications. On the other hand they do not believe the theory has received empirical support.

The RA model is usually estimated using, in part, an expression in the dependent variable along the lines of

$$Y_t = \alpha_0 + \alpha_L Y_{t-1} + \alpha_F Y_{t+1} \tag{1}$$

where Y_t is current consumption of the addictive commodity in question, Y_{t-1} is lagged, or past, consumption, with coefficient α_L and Y_{t+1} is leading, or future consumption, with coefficient α_F. Obviously, empirical implementation always includes a range of strictly exogenous variables, usually current and possibly past and future prices, but the primary focus of empirical implementation is the segment in equation (1) above, and it is the estimation of this expression with which we shall be concerned here.

The most commonly tested (but, as we shall see, not the only) empirical prediction of the RA model is that the coefficients on lead and lag consumption should be positive. This prediction is typically satisfied, but there seems to be no consensus in the literature as to the magnitude of the coefficients and whether the model yields testable predictions with regards to the absolute magnitude[1]. Further, it is not clear that positive and significant lead and lag coefficients should by themselves be taken as evidence of RA behavior - see the results of Auld and Grootendorst (2004) for example[2].

There is a second testable RA hypothesis that can be tested using equation (1), and that is

[1] For a review of panel data estimates of the RA model and a discussion of the wide range of results present in the literature see Baltagi (2007) and also Baltagi, B.H. and Geishecker, I. (2006).

[2] Auld and Grootendorst's estimation uses aggregate level data. Because the issues which arise when aggregate data are being used are quite different from those which arise with micro level data, with which we are concerned here, we discuss the Auld and Grootendorst (2004) results in the appendix.

$$\alpha_F = \beta\alpha_L \tag{2}$$

where $\beta = 1/(1+\rho)$ is the individual's discount factor, ρ being the individual's discount rate. This prediction (which strictly speaking is exact only under certain assumptions about functional forms) is tested reasonably often in the RA literature and the results have been described by Baltagi (2007), as "the fly in the ointment" of the RA model[3]. It is not uncommon for estimation at the individual level to yield wildly varying, and often implausible, estimates of β.

One obvious possible explanation for these results is that estimation of the RA model at the individual level involves the problems typically tackled in the Dynamic Panel Data (DPD) literature (see, for example, Arellano (2003)). In an earlier paper, however, we suggested that there might be a problem estimating RA-type equations even in the absence of DPD-type problems and even (and perhaps especially) when the theoretical model was in fact correct (Laporte et al. (2016)).

Our argument in that paper rested on the fact that the RA model is a particular case of an inter-temporal optimization problem (as indeed it was set out by Becker and Murphy (1988))[4] and that certain properties of the solution to a theoretically correct inter-temporal optimization problem might raise particular econometric problems in estimating equations like (1).

In that paper we focused on the general issue of estimating equations of the form of equation (1) above, while explicitly setting aside equation (2). In this paper we consider the implications of equation (2) in conjunction with the general econometric issues that arise from the nature of the solution to an inter-temporal optimization problem.

In the following section we set out the theoretical argument that places the RA model in the optimal control framework, and which will inform the interpretation of our simulation results. Section III considers the implications of hypotheses which are specific to the RA model, as set out in equation (2) above and also as set out in what Becker, Grossman and Murphy (1994) referred to as a stability condition (although it is perhaps better referred to as a uniqueness condition). Section IV uses Monte Carlo simulation to investigate the issue with which we are dealing here, first tackling the simulations in pure time series form and then, in Section V, in panel data form. The implications of the Monte Carlo experiments are discussed in Section VI.

[3] Baltagi and Grffin (2001) say: "In sum, we are optimistic that the rational-addiction model represents a significant improvement over models of myopic behavior. However, before it can be widely accepted, plausible and statistically significant estimates of the implied discount rate are needed. Based on BGM and our results, aggregate panel data do not seem likely to provide sharp estimates of the discount rate. The most promising approach appears to be microdata as discussed by Chaloupka (1991) but one hopes with samples much longer than T = 3."

[4] It is worth noting that, while Becker and Murphy (1988) and Becker Grossman and Murphy (1991) set the problem up in the form that we use below, in later papers, including Becker, Grossman and Murphy (1994) they put lagged consumption directly in the individual's utility function. This difference in formulations may help explain differences in the explanatory variables included in estimated equations in the literature, in particular whether lead and lagged prices should be included.

II A Brief Outline of the RA Model as an Optimal Control Problem

In setting the RA model up as an optimal control problem[5] we begin by defining the individual's lifetime utility function:

$$V = \int_0^T U(C(t), Y(t), A(t))e^{-\rho t} dt \tag{3}$$

where C is consumption of other, non-addictive commodities, Y is consumption of the addictive commodity and A is accumulated addiction capital which, as we noted above, is open to a number of interpretations, all of which are consistent with the notion that A yields disutility rather than utility, so the first and second derivatives of the instantaneous utility function, U, with respect to A are both negative. Both C and Y yield positive and diminishing marginal utility, and the individual discounts the future at the subjective rate ρ. The upper limit of integration, which represents the end of the individual's planning horizon - in most health economic applications, this refers to the end of life - is written T and, in terms of the formal analytics of an optimal control problem, can be either finite or infinite. It is not unusual for theoreticians to assume for simplicity that $T = \infty$ (as with for example, the infinitely-lived representative individual often found in micro-founded intertemporal macroeconomic models). However, the analytical differences between the case where T is infinite and that where T is finite- the latter obviously being the sensible case to assume in health economics applications- turns out to be critical to the argument we make in this paper.

Addiction capital accumulates according to the equation of motion:

$$\dot{A} = g(Y) - \delta A \tag{4}$$

where $\dot{A}$ is standard notation for $\partial A/\partial t$, the time derivative of A, g(Y) is a damage production function with $g(0) = 0$, $g'(Y) > 0$ and, typically in this application $g''(Y) > 0$, and δ, is the rate of depreciation of addiction capital, representing the rate at which the body can heal itself if the individual goes cold turkey on consumption of Y. We must also introduce a budget constraint relating C, the consumption of the non-addictive commodity and Y, the consumption of the addictive commodity. We can either introduce a lifetime budget constraint, by adding accumulation of financial assets to the problem, or an instantaneous budget constraint, requiring all income to be spent on C and Y at each point in time. Since this issue is not key to the point discussed in this paper we do not go into detail about it: implicitly in what follows we assume an instantaneous budget constraint.

We do not derive the intermediate steps for the theoretical analysis here - for the details see Ferguson (2000). We note simply that the application of optimal control techniques yields a pair of differential

[5]We work with continuous time at this stage so that we can illustrate an individual's time path using a phase diagram - the discrete time formulation which is analytically identical but notationally messier, will be used when we discuss empirical issues.

equations, one in A and one in Y, which incorporate the first order conditions for inter-temporal utility maximization and which can be used to map out the individual's lifetime utility maximizing trajectories for consumption of T and for addiction capital. These solution equations are typically used to draw a phase diagram, as in Figure 1 below:

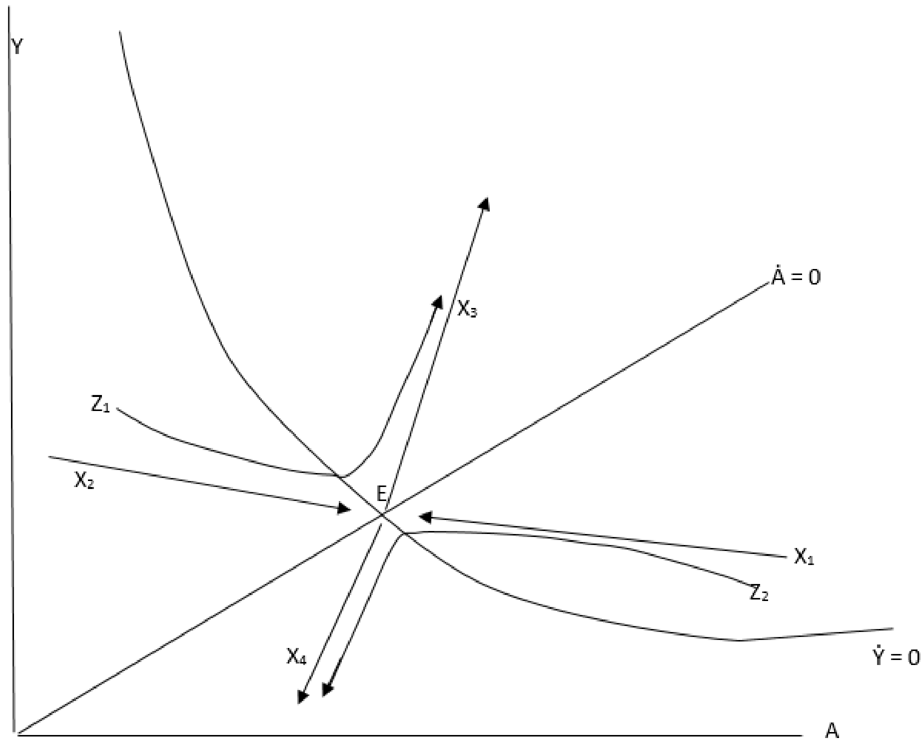

Figure 1: Individual's Phase Diagram for Rational Addiction Model

In Figure 1, Y, the consumption of the addictive commodity is plotted on the vertical axis and A, the accumulated addiction capital, is plotted on the horizontal. A phase diagram is a dynamic representation of the model, so we draw on it trajectories which show how Y and A evolve over time. The lines marked $\dot{A} = 0$ and $\dot{Y} = 0$ are called the stationary loci for the diagram. They are combinations of points along which the dynamics of the model mean that there is no intrinsic tendency for A or Y respectively to change. From equation (4), for example, $\dot{A} = 0$ when g(Y) - δ A $=$ 0. Thinking of A as a capital good, with δ as its depreciation rate, this says that there will be no intrinsic tendency for A to change if the amount of current Y yields a g(A) value just sufficient to balance the depreciation of the existing A. For different values of A we need different values of Y, and these pairings are marked out along the $\dot{A} = 0$ locus, whose shape will depend on the shape of the production function g(Y). We have drawn it as linear for expositional simplicity. The stationary locus for Y can be interpreted in a similar manner (see Ferguson (2000)). The intersection (E) of the two stationary loci is the equilibrium point for the system, the point at which neither Y nor A has any intrinsic tendency to change. The phase diagram shows the dynamics of the system when it is not at the equilibrium, and the stationary loci can be thought of as dividing the (Y, A) space into quadrants, in each of which Y and A can be shown to be either increasing or decreasing.

In Figure 1 we have drawn a number of trajectories. Two of them, referred to as the stable branches for the problem, and labeled X_1 and X_2, head directly towards the equilibrium and two more, the unstable branches, labeled X_3 and X_4, head directly away. In addition there are two other trajectories, labeled Z_1 and Z_2, each of which is initially heading towards the equilibrium, essentially tracking the nearest stable branch, but eventually curving away from it, tracking the nearest unstable branch and, it can be shown that in the long run it will asymptote on the relevant unstable branch. It is convenient to think of these Z trajectories as a form of weighted averages of the stable and unstable branches, along which initially the stable branch is the dominant influence but where eventually the unstable branch becomes the dominant influence. The optimizing individual must pick the best possible lifetime trajectory from the full set open to him (which is much larger than we have illustrated here) taking account of his initial value of A (usually zero in a RA application) and the length of his life. This type of phase diagram is referred to as having saddle-point dynamics and the equilibrium point as being a saddle-point equilibrium.

The issue of length of life is key to what follows. We noted above that it is not uncommon for theoretical exercises to assume an infinite horizon, for expositional simplicity. In empirical RA research, however, it clearly makes no sense to assume that our individual is fully rational in every way with the minor exception that she fully believes that she will live forever. If she is going to be a rational addict, she has to face up to the finiteness of her life.

