10ISSN 1183-1057

SIMON FRASER UNIVERSITY

Department of Economics

Working Papers

12-05

"A Behavioral Defense of Rational Expectations"

Ken Kasa

March, 2012

Economics

A BEHAVIORAL DEFENSE OF RATIONAL EXPECTATIONS

KENNETH KASA

Abstract. This paper studies decision making by agents who value optimism, but are unsure of their environment. As in Brunnermeier and Parker (2005), an agent's optimism is assumed to be tempered by the decision costs it imposes. As in Hansen and Sargent (2008), an agent's uncertainty about his environment leads him to formulate 'robust' decision rules. It is shown that when combined, these two considerations can lead agents to adhere to the Rational Expectations Hypothesis. Rather than being the outcome of the sophisticated statistical calculations of an impassive expected utility maximizer, Rational Expectations can instead be viewed as a useful approximation in environments where agents struggle to strike a balance between doubt and hope.

JEL Classification Numbers: D81, D84

Reason is, and ought only to be the slave of the passions, $\cdots$ – Hume (1739, p. 462)

1. Introduction

The Rational Expectations Hypothesis has always been controversial. Initially this was partly due to a perceived connection between Rational Expectations and conservative political views. It is now widely understood that no such connection exists. A more enduring source of skepticism has been its supposedly unrealistic assumptions. The Rational Expectations Hypothesis is based on two key assumptions: (1) Agents have common knowledge of the correct model of the economy, and (2) Given their knowledge of the model, agents make statistically optimal forecasts. A large and still active literature on learning has studied the consequences of relaxing the first assumption.[1] The results are rather negative. In general, adaptive learning converges to Self-Confirming Equilibria, not Rational Expectations Equilibria, and there can be important differences between these two equilibrium concepts (Sargent (2008)).[2] Until recently, the second assumption seemed relatively uncontroversial, or at least no more controversial than other maximization assumptions in economics. However, a number of recent papers have questioned the connection between utility maximization and statistically optimal forecasts. An implicit assumption of Rational Expectations is that beliefs are purely instrumental, i.e., agents derive no utility from expectations themselves. An abundance of experimental work casts doubt on this assumption. If agents value their own beliefs, then economists need to consider how agents go about choosing their expectations.[3]

This paper shows that when both assumptions are relaxed simultaneously, Optimal Expectations can in fact be Rational Expectations. Following Hansen and Sargent (2008), I assume agents have doubts about their models. To guard against potential model misspecification, agents formulate

Date: March, 2012.

[1]For useful reviews see Sargent (1993), Blume and Easley (1998), and Evans and Honkapohja (2001).

[2]A more recent literature has shifted attention from learning to selection. As long as markets are complete, and at least some agents' priors contain the true model, then eventually Rational Expectations will prevail (Blume and Easley (2006)). Of course, this just changes the question to the plausibility of complete markets.

[3]Akerlof and Dickens (1982) were perhaps the first to formally model belief selection.

'robust' decision rules. Agents accomplish this by imagining themselves to be immersed in a dynamic zero-sum game against a malevolent 'evil agent', who attempts to subvert their control and forecasting efforts. Essentially then, agents attain robustness by being pessimistic. To prevent agents from being excessively pessimistic, I again follow Hansen and Sargent (2008) and assume the evil agent's actions are bounded by a relative entropy constraint. The constraint is set so that agents only hedge against models that could have plausibly generated the observed data. At the same time, I assume agents derive utility from being optimistic. Following Brunnermeier and Parker (2005), I assume an agent's optimism is constrained by the decision costs it imposes. I show that these decision costs can also be related to relative entropy. To formulate Optimal Expectations, agents now invent an 'angelic agent', who selects *favorable* models subject to a relative entropy constraint. It turns out that if the evil and angelic agents are evenly matched, an agent's beliefs conform to the Rational Expectations Hypothesis, even though his knowledge of the true model is limited, and he is subject to the same psychological biases that have been highlighted in the recent experimental literature. Of course, the result that Robust Optimal beliefs conform *exactly* to the Rational Expectations Hypothesis is a non-generic, knife-edge result. However, the point here is not to argue that Rational Expectations is the uniquely correct way to model expectations. Instead, this paper merely attempts to point out that Rational Expectations can be a useful approximation even when its implausible assumptions are relaxed.

