OPERATIONS RESEARCH
Vol. 61, No. 4, July–August 2013, pp. 874–893
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2013.1180

CONTEXTUAL AREAS

Robust Portfolio Control with Stochastic Factor Dynamics

Paul Glasserman
Columbia Business School, Columbia University, New York, New York 10027, pg20@columbia.edu

Xingbo Xu
Industrial Engineering and Operations Research Department, Columbia University, New York, New York 10027, xx2126@columbia.edu

Portfolio selection is vulnerable to the error-amplifying effects of combining optimization with statistical estimation and model error. For dynamic portfolio control, sources of model error include the evolution of market factors and the influence of these factors on asset returns. We develop portfolio control rules that are robust to this type of uncertainty, applying a stochastic notion of robustness to uncertainty in model dynamics. In this stochastic formulation, robustness reflects uncertainty about the probability law generating market data, and not just uncertainty about model parameters. We analyze both finite- and infinite-horizon problems in a model in which returns are driven by factors that evolve stochastically. The model incorporates transaction costs and leads to simple and tractable optimal robust controls for multiple assets. We illustrate the performance of the controls on historical data. As one would expect, in-sample tests show no evidence of improved performance through robustness—evaluating performance on the same data used to estimate a model leaves no room to capture model uncertainty. However, robustness does improve performance in out-of-sample tests in which the model is estimated on a rolling window of data and then applied over a subsequent time period. By acknowledging uncertainty in the estimated model, the robust rules lead to less aggressive trading and are less sensitive to sharp moves in underlying prices.

Subject classifications: portfolio optimization; robustness; transaction costs; relative entropy.
Area of review: Financial Engineering.
History: Received November 2011; revisions received October 2012, January 2013; accepted March 2013. Published online in *Articles in Advance* July 31, 2013.

1. Introduction

Portfolio optimization, like most problems of optimization that rely on estimated quantities, is vulnerable to the error-amplifying effects of combining optimization with estimation. Any reasonable estimation procedure applied to multiple assets will overestimate the expected returns of some assets or underestimate their risk, and an optimization procedure that ignores this fact will drive a portfolio to overinvest in precisely these assets. Dynamic portfolio control introduces a further complication by requiring a model of the evolution of asset prices. Any practical model is likely to be misspecified, in addition to being subject to estimation error, and optimization will again amplify the effects of error, in this case, model error.

The goal of *robust* methods, as the term is used in the optimization and control literature, is to account for parameter or model uncertainty in the optimization procedure. Whereas Bayesian methods put a probability distribution on model parameters, robust optimization methods typically posit ranges over which parameters may vary (uncertainty sets) and then optimize against the worst-case configuration consistent with these ranges. This formulation leads to a min–max or max–min optimization problem. This approach to robustness has been developed extensively in the portfolio optimization setting in, e.g., Ben-Tal et al. (2000), Bertsimas and Pachamanova (2008), Goldfarb and Iyengar (2003), and, from a different perspective, in Lim et al. (2011). Robustness to uncertainty over a set of distributions in portfolio optimization is analyzed in, e.g., El Ghaoui et al. (2003), Natarajan et al. (2008, 2010), Delage and Ye (2010), and Goh and Sim (2010), primarily in single-period formulations.

Here we use a stochastic notion of robustness that allows model uncertainty in the law of evolution of the stochastic inputs to a model. This paper combines this approach to robustness with the following features: We study multiperiod (finite- and infinite-horizon) portfolio control problems. We work with a model in which returns are driven by factors that evolve stochastically; both the relationship between returns and factors and the evolution of the factors are subject to model error, and thus treated robustly. We incorporate transaction costs. We develop simple optimal controls that remain tractable for multiple assets. We demonstrate performance both in sample and out of sample on historical data.

It is convenient (and customary) to interpret the minimization in a max–min robust optimization problem as the work of a hypothetical adversary who chooses the least favorable set of parameters allowed by the uncertainty sets. In a stochastic formulation of robustness, the adversary can perturb the law of the stochastic inputs to the model, and not just a set of parameters. To get a sensible formulation, we must either constrain or penalize the magnitude of the perturbation; otherwise, the adversary's effect—reflecting the degree of model uncertainty—would be so great as to render optimization useless. This approach thus requires a way of quantifying deviations from a nominal or baseline model. Relative entropy turns out to be a particularly convenient way to measure the degree of change or error in the law of the stochastic inputs. This notion of stochastic robustness was introduced to the optimal control literature in Petersen et al. (2000) under a total cost criterion. Hansen and Sargent (2007) give a comprehensive treatment of the discounted case and expand the theory and applications of this approach in several directions; their monograph provides the main tool for the formulation of robustness we use, although, as we will see, the factor structure of asset returns leads to a model outside the scope of their results. Lim and Shanthikumar (2007) apply this notion of stochastic robustness in a point process model of pricing. El Ghaoui et al. (2003) apply a relative entropy constraint in a static setting.

For the factor model and factor dynamics, we start from the (nonrobust) model of Gârleanu and Pedersen (2013). Their model uses linear dynamics and a quadratic objective to achieve tractability with considerable flexibility and generality that lends itself to further study. Their analysis, motivated by realistic trading strategies, focuses on the impact of the speed of mean reversion in factor dynamics and how this affects portfolio control and, ultimately, equilibrium asset prices. By building on their framework, we retain a high degree of tractability, and we can study the effect of robustness in a current and independent model, rather than in a model introduced specifically for the comparison. As a by-product, we can also see the effect of model uncertainty on factor dynamics and the factor model of returns: the adversary in the robust formulation can perturb both, and the adversary's optimal choice points to the ways in which the investor is most vulnerable to model error. We test our portfolio rules on the same commodity futures as Gârleanu and Pedersen (2013). Briefly, we find that robustness leads to better performance in out-of-sample tests in which the model is reestimated on a rolling window; robust rules guard against model and estimation error by trading less aggressively on signals from the factors.

Maenhout (2004) extended the continuous-time formulation of robustness of Anderson et al. (2003) to solve a robust portfolio control problem and obtained closed-form solutions for an investor with power utility. His model has a single risky asset, a constant expected return, and no transaction costs; his focus is primarily on the impact of uncertainty on the aggregate demand for the risky asset and less on practical portfolio construction. Maenhout (2004) develops a connection between robustness and stochastic differential utility in the sense of Duffie and Epstein (1992), a connection also examined, from various perspectives, in Chen and Epstein (2003), Liu (2010), Skiadas (2003), and Uppal and Wang (2003). A parallel in the discrete-time setting is the emergence of a discounted recursive risk-sensitive control problem of the type studied by Hansen and Sargent (1995). (The total cost robust problem of Petersen et al. 2000 leads to an undiscounted risk-sensitive control problem of the type introduced and studied by Jacobson 1973 and Whittle 1981, Whittle and Whittle 1990.) A line of work that includes Bielecki and Pliska (1999), Bielecki et al. (2005), and Fleming and Sheu (2001) starts from a risk-sensitive objective and solves for the resulting optimal portfolio control. In contrast, in our setting the risk-sensitive control problem arises as a step in the solution to the robust problem.

In §6, we illustrate the performance of the controls on historical data. The robust controls do improve performance in out-of-sample tests. By acknowledging uncertainty in the estimated model, the robust rules lead to less aggressive trading and are less sensitive to sharp moves in underlying prices. We also compare the robust controls with three heuristic scaling and trimming methods similar to, e.g., DeMiguel et al. (2009). We find that risk scaling, as introduced in details in §6.4, is similar to isolating robustness in certain random sources.

The rest of this paper is organized as follows. Section 2 formulates the basic portfolio control problem and its robust extension. Section 3 solves the finite-horizon problem. Section 4 examines the effect of varying the degree of robustness and compares robust and nonrobust solutions. Section 5 solves the infinite-horizon control problem, and §6 presents numerical results. Most proofs are collected in the online e-companion (available as supplemental material at http://dx.doi.org/10.1287/opre.2013.1180).

2. Problem Formulation

2.1. Dynamics and Objective

We consider a portfolio optimization problem in which asset returns are driven by factors with stochastic dynamics. Examples of portfolio control problems with factor models of returns include, among many others, the work of Bielecki and Pliska (1999), Campbell and Viceira (2002), and Pesaran and Timmermann (2012). The formulation in Gârleanu and Pedersen (2013), which we now review, leads to particularly explicit solutions in both its original and robust form.