In terms of the phase diagram, the difference between infinite and finite horizon is key to the choice of the optimal trajectory. For most economists, it is instinctive to assume that the optimal trajectory will be one that converges on the equilibrium - i.e. one of the stable branches. It is well established in the optimal control literature, however, that the stable branch is the optimal trajectory only for an infinite horizon problem. For a finite horizon problem the optimal trajectory will be one of the ones that does not converge, typically one such as Z_1 and Z_2 above, which seems at first to be convergent but which, after sufficient time has elapsed, diverges from the equilibrium. Thus if we were to plot the individual's optimal trajectory against time it would look like one of those in Figure 2 below:

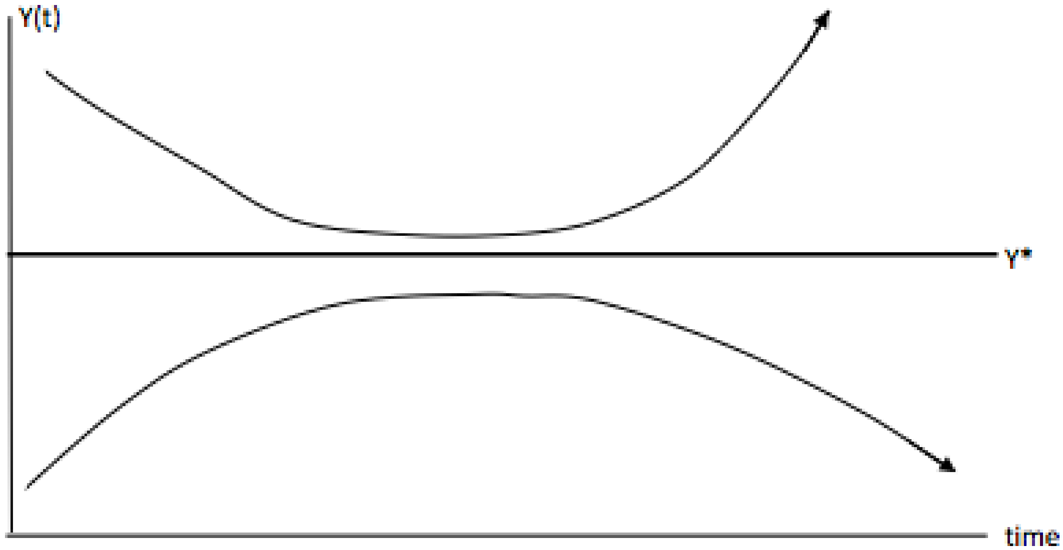

Figure 2: An Individual's time trajectories for Y

This highlights the fact that if we wish to test the RA model empirically we must use a functional form that has the potential to change direction as time passes.

A second empirical issue, common to many problems in health economics, is that while we can define A, addiction capital, we very seldom have actual measures of it which we can incorporate into empirical work. This is not such an insurmountable problem as it might at first seem.

It can be shown that, when we set the RA model up in discrete time terms, the solution, which yields the phase diagram, takes the general form

$$Y_t = \beta_{Y0} + \beta_{YY} Y_{t-1} + \beta_{YA} A_{t-1} \tag{5a}$$

$$A_t = \beta_{A0} + \beta_{AA} A_{t-1} + \beta_{AY} Y_{t-1} \tag{5b}$$

This is a pair of interrelated first order difference equations whose coefficients depend on the underlying parameters of the optimization problem, and it is the form we would like to estimate, if we could observe both A and Y. The fact that we cannot observe A is not fatal to our empirical work since it can be shown (see Ferguson and Lim (2003) and Jones et al. (2014)) that a pair of interrelated first order difference equations like (5a & b) can be combined into a single second order difference equation in either of the two variables: in our case, given the assumption that A is unobservable, this would be in Y.

Most commonly when this dynamic reduction is done we would write the second order difference equation (SODE) in backward looking form:

$$Y_t = \alpha_0 + \alpha_1 Y_{t-1} + \alpha_2 Y_{t-2} + \beta_X X_t + \epsilon_t \tag{6}$$

As we noted above, in the RA literature, because of interest in the forward-looking nature of the optimization problem, it is customary to write, and to estimate

$$Y_t = \alpha_0 + \alpha_L Y_{t-1} + \alpha_F Y_{t+1} + \beta_X X_t + \epsilon_t \tag{7}$$

which we shall refer to as a SODE in RA-form. We mentioned above that, for the most part, the focus of the empirical RA literature is on the issue of whether the coefficients on lead and lag consumption are positive. The theoretical optimal control framework, though, also raises the issue of how well the estimated equation tracks a nonlinear lifetime trajectory.

The non-linearity itself is not a fundamental problem, since a SODE, whether in backward looking or in RA-form, can map out a curved path. A key element in relating equation (7) to the form of trajectories set out in Figure 2 above is the notion of the solution form to a SODE.

It is shown in the theoretical dynamics literature that the solution to a SODE in either form yields an equation of the form:

$$Y(t) = A_1 \lambda_1^t + A_2 \lambda_2^t + Y^*(t) \tag{8a}$$

The terms λ_1 and λ_2 are what are referred to as the roots of the SODE, which are constants, and the t superscript is time, so the solution form (8) writes Y as a function of time. The roots of the system can be solved from either equation (6) or equation (7). Given the nonlinear form in which t enters the RHS of (8) it is clear that the values of the roots will be key to the shape of the time-trajectory of Y. In (8) Y*(t) is the equilibrium of the system, and corresponds to the value of Y at the intersection of the stationary loci in the phase diagram. Its position depends on the location of the stationary loci in the phase diagram, which in turn will depend on the exogenous variables in the problem. If there are exogenous variables, and if their value changes, the value of Y* will change and the trajectory generated by (8) will be adjusted accordingly. The terms A_1 and A_2 are constants whose values depend on the particular details of the model being analyzed. In equation (8) they serve as weights on the two roots.

Writing the solution to Y as a function of time as in (8) allows for a focus on the conditions for convergence to equilibrium. Since the As and the λs are constants, derived from the parameters of the particular problem being analyzed, the behavior of Y over time can be mapped out relative to Y^*. Suppose, for example, that both λ_1 and λ_2 are positive fractions. Then as time passes, regardless of the magnitude of the two A terms, eventually both λt terms will converge on zero and Y(t) will converge on Y^*. This is the case of a dynamically stable equilibrium. Similarly if both λ_1 and λ_2 are positive but larger than 1 (negative roots are very rare in economic applications, so we do not discuss them here), as time passes both λ^t terms will become extremely large and no matter how small the A terms might be, Y will tend to diverge from Y^*.

The saddle-point equilibrium as set out in Figure 1 above has the property that while both roots are positive, one of the roots is larger than one and the other smaller than one. The root that is larger than one (we shall assume that λ_1 is the larger root and λ_2 the smaller one) we refer to as the unstable root and the smaller one as the stable root. Clearly, as time passes and t increases, the unstable root, which is greater than 1, will tend to have more and more influence on the behavior of Y(t) while the stable root, which is a positive fraction, will have a diminishing influence. The terms A_1 and A_2 in effect pick which possible trajectory on the phase diagram the individual is on, and hence what the time-trajectory of Y (i.e. the shape of Y(t) plotted against t) looks like. For example, if the conditions of the problem are such that A1, the weight on the unstable root, is zero, expression (8) will collapse to the solution form for a stable first order difference equation, $Y(t) = A_2\lambda_2^t + Y^*(t)$. In this case, since λ_2 is a positive fraction, as t increases, λ_2^t will go to zero and the actual value of Y will converge, as time passes, on its equilibrium value. In terms of the phase diagram this is the case where the individual's optimal trajectory is the stable branch to the equilibrium. As we have already noted, the results of optimal control theory tell us that this will be the optimal solution only in the case of an infinite horizon problem. For a finite horizon problem, the optimal trajectory for an individual will be one along which both roots have non-zero weights. If A_1 is very small relative to A_2, it will be the case that initially the stable root will dominate the trajectory and the system will seem to be converging to the equilibrium, but eventually t will reach a value large enough that the unstable root will come to dominate and the time-trajectory of an individual's Y will start to swing away from the equilibrium. For t sufficiently large the influence of the stable root will be swamped by the influence of the unstable root and the individual's time trajectory of Y will be very close to that for an unstable FODE: $Y(t) = A_1\lambda_1^t + Y^*(t)$.

As noted above, (8) writes the solution to a SODE strictly as a function of time. The solution relation must hold at each value of t, so we have

$$Y(t-1) = A_1\lambda_1^{t-1} + A_2\lambda_2^{t-1} + Y^*(t-1) \tag{8b}$$

and

$$Y(t+1) = A_1\lambda_1^{t+1} + A_2\lambda_2^{t+1} + Y^*(t+1) \tag{8c}$$

Given that we can solve for the roots of the estimated RA-form SODE (7), the upshot of this is that, if the RA model is in fact the solution to an individual's inter-temporal optimization problem, when we estimate equation (7) and solve for its roots we should find one stable and one unstable root (i.e., because we are using the difference equation form for our empirical work, one root less than one and one root greater than one).

In Laporte et al. (2016) we presented Monte Carlo results that suggest that, in the case of a standard, backward-looking SODE such as (6) above, the fact of saddle-point dynamics might cause problems for estimation. Thus consider a SODE of the form:

$$Y_t = -1000 + 2.055Y_{t-1} - 1.05Y_{t-2} \tag{9}$$

(where for simplicity we have omitted exogenous variables since our focus is on the SODE itself). This equation has roots 1.10 and 0.952, so it satisfies the saddle-point condition. In Laporte et al. (2016) the Monte Carlo experiments were conducted with 500 replications on 500 observations, using recursive regression. Within each experiment the deterministic part of the SODE was held unchanged and a zero mean NID disturbance term was added. The observations were ordered so that a time series graph of the data would initially seem to be converging on the equilibrium but then, as the unstable root came to dominate, turn and diverge from it, as in the case of trajectory Z_1 in Figure 1 above. For this equation we report the means of the Monte Carlo estimates on the two lag coefficients in Figure 3 below, where the first 20 observations were reserved for initialization of the recursive estimation.

Figure 3: Means of Monte Carlo Coefficients on two lags of Y, case with λ_H = 1.1

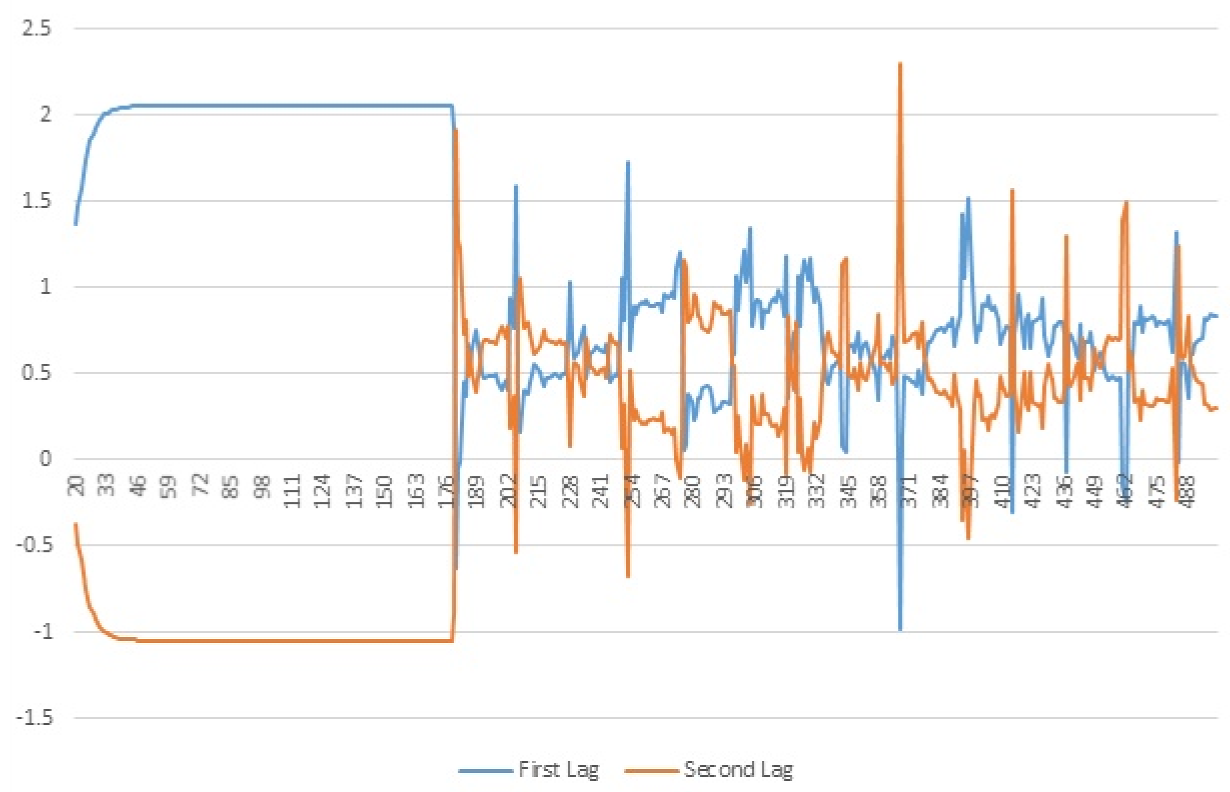

In Figure 3 we see that through the first part of the data set both the coefficients on the first and the second lag are well estimated, but after about 175 observations the estimated values suddenly become wildly variable. We can get a sense of what is going on if we consider the values of the roots of SODE (9) as derived from the estimated coefficients:

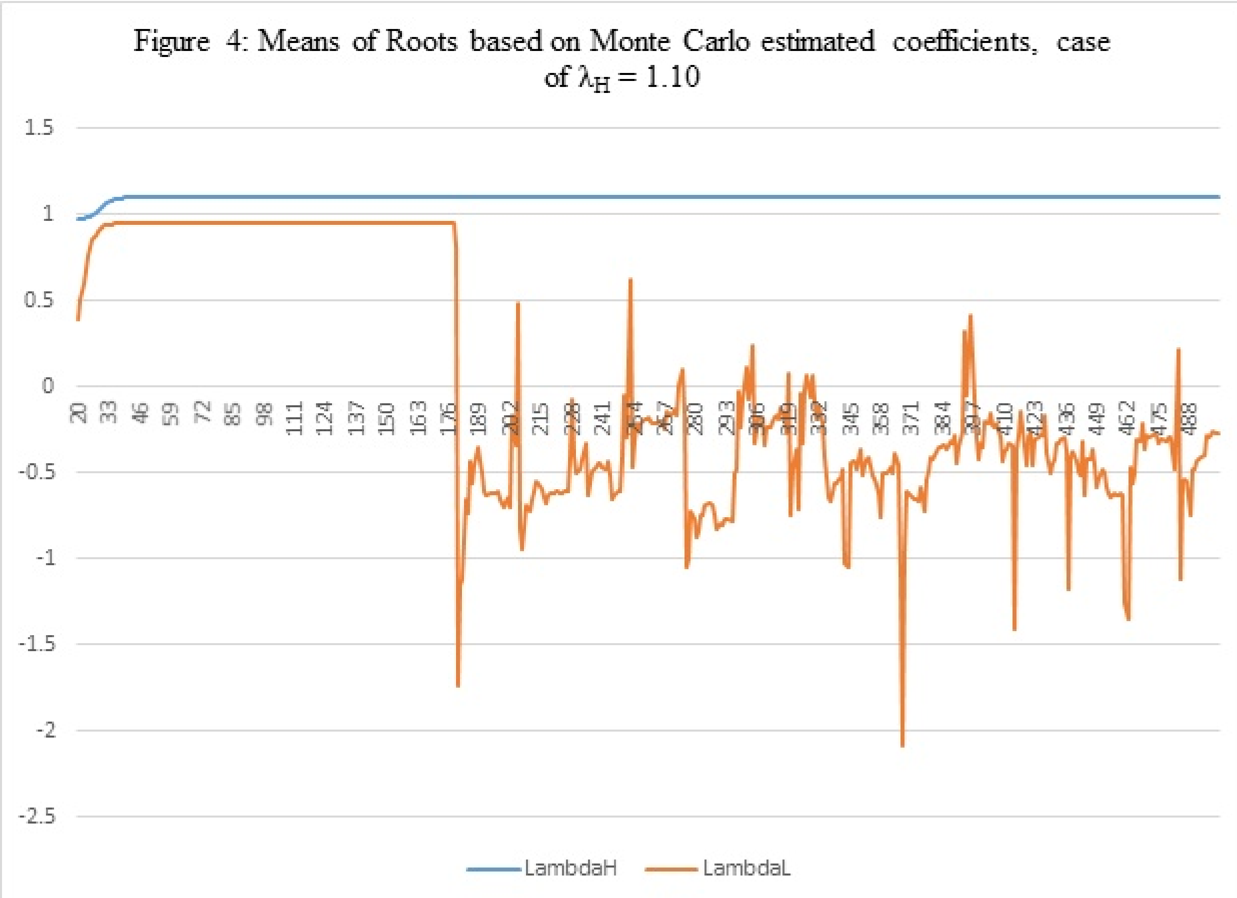

In Figure 4 we see that the unstable root is well estimated throughout our time series, including in the convergent region, where the stable root is dominant, but that the stable root is well estimated only in that first segment - once the time series reaches a certain length, with the critical value apparently tied to the unstable root's acquiring a minimum degree of dominance over the dynamics of the dependent variable, our calculated value for the stable root goes astray. In Laporte et al. (2016) we suggest that in the first part of the time series data both roots matter in the determination of the behavior of Y, and that that is sufficient to tie down the value of both αs, but that once the unstable root becomes dominant and the stable root has minimal impact on the behavior of the dependent variable, the system is essentially underdetermined, in the sense that any pair of α values will do, so long as they yield the unstable root, but that the stable root is not having enough of an effect on the time series behavior of Y to be able to tie down both coefficients.

Our interest in this paper is whether a similar effect could be at play in the estimation of RA models that, despite the forward/backward difference equation form typically estimated, should display saddle-point dynamics.

Issues arising from the presence of unstable roots have been discussed in the macro-econometrics literature: Nielsen and Reade (2007), for example, discuss how numerical instability arising from the presence of an unstable root affects the development of critical values for tests for the presence of a unit root in a variable which also has an unstable root, and Nielsen (2008) discusses how the presence of an explosive root affects testing models of Yugoslav hyperinflation. In the macroeconomic context, the presence of an explosive root is typically an aberration - a hyperinflation, or possibly a bubble in the price of an asset. This focus of the literature on macro-econometric issues presumably arises from the availability of macro data sets which at least seem to display explosive behaviour. In microeconomics, if we take the models of inter-temporal optimization seriously, the presence of an unstable root isn't a bug, it's a feature. As the number of individual-level longitudinal data sets available increases, investigating unstable roots in the context of saddle-point dynamics is likely to become more of an issue.

III The Predictions of the RA Model

We noted above that, while it is common in empirical RA applications to find positive values of the α_F and α_L coefficients, condition (2), that $\alpha_F = \beta\ \alpha_L$ where β is the discount factor, is much more problematic, in the sense that the estimated values of α_F and α_L very frequently yield implausible values of β. It is worth noting that, as usually set up, the theoretical RA model yields two other conditions on the coefficients on lead and lag consumption. One is the Becker-Murphy stability condition:

$$\alpha_L \alpha_F < 0.25 \tag{10}$$

While the other is the condition that the RA model display saddle-point dynamics

$$\alpha_L + \alpha_F < 1 \tag{11}$$

We are now in a position to set out all of our hypotheses (which we shall maintain in the design of our Monte Carlo experiments) in diagrammatic form, in Figure 5 below:

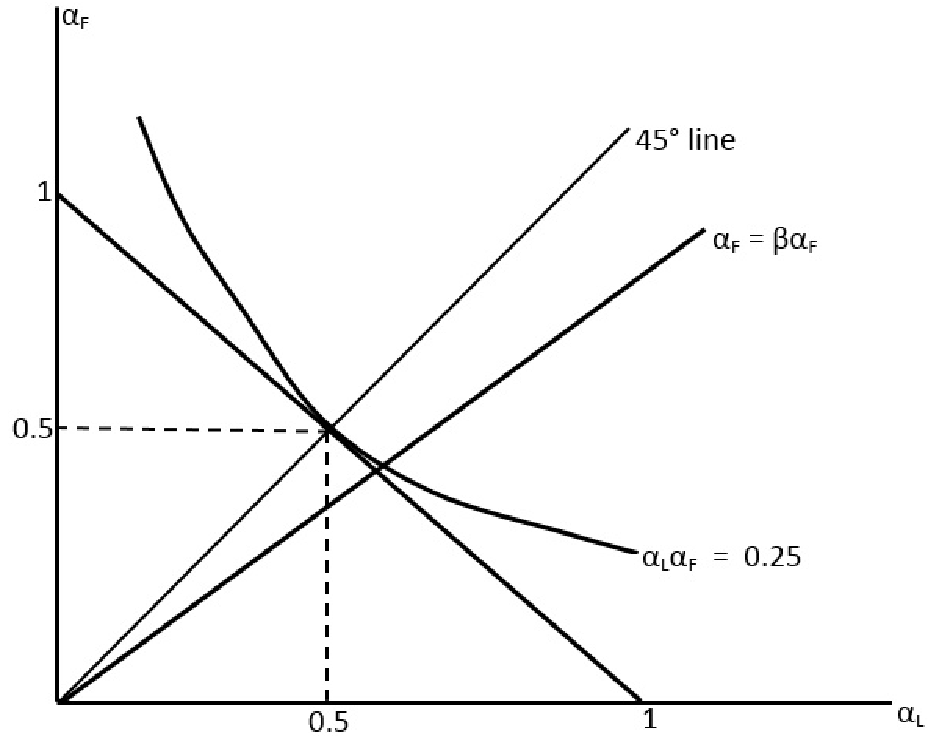

Figure 5: Theoretical restrictions on the RA coefficients

In Figure 5 above the fact that we are looking only at the positive quadrant means that the αs are both positive and the negatively sloped line is $\alpha_L + \alpha_F = 1$, so the condition for saddle-point dynamics is that the coefficients on lead and lag consumption lie inside that line. The negatively sloped curve represents the Becker-Murphy stability condition, in that the αs must lie inside it. The curve coincides with the line only when both αs $= 0.5$, so satisfaction of the saddle-point condition implies satisfaction of the BM condition[6]. The ray from the origin represents condition (2): because the discount factor is less than one, the ray is flatter than the 45° line. Thus if all of the RA hypotheses are true: the values of the coefficients on the lead and lag consumption terms must lie on the ray from the origin, inside the negatively sloped line.

Condition (2), on the relation between the coefficients, imposes a condition on the roots of the RA estimating equation. It can be shown that, if condition (2) holds,

$$\lambda_H \lambda_L = 1/\beta \tag{12}$$

[6]This is because the Becker-Murphy condition is actually a condition for the roots of the equation to be real - if it is violated they are complex, meaning that the optimal trajectory displays cyclical behavior, and cyclical behavior is not consistent with saddle-point dynamics. Violation of the BM condition would be a rejection of the RA model.

where the two λs are the larger and smaller roots respectively. By the definition of β, this gives

$$\lambda_H \lambda_L = 1 + \rho \tag{13}$$

where, as, above, ρ is the individual's subjective discount rate.

This would not, in principle, seem to be too demanding a condition - given ρ it is not that difficult to define a rectangular hyperbola in the two roots. It is important to remember, however, that the saddle-point condition requires that one root be stable - i.e. less than 1 - and the other unstable - i.e. greater than 1. Thus the unstable root must be larger than $1+\rho$, and the more heavily the individual discounts the future (i.e. the larger ρ), the larger the unstable root must be. This relation between the individuals' discount rate and λ_H is, of course, consistent with the interpretation of $1/\lambda_H$ as the strength of the forward looking effect in the RA equation: i.e.

$$1/\lambda H = \partial Y_t / \partial Y_{t+1} \tag{14}$$

(see Laporte, Karimova and Ferguson (2010) for the derivation of this result - it is also shown there that the strength of the backward looking effect is $\partial Y_t / \partial Y_{t+1} = \lambda_L$)

Clearly this result makes intuitive sense: the larger the rate at which the individual discounts the future, the smaller we would expect the influence of the future - i.e. the forward looking effect - to be in his consumption decisions and hence, given the definition of the forward looking effect, the larger the unstable root.

While we shall consider all three of the conditions on the αs in what follows, most of the attention in the literature has been focused on the discount factor condition. Considering Figure 5 we see that, for what we would probably regard as reasonable values of β, (0.95 on annual data, for example) the ray representing the discount factor condition will tend to lie very close to the 45° line, meaning that α_F and α_L will take on very similar values, which may be expected to have implications for the estimation of β.

IV Panel Data Monte Carlo Results

In this section we report results of Monte Carlo estimation of RA type equations. Because the RA model is typically estimated using panel data (when individual level data are available - much of the earlier empirical work on RA was done using aggregate level data, which raised its own econometrics issues distinct from those which we are considering here). To do this we first generate pure time series data sets of 500 observations, following trajectories like Z_1 and Z_2 in Figure 1. Within each experiment we hold the deterministic SODE unchanged and add a zero mean NID disturbance term[7]. Each of our data sets satisfies the RA hypotheses. We assume that each trajectory represents

[7] There is always a question as to how the disturbance term should be regarded in a model of addictive consumption. A myopic addict, on experiencing a random upward shock in consumption, might be expected to remain on the

a full lifetime of consumption behavior for a rational consumer of an addictive commodity. Because we never actually observe a full lifetime of data for any individual, we next divide the full lifetime trajectory into 50 non-overlapping panels, each of 10 observations, each taken to represent a different individual at different stages along the life course. This means that we are working with a very artificial age distribution - we leave the question of the effects of different age distributions on the estimated coefficients for later work. This also means that the individuals in our panel will share the same intercept, as we have not introduced individual specific effects in our DGP, so every individual follows the same lifetime trajectory. Again the purpose of this assumption is to allow us to focus on the econometric implications of saddle-point dynamics, ignoring the other econometric issues that arise in dynamic panel data estimation. As in our pure time series illustration above we do not include any exogenous variables other than the constant term. Our individuals are ordered from youngest to oldest, meaning that the early, youngest, panel entrants are drawn from that part of the lifetime trajectory in which both roots are operative and the stable root is dominant while the later individuals are older and are drawn from the portion of the lifetime trajectory whose behavior is dominated by the unstable root.

Consider first an RA-type equation, with $\alpha_F = 0.4866$ and $\alpha_L = 0.5109$, and $\beta = 0.9524$. We chose this as the first RA illustration because the roots of the RA-type SODE are the same as the roots of the pure time series illustration above, 1.10 and 0.952. Figure 6 below shows the means of the Monte Carlo coefficient estimates for the recursive panel estimation of this equation:

permanently higher path - i.e. myopic addicts might display unit root behavior in individual-level consumption data. Since RA is our maintained assumption for these experiments, we chose to treat the error term as noise around a single chosen consumption trajectory, assuming that, if the rational addict is temporarily knocked off her trajectory she corrects back on to it. Since the econometric issue here arises from the characteristic roots of the chosen trajectory, our fundamental point should not be affected by the variance of the disturbance terms. Experiments with differing variances, early in our research program, did not indicate sensitivity to the choice of variance.