The remainder of the paper is organized as follows. The next section develops a simple static, Linear-Quadratic model of Optimal Expectations. An LQ setting is convenient, since decisions only depend on first moments. The key result here is to show that the decision costs arising from optimistically biased beliefs can be related to the relative entropy between the true model and the agent's biased model. Section 3 does the same thing for the case of model uncertainty and robust control. Section 4 combines Optimal and Robust Expectations and derives conditions under which they lead to Rational Expectations. Section 5 shows that the results extend naturally to a dynamic, infinite-horizon setting. This is of some independent interest given the fact that Optimal Expectations dissipate in the model of Brunnermeier and Parker (2005). Finally, section 6 concludes by speculating how recent developments in the robust control literature could be used to allow agents to be optimistic over some domains, while at the same time being pessimistic over other domains.

2. Optimal Expectations

The chance of gain is by every man more or less overvalued, and the chance of loss is by most men undervalued, and by scarce any man, $\cdots$, valued more than it is worth. – Smith (1776, p. 120)

Consider an agent who has the following model relating a control action, u, to a state variable, x:

$$x = u + \epsilon \tag{2.1}$$

where ϵ is an exogenous random variable. Suppose the agent has the following preferences,

$$\mathcal{U} = \max_u -\frac{1}{2}\left\{u^2 + E(x-b)^2\right\} \tag{2.2}$$

which reflects a trade-off between control effort and the costs of having the state deviate from its target value, b. As a benchmark, let's first suppose the agent has Rational Expectations. Specifically, suppose the model in eq. (2.1) is known to be the correct model, and furthermore, suppose the agent knows that $\epsilon \sim N(0,1)$. By definition, Rational Expectations means that the subjective expectations appearing in (2.2) are evaluated using this objective distribution. Given all this information, the agent can easily solve for the optimal action, $u = b/2$.

Now suppose, following Brunnermeier and Parker (2005), that the agent values optimism.[4] Left unconstrained, it is clear what Optimal Expectations would be here. In this LQ setting, variance only influences payoffs (negatively), so clearly the agent is going to believe that ϵ is degenerate. Moreover, since effort is costly, he's going to believe that $\epsilon = b$. In this case, no effort is required. He simply chooses $u = 0$, and achieves the bliss point (anticipatory) utility of $\mathcal{U} = 0$. Ex post, of course, realized utility will on average be less than this. It will be $\mathcal{U}^{oe} = -\frac{1}{2}(1 + b^2)$, which in turn is less than expected utility under Rational Expectations, $\mathcal{U}^{re} = -\frac{1}{2}(1 + b^2/2)$. The difference between $\mathcal{U}^{re}$ and $\mathcal{U}^{oe}$ can be interpreted as a 'cost of optimism'. A natural way to define Optimal Expectations is to weigh the ex ante benefits of optimism against these ex post realized costs. This is essentially what Brunnermeier and Parker (2005) do. They compute Optimal Expectations in two steps. First, actions are chosen to maximize 'felicity', which is the discounted sum of past, current, and expected future utility. Optimal actions will in general depend on beliefs about the future. Since current felicity depends on expected future utility, agents have an incentive to be optimistic. Second, agents are assumed to be sophisticated, and are aware that optimistically biased beliefs will yield inferior actions. To account for this, agents are presumed to choose beliefs so as to maximize 'welfare', which is the time-0 expectation of lifetime felicity, computed using the true *objective* probabilities. The idea here is that agents know that outcomes over their lifetimes will in fact be generated by the objective probability distribution. Choosing beliefs to maximize ex ante welfare is a way of constraining the agent's optimism.[5]

Note that in the Brunnermeir-Parker approach, Optimal Expectations are chosen once-and-for-all, at time-0. They are then revised according to Bayes Rule as events unfold. In contrast, I assume Optimal Beliefs are chosen recursively. Rather than constrain optimism in reference to an ex ante notion of 'welfare', which produces dynamically inconsistent beliefs and requires commitment, I assume optimism is constrained by a recursively updated, forward-looking estimate of the (discounted) relative entropy between the agent's subjective beliefs and the objective probabilities. This is especially convenient in Linear-Quadratic-Gaussian environments, since here relative entropy turns out to be closely linked to the decision costs of optimism.[6] It is also convenient since, as we shall see in section 5, it produces stationary, dynamically consistent, decision rules in infinite-horizon settings. In contrast, Brunnermeir-Parker Optimal Expectations converge to Rational Expectations as the horizon grows.[7]

To set the stage for these later developments, let's see how relative entropy costs work in this simple 1-period example. Suppose the agent's subjective beliefs are $\epsilon \sim N(w, \sigma^2)$, rather than the objective distribution, $N(0, 1)$. Denote the $N(0, 1)$ density $f_o(\epsilon)$ and the distorted $N(w, \sigma^2)$ density

---

[4]Brunnermeier and Parker (2005) discuss evidence supporting this assumption. Note that optimism has no instrumental value here. It is valued for its own sake. A separate literature has argued that under certain conditions optimism can have instrumental value in helping the agent to achieve his objectives. See, e.g., Benabou and Tirole (2002) and Compte and Postlewaite (2004).