The investor has access to n_x underlying assets evolving in discrete time. The changes in prices of the assets from time t to time $t+1$ are indicated by a vector $r_{t+1} \in \mathbb{R}^{n_x}$, specified by

$$r_{t+1} = \mu + Bf_t + u_{t+1}, \tag{1}$$

where $\mu \in \mathbb{R}^{n_x}$ represents an expected or "fair" return, $f_t \in \mathbb{R}^{n_f}$ is a vector of factors influencing price changes and known to the investor at time t, $B \in \mathbb{R}^{n_x \times n_f}$ is a factor-loading matrix, and $u_1, u_2, \ldots$, are i.i.d. random vectors in $\mathbb{R}^{n_x}$ following a multivariate normal distribution with mean zero and covariance matrix Σ_u. The factors are mean reverting, evolving according to the equation

$$f_{t+1} = Cf_t + v_{t+1}, \tag{2}$$

with coefficient matrix $C \in \mathbb{R}^{n_f \times n_f}$ and i.i.d. noise vectors $v_1, v_2, \ldots$ in $\mathbb{R}^{n_f}$ following a multivariate normal distribution with mean zero and covariance matrix Σ_v. We assume that the v_t are independent of the u_t. To make the factors stable, we assume throughout that $\sigma(C) < 1$, where $\sigma(\cdot)$ gives the spectral radius of a square matrix. Equations (1) and (2) allow the possibility that prices become negative. However, we measure performance based on price changes, so this does not present a problem. The probability of this occurring can also be made very small through parameter choices.

Denote by $x_t \in \mathbb{R}^{n_x}$ the vector of shares of underlying assets held in the portfolio just after any transactions made at time t; in other words, at time t the portfolio's holdings are rebalanced from x_{t-1} to x_t. Rebalancing the portfolio imposes transactions costs modeled as

$$\tfrac{1}{2}\Delta x_t^\top \Lambda \Delta x_t, \quad \text{where } \Delta x_t = x_t - x_{t-1}, \tag{3}$$

with cost matrix Λ symmetric and positive definite. For a square matrix A, we write $A > 0$ (or $\geqslant$, $<$, $\leqslant$) if A is positive definite (or positive semidefinite, negative definite, negative semidefinite, respectively). If a small transaction of dx shares temporarily moves the market price unfavorably by the amount Λdx, then a transaction of size Δx_t results in a total cost of $\Delta x_t^\top \Lambda \Delta x_t / 2$, compared to executing the transaction at the original price. The simple model penalizes large trades and provides tractability.

In a mild abuse of notation, we use Δx to denote an investment *policy*—that is, a rule for determining transactions given the information available. With this convention, we write the objective introduced by Gârleanu and Pedersen (2013), which models an investor seeking to maximize the present value of risk-adjusted excess gains, net of transaction costs, as follows:

$$\begin{aligned}
&\sup_{\Delta x} E\Bigg[\sum_{t=0}^{T}\beta^t\bigg((x_{t-1}+\Delta x_t)^\top(r_{t+1}-\mu) \\
&\qquad\qquad -\frac{\gamma}{2}\mathrm{Var}_t[(x_{t-1}+\Delta x_t)^\top r_{t+1}] - \frac{1}{2}\Delta x_t^\top \Lambda \Delta x_t\bigg)\Bigg] \\
&\quad = \sup_{\Delta x} E\Bigg[\sum_{t=0}^{T}\beta^t\bigg(x_t^\top(Bf_t+u_{t+1}) - \frac{\gamma}{2}x_t^\top \Sigma_u x_t \\
&\qquad\qquad - \frac{1}{2}\Delta x_t^\top \Lambda \Delta x_t\bigg)\Bigg].
\end{aligned} \tag{4}$$

The objective (4) consists of three terms. The first term is the discounted sum of future excess returns with discount factor $\beta \in (0, 1)$. The third term measures discounted transaction costs. The difference between these two terms measures the discounted net cash flow to the investor. The middle term is a risk penalty, in which $\gamma > 0$ measures the investor's risk aversion. The notation Var_t denotes conditional variance of excess return, given information up to time t, including position x_t, and we use E_t analogously. Measuring risk through the expectation of the discounted sum of conditional variances is a compromise made for tractability, as is the case with the quadratic measure of transaction costs. This also makes the objective time consistent, which is not typically true for dynamic mean-variance problems (e.g., Basak and Chabakauri 2010, Ruszczyński 2010). Interestingly, even if we drop the risk penalty (setting $\gamma = 0$), a term of exactly this form appears in the solution to a robust formulation. We view (4) as a guide to selecting sensible strategies rather than as a precise representation of an investor's preference. In our numerical tests, we therefore evaluate performance through a Sharpe ratio as well as directly through (4).

Given the Markov structure of the problem, it suffices for the investor to consider policies under which Δx_t is a deterministic function of (x_{t-1}, f_t), and the supremum in (4) is taken over such policies. In choosing an optimal policy, the investor must, as usual, balance risk and reward. In addition, the combination of the factor structure in (1) and the mean reversion in (2) requires the investor to balance the benefits of acting on a signal before the factors decay against the costs of large transactions.

Remark 1. Our model of transaction costs can be generalized to incorporate more features while preserving tractability. One generalization is to incorporate permanent price impact by adding $x_t^\top \tilde{\Lambda} \Delta x_t$ to the transaction costs, where $\tilde{\Lambda}\Delta x_t$ is the permanent price impact caused by transaction Δx_t. Because this term is linear in Δx_t, we still get an explicit iteration similar to Proposition 3. Moallemi and Saglam (2012) consider more general models with transaction costs and more general performance objectives than (4); they forgo explicit solutions and instead optimize numerically within the class of linear rebalancing rules.

2.2. Robust Formulation

We now introduce model uncertainty by allowing perturbations in the stochastic dynamics of the model. The stochastic input to the model is the sequence $\{(u_t, v_t), t = 1, 2, \ldots\}$ of noise terms, so we will translate uncertainty about the model into uncertainty about the law of this sequence. As is usually the case in discussions of robustness, it is convenient to describe uncertainty through the possible actions of a hypothetical adversary who changes the model to maximize harm to the original agent—in our setting, the investor. We will constrain the actions of the adversary shortly, but first we briefly illustrate the effect the adversary can have by changing the law of the noise sequence.

The noise vectors all have mean zero under the original model. If the adversary changes the conditional mean of

v_{t+1} to $-Df_t$, for some $n_f \times n_f$ matrix D, then (2) becomes

$$f_{t+1} = (C-D)f_t + \tilde{v}_{t+1},$$

where $\tilde{v}_{t+1}$ has the law of the original v_{t+1}. Thus, the adversary can change the dynamics of the factors and, for example, accelerate the speed of mean reversion and potentially reduce the value of the factors to the investor. By instead setting the conditional mean of v_{t+1} to $(I-C)\bar{f}$, for some fixed $\bar{f}$, the adversary moves the long-run mean of the factors from the origin to $\bar{f}$. Changing the conditional mean of u_{t+1} to $-\bar{B}f_t$ changes (1) to

$$r_{t+1} = \mu + (B-\bar{B})f_t + \tilde{u}_{t+1},$$

$\tilde{u}_{t+1}$ having the original distribution of u_{t+1}, and thus allows the adversary to change the factor loadings. By changing the covariance of u_{t+1}, the adversary changes the covariance of the price changes r_{t+1}. The adversary can also introduce correlation between u_{t+1} and v_{t+1}. We will see that the adversary's optimal controls take advantage of dynamic information about both portfolio holdings and factor levels and thus go beyond robustness to uncertainty about static parameters like C and B.

These examples also serve to illustrate that alternative but equivalent formulations of the original control problem (4) can lead to distinct formulations when we consider robustness. For instance, replacing $x_t^\top(r_{t+1}-\mu)$ on the left side of (4) with its conditional expectation $E_t[x_t^\top(r_{t+1}-\mu)] = x_t^\top Bf_t$ clearly has no effect on the investor's portfolio choice or its performance as measured by (4). However, by including u_{t+1} in the objective, we allow the adversary to influence performance by changing the distribution of this term. Put differently, including this term leads the investor to a strategy that is robust to errors in both the return model (1) and the factor dynamics (2), whereas omitting u_{t+1} focuses robustness exclusively on the factor dynamics. We solve and test both formulations.

We now formulate the adversary's actions more precisely. Let g_u and g_v denote the probability densities of u_t and v_t. The adversary may choose a new joint density $\tilde{g}_t$ for (u_t, v_t), which could, in the most general formulation, depend on past values (u_s, v_s), $s<t$, of the noise sequence. However, we will restrict our analysis to the Markovian case in which any dependence of the density function $\tilde{g}_t$ on the past is captured through dependence on the state (x_{t-1}, f_t). If we set

$$m_t = \frac{\tilde{g}_t(v_t, u_t \mid x_{t-1}, f_t)}{g_v(v_t)g_u(u_t)}, \qquad M_t = \prod_{s=1}^{t} m_s, \tag{5}$$

then M_t is the likelihood ratio relating the distribution of $(u_1, v_1, \ldots, u_t, v_t)$ selected by the adversary to the original distribution. Because g_u and g_v are multivariate normal densities, so the denominator of m_t is supported on all of $\mathbb{R}^{n_x} \times \mathbb{R}^{n_x}$ and is never zero.