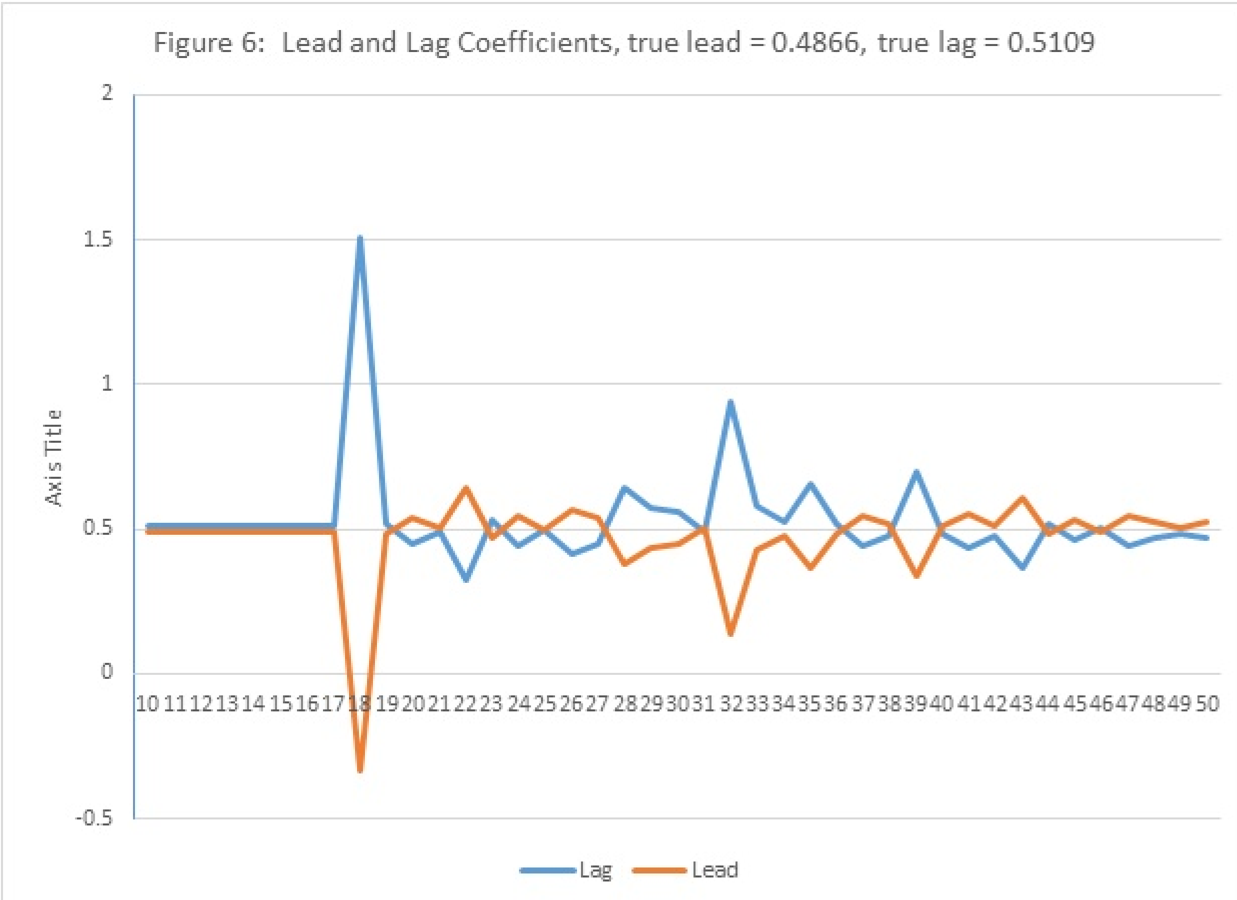

The numbers along the horizontal axis represent individuals in the data set: our 500 observations have been allocated to 50 individuals who occupy successive, non-overlapping segments of the same lifetime trajectory. We see that, as in the case of the pure time series estimation with the same roots, the coefficient estimates drawn from the earlier part of the life course are accurate but that around the 17th individual, meaning at about the same point along the trajectory as in the time series case, the coefficient estimates suddenly become extremely variable.

To see what this means for the RA conditions as set out in Section III above, we first consider Equation (9), $\alpha_L \alpha_F < 0.25$, which we graph in the form $\alpha_L \alpha_F - 0.25$, meaning that, if the condition in equation (9) is satisfied our plotted values should all be negative:

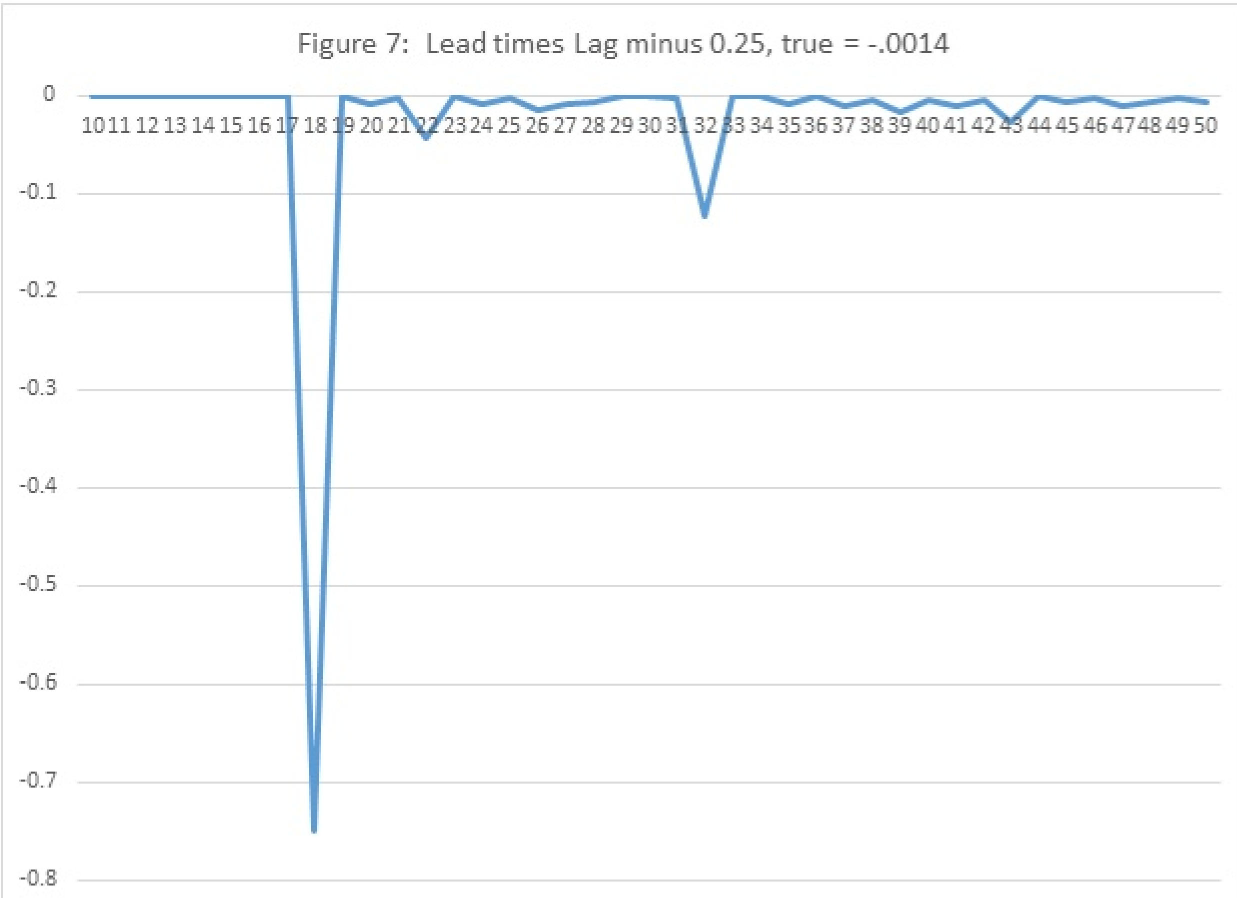

Given the α values the plot should be a straight line at -0.0014. While we have some deviation from this, the values reported in Figure (7) are all negative.

Next we consider the condition set out in Equation (10): $\alpha_L + \alpha_F < 1$, which we plot in Figure 8 below in the form $\alpha_L + \alpha_F - 1$, so again all of the values should be negative:

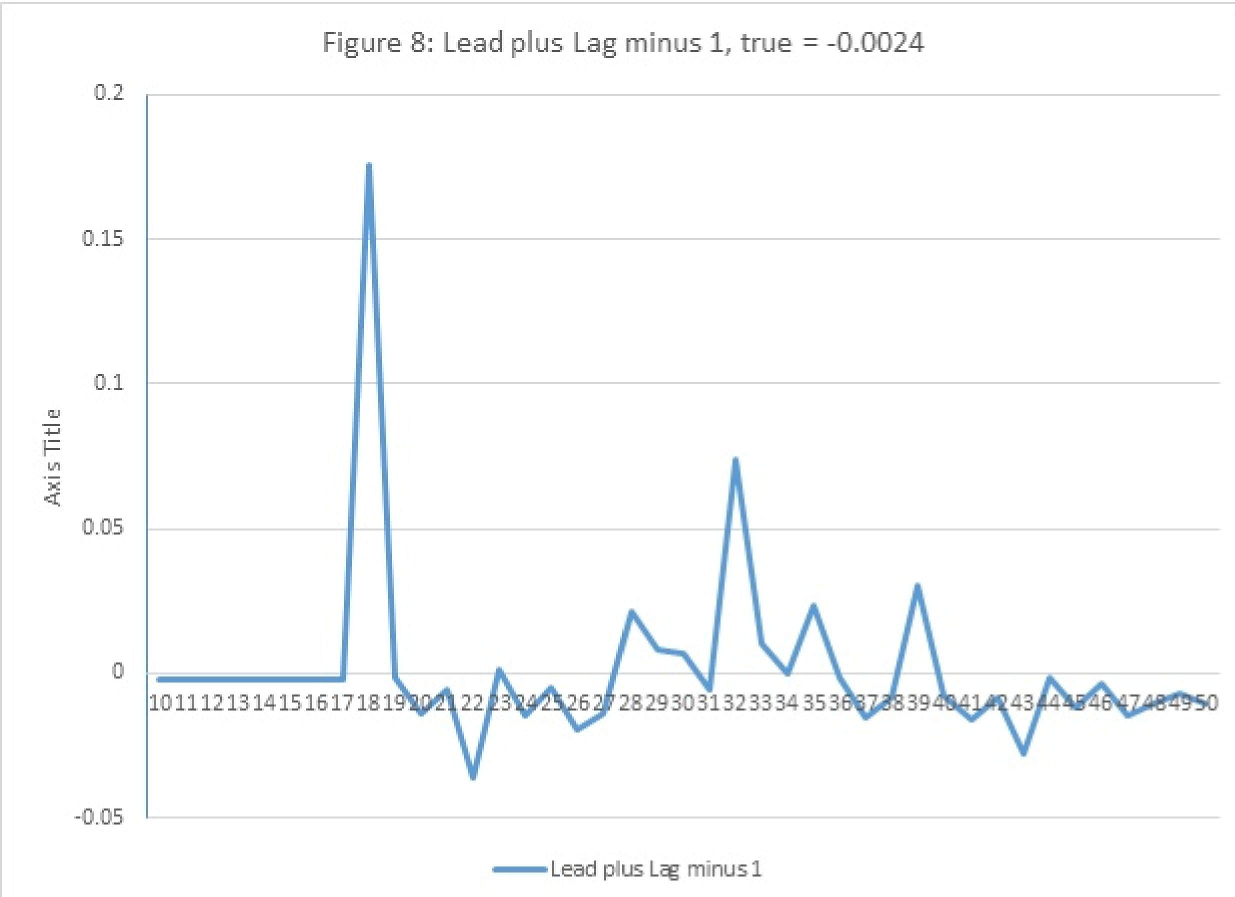

Here, even though the condition is satisfied in the DGP, the estimation results are much less well behaved.