[5]Note that as defined by Brunnermeir-Parker there can be no meaningful notion of Optimal Expectations in a static 1-period model, since there is no future to be optimistic about. Strictly speaking then, Optimal Expectations would be Rational Expectations in the above example.

[6]As noted by Arrow (1986), when agents have log utility relative entropy captures the cost of having subjective beliefs differ from objective probabilities even in non-Gaussian settings.

[7]Sarver (2011) proposes an alternative approach to obtaining stationary decision rules when agents value optimism in infinite-horizon settings. In his approach, optimism is constrained by loss aversion, rather than by decision costs. Agents take actions using objective probabilities, but value future utility directly. However, agents pay a utility cost if future utility falls short of this ex ante expected value.

$f(\epsilon)$. The relative entropy, $\mathcal{E}(f; f_o)$, between f and f_o is the following expected log-likelihood ratio,

$$\mathcal{E}(f; f_o) = \int \log\left(\frac{f(\epsilon)}{f_o(\epsilon)}\right) f(\epsilon) d\epsilon = \frac{1}{2}\left(w^2 + \sigma^2 - 1 - \log(\sigma^2)\right)$$

This can be interpreted as a measure of the 'distance' between the subjective and objective probability distributions.[8] Note that $\mathcal{E}(f) \geq 0$ and that $\mathcal{E}(f) = 0$ if and only if $w = 0$ and $\sigma^2 = 1$. Also note that in this Gaussian setting, relative entropy is a quadratic function of the mean distortion. As it turns out, decision costs are also a quadratic function of the mean distortion.

To see this, observe that given the subjective beliefs the optimal action is $u = \frac{1}{2}(b - w)$, and anticipated utility is $\mathcal{U} = -\frac{1}{2}\left[\sigma^2 + \frac{1}{2}(b - w)^2\right]$. Realized expected utility, computed using the objective $N(0, 1)$ distribution, is then $\mathcal{U}^{oe} = -\frac{1}{2}[1 + \frac{1}{2}(b^2 + w^2)]$. As noted above, if instead the agent had Rational Expections, his expected utility would be $\mathcal{U}^{re} = -\frac{1}{2}(1 + b^2/2)$. We can now define Optimal Expectations to be the solution of the problem of maximizing ex ante expected utility subject to the constraint $\mathcal{U}^{re} - \mathcal{U}^{oe} \leq \eta$, where η is an exogenous bound on decision costs. Written out, Optimal Expectations is the solution to

$$\max_{w,\sigma^2}\left\{-\frac{1}{2}\left[\sigma^2 + \frac{1}{2}(b - w)^2\right] + \theta\left(\eta - \frac{1}{4}w^2\right)\right\}$$

where θ is a Lagrange multiplier. By inspection, it is clear that $\sigma^2 = 0$. It is also clear that if $\eta \geq \frac{1}{4}b^2$, the constraint is slack and the optimal solution is $w = b$, as we saw before. On the other hand, if $\eta < \frac{1}{4}b^2$ the constraint binds and the optimal solution is $w = 2\sqrt{\eta}$.

As stated, the previous analysis views η as exogenous and θ as endogenous. Following Hansen and Sargent (2008), it proves more convenient to *penalize* decision costs rather than to impose a hard constraint on them. In this case, we can regard the Lagrange multiplier, θ, as an exogenous free parameter. At the same time, since η is exogenous, we can absorb variance distortions into it and just focus on the mean distortion, w. This gives us the following 2-agent problem characterizing Optimal Expectations.

$$\max_u \max_w \left\{-\frac{1}{2}u^2 - \frac{1}{2}(u + w - b)^2 - \frac{1}{4}\theta w^2\right\}$$

Interestingly, in the case of Optimal Expectations, the agent employs the device of an *angelic* agent, who chooses shocks so at to *help* the agent. As in robust control, the angelic agent's actions are penalized by a relative entropy constraint. The first-order conditions are

$$w = \frac{b - u}{\frac{1}{2}\theta + 1} \qquad u = \frac{b - w}{2}$$

Solving these yields the pair of optimal actions and beliefs

$$w = \frac{b}{1 + \theta} \qquad u = \frac{1}{2}b\left(\frac{\theta}{1 + \theta}\right)$$

Note that as decision costs are penalized more heavily, so that $\theta \to \infty$, actions and beliefs converge to those of the Rational Expectations Equilibrium.

One final point is worth noting here before turning to model uncertainty and Robust Expectations. A natural question to ask is whether we are imposing any unjustifiable restrictions on the agent's subjective beliefs by supposing they are Gaussian. Could the agent do better with a more general distribution? It turns out that in LQG environments the answer is no. (See, e.g., Hansen and Sargent (2008) for details).