As in Hansen and Sargent (2007), Hansen et al. (2006), and Petersen et al. (2000), we limit the adversary by constraining or penalizing the relative entropy of the change of measure. How tightly we constrain or penalize the relative entropy determines the degree of model uncertainty by limiting how far the adversary can change the stochastic evolution of the data away from the investor's model. The relative entropy at time t is $E[M_t \log M_t]$, which is always positive and is equal to zero only when the adversary leaves the original measure unchanged by taking $M_t \equiv 1$. Given M_0 and a sequence of one-period likelihood ratios $m = \{m_t, t = 1, 2, \ldots\}$ as in (5), let

$$\begin{aligned}\mathbf{R}_\beta(m) &= (1-\beta)\sum_{t=0}^{\infty}\beta^t E[M_t \log M_t] \\ &= \sum_{t=0}^{\infty}\beta^{t+1}E\big[M_t E_t[m_{t+1}\log m_{t+1}]\big]\end{aligned} \tag{6}$$

denote the infinite-horizon discounted sum of relative entropy, where the term $1-\beta$ is introduced to simplify the final expression. We can give the adversary a budget $\eta > 0$ and constrain the measure change to satisfy $\mathbf{R}_\beta(m) \leqslant \eta$. When truncated at a finite upper limit T, the two sums in (6) no longer coincide—the discount factor on the right would need to be replaced with $(\beta^{t+1} - \beta^{T+1})$, leading to a control problem that depends on both t and T, and not just on the time-to-go $T-t$. To avoid this feature and to preserve consistency with the infinite-horizon case, we use the rightmost sum in (6), truncated at T, as our finite-horizon measure of discounted entropy.

Constraining the adversary's measure change to satisfy $\mathbf{R}_\beta(m) \leqslant \eta$ results in the robust control problem (for either a finite or infinite horizon)

$$\sup_{\Delta x}\inf_{m:\mathbf{R}_\beta(m)<\eta} E\bigg[\sum_t \beta^t M_t\bigg(x_t^\top(Bf_t + u_{t+1}) - \frac{\gamma}{2}x_t^\top\Sigma_v x_t - \frac{1}{2}\Delta x_t^\top \Lambda \Delta x_t\bigg)\bigg]. \tag{7}$$

Here, the investor seeks to optimize performance in the face of model uncertainty by maximizing performance against the worst-case stochastic perturbation to the original model, considering only perturbations that are sufficiently close to the original model to satisfy the relative entropy constraint. As before, the supremum is taken over policies under which each Δx_t is a deterministic function of (x_{t-1}, f_t); the infimum is taken over measure changes satisfying the relative entropy constraint and having the form in (5) in which each new density $\tilde{g}_t$ is determined by (x_{t-1}, f_t). Thus, (x_{t-1}, f_t) remains Markovian under any policy pair $(\Delta x, m)$.

The Lagrangian of the constrained problem (7) is a penalty problem with parameter $\theta > 0$

$$\sup_{\Delta x}\inf_m E\bigg[\sum_{t=0}^{T}\beta^t M_t\bigg(x_t^\top(Bf_t + E_t[m_{t+1}u_{t+1}]) - \frac{\gamma}{2}x_t^\top\Sigma_v x_t - \frac{1}{2}\Delta x_t^\top\Lambda\Delta x_t + \theta\beta E_t[m_{t+1}\log m_{t+1}]\bigg)\bigg], \tag{8}$$

with the constraint $\mathbf{R}_\beta(m) < \eta$ replaced by an admissibility condition $\mathbf{R}_\beta(m) < \infty$. Hansen et al. (2006, Claim 5.4), establish the equivalence of constrained and penalized formulations through convex duality under mild conditions, complementing a similar result in Petersen et al. (2000). We work directly with (8) because it is more amenable to explicit solution. This formulation also has an interpretation in terms of dynamic risk measures, in which case θ measures the investor's aversion to ambiguity; see Ruszczyński (2010).

Our restriction to investment policies and measure changes that are Markovian in the sense that their dependence on the past is fully captured by dependence on the state (x_{t-1}, f_t) does not change the value of (8) in the finite-horizon case, provided the measure changes satisfy a more general rectangularity condition; this is shown in Theorems 2.1–2.2 of Iyengar (2005) for discrete state spaces, but his argument applies here as well. The infinite-horizon problem raises stability issues, but §7.6 of Hansen and Sargent (2007) shows an analogous reduction to Markovian strategies, under modest technical conditions, for problems of the type we consider. We avoid a digression into these issues by limiting ourselves to Markovian strategies throughout.

3. Finite-Horizon Robust Problem

3.1. Robust Bellman Equation for *W*

To lighten notation, we define

$$Q(x, \Delta x, f) = x^\top Bf - \frac{\gamma}{2} x^\top \Sigma_u x - \frac{1}{2}\Delta x^\top \Lambda x.$$

For a fixed horizon $T < \infty$, and any $0 \leqslant t \leqslant T$, the finite-horizon value function $W_{t,T}$ for the penalty problem (8) is given by

$$\begin{aligned} W_{t,T}&(M_t, x_{t-1}, f_t) \\ &= \sup_{\Delta x} \inf_{m} E_t\bigg[\sum_{s=t}^{T} \beta^s M_s\big(Q(x_s, \Delta x_s, f_s) + E_s[m_{s+1}u_{s+1}^\top]x_s \\ &\qquad + \theta\beta E_s[m_{s+1}\log m_{s+1}]\big)\bigg], \end{aligned} \tag{9}$$

with (x_{-1}, f_0) fixed and $W_{T+1,T} = 0$. That $W_{t,T}$ is indeed a function of only (M_t, x_{t-1}, f_t) follows from our restriction to Markovian strategies $(\Delta x, m)$ for the investor and the adversary. In particular, under a fixed pair of policies, the conditional expectations inside the summation reduce to functions of (x_s, f_s), and x_s is a function of (x_{s-1}, f_s).

Define the one-step robust dynamic programming operator $\mathcal{T}$ acting on functions h: $\mathbb{R} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_f} \to \mathbb{R}$ by

$$\begin{aligned} \mathcal{T}&(h)(M, x, f) \\ &= \sup_{\Delta x_+} \inf_{m_+} M\big(Q(x_+, \Delta x_+, f) + E[m_+ u_+^\top]x_+ + \beta E[\theta m_+ \log m_+] \\ &\qquad + \beta E[h(Mm_+, x_+, f_+)]\big), \end{aligned} \tag{10}$$

where $x_+ = x + \Delta x_+$, $f_+ = Cf + v_+$, $v_+ \sim N(0, \Sigma_v)$, and $u_+ \sim N(0, \Sigma_u)$. The supremum is over $\Delta x_+ \in \mathbb{R}^{n_x}$ and the infimum is over m_+ of the form in (5). It is always feasible (and optimal) for the adversary to choose m_+ with finite relative entropy.

Proposition 1. *$W_{t,T}$ satisfies, for $0 \leqslant t \leqslant T$, the robust Bellman equation*

$$W_{t,T} = \mathcal{T}(W_{t+1,T}). \tag{11}$$

Proof. The proof is the same as that of Theorem 2.2 of Iyengar (2005), even though the setting there is a discrete state space. In particular, the rectangularity condition required there holds in our setting. (We detail the argument for the infinite-horizon case in the proof of Proposition 6.) □

We have taken $W_{T+1,T} = 0$ as our terminal condition for simplicity. If the underlying assets are futures contracts (the focus of §6), then we can interpret this condition as having all contracts mature at $T+1$. Alternatively, we could assign $W_{T+1,T}$ a liquidation value for the portfolio, considering both asset prices at time $T+1$ and the transactions costs incurred in selling off the portfolio's holdings. This formulation would require recording price levels (the cumulative sum of the price differences r_t) in the state vector, which could be done quite easily. The final portfolio is just the scalar product of x_{T+1} and the price vector, so this formulation remains within the linear-quadratic framework. We omit this extension for simplicity, particularly since it does not apply to the infinite-horizon problem.

