Journal of Economic Dynamics & Control 42 (2014) 1–12

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

An escape time interpretation of robust control

In-Koo Cho [a,c], Kenneth Kasa [b,*]

[a] Department of Economics, University of Illinois, Urbana, IL 61801, USA
[b] Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, BC, Canada V5A 1S6
[c] College of Economics and Finance, Hanyang University, Seongdong-gu, Seoul 133-1791, Korea

ARTICLE INFO

Article history:
Received 28 August 2013
Received in revised form
23 February 2014
Accepted 28 February 2014
Available online 12 March 2014

JEL classification:
C61
D81

Keywords:
Robust control
Large deviations

ABSTRACT

This paper studies the problem of an agent who wants to prevent the state from exceeding a critical threshold. Even though the agent is presumed to know the model, the optimal policy is computed by solving a conventional robust control problem. That is, robustness is induced here by objectives rather than uncertainty, and so is an example of the duality between risk-sensitivity and robustness. However, here the agent only incurs costs upon escape to a critical region, not during 'normal times'. We argue that this is often a more realistic model of macroeconomic policymaking.

1. Introduction

Robust control methods have become popular in economics.[1] This popularity stems from their ability to provide convenient formalizations of ambiguity and Knightian Uncertainty.[2] The presumption in typical applications of robust control is that agents want to maximize expected utility, but confront forms of model uncertainty that are difficult to capture with conventional finite-dimensional Bayesian priors. In response, agents construct a *set* of plausible priors, and optimize against the worst-case prior. This can be interpreted as maximizing expected utility with respect to a pessimistically distorted prior.

This paper argues that robust control methods are useful even when agents know the correct model, or are able to construct a unique prior about it. In many policy settings, what policymakers really care about are extreme events, e.g., crises or 'market meltdowns'. Using results from Dupuis and McEneaney (1997), we show that robust control policies can be interpreted as minimizing the expectation of an exponential function of the escape time to a critical loss threshold. This implies agents are especially averse to rapid escapes. The fact that expectations are evaluated using the true (undistorted) probability measure implies agents trust their models and confront no model uncertainty. Instead, robust policies are useful because they reduce the likelihood that a known and trusted system will rapidly hit an undesirable threshold.[3]

* Corresponding author. Tel.: +1 778 782 5406.
E-mail addresses: inkoocho@uiuc.edu (I.-K. Cho), kkasa@sfu.ca (K. Kasa).
URL: http://www.sfu.ca/kkasa (K. Kasa).

[1] See Hansen and Sargent (2008) for a textbook treatment. For an up-to-date survey see Hansen and Sargent (2011).
[2] Strzalecki (2011) provides axiomatic foundations linking robust control, ambiguity, and Knightian Uncertainty.
[3] This class of problems has been studied in the engineering literature. Meerkov and Runolfsson (1988) and Clark and Vinter (2012) provide examples.

http://dx.doi.org/10.1016/j.jedc.2014.02.014
0165-1889/

That robust control, motivated by model uncertainty, can be reinterpreted as an aspect of preferences, motivated by 'risk-sensitivity', is already well known. Mathematically, this equivalence derives from a Legendre transform duality between the two.[4] What is new here is our specification of the costs experienced by our risk-sensitive agent. Normally, these costs are determined by the instantaneous loss function of the agent, which depends on the current values of the state and control variables. In contrast, we suppose our agent experiences no state costs whatsoever during 'normal times', i.e., when instantaneous losses are below the threshold. What matters to him is the frequency of disasters. Despite this change in the cost function, we show there is still a connection between robustness and risk-sensitivity. However, from a mathematical standpoint, our result is based more on the Feynman–Kac formula than on the Legendre transform.

We make no effort here to derive this boundary cost objective function from first principles. A natural question, of course, is why the policymaker simply does not maximize the welfare of the agents who elected or appointed him. One possibility is that policymakers lack the sort of detailed information on individual preferences that would allow them to perform this sort of optimization, and so resort to simpler and more robust rules of thumb, like avoiding disasters. Another possibility is that agents themselves experience 'zones of indifference', perhaps based on information processing constraints and 'rational inattention', which make them primarily sensitive to thresholds. However, we leave the exploration of these possibilities to future research. It should be noted, however, that our analysis does not *require* the assumption that costs only depend on hitting times. Neglecting running costs simply makes the math easier, and in our opinion, makes the demonstrated correspondence more interesting.

The remainder of the paper is organized as follows. Section 2 lays out a conventional continuous-time Linear-Quadratic Regulator (LQR) problem, and shows that the optimally controlled state evolves as an Ornstein–Uhlenbeck process. Section 3 does the same thing for the robust control case. Section 4 first compares the stationary and hitting-time distributions implied by conventional and robust control policies. The stationary distributions are both Gaussian, but robust policies produce distributions with a smaller variance. Using results from Day (1983), we show that first hitting-time densities are asymptotically exponential (in the small noise limit), and that robust policies produce densities that are skewed toward longer escape times. We then discuss the sense in which conventional robust control policies are implied by an objective that penalizes rapid hitting times. Section 5 applies the results to an example in which a firm attempts to maximize the expected present value of dividend payments out of a stochastic net cash flow process subject to a bankruptcy constraint. We show that an increase in bankruptcy costs and an increased preference for robustness both lead to an increase in the dividend payment threshold. Section 6 concludes and offers a few suggestions for future research.

2. The conventional LQR

Given our focus on hitting and escape time issues, it proves to be convenient to work in continuous time. So consider an agent with the following objective function:

$$V(x_0) = \min_u E_0 \int_0^\infty \frac{1}{2}\left[x(t)^2 + \lambda u(t)^2\right] e^{-\rho t}\, dt \tag{2.1}$$

where $x(t)$ is the state, $u(t)$ is a control variable, and λ is a parameter representing the relative cost of control. Without loss of generality, we assume that $x(t)$ and $u(t)$ are both scalars. The state transition equation is given by the following linear diffusion process:

$$dx = (-ax + bu)\, dt + \sigma\, dW \tag{2.2}$$

where $W(t)$ is a Brownian motion process. Employing Ito's Lemma, the Hamilton–Jacobi–Bellman (HJB) equation for this problem is given by the following second-order ordinary differential equation:

$$\rho V(x) = \min_u \left\{ \frac{1}{2}\left[x^2 + \lambda u^2\right] + V'(x)(-ax + bu) + \frac{1}{2}\sigma^2 V''(x) \right\} \tag{2.3}$$

which yields the optimal policy function, $u(x) = -(b/\lambda)V'(x)$. Substituting this back in, and conjecturing that $V(x) = \frac{1}{2}Px^2 + C$ produces the following algebraic Riccati equation for P

$$\rho P = 1 + \frac{b^2}{\lambda}P^2 - 2P\left(a + \frac{b^2}{\lambda}P\right) \tag{2.4}$$

which has the following unique positive solution[5]

$$P = \frac{-[\rho + 2a] + \sqrt{[\rho + 2a]^2 + 4b^2\lambda^{-1}}}{2b^2\lambda^{-1}} \tag{2.5}$$

[4] Hansen and Sargent (2008) provide a detailed discussion of the connections between robustness and risk-sensitivity.

[5] Given P, one can readily verify that the constant is given by $C = \sigma^2 P/(2\rho)$. Note, this constant does not influence behavior.