[8]Strictly speaking, it is not a metric since it is not symmetric, nor does it satisfy the triangle inequality.

3. Robust Expectations

I will suppose therefore $\cdots$ some malicious demon of the utmost power and cunning has employed all his energies in order to deceive me. – Descartes (1641, p. 79)

Now suppose the agent no longer values optimism. Instead, suppose he is not sure the true model is given by eq. (2.1). He merely views it as a useful approximation. To cope with this model uncertainty, suppose he follows the robust control procedures advocated by Hansen and Sargent (2008). He begins by surrounding (2.1) with a cloud of alternative candidate models

$$x = u + (\epsilon + w)$$

Although he knows $\epsilon \sim N(0,1)$, he is presumed to know very little about w. If the agent were a Bayesian, and adhered to the Savage axioms, the presence of w would not be a problem. He would simply formulate a well defined subjective prior distribution over w, and maximize expected utility as usual. In the context of model uncertainty, this strategy is often called 'model averaging'.

The assumption here is that agents *cannot* formulate a conventional finite-dimensional prior over the space of alternative models. In fact, doing so essentially eliminates the problem of model uncertainty altogether, by converting it to a problem of parameter uncertainty.[9] Instead of optimizing against a weighted average of alternative models, robust control is based on the assumption that agents optimize against a single *worst case* model. By doing this, the agent can enforce a lower bound on his payoffs. This makes sense when agents 'distrust' their priors. If instead an agent optimizes against a particular weighted average, he might worry that outcomes would change drastically if actions were based on a different prior.[10]

In general, this worst-case model depends on the decisions of the agent. It is endogenous. To formulate a robust policy, the agent therefore enlists the services of an 'evil agent', who is imagined to be playing a (dynamic) zero-sum game againts the agent. A robust policy is a Nash equilibrium of this game. In contrast to the supernatural powers possessed by Descartes' malicious demon, the evil agent in robust control faces constraints on his actions. As discussed earlier, rather than impose a hard bound on the actions of the evil agent, it is more convenient to penalize them. In general, what is penalized is the relative entropy between the agent's benchmark model and the evil agent's worst-case model. In LQG environments, however, this takes the form of a simple quadratic cost term. Hence, a robust policy turns out to be a Nash equilibrium of the following zero-sum game,

$$\max_u \min_w \left\{ -\frac{1}{2}u^2 - \frac{1}{2}E\left(u + \epsilon + w - b\right)^2 + \frac{1}{2}\theta w^2 \right\}$$

[9]Bayesians are unapologetic about this. For Bayesians there is no meaningful distinction between model uncertainty and parameter uncertainty.

[10]A few comments: (1) Obviously, the idea that agents distrust their priors only makes sense outside of the Savage realm, (2) One could argue that picking the worst case model is equivalent to picking a (degenerate) weighted average, so there is no difference between Bayesian model averaging and robust control (Sims (2001)). See Hansen and Sargent (2008) for a critique of this argument. (3) An attractive aspect of Bayesian decision making is that it possesses axiomatic foundations (courtesy of Savage and others). There are now formal axiomatizations of robust control as well. See, e.g., Strzalecki (2011). (4) It is noteworthy that (good) practicing Bayesians are careful to check the robustness of their conclusions to alternative prior specifications. Bayesians rationalize this practice by claiming they are presenting their results for the benefit of readers with different priors. One might wonder whether they distrust their own priors.

where expectations are computed using the known $N(0,1)$ distribution of ϵ. Following Hansen and Sargent (2008), the penalty parameter, θ, can be calibrated to 'detection error probabilities', so that agents only worry about empirically plausible alternative models.[11]

The first-order conditions are now

$$w = \frac{u-b}{\theta-1} \qquad u = \frac{b-w}{2}$$

Note that for an interior solution to exist, it must be the case that $\theta > 1$. This ensures that the evil agent's minimization problem is strictly convex. Solving these yields a robust control policy and a worst-case model,

$$w = -b\left(\frac{1}{2\theta-1}\right) \qquad u = b\left(\frac{\theta}{2\theta-1}\right)$$

Notice that in contrast to the case of Optimal Expectations, here the agent believes that w is working against him, pushing him away from his target. As before, notice that as $\theta \to \infty$ the agent's policy and beliefs converge to the Rational Expectations Equilibrium.