3.2. Bellman Equation for *V*

To solve the Bellman equation (11), we will follow the approach in Hansen and Sargent (2007) and prove that $W_{t,T}$ can be decomposed as a product of M_t and a function of (x_{t-1}, f_t), which will simplify (11) so that we can solve it analytically. First, we write W in the form

$$W_{t,T}(M_t, x_{t-1}, f_t) = M_t V_{t,T}(M_t, x_{t-1}, f_t), \tag{12}$$

taking this as the definition of $V_{t,T}$, since $M_t \in (0, \infty)$. Then (11) becomes

$$\mathcal{T}(M_t m_{t+1} V_{t+1,T})(M_t, x_{t-1}, f_t) = M_t V_{t,T}(M_t, x_{t-1}, f_t). \tag{13}$$

Recalling the definition of $\mathcal{T}$, we can divide both sides of (13) by M_t to get

$$\begin{aligned} V_{t,T}&(M_t, x_{t-1}, f_t) \\ &= \sup_{\Delta x_t} \inf_{m_{t+1}} \big\{Q(x_t, \Delta x_t, f_t) + E_t\big[m_{t+1}u_{t+1}^\top x_t + \beta\theta m_{t+1}\log m_{t+1} \\ &\qquad + \beta m_{t+1} V_{t+1,T}(M_{t+1}, x_t, f_{t+1})\big]\big\}. \end{aligned} \tag{14}$$

Set $V_{T+1,T} = 0$. It now follows by induction that $V_{t,T}$ does not depend on M_t: Suppose this is true of $V_{t+1,T}$; then, under our Markovian restriction on strategies, the conditional expectations in (14) are functions of (x_{t-1}, f_t), and thus so is $V_{t,T}$.

Since $\theta > 0$, the term in the conditional expectation of (14) is convex in m_{t+1}, so we can solve the minimization problem through the first-order conditions, which leads to the optimal choice

$$m^*_{t+1} = \frac{\exp\{-(1/\theta)(V_{t+1,T}(x_t, f_{t+1}) + (1/\beta)x_t^\top u_{t+1})\}}{E_t[\exp\{-(1/\theta)(V_{t+1,T}(x_t, f_{t+1}) + (1/\beta)x_t^\top u_{t+1})\}]}. \quad (15)$$

This is a positive function of x_t, f_{t+1}, u_{t+1} and v_{t+1}, normalized to integrate to 1, so it has the form required by (5). (We verify that the normalization in the denominator is finite in the case of interest as part of Theorem 1.) Substituting expression (15) into (14), we get the following recursion for V:

$$V_{t,T}(x_{t-1}, f_t) = \sup_{\Delta x_t \in \mathbb{R}^d} \Bigg\{ Q(x_t, \Delta x_t, f_t) - \beta\theta \log E_t \cdot \left[\exp\left\{-\frac{V_{t+1,T}(x_t, f_{t+1}) + (1/\beta)x_t^\top u_{t+1}}{\theta}\right\}\right]\Bigg\}, \quad (16)$$

which we abbreviate as $V_{t,T} = \mathcal{T}^V(V_{t+1,T})$ by defining $\mathcal{T}^V$ as the operator on the right. We summarize this transformation as follows:

Proposition 2. *For any $T < \infty$, any solution to* (16) *gives a solution $W_{t,T}(M_t, x_{t-1}, f_t) = M_t V_{t,T}(x_{t-1}, f_t)$ to* (11), *where the adversary's choice is given by* (15).

The recursion in (16) has the form of a risk-sensitive optimal control problem. With $\beta = 1$, the recursion can be unwound and V expressed as the value function of a control problem; this case is treated extensively in Whittle (1981), Whittle and Whittle (1990). With $\beta < 1$, there is no nonrecursive expression for V: V cannot be expressed as the value function for a control problem with a time-separable objective, nor is it equivalent to specifying an exponential utility function or any other standard utility function. This discounted case is treated in Hansen and Sargent (1995), although their convexity condition does not hold in our setting; see also Skiadas (2003). Portfolio optimization problems with risk-sensitive criteria are solved in, e.g., Bielecki and Pliska (1999), Bielecki et al. (2005), and Fleming and Sheu (2001), with the risk-sensitive objective posited from the outset. It should be stressed that, in our setting, the risk-sensitive problem (16) emerges only as an intermediate step in solving the robust control problem, in response to the adversary's optimal strategy, and not as the primary objective.

3.3. Saddlepoint Condition and Solution to the Bellman Equation

Building on Whittle (1981) and Whittle and Whittle (1990), we will give conditions leading to a quadratic solution for $V_{t,T}$. To motivate the argument, we first observe that starting with $V_{T+1,T} = 0$ and taking one step backward in (16), we find that $V_{T,T}$ is quadratic in (x_{T-1}, f_T). If, for some $t \leqslant T$, $V_{t+1,T}$ is quadratic in (x_t, f_{t+1}) with no dependence on M_t, then the last term in (16) becomes

$$-\theta\beta\log\left(E_t\left[\exp\left\{-\frac{1}{\theta}G\right\}\right]E_t\left[\exp\left\{-\frac{1}{\theta\beta}x_t^\top u_{t+1}\right\}\right]\right)$$
$$= -\theta\beta\log E_t\left[\exp\left\{-\frac{1}{\theta}G\right\}\right] - \frac{1}{2\beta\theta}x_t^\top \Sigma_u x_t, \quad (17)$$

where G is a quadratic function of (x_t, f_{t+1}) (and thus of v_{t+1}), and the conditional expectation factors because of the independence of u_{t+1} under the original probability measure. Under the saddlepoint conditions given below, (17) then reduces to a quadratic function of (x_t, f_t). So, the right side of (16) is a quadratic function of x_t, f_t, and x_{t-1}. Maximizing over x_t under the saddlepoint conditions, the right side of (16) becomes a quadratic function of (x_{t-1}, f_t). Thus, $V_{t,T}$ is quadratic in (x_{t-1}, f_t) and has no functional dependence on M_t.

A consequence of these properties of $V_{t,T}$ is that it is quadratic in v_{t+1}—implying that v_{t+1} continues to be normally distributed under the change of measure, although with a different mean and covariance—and linear in u_{t+1}—implying that u_{t+1} continues to be normally distributed but with a different mean. The absence of a cross term multiplying v_{t+1} and u_{t+1} in m^*_{t+1} preserves the independence of the two vectors.

In light of the foregoing discussion, we posit the representation

$$V_{t,T}(x_{t-1}, f_t) = x_{t-1}^\top A_{xx}^{(t,T)} x_{t-1} + x_{t-1}^\top A_{xf}^{(t,T)} f_t + f_t^\top A_{ff}^{(t,T)} f_t + A_0^{(t,T)},$$

and set

$$A^{(t,T)} = \begin{bmatrix} A_{xx}^{(t,T)} & \frac{1}{2}A_{xf}^{(t,T)} \\ \frac{1}{2}(A_{xf}^{(t,T)})^\top & A_{ff}^{(t,T)} \end{bmatrix}.$$

Without loss of generality, we take $A^{(t,T)}$ to be symmetric. We introduce two conditions to ensure that this structure is preserved by the recursion (16). To state the conditions generically, we drop the superscript (t, T).

Condition 1. $\Sigma_v^{-1} + (2/\theta)A_{ff} > 0$.

Condition 2. $J_1 > 0$, *with $\gamma_\theta = \gamma + (1/\theta\beta)$ and*

$$J_1 = \gamma_\theta \Sigma_u + \Lambda - 2\beta A_{xx} + \frac{\beta}{\theta}A_{xf}\left(\Sigma_v^{-1} + \frac{2}{\theta}A_{ff}\right)^{-1} A_{xf}^\top > 0. \quad (18)$$

By analogy with Whittle and Whittle (1990, pp. 81–83), we call these saddlepoint conditions.

The following lemma provides sufficient conditions for the required properties.

LEMMA 1. (i) *If $A_{xx} \leqslant 0$ and Condition* 1 *holds, then Condition* 2 *holds.* (ii) *If $A_{ff} \geqslant 0$, then Condition* 1 *holds.*

The lemma helps explain the name "saddlepoint" and shows that the conditions we have are weaker than concavity in x and convexity in f. We now apply our conditions to the Bellman equation (16). The result is similar to that of Hansen and Sargent (1995); however, a separate argument is needed because we do not have joint concavity. Given a matrix A, define

$$J_2 = Bf_t + \Lambda x_{t-1} + \beta A_{xf} C f_t - 2\frac{\beta}{\theta} A_{xf}\left(\Sigma_v^{-1} + \frac{2}{\theta}A_{ff}\right)^{-1} A_{ff} C f_t, \tag{19}$$

$$J_3 = B + \beta A_{xf} C - 2\frac{\beta}{\theta} A_{xf}\left(\Sigma_v^{-1} + \frac{2}{\theta}A_{ff}\right)^{-1} A_{ff} C. \tag{20}$$

The matrix J_3 appears when we substitute a quadratic function into (16); the vector J_2 depends on the state (x_{t-1}, f_t) and will be used to describe the investor's optimal portfolio. The proof of the following result, and that of most results throughout the paper, appears in appendices.