Finally, for this experiment, we plot the values of β generated using the means of the recursive Monte Carlo coefficient estimates:

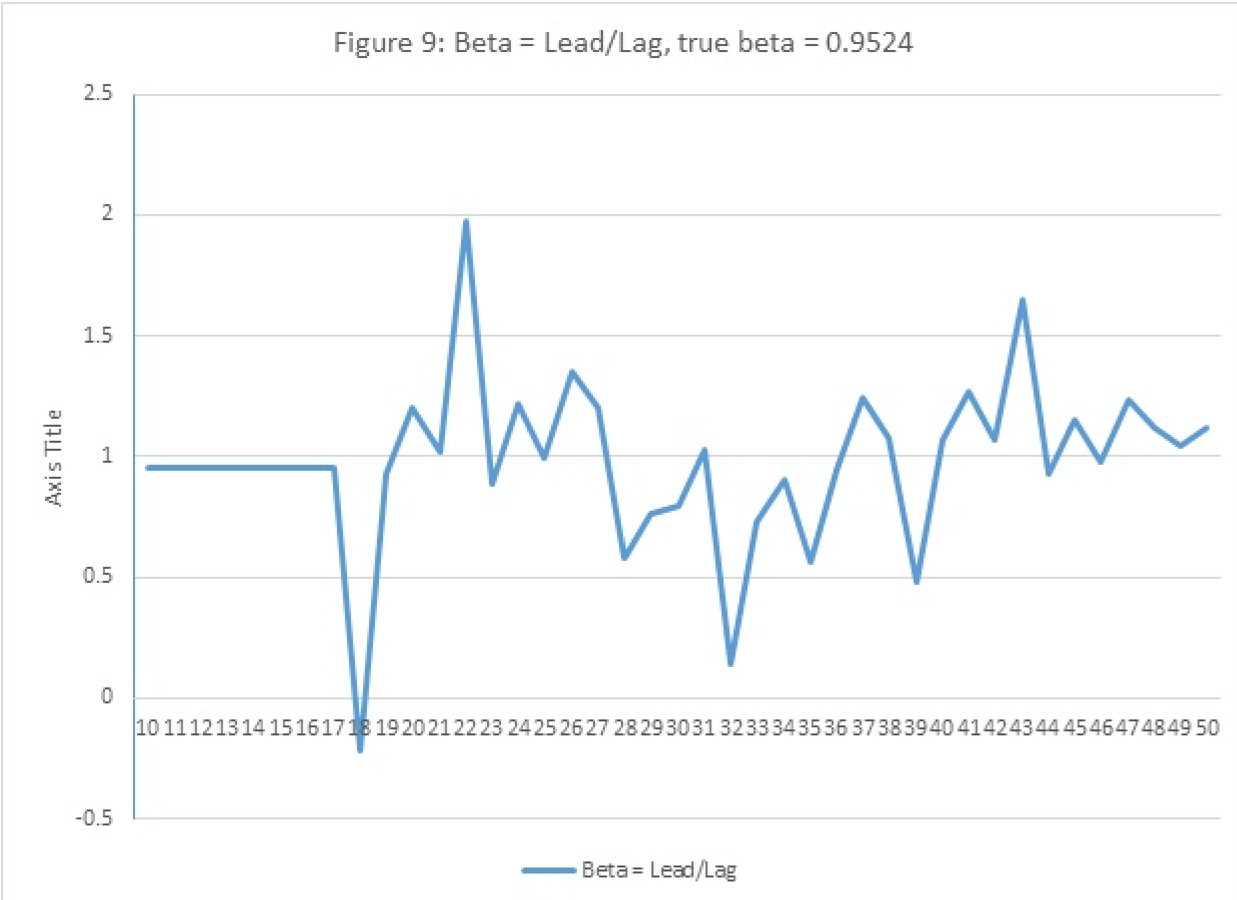

Figure 9: Beta = Lead/Lag, true beta = 0.9524

In Figure 9 we see that, as we would expect from the graph of the MC coefficient means, β is accurately estimated in the first part of the panel data set, i.e. among the observations where both roots are influential, but badly estimated later on.

Next we consider the case of a RA-type SODE where $\alpha_F = 0.4819$ and $\alpha_L = 0.5060$, with $\beta = 0.9524$, giving roots of 1.20 and 0.875. The means of the Monte Carlo coefficients are shown in Figure 10 below:

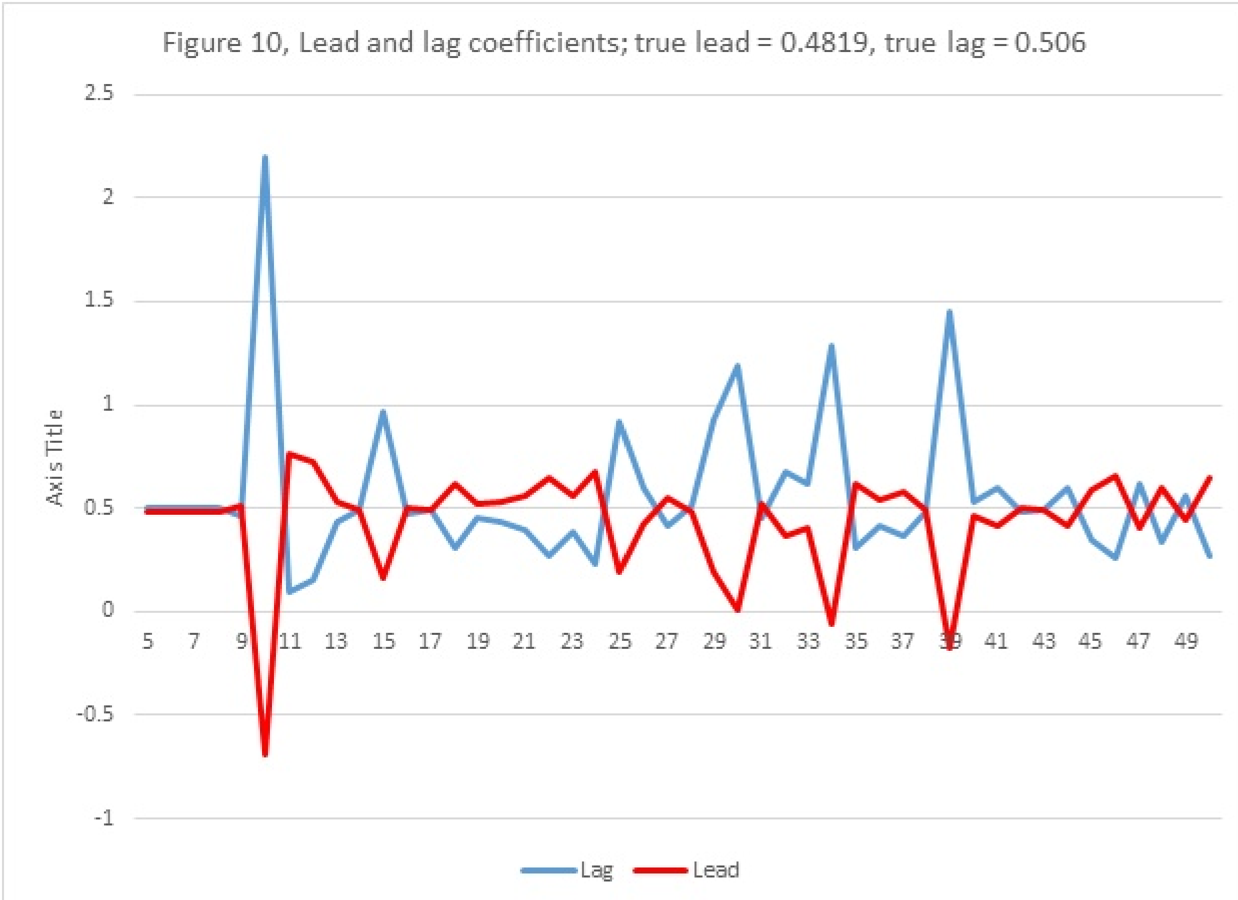

Here we see that the general pattern is as in the previous case, but the well-behaved range is smaller, presumably because the unstable root is larger in this case than in the previous one. Looking at the condition of the product of the coefficients we have

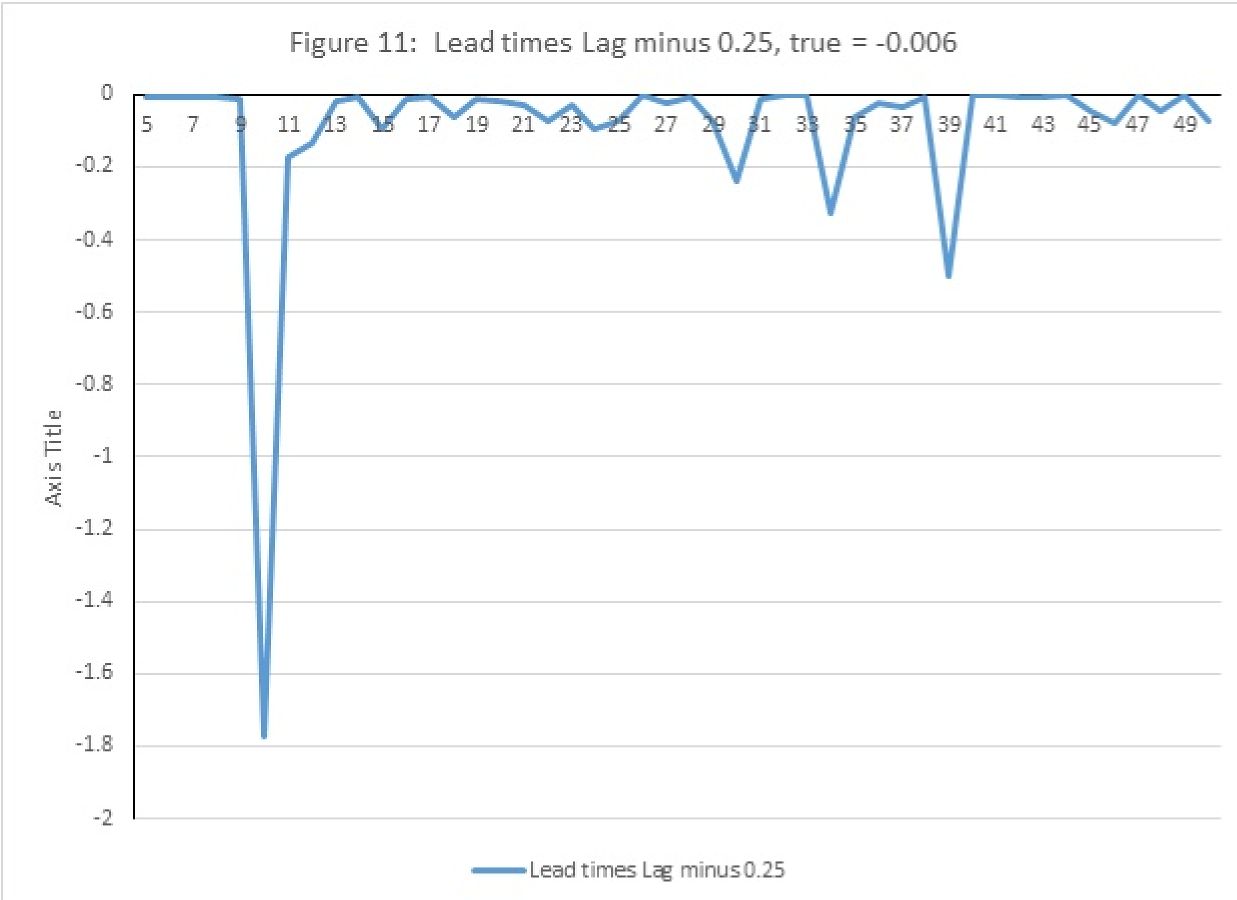

and the sum of the lead and lag coefficients:

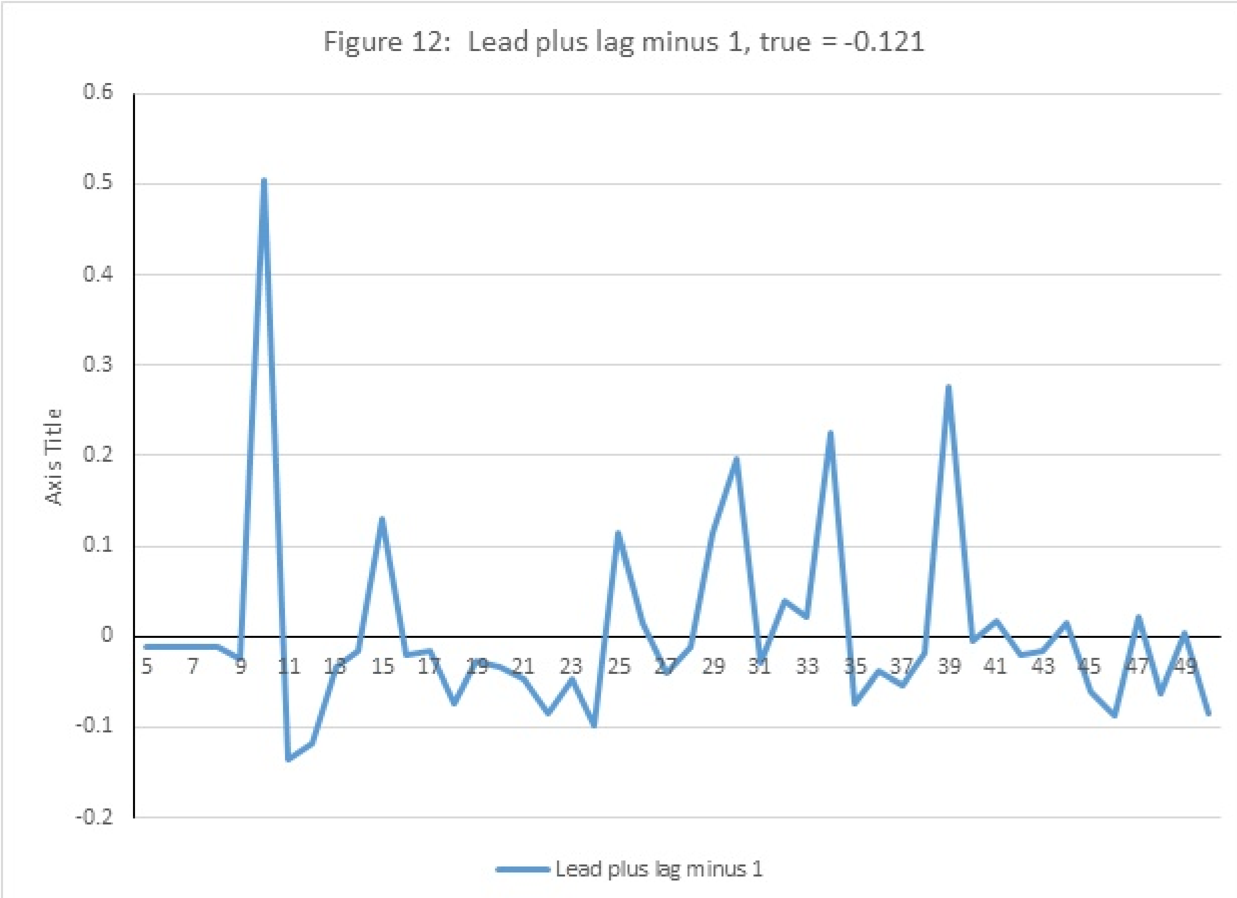

and the panel-recursive β values for this case are:

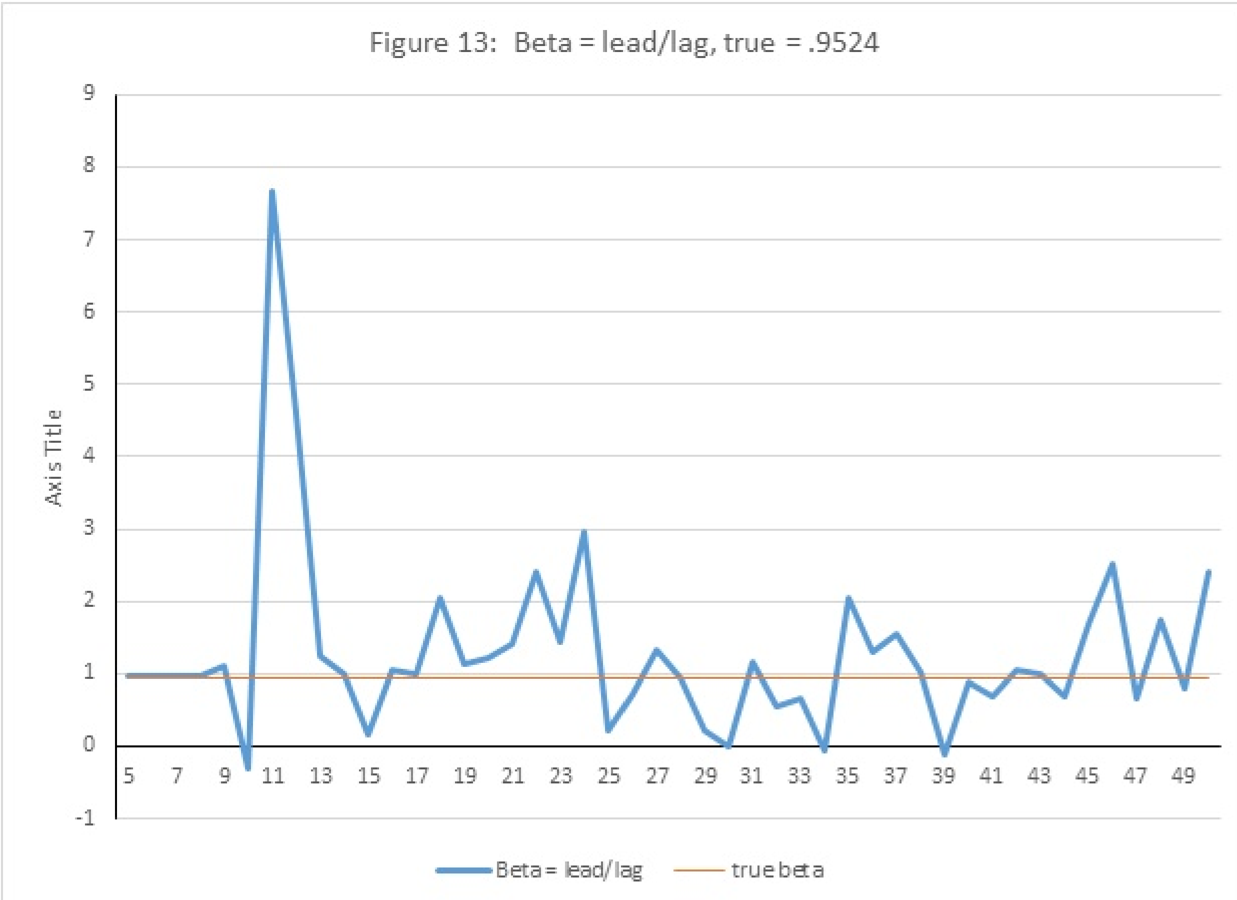

Finally, we consider the case where $\alpha_F = 0.4785$ and $\alpha_L = 0.5024$, with $\beta = 0.9524$, giving roots of 1.25 and 0.84:

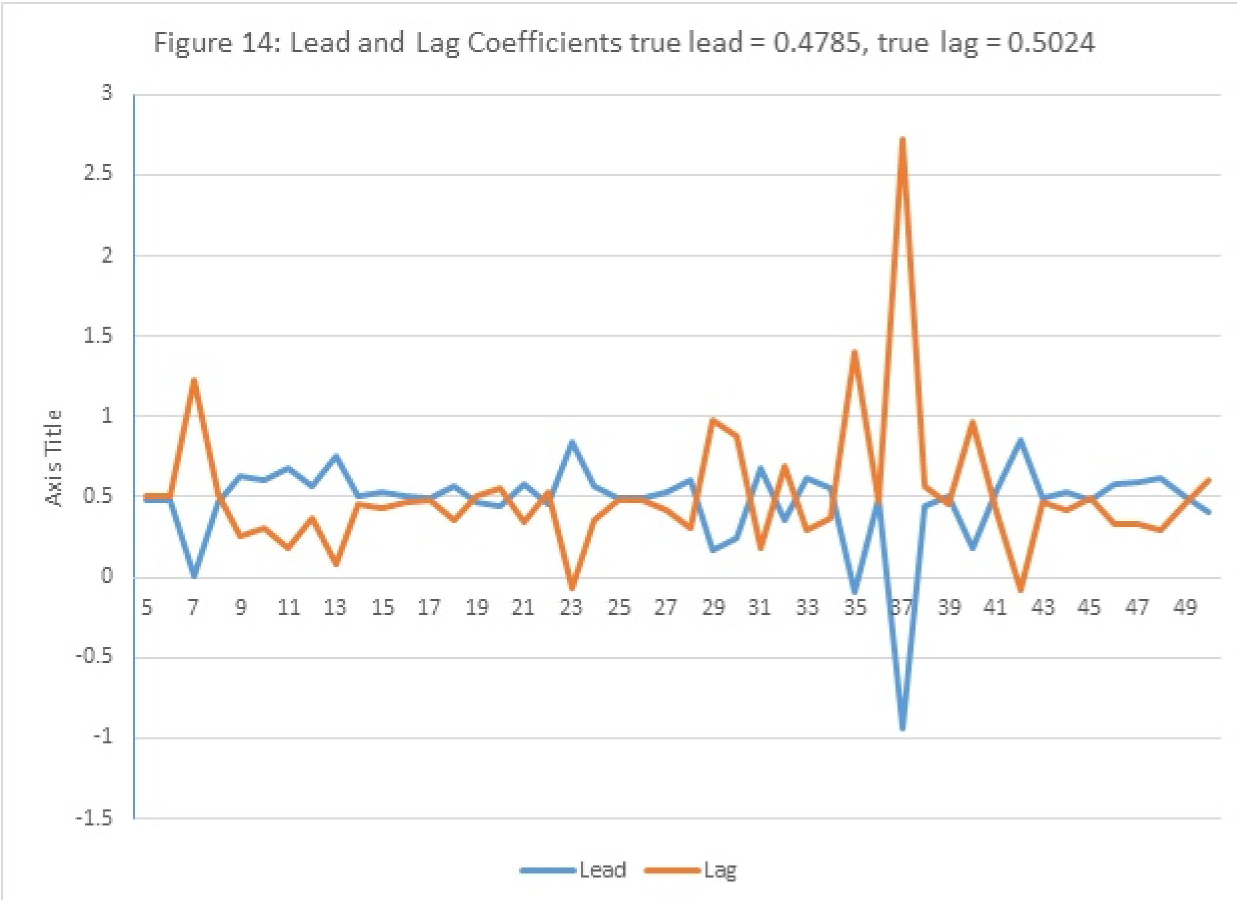

Here the higher unstable root has (presumably) shortened the well-behaved region even further. For this case we have:

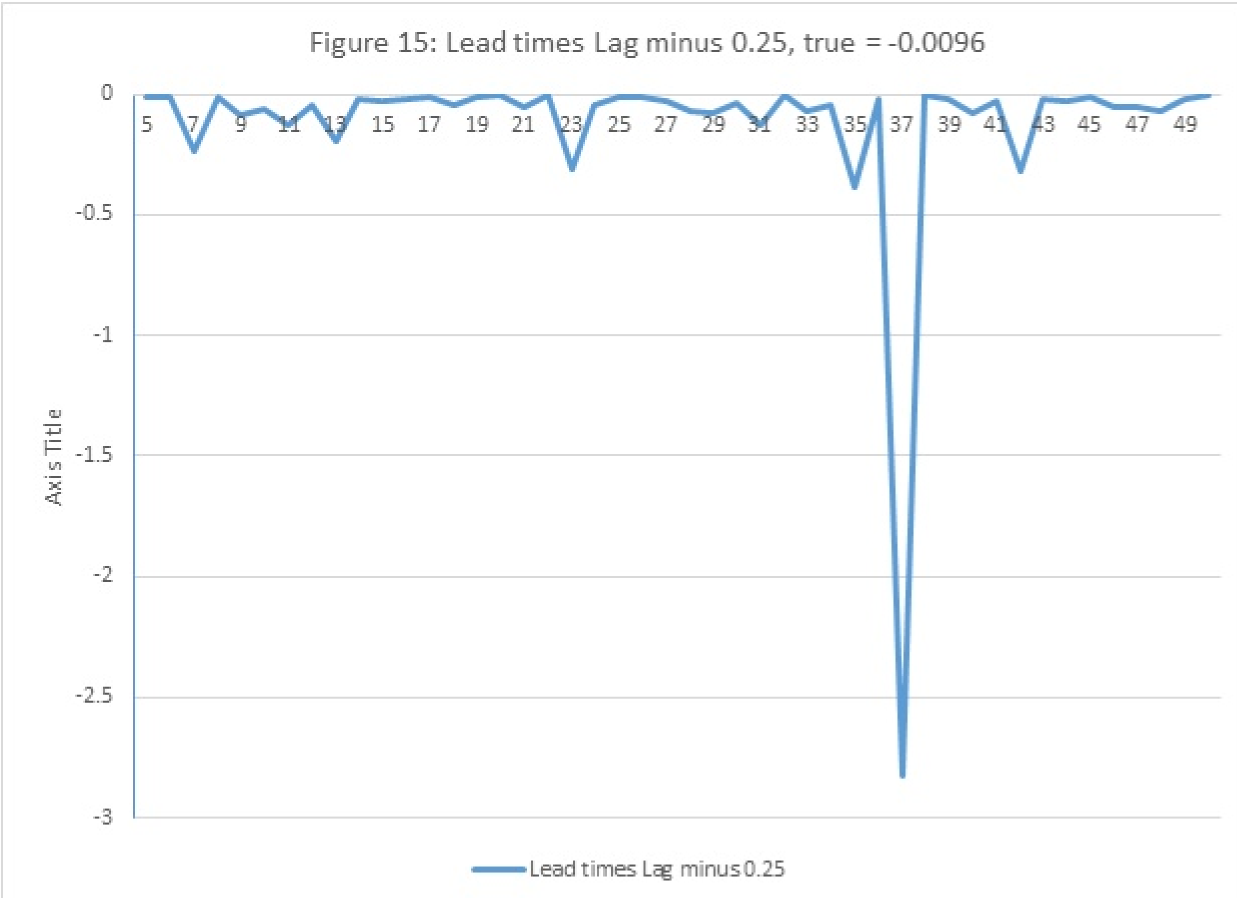
Figure 15: Lead times Lag minus 0.25, true = -0.0096
0
-0.5
-1
-1.5
-2
-2.5
-3
Axis Title
5
7
9
11
13
15
17
19
21
23
25
27
29
31
33
35
37
39
41
43
45
47
49
Lead times Lag minus 0.25