4. Robust Optimal Expectations

Third, rationality is an assumption that can be modified. Systematic biases, $\cdots$, can be examined with analytical methods based on rationality. – Muth (1961, p. 330)

Finally, consider an agent who values optimism but at the same time has fears of model misspecification, fears which are difficult to summarize in a unique, finite-dimensional prior. How will such an agent behave? He begins by assuming the model takes the following form:

$$x = u + (\epsilon + w_1 + w_2)$$

where as before ϵ is known to be $N(0,1)$. However, now there are imagined to be two endogenously selected disturbances. Doubts about model specification are reflected by w_1, which is chosen by an evil agent. Hopes for the future are reflected by w_2, which is chosen by an angelic agent. These two hypothetical agents operate independently and noncooperatively, each taking the actions of the other as given. The underlying tensions in the agent's environment are depicted in the following zero-sum game,

$$\max_u \max_{w_2} \min_{w_1} \left\{ -\frac{1}{2}u^2 - \frac{1}{2}E\left(u+\epsilon+w_1+w_2-b\right)^2 + \frac{1}{2}\theta_1 w_1^2 - \frac{1}{2}\theta_2 w_2^2 \right\}$$

where θ_1 penalizes the actions of the evil agent, and θ_2 penalizes the actions of the angelic agent. Note that as long as appropriate concavity and convexity conditions are satisfied, the timing of the three agents' actions makes no difference. (This irrelevance follows from the zero-sum structure of the problem). The first-order conditions are,

$$\begin{aligned} 2u + w_1 + w_2 - b &= 0 \\ u + (1-\theta_1)w_1 + w_2 - b &= 0 \\ u + w_1 + (1+\theta_2)w_2 - b &= 0 \end{aligned}$$

The solution can be expressed as follows,

[11] Strzalecki's (2011) recent axiomatization suggests an alternative calibration strategy, based on eliciting certainty equivalents from bets on Ellsberg urns.

$$
\begin{aligned}
u &= \left(\frac{1}{2+\theta_2^{-1}-\theta_1^{-1}}\right) b \\
w_1 &= \left(\frac{-\theta_1^{-1}}{2+\theta_2^{-1}-\theta_1^{-1}}\right) b \\
w_2 &= \left(\frac{\theta_2^{-1}}{2+\theta_2^{-1}-\theta_1^{-1}}\right) b
\end{aligned}
$$

Notice first that as $\theta_1 \to \infty$ and $\theta_2 \to \infty$, the solution collapses to the Rational Expectations solution, as one would expect. Second, observe that the evil agent tries to push the agent away from his target, while at the same time the angelic agent tries to pull him toward it. Third, notice that because the agent is receiving help from the angelic agent, the robust control 'breakdown point' for θ_1 is smaller than before. Rather than $\theta_1 > 1$, we now have the existence condition, $\theta_1 > \theta_2/(1+\theta_2)$. Fourth, and most importantly, notice that the agent might choose the Rational Expectations policy, $u = b/2$, even when he doubts his model, as long as he possesses the right degree of optimism. That is, as long as the knife-edge condition, $\theta_1 = \theta_2$, is satisfied. One might suspect that a finely tuned balance between doubt and hope would be a recipe for schizophrenia. Quite the contrary. The results here suggest 'rationality' might simply reflect a compromise between competing 'irrational' forces.

5. A Dynamic Multivariate Extension

The previous analysis is obviously too simple to be applied. Fortunately, it is straightforward to extend it to dynamic, multivariate settings. The only conceptual subtlety concerns the nature of the penalties imposed on the evil and angelic agents. To ensure stationary, dynamically consistent, decision rules for the actions of the evil and angelic agents, it turns out that we must recast their relative entropy constraints in a recursive, forward-looking manner, so that they resemble the sort of promise-keeping constraints that arise in the recursive contracts literature. This implies that time-t increments to relative entropy must be discounted at the same rate the agent discounts increments to his utility. Otherwise, the evil and angelic agents would want to front-load their distortions in a way that would cause optimism and specification doubts to dissipate over time. Writing the relative entropy constraints in this way also effectively constrains the ability of the evil/angelic twins to reconsider their earlier planned distortions in response to subsequently unrealized states. Otherwise, as events unfold, they could 'save entropy' by reallocating it away from unrealized states. (See Hansen and Sargent (2008) for a detailed discussion).