PROPOSITION 3. *If a symmetric matrix A satisfies Conditions* 1 *and* 2, *then $\mathcal{T}^V$ maps a quadratic function with coefficients (A, A_0) to a quadratic function with coefficients $(\mathcal{S}(A), \mathcal{U}(A, A_0))$, where*

$$\mathcal{S}(A)_{xx} = -\tfrac{1}{2}\Lambda + \tfrac{1}{2}\Lambda J_1^{-1}\Lambda^\top, \tag{21}$$

$$\mathcal{S}(A)_{xf} = \Lambda J_1^{-1}\left[B + \beta A_{xf} C - 2\frac{\beta}{\theta} A_{xf}\left(\Sigma_v^{-1} + \frac{2}{\theta}A_{ff}\right)^{-1} A_{ff} C\right], \tag{22}$$

$$\mathcal{S}(A)_{ff} = \beta C^\top A_{ff} C - 2\frac{\beta}{\theta} C^\top A_{ff}\left(\Sigma_v^{-1} + \frac{2}{\theta}A_{ff}\right)^{-1} A_{ff} C + \frac{1}{2} J_3^\top J_1^{-1} J_3, \tag{23}$$

$$\mathcal{U}(A, A_0) = \frac{\beta\theta}{2}\log\left|I + \frac{2}{\theta}\Sigma_v A_{ff}\right| + \beta A_0, \tag{24}$$

and $|\cdot|$ denotes the determinant of a matrix.

We now show that the properties we need for a quadratic representation of $V_{t,T}$ are indeed preserved by (16). We write $\mathcal{S}^n$ for the n-fold iteration of the mapping $\mathcal{S}$.

PROPOSITION 4. (i) *If A is symmetric, then $\mathcal{S}^n(A)$ is symmetric for all $n > 0$.* (ii) *If A is symmetric and satisfies $\Lambda^{1/2} J_1^{-1} \Lambda^{1/2} < I$, $A_{ff} \geqslant 0$, and Condition* 2, *then*

$$-\Lambda/2 < \mathcal{S}^n(A)_{xx} < 0 \quad \text{and} \quad \mathcal{S}^n(A)_{ff} \geqslant 0, \quad \text{for } n > 0.$$

Hence, Conditions 2 *and* 1 *hold for all $n > 0$.*

3.4. Optimal Controls

We can now summarize the optimal controls for the adversary and the investor. We use $\tilde{E}$ to denote expectation under the change of measure selected by the adversary. The conditions on $A^{(T,T)}$ in the following result hold, in particular, for the terminal condition $A^{(T+1,T)} = 0$ corresponding to $W_{T+1,T} = 0$, but they hold more generally as well.

THEOREM 1. *Suppose $A^{(T,T)}$ satisfies the conditions in Proposition* 4 *so that Conditions* 2 *and* 1 *hold for all $A^{(t,T)}$, $t \leqslant T$.*

(i) *Under the adversary's optimal change of measure, the conditional distribution of (u_{t+1}, v_{t+1}) given $(u_1, v_1), \ldots, (u_t, v_t)$ is normal with conditional covariance $\tilde{\Sigma}_{t+1}$,*

$$\tilde{\Sigma}_{t+1} = \begin{bmatrix} \Sigma_u & 0 \\ 0 & (\tilde{\Sigma}_v)_{t+1} \end{bmatrix}, \quad \textit{where}$$

$$(\tilde{\Sigma}_v)_{t+1} = \left(\Sigma_v^{-1} + \frac{2}{\theta}A_{ff}^{(t+1,T)}\right)^{-1},$$

and conditional mean $(z_{u,t+1}, z_{v,t+1})$, with

$$z_{u,t+1} = \frac{-\Sigma_u x_t}{\theta\beta}$$

and

$$z_{v,t+1} = -\frac{1}{\theta}(\tilde{\Sigma}_v)_{t+1}\left((A_{xf}^{(t+1,T)})^\top x_t + 2A_{ff}^{(t+1,T)} C f_t\right) = -\frac{1}{\theta}(\tilde{\Sigma}_v)_t \frac{\partial V_{t+1,T}(x_t, \tilde{E}_t[f_{t+1}])}{\partial f_t}. \tag{25}$$

(ii) *The investor's optimal choice is*

$$\Delta x_t^* = 2\Lambda^{-1} A_{xx}^{(t,T)}\left(x_{t-1} + \tfrac{1}{2}(A_{xx}^{(t,T)})^{-1} A_{xf}^{(t,T)} f_t\right) = \Lambda^{-1}\frac{\partial V_{t,T}}{\partial x}, \tag{26}$$

so

$$x_t = J_1^{-1} J_2 = (I + 2\Lambda^{-1} A_{xx}^{(t,T)}) x_{t-1} + \Lambda^{-1} A_{xf}^{(t,T)} f_t. \tag{27}$$

The effect of the adversary's control is to change the evolution of the factors from (2) to

$$f_{t+1} = C f_t + z_{v,t+1} + \tilde{v}_{t+1}, \qquad \tilde{v}_{t+1} \sim N(0, (\tilde{\Sigma}_v)_{t+1});$$

in particular, this makes

$$\tilde{E}_t[f_{t+1}] = C f_t + z_{v,t+1} = \left(I - \frac{2}{\theta}(\tilde{\Sigma}_v)_{t+1} A_{ff}^{(t+1,T)}\right) C f_t - \frac{1}{\theta}(\tilde{\Sigma}_v)_{t+1}(A_{xf}^{(t+1,T)})^\top x_t = (\tilde{\Sigma}_v)_{t+1}\left(\Sigma_v^{-1} C f_t - \frac{1}{\theta}(A_{xf}^{(t+1,T)})^\top x_t\right), \tag{28}$$

where $\tilde{E}$ denotes expectation under the change of measure selected by the adversary. Because $(\tilde{\Sigma}_v)_{t+1} \leqslant \Sigma_v$, we interpret the first term in (28) as shrinking the persistence of the factors, and thus potentially reducing their value to the investor; the second term in (28) indicates that the adversary also exploits the investor's current portfolio in setting the conditional mean of the factors, as suggested by the expression following (25).

Corollary 1. *We can write the investor's optimal choice* (27) *as*

$$\begin{aligned} x_t &= (I + 2\Lambda^{-1} A_{xx}^{(t,T)}) x_{t-1} \\ &\quad - 2\Lambda^{-1} A_{xx}^{(t,T)} \big(-\tfrac{1}{2} (A_{xx}^{(t,T)})^{-1} A_{xf}^{(t,T)} f_t\big) \qquad (29) \\ &= \big(\gamma_\theta \Sigma_u + \Lambda - 2\beta A_{xx}^{(t+1,T)}\big)^{-1} \\ &\quad \times \big(\Lambda x_{t-1} + \gamma_\theta \Sigma_u (\gamma_\theta \Sigma_u)^{-1} B f_t - 2\beta A_{xx}^{(t+1,T)} \\ &\quad \cdot \big(-\tfrac{1}{2} (A_{xx}^{(t+1,T)})^{-1} A_{xf}^{(t+1,T)} \tilde{E}_t[f_{t+1}]\big)\big). \qquad (30) \end{aligned}$$

For (29), as in Gârleanu and Pedersen (2013), we can interpret this choice as a weighted average of the current portfolio x_{t-1} and a target portfolio given by

$$\mathit{target} = -\tfrac{1}{2} (A_{xx}^{(t,T)})^{-1} A_{xf}^{(t,T)} f_t. \qquad (31)$$

The target portfolio maximizes the quadratic function $V_{t,T}$, given the factor level f_t. If transaction costs were waived for one period, the investor would move immediately to the target; otherwise, the investor's optimal trade (26) is proportional to the difference between the current portfolio and the target, the proportion depending on the cost matrix Λ. Recall that the A matrix depends on the robustness parameter θ through the recursions in Proposition 3.

In (30), the term $(\gamma_\theta \Sigma_u)^{-1} B f_t$, which we call the *myopic* portfolio, maximizes the single-period mean-variance objective

$$x_t^\top B f_t - \frac{\gamma_\theta}{2} x_t^\top \Sigma_u x_t.$$

Thus, (30) represents x_t as a weighted average of the current portfolio x_{t-1}, the myopic portfolio, and the conditional expectation of the target portfolio one step ahead. Comparing this expression to Equation (16) of Gârleanu and Pedersen (2013), we can interpret the effect of the robust solution as replacing the original conditional expectation of the factors with their conditional expectation under the adversary's change of measure and implicitly increasing the investor's risk-aversion parameter from γ to γ_θ. Interestingly, if we omitted the variance penalty $\gamma x_t^\top \Sigma_u x_t / 2$ from the original objective (4), it would still appear in the robust formulation, because $\gamma_\theta > 0$ even if $\gamma = 0$. Uncertainty in the linear term $u_{t+1}^\top x_t$ has the effect of increasing risk aversion.

4. Comparison with the Nonrobust Case

In this section, we examine the effect of varying the robustness parameter θ, including the nonrobust formulation $\theta = \infty$ as a limiting case. We affix θ as a subscript or superscript to indicate functions and quantities tied to a specific value of the parameter. The nonrobust version is indicated by a subscript or superscript ∞.