and

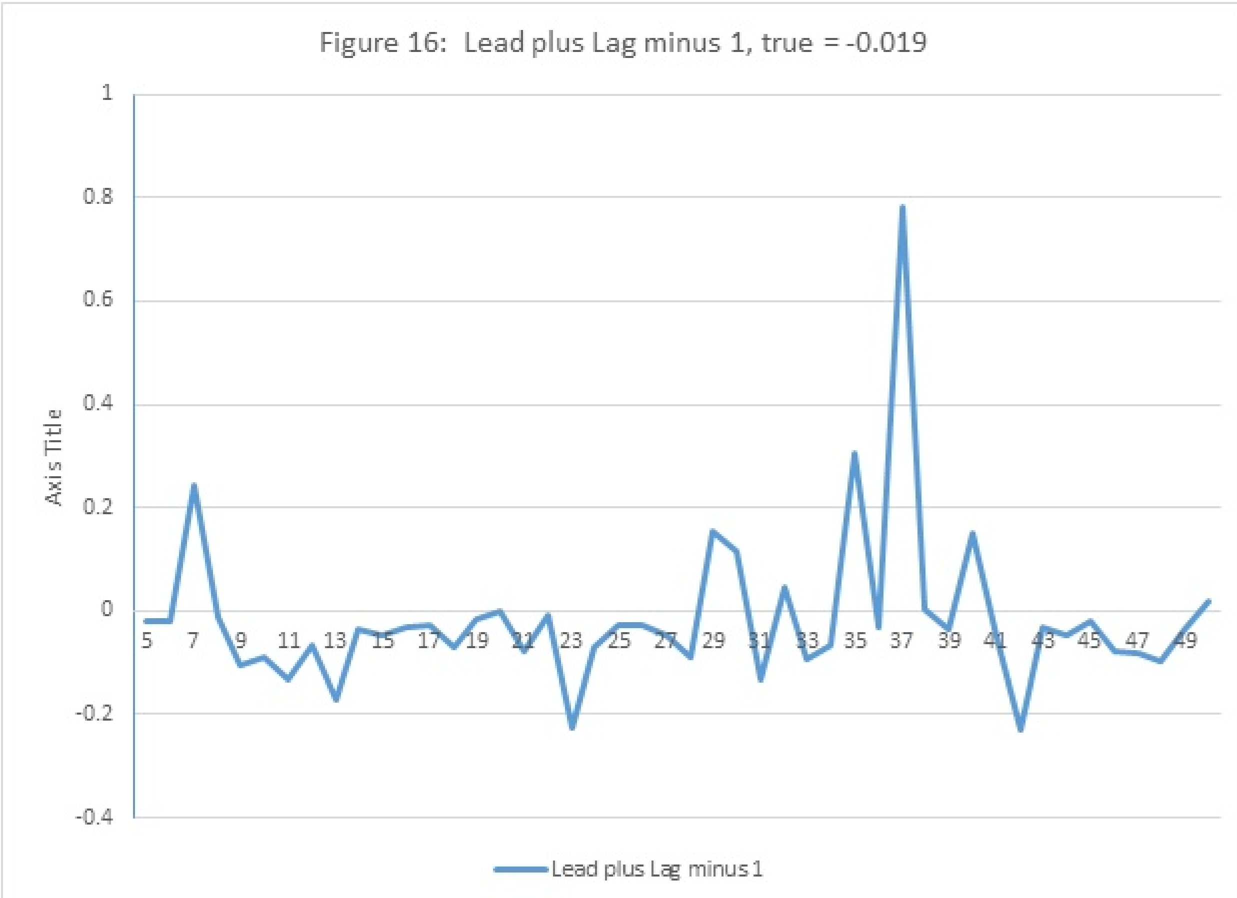

For the discount factor β we have:

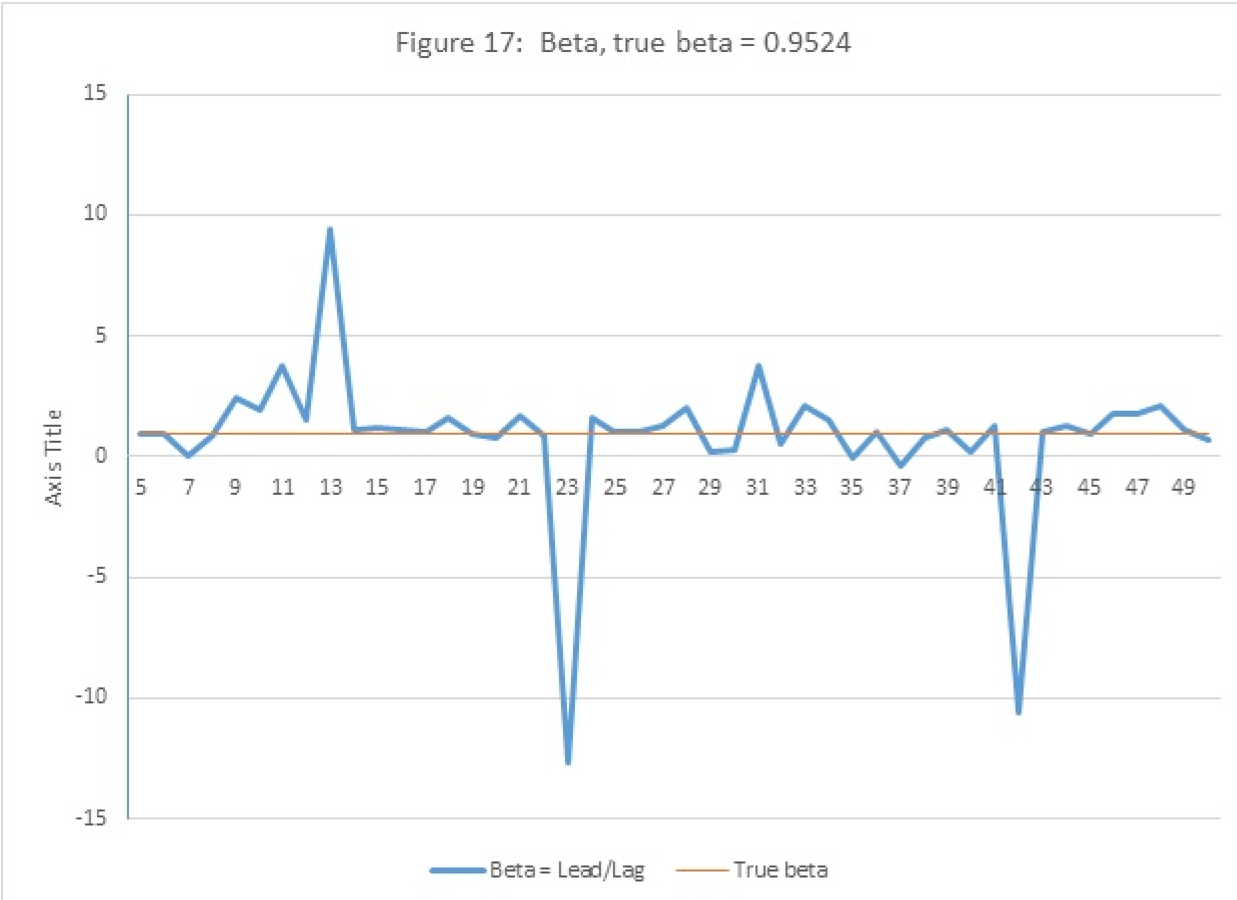

Where we see the same general result as before, but with a couple of much more extreme outlier values for β.

V Discussion

It is important to keep in mind what we have not been trying to do in this paper. We have not been trying to test the Rational Addiction model. Rather, we have discussed a Monte Carlo experiment in which the DGP satisfies the key properties of the theoretical RA model, including the property, based on the usual results of optimal control theory as applied to models of inter-temporal optimization in economics, that the solution to such a model should display saddle-point dynamics, i.e. that its solution equation, writing Y as a function of elapsed time, should contain one stable and one unstable root. We note that we have not added exogenous variables and that the standard deviation of the disturbance term is small relative to many of the values in the data set.

Our proposition is that, even when the data set is known to satisfy all of the properties of the theoretical RA model, the presence of the unstable root may make it very difficult to estimate the coefficients accurately, with the result that, looking at the estimated value of β, the discount factor, we may be led to reject the RA model even when it is true. Again, we emphasize that we are not saying that this necessarily explains the “fly in the ointment” which Baltagi(2007) identified, but rather that it may greatly complicate the problem of testing the RA hypothesis. Even when the individual is a rational addict, the properties of the data may lead us to reject a true hypothesis.

We noted above that some of the Monte Carlo DGPs were better suited to finding the RA coefficients than others, although even in the best of the cases the results went off relatively quickly. The precise pattern which we observe in practice will, we hypothesize, depend on two factors - how large is the unstable root and where in our hypothesized lifetime trajectory do the bulk of the individuals in the panel come from. In this paper we have, as we noted above, chosen an unrealistically uniform distribution of ages of our individuals. We hypothesize that in practice the estimated coefficients will be forms of weighted averages, with the weights on various parts of the trajectory depending on how many of the individuals in the particular sample happen to lie at various points along it. We emphasize here that because all of our individuals lie on the same lifetime trajectory, none of the problems associated with different fundamental preferences, which are usually associated with the DPD literature, should affect our results here. Those problems would presumably present an additional issue to the issue with which we are concerned here.

The hypothesized properties of the RA model have been translated here into their implications for the roots of the solution to the RA-form SODE. We should note that the particular values of the roots have not been calibrated to any particular numbers, since we have no solid information on what values to calibrate them to. We have chosen values within a very small absolute range, because we have in general tried to use values that seem sensible for RA models. The relative values of α_F and α_L have been dictated by the choice of a β which does not seem out of place for a RA model based on annual data and the absolute values have been chosen with an eye to having one stable and one unstable root, but not making the unstable root dominate too early, so that we could illustrate the econometric effect in which we are interested, and because we do not observe extreme instability in consumption behavior in the most commonly studied addictive commodities - cigarettes and alcohol. For narcotics, of course, the story might be different. Beyond that, however, our results raise the possibility that even in what appears to be a well-studied case, the presence of the unstable root may be affecting empirical results.

VI Conclusion

The starting point for this paper was the observation that, while the Becker-Murphy model of Rational Addiction is widely used in the economics literature on the consumption of addictive substances, and widely accepted among economists, there is controversy in the literature as to how strongly the empirical results actually tend to support the theoretical model. In particular, it is not uncommon for the estimates of β, the individual's time discount factor, to seem implausible, and indeed, for estimates of the value of β to range widely across the literature.

In this paper we have not attempted to test the RA model. Rather, our interest has been in the question of whether the unstable root which theory predicts will be present in the individual's optimal solution trajectory might complicate the estimation of the RA model even when the true DGP possesses the key characteristics of the RA model. Our Monte Carlo results suggest that this may well be the case. The results of the experiments presented here are consistent with the proposition that, the stronger the unstable root the more unreliable are estimates of the coefficients of an RA-form SODE. This implication is complicated by the fact that the key issue is not just the magnitude of the unstable root, although that is a powerful factor, but the relative importance of the stable and unstable roots in characterizing the segment of the lifetime trajectory which an individual happens to lie on. As we saw in the Monte Carlo results for our best-behaved RA-form SODE, there will exist a segment of the lifetime trajectory along which it is possible to estimate the RA-form SODE accurately, and there will exist segments along which it is not. We emphasize that these segments all lie on the same lifetime trajectory: it is not a matter of different individuals having different characteristic roots. This effect will be complicated in estimation on real-world data sets in that, even if everyone in the longitudinal data set did happen to be following exactly the same lifetime trajectory, we would expect to have clusters of individuals at different stages along it, so the estimated coefficients would be, in a sense, weighted averages of the types of coefficient values we have found in our experiments.

We note that there are values of the unstable root for which well-behaved segments do exist. We also note that, at least in our panel simulations, the unstable root could be recovered with considerable accuracy, although the same could not be said for the stable root. The unstable root by itself, however, will not allow for identification of the true values of the coefficients on lead and lag consumption.

In addition, we have the problem that we do not, at this point, know what constitutes a reasonable value for λ_H. In our experiments, we simply generated data series that were driven by certain values of the stable and unstable roots. While the roots are raised to the power t in the solution form for a SODE, in our experiments we have not defined the length of a period, t: there is no reason for it to be in years, it could perfectly well be in months. While there must be some form of internal consistency between the dynamics of a model set out in terms of months and the same model set out in terms of years, that fact by itself does not tell us what a well-behaved value of λ_H should be in any real-world application. Indeed, the only basis we have for assuming that the unstable root will be small enough to give us what we have been calling a well-behaved data set is that we do not seem to observe, in real world applications, the explosive dynamics that would be associated with a large unstable root.

Our problem, then, appears to be that, when the RA model is the true DGP, the presence of the unstable root makes it extremely difficult to estimate the true coefficients and to extract the true value of the discount factor. We do not assert that our results must absolutely be the explanation for the wide range of estimated values, especially for β, which have come out of the empirical RA literature. It is, after all, always possible that the RA hypothesis is simply wrong[8]. We do, however, suggest that they may be part of the explanation for the variability of the estimates, that they may well represent a serious complicating factor in empirical RA research, and that they merit further investigation if we are to hope to understand the empirical behavior of the RA model[9].