The agent's model now takes the form:

$$x_{t+1} = Ax_t + Bu_t + C(\epsilon_{t+1} + w_{1,t} + w_{2,t}) \tag{5.3}$$

where x_t is a state vector, u_t is a control vector, ϵ_{t+1} is an exogenous disturbance with a known $N(0,1)$ distribution, and $(w_{1,t}, w_{2,t})$ are endogenous disturbance vectors chosen by the evil and angelic agents. A, B, and C are appropriately dimensioned coefficient matrices. The Bellman equation for the agent's dynamic optimization problem takes the following form[12]

[12] Two comments: (1) I have taken advantage of a modified Certainty Equivalence principle to eliminate stochastic elements from (5.4) and the resulting constant term. The constant term in the value function does not influence behavior, although it would be relevant for welfare calculations, (2) For notational simplicity I have omitted cross-products from the utility function and normalized the targets to zero. This is without loss of generality. (See Hansen and Sargent (2008) for a discussion of both these features).

$$-x'Px = \max_u \max_{w_2} \min_{w_1} \{-(x'Qx + u'Ru) + \beta\theta_1 w_1' w_1 - \beta\theta_2 w_2' w_2 - \beta x^{*\prime} P x^*\} \tag{5.4}$$

where x^* denotes next period's state, and P is a positive definite matrix representing the value function. The first-order conditions for the inner two extrema problems are

$$\begin{aligned} (\theta_1 - C'PC)w_1 - C'PCw_2 &= C'P(Ax + Bu) \\ -(\theta_2 + C'PC)w_2 - C'PCw_1 &= C'P(Ax + Bu) \end{aligned}$$

Solving these yields the following two reaction functions,

$$\begin{aligned} w_1 &= \theta_1^{-1}[I - (\theta_1^{-1} - \theta_2^{-1})C'PC]^{-1}C'P(Ax + Bu) \\ w_2 &= -\theta_2^{-1}[I - (\theta_1^{-1} - \theta_2^{-1})C'PC]^{-1}C'P(Ax + Bu) \end{aligned}$$

If we then substitute these back into (5.4) we get the following 'indirect' Bellman equation,

$$-x'Px = \max_u \{-(x'Qx + u'Ru) - \beta(Ax + Bu)'\mathcal{D}(P)(Ax + Bu)\} \tag{5.5}$$

where the nonlinear operator $\mathcal{D}(P)$ is given by

$$\mathcal{D}(P) = P + (\theta_1^{-1} - \theta_2^{-1})PC[I - (\theta_1^{-1} - \theta_2^{-1})C'PC]^{-1}C'P$$

Notice that the Bellman equation in (5.5) assumes a distortionless state transition equation. The agent's conflicting tensions between doubt and hope are fully captured by the operator $\mathcal{D}(P)$. If $\theta_1 < \theta_2$, the value function undergoes a pessimistic adjustment, whereas if $\theta_1 > \theta_2$ the value function is adjusted optimistically.

Finally, if we now define the conventional Bellman operator as follows,

$$T(P) = Q + \beta A'PA - \beta^2 A'PB(R + \beta B'PB)^{-1}B'PA$$

we can represent the agent's value function as the fixed point of the following composite operator, $P = T \circ \mathcal{D}(P)$. The existence condition is that $I - (\theta_1^{-1} - \theta_2^{-1})C'PC$ be positive definite. If this condition is satisfied, then P can be computed by iterating until convergence on $P_{j+1} = T \circ \mathcal{D}(P_j)$, starting from $P_0 = 0$ (making sure at each step that the existence condition remains satisfied).

Once we have P, we can then write the Markov perfect policy functions as $u_t = -Fx_t$ and $w_{i,t} = K_i x_t$, where the feedback matrices are given by

$$\begin{aligned} F &= \beta(R + \beta B'\mathcal{D}(P)B)^{-1}B'\mathcal{D}(P)A \\ K_i &= \theta_i^{-1}[I - (\theta_1^{-1} - \theta_2^{-1})C'PC]^{-1}C'P(A - BF) \end{aligned}$$

Once again, as $\theta_1 \to \infty$ and $\theta_2 \to \infty$, the solution converges to the Rational Expectations solution, since in this case $\mathcal{D}(P) \to P$. More interestingly, notice that this is also the case when $\theta_1 = \theta_2$. Thus, the agent appears to adhere to the Rational Expectations Hypothesis when doubts about model specification are exactly offset by hopes for the future.

More generally, it is clear by inspection that only $(\theta_1 - \theta_2)$ is identified from observed behavior. In fact, the identification problem is even worse than this. First, we know from Hansen and Sargent (2008, chpt. 10) that in this class problems robustness induces a form of precautionary savings that can be mimicked by an increase in β. At the same time, Brunnermeier and Parker (2005) note that Optimal Expectations induce a downward drift in consumption, which can be mimicked by a decrease in β. When combined, an agent with $\theta_1 > \theta_2$ would be observationally equivalent to a Rational Expectations agent with a smaller β. Conversely, an agent with $\theta_1 < \theta_2$ would be observationally equivalent to a Rational Expectations agent with a larger β. However, Hansen and Sargent (2008) go on to show that identification can be achieved by looking at asset prices. Robustness has effects on risky asset returns that cannot be replicated by changing the rate of time preference.