We denote by $W_{t,T}^\infty$ and W^∞, respectively, the finite-horizon and infinite-horizon value functions for the (nonrobust) objective (4) and define a dynamic programming operator acting on functions h: $\mathbb{R}^{n_x} \times \mathbb{R}^{n_f} \to \mathbb{R}$ by

$$\mathcal{T}^\infty(h)(x, f) = \sup_{\Delta x_+ \in \mathbb{R}^{n_x}} \left\{ x_+^\top B f_+ - \frac{\gamma}{2} x_+^\top \Sigma_v x_+ - \frac{1}{2} \Delta x_+^\top \Lambda \Delta x_+ + E[h(x_+, f_+)] \right\},$$

with $x_+ = x + \Delta x_+$, and the expectation taken over $f_+ = Cf + v$, $v \sim N(0, \Sigma_v)$. Then $W_{t,T}^\infty$ satisfies the recursion $\mathcal{T}^\infty(W_{t+1,T}^\infty) = W_{t,T}^\infty$. Gârleanu and Pedersen (2013) show that this dynamic programming equation maps a quadratic function backward to another quadratic function. We can therefore write

$$\begin{aligned} W_{t,T}^\infty(x_{t-1}, f_t) &= x_{t-1}^\top \mathcal{S}_\infty^{T-t}(A)_{xx} x_{t-1} + x_{t-1}^\top \mathcal{S}_\infty^{T-t}(A)_{xf} f_t \\ &\quad + f_t^\top \mathcal{S}_\infty^{T-t}(A)_{ff} f_t + \mathcal{U}_\infty^{T-t}(A, A_0), \qquad (32) \end{aligned}$$

with $\mathcal{S}_\infty^{T-t}(A)$ and $\mathcal{U}_\infty^{T-t}(A, A_0)$ the coefficients of $W_{t,T}^\infty$ at time t when $W_{T,T}^\infty$ is quadratic with coefficient matrix A. Here, $\mathcal{S}_\infty^n$ is the n-fold iteration of $\mathcal{S}_\infty$, but $\mathcal{U}_\infty^n$ is defined recursively by setting $\mathcal{U}_\infty^1 = \mathcal{U}_\infty$ and $\mathcal{U}_\infty^n(A, A_0) = \mathcal{U}_\infty(\mathcal{S}_\infty^{n-1}(A), \mathcal{U}_\infty^{n-1}(A, A_0))$; see the analogous dependence on A and A_0 in (24).

The nonrobust case can be considered a special case of the robust formulation. Condition 1 holds automatically when $\theta = \infty$, and if Condition 2 holds for some matrix A for $\theta = \infty$, then it also holds for any $\theta \in (0, \infty)$. This is because the last term in (18) is positive definite for $\theta \in (0, \infty)$, but vanishes when $\theta = \infty$, and γ_θ is decreasing in θ, so $J_1^\theta \geqslant J_1^\infty$. It is also easy to verify that Proposition 4 holds at $\theta = \infty$. As we vary θ (smaller θ indicating greater robustness), the coefficient matrices are ordered as follows:

Lemma 2. *If* (A, A_0) *satisfies* $A_{ff} \geqslant 0$, $\Lambda^{1/2} J_1^{-1} \Lambda^{1/2} < I$, *and Condition* 2 *for some* $0 < \theta_1 < \theta_2 \leqslant \infty$, *then for any* $n \geqslant 0$, $\mathcal{S}_{\theta_1}^n(A) \leqslant \mathcal{S}_{\theta_2}^n(A)$ *and* $\mathcal{U}_{\theta_1}^n(A, A_0) \leqslant \mathcal{U}_{\theta_2}^n(A, A_0)$.

To illustrate, suppose we start the recursions for two parameter levels $0 < \theta_1 < \theta_2 \leqslant \infty$ from the same terminal condition with coefficients (A, A_0) (including $A = A_0 = 0$ as a special case). Suppose the conditions of Lemma 2 hold. We make the following observations:

(a) In the portfolio decomposition (29), the weight on the previous position satisfies $1 + 2\Lambda^{-1} \mathcal{S}_{\theta_1}^n(A)_{xx} < 1 + 2\Lambda^{-1} \mathcal{S}_{\theta_2}^n(A)_{xx}$, so more robustness (smaller θ) leads to less weight on the previous position x_{t-1} and more weight on the target portfolio (31). The less-robust investor puts greater trust in the persistence of the factors described by the model and thus attaches greater value to the previous portfolio. However, the coefficient of f_t in (29) can either increase or decrease with θ because $\mathcal{S}_\theta^n(A)_{xf}$ can increase or decrease or change in a more complicated way.

(b) From the decomposition (30), we find similarly that increasing robustness decreases weight on the myopic portfolio $(\gamma_\theta \Sigma_u)^{-1} B f_t$, and it also decreases the size of the myopic portfolio because γ_θ increases with θ. If we remove u_{t+1} from (4) and limit robustness to the factor dynamics (2) only, then $\gamma_\theta \equiv \gamma$ and the myopic portfolio does not vary with θ.

(c) Also from (30), we see that increasing robustness puts more weight on the conditional expectation of the target portfolio while decreasing the coefficient on the conditional expectation of the factors. In numerical examples, we find that the conditional expectation of the target portfolio is very sensitive to θ.

We can also interpret the effect of robustness from the optimal controls

$$\Delta x_t^* = \Lambda^{-1} \frac{\partial V_{t+1,T}}{\partial x} \quad \text{and}$$

$$z_{v,t+1} = -\frac{1}{\theta} (\tilde{\Sigma}_v)_t \frac{\partial V_{t+1,T}(x_t, \tilde{E}_t[f_{t+1}])}{\partial f_t}.$$

These expressions are already very suggestive, as they show the investor and the adversary using their controls to increase and decrease V, respectively. Also, the quadratic function $V_{t,T}$ is concave in x_{t-1} and convex in f_t, making it a hyperbolic paraboloid. The cross term $x_t^\top B f_t$ in the objective function leads to the cross-term coefficient $A_{xf} \neq 0$ in the value function. The presence of this term means that the axes of the hyperbolic paraboloid are twisted and not orthogonal to each other. As a result, the minimum point for f is linear in x, and maximum point for x is linear in f, properties exploited by both players. This can also be seen in (25) and (27).

If there were no cross term in the objective and we had $A_{xf}^{t+1,T} = 0$, then the coefficient of Cf_t in (25) would be a negative definite matrix, and the effect of the adversary's choice of $z_{v,t+1}$ would thus be to accelerate the mean reversion of the factors in (2) and reduce their value to the investor. In fact, in the limit as θ approaches zero, the coefficient of f_t in (25) becomes $-C$, which eliminates any persistence in the factor dynamics (2). With a nonzero cross term, the adversary can do further harm by moving the factors in a direction that depends on the investor's current portfolio.

We conclude this section by verifying that value iteration for the nonrobust problem converges; this is needed to confirm that the solution to the Bellman equation found in Gârleanu and Pedersen (2013) is in fact the value function for the infinite-horizon problem and that the corresponding control is optimal. In the following, J_1 is evaluated with $\theta = \infty$.

Proposition 5. *If A is such that $J_1 > 0$, $A_{ff} \geqslant 0$, and*

$$A - \begin{bmatrix} 0 & 0 \\ 0 & \dfrac{1}{2\gamma} B^\top \Sigma_u^{-1} B \end{bmatrix} \leqslant 0,$$

then the iteration of (32) converges; i.e., $\lim_{n\to\infty}(\mathcal{S}_\infty^n(A), \mathcal{U}_\infty^n(A, A_0))$ exists. The control (26) obtained from the limit is optimal, and the quadratic function defined by the limit solves the Bellman equation and is the value function for the infinite-horizon problem.

5. Infinite-Horizon Robust Problem

5.1. Formulation and Bellman Equation

For the robust infinite-horizon problem, define W by setting $t = 0$ and $T = \infty$ on the right side of (9). This robust value function is bounded above by the nonrobust value function (corresponding to $\theta = \infty$ in Lemma 2), and it is bounded below because the investor can choose $x_t \equiv 0$.

Proposition 6. *With $\mathcal{T}$ the operator defined in (10), W satisfies*

$$W = \mathcal{T}(W). \tag{33}$$

Similarly, by arguing as in Proposition 2 and the subsequent discussion, we arrive at the following result.

Proposition 7. *Suppose $V(x_{t-1}, f_t)$ satisfies $V = \mathcal{T}^V(V)$, with $\mathcal{T}^V$ as defined by (16). Then $W(M_t, x_{t-1}, f_t) = M_t V(x_{t-1}, f_t)$ satisfies (33) with*

$$m_{t+1}^* = \frac{\exp\{-(1/\theta)(V(x_t, f_{t+1}) + (1/\beta) x_t^\top u_{t+1})\}}{E_t[\exp\{-(1/\theta)(V(x_t, f_{t+1}) + (1/\beta) x_t^\top u_{t+1})\}]}, \tag{34}$$

provided the normalization in (34) is finite.