[8] A referee has raised the issue of whether the argument made here might also shed light on the exponential versus quasi-hyperbolic discounting debate. See, for example, Gruber and Koszegi (2001).

[9] We do not, in this paper, consider the possibility of using pseudo panels constructed from repeated cross section surveys. In what must be one of the most meticulous papers in the RA literature, in the sense that they consider all of the theoretical predictions of the theoretical model, including the roots, Pierani and Tiezzi (2011) construct a

Appendix: Macro versus Micro estimation

In discussing issues in the estimation of the RA model, it is important to note that, while the theoretical model is one of inter-temporal optimization by the individual, past limitations on the availability of individual level panel data sets has meant that most empirical estimation has been done at the aggregate, or macro, level. American studies[10] have often used national or state-level data and studies from other countries have used national level data. Among the few early exceptions were Chaloupka (1991) that used data from the National Health and Nutrition Examination Survey, and Grossman, Chaloupka and Sirtalan (1998), and Grossman and Chaloupka (1998) both of which used data from the Monitoring the Future program. This means that most of the estimates of the coefficients on lead and lag consumption, and therefore of the discount factor, which are at the core of the debate about the empirical validity of the theoretical model, have been based on aggregate level studies. The usual assumption is that macro level responses can be derived, with suitable attention to possible distributional effects (differences in coefficients across age groups and between males and females, for example) from individual level coefficients. This is true of factors such as price elasticities - the market level response to a price change should be a weighted average of individual level responses - but it is not true of the intrinsic dynamics of consumption and therefore it is not true of the discount factor. Since estimates of the discount factor depend on estimates of the coefficients on lead and lag consumption (as do presumed tests of rationality in addiction) and the intrinsic dynamics of individual level consumption do not carry over to the aggregate level, we cannot rely on aggregate level estimates of the discount factor to tell us about the true individual-level discount factor.

The reason for this is that the lead and lag consumption terms in the estimated RA equation reflect the evolution of optimal consumption through the individual's life course. In terms of the phase diagram in Figure 1 above, the individual's optimal lifetime consumption rule has her following a trajectory like Z_1, so that her consumption changes as time passes, even if none of the exogenous variables of the problem change. (A change in an exogenous variable, such as price, would show up as a shift of the entire trajectory.) At the aggregate level, however, even if each individual in society is following a U-shaped trajectory such as Z_1, there is no reason for aggregate consumption to behave the same way. Aggregate consumption, and in particular per capita consumption (the variable typically used in empirical rational addiction studies) will be a weighted sum of individual consumption. To see this, assume that every individual in the society follows the same lifetime consumption path, as in Z_1 in Figure 1. Individuals of different ages will be at different points along the trajectory. Aggregate per capita consumption at any point in time will therefore depend on what proportion of the population are at each point along the trajectory at that point in time. If the population age distribution is stable, so that the people flow in and out in such a manner that the proportions at each age remain unchanged over many years, then aggregate per capita

pseudo panel using Italian monthly cross section data on wine consumption from January 1999 to December 2006 and estimate the RA model on it. Their results are consistent with the RA theory. Pseudo panel methods should, in principle, allow us to track cohorts which are, in the first year of the surveys, at different points along the lifetime consumption trajectory. While there will be different individuals in the sample each year, and we would be looking at mean levels of each age cohort's consumption in each year, the argument would be that, so long as we include a large enough number of trajectory-shifting factors, including, ideally, measures of distributions as well as means of those shifters, age cohort average consumption data should allow us to map out the typical lifetime trajectory. However, since the trajectory being mapped out is in effect an average of trajectories obeying saddle-point dynamics, the problem which we have been discussing here may well still arise.

[10] For example, Becker, Grossman and Murphy (1994)

consumption will remain unchanged over time (so long as none of the exogenous variables change) even though each individual's consumption is changing as time passes. Per capita consumption by a population of rational addicts can remain unchanged over the years, even though each individual is following a trajectory such as Z1, with consumption (in this case) initially falling and then rising. At the very least, there would seem to be a case for arguing that any aggregate level RA estimation should include population age distribution among its explanatory variables.

The question then arises as to why so many aggregate level studies seemed to find support for rational addiction behavior. Auld and Grootendorst (2004) argue, correctly, in our view, that this is explained by time series econometrics issues at the aggregate level. They argue that the standard lead-lag form of estimating equation can be expected to yield spurious results at the aggregate level, and illustrate their case by estimating RA-type SODEs on Canadian data for milk, eggs, oranges, apples and cigarettes, using OLS and three different 2SLS methods on each. They find positive, though not always significant, coefficients on all of their variables but generally implausible values of the discount factor. They conclude that, if the RA equations are to be believed, milk is more addictive than cigarettes.

Auld and Grootendorst's argument is to some degree weakened by their empirical illustrations. While it is true that with one exception (one of the 2SLS estimates for oranges) they find positive coefficients on lead and lag consumption, the roots of their RA-type SODEs (which are not discussed in their paper but relate to the discussion here) are not consistent with individual level optimization. For milk, for example, which seems to yield strong RA results, all of the roots are complex with modulus greater than one, meaning that milk consumption in Canada is apparently characterized by unstable spirals, a result which does not seem likely to fall out of an inter-temporal optimization problem[11]. Similar issues arise in most of their other examples. The only set of results which do not yield at least some implausible roots are those for cigarettes. Thus if we go beyond simply looking at the signs on the lead and lag consumption terms, and look at the implied dynamics, only cigarettes seem to be consistent with RA theory. On the basis of the A&G empirical examples it would appear that an important implication of their results is that simply looking at the signs on the lead and lag coefficients is not sufficient to test RA.

A&G's key result about the use of the RA methodology on certain types of macro data is undoubtedly correct. Consider the case of a FODE whose true DGP is stable:

$$Y_t = -1.05 + 0.8Y_{t-1} + \epsilon_t$$

When we, following A&G, run a Monte Carlo experiment estimating an RA form SODE on this equation we obtain

$$\begin{array}{rllll} & Y_t = & -\,0.14 + & 0.486Y_{t-1} + & 0.485Y_{t+1} + \epsilon_t \\ MCSD & & (0.09) & (0.01) & (0.0099) \end{array}$$

[11]Dockner and Feichtinger (1993) propose a model which, they argue, can yield a cyclical consumption pattern as a solution to the inter-temporal optimization problem, but it seems unlikely that the particular conditions required by their model are satisfied in the case of milk and even more unlikely that the optimal cycle would be explosive

which would appear to be a well-behaved RA equation, with roots 1.27 and 0.79 and with discount factor 0.997.

When we go a step further and estimate a more standard backward looking SODE on the same data, however, we find

$$\begin{array}{rllll} & Y_t = & -\,0.14 + 0.785Y_{t-1} + 0.004Y_{t-2} + \epsilon_t \\ MCSD & & (0.21) \;\; (0.068) \qquad (0.073) \end{array}$$

Thus it appears that, when the true DGP is a stable FODE, as would be the case in aggregate where the dynamics of the dependent variable arose from a lagged adjustment process, if we estimate a RA-type forward looking SODE we seem to find well-behaved RA behavior but when we estimate a backward-looking SODE, the second lag falls well short of significance. Other preliminary Monte Carlo results suggest that when the DGP is in fact second order, both backward-looking and RA-type SODEs can be fitted to it, which suggests a fairly simple possible check for spurious RA, although we have not yet done enough work to propose this as a formal test.

Taken in all, then, there is a strong theoretical argument against estimating RA models on aggregate level data, and our empirical results, consistent with those of A&G, suggest that there are also serious dynamic econometric issues[12] associated with the use of aggregate data[13].

[12]Note that the econometric issues are quite separate from the theoretical issues: the theoretical issues argue in favour of modifying the form of the estimating equation if aggregate data is to be used, while the econometric issues focus on investigating the dynamics implied by the estimates in much more detail than has usually been the case.

[13]Just to add to the complications, when working with aggregate data, unit root problems may arise.

References

Arellano, M. (2003) Panel Data Econometrics, Oxford University Press: Advanced Texts in Econometrics, Oxford.

Auld, M. C. and Grootendorst P. (2004), "An empirical analysis of milk addiction", *Journal of Health Economics*, 23: 1117-1133.

Baltagi, B.H. (2007) "On the use of panel data methods to estimate rational addiction models", *Scottish Journal of Political Economy*, 54 (1):1-18.

Baltagi, B.H. and Geishecker, I. (2006) "Rational alcohol addiction: evidence from the Russian longitudinal monitoring survey," *Health Economics*, 15(9): 893-914.

Baltagi, B. H. and Griffin, J.M. (2001) "The Econometrics of Rational Addiction: the Case of Cigarettes" *Journal of Business and Economic Statistics* 19(4), October, 449-454.

Becker, G.S. and Murphy, K.M. (1988) "A Theory of Rational Addiction", *Journal of Political Economy*, 96(4): 675-700.

Becker, G.S., Grossman, M. and Murphy, K.M. (1991) "Rational Addiction and the Effect of Price on Consumption" *American Economic Review* 81(2), Papers and Proceedings of the Hundred and Third Meeting of the American Economic Association, May, 237-241.

Becker, G. S., Grossman, M. and Murphy, K.M. (1994) "An Empirical Analysis of Cigarette Addiction" *American Economic Review* 84(3), June, 396-418.

Chaloupka, F. (1991) "Rational Addictive Behavior and Cigarette Smoking", *Journal of Political Economy*, 99(4) August, pp. 722-742

Dockner, E. J. and Feichtinger, G. (1993): "Consumption Patterns and Rational Addiction" *American Economic Review*, 83(1) March, pp. 256-263

Doornik, Jurgen A. and David F. Hendry (2013a): Empirical Econometric Modelling - PcGive 14 pub. Timberlake Consultants

Doornik, Jurgen A. and David F. Hendry (2013b): Interactive Monte Carlo Experimentation in Econometrics - PcNaive 6 pub. Timberlake Consultants

Ferguson, B. S. (2000) "Interpreting the rational addiction model", *Health Economics*, 9(7): 587-598.

Ferguson, B.S. and Lim, G.C. (1988), Discrete Time Dynamic Economic Models: Theory and Empirical Applications, Routledge.

Grossman, M., Chaloupka F.J., and Sirtalan, I. (1998) "An Empirical Analysis of Alcohol Addiction: Results from the Monitoring the Future Panels" *Economic Inquiry* 36, January, 39-48.

Grossman, M. and Chaloupka, F. J. (1998) "The demand for cocaine by young adults: a rational addiction approach" *Journal of Health Economics* 17, 427-474.

Gruber, J. and Köszegi, B. (2001) "Is Addiction 'Rational'? Theory and Evidence" *The Quarterly Journal of Economics* 116 (4): 1261-1303.

Jones, Andrew M., Audrey Laporte, Nigel Rice, Eugenio Zucchelli (2014): "A Synthesis of the Grossman and Becker-Murphy Models of Health and Addiction: Theoretical and Empirical Implications" Working Papers 140007, Canadian Centre for Health Economics.

Laporte, A., Karimova A., and Ferguson B.S. (2010) "Quantile regression analysis of the rational addiction model: investigating heterogeneity in forward-looking behavior", *Health Economics*,19(9): 1063-1074.

Laporte, A., Rohit Dass, A., and Ferguson, B.S. (2016) "On a Possible Non-Stationarity Problem in the Estimation of Rational Addiction-type Equations", Working Papers 160004, Canadian Centre for Health Economics.

Nielsen, B. (2008) "On the Explosive Nature of Hyper-Inflation Data" *Economics: The Open-Access, Open-Assessment E-Journal*, 2 (2008-21): 1-29.

Nielsen, B. and Reade, J.J. (2007) "Simulating Properties of the Likelihood Ratio Test for a Unit Root in an Explosive Second-Order Autoregression," *Econometric Reviews*, vol. 26(5): 487-501.

Melberg, H. O. (2008) "Rational addiction theory - a survey of opinions", University of Oslo, Health Economics Research Programme, Working Paper 2008:7

Olekalns, N. and Bardsley, P. (1996) "Rational Addiction to Caffine: An Analysis of Coffee Consumption" *Journal of Political Economy* 104(5), October, 1100-1104.

Pierani, P. and Tiezzi, S. (2011) "Infrequency of Purchase, Individual Heterogeneity and Rational Addiction in Single Households' Estimates of Alcohol Consumption", American Association of Wine Economists AAWE Working Paper No. 78.

Shen, K. and Giles, D.E. (2006): "Rational exuberance at the mall: addiction to carrying a credit card balance" *Applied Economics* 38, 587-592.

Sisto, A. and Zanola, R. (2010): "Cinema attendance in Europe" *Applied Economics Letters* 17, 515-517

Spenner, E. L., Fenn, A. J., Crooker, J.R. (2010) "The Demand for NFL attendance: A Rational Addiction Model" *Journal of Business and Economics Research* 8(12), December, 21-42.