A second source of identification problems arise once we relax the highly restrictive/symmetric manner in which the w_1 and w_2 disturbances enter the benchmark model in (5.3). As written, the model presumes the agent has equal degrees of optimism and pessimism in all parts of his model. This seems overly restrictive. In practice, agents are likely to doubt some parts of their models more than others, while at the same time having greater optimism over some aspects of the future than others. In fact, Haselton and Nettle (2006) provide evidence that agents appear to be pessimistically biased in some domains, while at the same time appearing to be be optimistically biased in others. They argue that evolution can explain this. Sometimes evolutionary forces select in favor of pessimism, while in others optimism is favorably selected. In most of their examples, fitness is defined biologically, as reproductive success, but the point seems more general. Allowing agents to have differing degrees of doubt about different aspects of their models is also at the forefront of the robust control literature. For example, Hansen and Sargent (2011) discuss problems of robust control and filtering, where one distortion operator (T_1) can be used to express doubts about model specification, while a different distortion operator (T_2) can be used to express doubts about which of a finite collection of models is actually generating the data.

Although Hansen and Sargent's T_1/T_2 operators could likely be adapted to the case of Robust Optimal Expectations, I will instead use a more structured example to illustrate the identification difficulties that arise in asymmetric specifications. Suppose x_t is n-dimensional, and we partition the model's n equations into n_1 'doubtful' equations and n_2 'hopeful' equations, where $n = n_1+n_2$. If we place the doubtful equations first, we can write the model as follows

$$x_{t+1} = Ax_t + Bu_t + C\epsilon_{t+1} + \tilde{I}_1 w_{1,t} + \tilde{I}_2 w_{2,t}$$

where now $w_{1,t}$ is an $n_1 \times 1$ vector of pessimistic distortions, $w_{2,t}$ is an $n_2 \times 1$ vector of optimistic distortions, $\tilde{I}_1$ is an $n \times n_1$ matrix consisting of an $n_1 \times n_1$ identity matrix stacked onto n_2 rows of zeros, and $\tilde{I}_2$ is an $n \times n_2$ matrix consisting of n_1 rows of zeros stacked on top of a $n_2 \times n_2$ identity matrix. The agent's Bellman equation now becomes

$$-x'Px = \max_u \max_{w_2} \min_{w_1} \{-(x'Qx + u'Ru) + \beta\theta_1 w_1'\Theta_1 w_1 - \beta\theta_2 w_2'\Theta_2 w_2 - \beta x^{*\prime}Px^*\}$$

where Θ_1 is an $n_1 \times n_1$ diagonal matrix with equation specific penalty terms applied to the evil agent, and Θ_2 is an $n_2 \times n_2$ diagonal matrix of angelic agent penalties. Optimizing over w_1 and w_2 implies the following,

$$\tilde{I}_1 w_{1,t} + \tilde{I}_2 w_{2,t} = (\tilde{I}_1, \tilde{I}_2)[\Theta\tilde{I} - P]^{-1}P(Ax + Bu)$$

where Θ is an $n \times n$ diagonal matrix formed from Θ_1 and Θ_2, and $\tilde{I}$ is an $n \times n$ block diagonal matrix with an $n_1 \times n_1$ diagonal matrix in the top left, and a negative $n_2 \times n_2$ diagonal matrix on the bottom right. Evidently, if the penalty parameters satisfy

$$(\tilde{I}_1, \tilde{I}_2)[\Theta\tilde{I} - P]^{-1} = 0$$

where P is the solution of the Rational Expectations Riccati equation, then the agent will adhere to the Rational Expectations Hypothesis, despite his varying degrees of doubt and hope. What's going on here is that optimism and pessimism are interacting in just the right way across equations so as to nullify each other within each equation. Clearly, this is a highly non-generic example, but it does illustrate the new possibilities that arise once we entertain the possibility that doubts and hopes can interact in a general, unrestricted manner.

6. Conclusion

Friedman (1953) famously argued that we should not worry about a theory's assumptions. Lucas (1986, p. 241) as well was not sympathetic to arguments claiming that the Rational Expectations Hypothesis is based on implausible assumptions:

> *To observe that economics is based on a superficial view of individual and social behavior does not $\cdots$ seem to me to be much of an insight. I think it is exactly this superficiality that gives economics much of the power that it has: its ability to predict human behavior without knowing very much about the makeup and the lives of the people whose behavior we are trying to understand.*

This paper has been devoted to relaxing the assumptions of the Rational Expectations Hypothesis, so one might reasonably ask whether it's been worth the effort. By revealed effort, I think it has. In the very next sentence following the above quote, Lucas states: "*Yet an ability such as this necessarily has its limits, and I have spent most of this essay on cases that lie close to these limits,* $\cdots$". The cases to which he was referring concerned questions where equilibrium reasoning by itself was insufficient to make unambiguous predictions, due to the presence of multiple equilibria. Rational Expectations models have also encountered limits along other dimensions, particularly with respect to asset market phenomena, and the robust control literature has already demonstrated the benefits of a disciplined retreat from the Rational Expectations Hypothesis, even for those who follow Friedman's advice (see Hansen and Sargent (2011) for a survey). I would argue that incorporating elements of hope and optimism can yield additional empirical payoffs (see Brunnermeier and Parker (2005) for a catalogue of examples).