This reduces the problem of finding a solution to the robust Bellman equation (33) to one of solving $V = \mathcal{T}^V(V)$. In solving for the infinite-horizon V, the finite-horizon recursions for $V_{t,T}$ in Proposition 3 become simultaneous equations. Given coefficients (A, A_0) of a quadratic function, define J_1, J_2, and J_3 as in (18)–(20).

Proposition 8. *If A is symmetric and (A, A_0) satisfy $A_{ff} \geqslant 0$, $J_1 > 0$, $\Lambda^{1/2} J_1^{-1} \Lambda^{1/2} < I$, and*

$$A_{xx} = -\tfrac{1}{2}\Lambda + \tfrac{1}{2}\Lambda J_1^{-1} \Lambda^\top,$$

$$A_{xf} = \Lambda J_1^{-1} \left[B + \beta A_{xf} C - 2\frac{\beta}{\theta} A_{xf} \left(\Sigma_v^{-1} + \frac{2}{\theta} A_{ff} \right)^{-1} A_{ff} C \right],$$

$$A_{ff} = \beta C^\top A_{ff} C - 2\frac{\beta}{\theta} C^\top A_{ff} \left(\Sigma_v^{-1} + \frac{2}{\theta} A_{ff} \right)^{-1} A_{ff} C + \tfrac{1}{2} J_3^\top J_1^{-1} J_3,$$

$$(1-\beta) A_0 = -\frac{\beta\theta}{2} \log\left| I + \frac{2}{\theta} \Sigma_v A_{ff} \right|,$$

then the quadratic function V defined by (A, A_0) is a fixed point of $\mathcal{T}^V$.

This is a direct consequence of Proposition 4. With this result, we can solve the equations for (A, A_0) and check the conditions in the statement of the proposition (which ensure the saddlepoint conditions we need for optimality). If these are satisfied, then we have a fixed point V, from which we get a solution W to the Bellman equation $\mathcal{T}(W) = W$ by setting $W(M, x, f) = MV(x, f)$. Such a solution provides candidate optimal controls for both the investor and the adversary—controls that attain the supremum and the infimum in the one-step operator $\mathcal{T}$. The calculation of these controls is similar to that in Theorem 1, but simpler because of that stationarity implicit in the infinite-horizon setting. We summarize the calculation as follows:

Lemma 3. *Suppose the conditions of Proposition* 8 *hold and define V from (A, A_0) accordingly.*

(i) *Under the change of measure* (34), *the conditional distribution of (u_{t+1}, v_{t+1}) given $(u_1, v_1), \dots, (u_t, v_t)$ is normal with conditional covariance $\tilde{\Sigma}$,*

$$\tilde{\Sigma} = \begin{bmatrix} \Sigma_u & 0 \\ 0 & \tilde{\Sigma}_v \end{bmatrix}, \quad \text{where} \quad \tilde{\Sigma}_v = \left(\Sigma_v^{-1} + \frac{2}{\theta} A_{ff} \right)^{-1},$$

and conditional mean $(z_{u,t+1}, z_{v,t+1})$, with

$$z_{u,t+1} = -\Sigma_u x_t / (\theta\beta) \quad \text{and}$$

$$z_{v,t+1} = -\frac{1}{\theta} \tilde{\Sigma}_v (A_{xf}{}^\top x_t + 2A_{ff} C f_t)$$
$$= -\frac{1}{\theta} \tilde{\Sigma}_v \frac{\partial V(x_t, \tilde{E}_t[f_{t+1}])}{\partial f_t}.$$

(ii) *The supremum over Δx in the Bellman equation $V = \mathcal{T}^V(V)$ is given by the investment choice*

$$\Delta x_t^* = 2\Lambda^{-1} A_{xx} \big(x_{t-1} + \tfrac{1}{2} (A_{xx})^{-1} A_{xf} f_t \big)$$
$$= \Lambda^{-1} \frac{\partial V(x_{t-1}, f_t)}{\partial x},$$

under which

$$x_t = J_1^{-1} J_2 = (I + 2\Lambda^{-1} A_{xx}) x_{t-1} + \Lambda^{-1} A_{xf} f_t.$$

5.2. Stability

Lemma 3 provides explicit expressions for the controls obtained by solving the robust Bellman equation. As is often the case in infinite-horizon problems, we need additional conditions to verify that a solution to the Bellman equation is in fact the value function (9) (with $T = \infty$) and that the corresponding controls are optimal. For these properties, we need to impose stability properties on the evolution of the controlled system. The key property is the admissibility condition $\mathbf{R}_\beta(m) < \infty$, with $\mathbf{R}_\beta$ as defined in (6).

Although it refers only to the adversary's control, this property is best viewed as a condition on the controls of both players because the adversary's choice of m_t may depend on the investor's choice of portfolio. Define the state y_t and the extended state y_t^e by setting

$$y_t = \begin{bmatrix} x_{t-1} \\ f_t \end{bmatrix} \quad \text{and} \quad y_t^e = \begin{bmatrix} u_t \\ x_{t-1} \\ f_t \end{bmatrix}.$$

The state evolution depends on the chosen pair of policies $(\Delta x, m)$. Details of the state dynamics can be found in the beginning of §8 in online e-companion (available as supplemental material at http://dx.doi.org/10.1287/opre.2013.1180). We use $\|\cdot\|$ to denote the usual Euclidean vector norm. The full stability condition we use is as follows.

Definition 1 (β-Stability). We call a policy pair $(\Delta x, m)$ and the resulting extended state evolution β-stable if $\mathbf{R}_\beta(m) < \infty$ and if $\alpha^t \tilde{E}[\|y_t^e\|^2] \Rightarrow 0$ for some $\alpha \in (\beta, 1)$, for all y_0^e.

The mean square convergence to zero of $\alpha^{t/2} y_t^e$ under the change of measure is sufficient to ensure that the infinite discounted sum (with discount factor β) of a quadratic function of the extended state is finite. We thus interpret β-stability as ensuring that the adversary cannot drive the investor reward to $-\infty$ and that the investor cannot drive the relative entropy penalty to $+\infty$; in particular, the condition avoids the possibility of getting $\infty - \infty$ in the robust value function. For the controls obtained from the Bellman equation (the controls in Lemma 3), a simpler condition characterizes β-stability. We use the condition

$$\mathbf{R}_\alpha(m) < \infty \quad \text{for some } \alpha \in (\beta, 1). \tag{35}$$

Lemma 4. *Suppose U is a solution to the robust Bellman equation* (33) *for which $U(M, x, f)$ is the product of M and a quadratic function of (x, f). If the quadratic function satisfies Conditions* 2 *and* 1, *then the resulting controls $(\Delta x, m)$ are β-stable if and only if* (35) *holds.*

5.3. Optimality

We now verify that that the policies provided by the robust Bellman equation through Lemma 3 do indeed solve the robust control problem in a suitable sense. Suppose both the investor and the adversary choose their policies (Markov, as we assume throughout), and let x_t^* and m_t^* be the resulting portfolios and likelihood ratios. Then the value attained by this pair of policies, starting from (M_0, x_{-1}, f_0), whenever this expression is well defined, is given by

$$W^*(M_0, x_{-1}, f_0) = E\bigg[\sum_{t=0}^{\infty} \beta^t M_t \big(Q(x_t^*, \Delta x_t^*, f_t) + \theta\beta m_{t+1}^* (\log m_{t+1}^* + u_{t+1}^\top x_t^*) \big) \bigg] \tag{36}$$

$$= \tilde{E}\bigg[\sum_{t=0}^{\infty} \beta^t \big(Q(x_t^*, \Delta x_t^*, f_t) + \theta\beta m_{t+1}^* (\log m_{t+1}^* + u_{t+1}^\top x_t^*) \big) \bigg], \tag{37}$$

where, as before, $\tilde{E}$ denotes expectation under the adversary's change of measure. We show that a solution to the robust Bellman equation (33) is indeed the value attained under the corresponding policies, and the policy forms an equilibrium. Once one player has selected a policy, we call a policy selected by the other player a β-stable response if the resulting policy pair is β-stable.

THEOREM 2. *Suppose U^* is a solution to the robust Bellman equation* (33), *and suppose $U^*(M, x, f)$ is the product of M and a quadratic function of (x, f) satisfying Conditions* 1–2. *Suppose the corresponding policy pair $(\Delta x^*, m^*)$ satisfies $\mathbf{R}_\alpha(m^*) < \infty$ for some $\alpha \in (\beta, 1)$. Then*

(i) *U^* is the value attained under the policy pair $(\Delta x^*, m^*)$.*

(ii) *The investor's best β-stable response to m^* is Δx^*. The adversary's best β-stable response to Δx^* is m^*.*

This result justifies the controls that come out of Lemma 3. It is worth noting that a violation of β-stability entails either a portfolio size that grows exponentially or an infinite relative entropy penalty. The restriction to β-stable policy pairs is therefore sensible, and it is appropriate to view the policy Δx^* derived from the robust Bellman equation as the investor's optimal choice in the face of the model uncertainty captured by the robust formulation.