Another contribution of this essay is to sound a note of caution concerning the interpretation of experimental results in economics. Lab experiments are useful for isolating specific aspects of human behavior, and there have been many studies claiming that individuals are optimistic and over-confident, and an equal number of studies claiming that individuals are pessimistic and ambiguity-averse. Of course, these studies need not be contradictory. In fact, this paper shows how both phenomena could be captured within a single model. However, the main result here potentially calls into question the external validity of experimental results that control for too much. Optimism and pessimism that may be present within certain ceteris paribus domains can interact in ways that make individuals behave 'rationally'.

One question that might have occurred to some readers is whether the results here are driven by the specialized Linear-Quadratic-Gaussian environment. In a way they are, and in a way they aren't. The tight connection between decision costs and relative entropy is certainly dependent on the LQG structure. So is the (modified) Certainty Equivalence result that allowed us to focus on mean distortions and abstract from distortions in higher moments. However, I would argue that the key result concerning the interactions between doubt and hope are far more general. One strategy for generalizing the results would be to consider a continuous-time version. Here relative entropy constraints necessarily consist of only mean distortions, even in general nonlinear models. However, normality in some sense remains important, since the underlying model must consist of diffusion processes.

References

Akerlof, G. A., and W. T. Dickens (1982): "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, 72, 307–319.

Arrow, K. J. (1986): "The Value of and Demand for Information," in *Decision and Organization*, ed. by C. McGuire, and R. Radner, pp. 131–139. University of Minnesota Press, 2 edn.

Benabou, R., and J. Tirole (2002): "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 117, 871–915.

Blume, L., and D. Easley (1998): "Rational Expectations and Rational Learning," in *Organizations with Incomplete Information*, ed. by M. Majumdar, pp. 61–109. Cambridge University Press.

——— (2006): "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets," *Econometrica*, 74, 929–66.

Brunnermeier, M. K., and J. A. Parker (2005): "Optimal Expectations," *American Economic Review*, 95, 1092–1118.

Compte, O., and A. Postlewaite (2004): "Confidence-Enhanced Performance," *American Economic Review*, 94, 1536–57.

Descartes, R. (1998): "Meditations on First Philosophy, I," in *Descartes: Selected Philosophical Writings*, ed. by J. Cottingham, and R. Stoothoff, pp. 80–86. Cambridge University Press.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

Friedman, M. (1953): "The Methodology of Positive Economics," in *Essays in Positive Economics*, ed. by M. Friedman, pp. 3–43. University of Chicago Press.

Hansen, L. P., and T. J. Sargent (2008): *Robustness*. Princeton University Press, Princeton.

——— (2011): "Wanting Robustness in Macroeconomics," in *Handbook of Monetary Economics, 3B*, ed. by B. M. Friedman, and M. Woodford, pp. 1097–1157. North-Holland.

Haselton, M. G., and D. Nettle (2006): "The Paranoid Optimist: An Integrative Evolutionary Model of Cognitive Biases," *Personality and Social Psychology Review*, 10, 47–66.

Hume, D. (1739): *A Treatise of Human Nature*. Penguin Classics (1985 Edition).

Lucas, R. E. (1986): "Adaptive Behavior and Economic Theory," in *Rational Choice: The Contrast Between Economics and Psychology*, ed. by R. M. Hogarth, and M. W. Reder, pp. 217–242. University of Chicago Press.

Muth, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29, 315–35.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press.

——— (2008): "Evolution and Intelligent Design," *American Economic Review*, 98, 5–37.

Sarver, T. (2011): "Optimal Anticipation," mimeo, Northwestern University.

Sims, C. A. (2001): "Pitfalls of a Minimax Approach to Model Uncertainty," *American Economic Review*, 91, 51–54.

Smith, A. (1776): *An Inquiry into the Nature and Causes of The Wealth of Nations*. University of Chicago Press (1976 Edition).

Strzalecki, T. (2011): "Axiomatic Foundations of Multiplier Preferences," *Econometrica*, 79, 47–73.

Kenneth Kasa
Department of Economics
Simon Fraser University
8888 University Drive
Burnaby, BC, V5A 1S6 CANADA
email: kkasa@sfu.ca