5.4. Convergence of Value Iteration

From Lemma 3, we see that the key step in solving the infinite-horizon robust control problem is solving the equations in Proposition 8, which restate the condition $V = \mathcal{T}^V(V)$ for quadratic V. A natural approach is to start from some initial (A, A_0) and apply the equations iteratively. Each application of the equations is an application of the operator $\mathcal{T}^V$, so the question of convergence of this iterative approach is equivalent to the question of convergence of the finite-horizon function $V_{t,T}$ as $t \to -\infty$ with $V_{T,T}$ the quadratic function determined by the starting point (A, A_0). Hansen and Sargent (1995) consider the case where the objective function is concave in the state variable, which allows a simple proof through a monotone convergence argument, but our setting is beyond the scope of their result.

Over a finite horizon T, each $(A^{t+1,T}, A_0^{t+1,T})$ determines candidate controls through the prescription in Theorem 1. Under these controls, the state y_t evolves as in (5) in online e-companion (available as supplemental material at http://dx.doi.org/10.1287/opre.2013.1180), but with a time-dependent transition matrix

$$\bar{\Psi}_{t,T} = \begin{bmatrix} I & 0 \\ -\dfrac{1}{\theta}(\tilde{\Sigma}_v)_{t+1} A_{xf}^{(t+1,T)} & (\tilde{\Sigma}_v)_{t+1}\Sigma_v^{-1} C \end{bmatrix} \cdot \begin{bmatrix} I + 2\Lambda^{-1} A_{xx}^{(t+1,T)} & \Lambda^{-1} A_{xf}^{(t+1,T)} \\ 0 & I \end{bmatrix}.$$

If both factors in this representation have norm less than $\beta^{-1/2}$, then we have convergence of $\mathcal{S}^n(A)$ $\mathcal{U}^n(A, A_0)$ for any initial (A, A_0) that satisfies $\Lambda^{1/2} J_1^{-1} \Lambda^{1/2} < I$ and $A_{ff} \geqslant 0$. The norm here can be any matrix norm for which $\|M^n\|^{1/n} \to \sigma(M)$, such as any p-norm.

We have not found simple sufficient conditions that ensure this uniform stability condition. The condition can easily be checked for each $\bar{\Psi}_{t,T}$ at each iteration as part of an iterative algorithm, but given the difficulty of verifying the condition in advance we omit the details of the result. In our numerical experiments, we have never observed a failure to converge, starting either from zero or the solution of the nonrobust case and, indeed, the convergence appears to be quite fast. An alternative to iteration is the decomposition method covered in Hansen and Sargent (2007). This approach leads to conditions that guarantee a solution, but it requires a lengthy and technical digression, so we omit it.

6. Numerical Results

6.1. Data Description and Model Estimation

In order to test the effect of the robust formulation, we work with the application to commodity futures in Gârleanu and Pedersen (2013), using futures prices on the following commodities: aluminum, copper, nickel, zinc, lead, and tin from the London Metal Exchange; gas oil from the Intercontinental Exchange; WTI crude, RBOB unleaded gasoline, and natural gas from the New York Mercantile Exchange; gold and silver from the New York Commodities Exchange; and coffee, cocoa, and sugar from the New York Board of Trade. For consistency with Gârleanu and Pedersen (2013), we use daily prices for the period 01/01/1996–01/23/2009 for our in-sample tests; we use data through 04/09/2010 for out-of-sample tests. As discussed in Gârleanu and Pedersen (2013), extracting price changes from futures prices requires some assumptions on how contracts are rolled, and this makes it difficult to reproduce exactly the same time series of price changes. We choose the contract with the largest volume on each day. In some early samples when volumes for some commodities are not available, we choose the contract closest to maturity that does not expire in the current month. Our estimates and results are quite close and adequate for the purpose of examining the effect of robustness.

For each commodity, Gârleanu and Pedersen (2013) introduce factors f^{5D}, f^{1Y}, and f^{5Y}, which are the moving averages of price changes over the previous five days, one year, and five years, normalized by their respective standard deviations. Using these factors, we estimate the following model of price changes for each commodity:

$$r_{t+1}^s = \underset{(0.54)}{0.004} + \underset{(2.41)}{11.43}\, f_t^{5D,s} + \underset{(2.42)}{107.55}\, f_t^{1Y,s} - \underset{(-1.65)}{218.76}\, f_t^{5Y,s} + u_{t+1}^s,$$

the superscript s indexing the 15 commodities. This is a pooled panel regression—the coefficients are the same across all commodities—with parameters estimated using

Table 1. In-sample performance comparisons using the full data series with robustness in returns (u_t) and factor dynamics (v_t).

		Obj × 10^{-6}				Mean/std			
	θ	Gross	t-stat	Net	t-stat	Gross	t-stat	Net	t-stat
No TC		1.22		−159.68		0.82		−11.67	
Myopic		−0.21		−0.22		0.08		0.08	
Nonrobust		0.70		0.62		0.60		0.56	
Robust	10^{10}	0.70	−0.04	0.62	0.12	0.61	−0.08	0.57	−0.13
	10^9	0.55	−0.65	0.51	−0.49	0.61	−0.29	0.57	−0.48
	10^8	0.15	−1.24	0.13	−1.09	0.63	0.09	0.57	−0.35
	10^7	0.02	−1.37	0.02	−1.21	0.72	0.89	0.66	0.70
	10^6	0.00	−1.40	0.00	−1.23	0.79	1.04	0.77	1.18
	10^5	0.00	−1.40	0.00	−1.23	0.82	1.13	0.82	1.40

Note. For each θ, the t-stats compare performance of the robust rule at that θ with the nonrobust case, based on grouping the data into 40 batches.

feasible generalized least squares. The numbers reported under the coefficients are t-statistics. Similarly, for the factor dynamics we get the following estimates:

$$\Delta f_{t+1}^{5D,s} = \underset{(-0.67)}{-0.2510}\, f_t^{5D,s} + v_{t+1}^{5D,s},$$

$$\Delta f_{t+1}^{1Y,s} = \underset{(-0.64)}{-0.0039}\, f_t^{1Y,s} + v_{t+1}^{1Y,s},$$

$$\Delta f_{t+1}^{5Y,s} = \underset{(-0.78)}{-0.0010}\, f_t^{5Y,s} + v_{t+1}^{5Y,s}.$$

The matrix C is thus diagonal, and, in light of the t-statistics, a potential source of model error to be captured in the v_{t+1} terms. With $f = (f^{5D,1}, f^{1Y,1}, f^{5Y,1}, \ldots, f^{5D,15}, f^{1Y,15}, f^{5Y,15})^\top$, the form of the loading matrix B follows from the regression equation for r_t^s. Erb and Harvey (2006) documented the 1-year momentum factor in commodity futures prices. Asness et al. (2013) and Moskowitz et al. (2012) documented 1-year and 5-year many asset classes.

We adopt the choices in Gârleanu and Pedersen (2013) in estimating Σ_v and Σ_u, and in setting the risk-aversion parameter to $\gamma = 10^{-9}$, the one-day discount factor $\beta = \exp(-0.02/260)$ corresponding to a 2% annual rate, and the transaction cost matrix to $\Lambda = \lambda\Sigma_u$, with $\lambda = 3\times 10^{-7}$.

6.2. In-Sample Tests

This section reports results of in-sample tests in which we evaluate portfolio performance on the same price data used to estimate the model. We compare performance at various levels of the robustness parameter θ, including the nonrobust case $\theta = \infty$. The "No TC" case is a strategy that ignores transaction costs and thus reduces to the mean-variance optimal portfolio $x_{t+1} = (\gamma\Sigma_u)^{-1}Bf_t$. With $\Lambda = \lambda\Sigma_u$, the myopic portfolio corresponds to taking $\beta = 0$, and it evolves as $x_{t+1} = (\lambda/(\lambda+\gamma))x_t + (1/(\lambda+\gamma))\Sigma_u^{-1}Bf_t$.

Tables 1 and 2 summarize performance results. The robust results in Table 1 are based on allowing the changes in both returns (through u_t) and factor dynamics (through v_t); in Table 2, robustness is limited to v_t by omitting u_{t+1} from the original problem (4). As we discussed in §2.2, alternative but equivalent nonrobust objectives can lead to different robust problems.

The columns labeled "mean/std" report annualized performance ratios computed as

$$\sqrt{260} \times \text{Mean(daily \$ profit)}/$$
$$\text{Standard deviation(daily \$ profit)}.$$

We refer to these loosely as Sharpe ratios although they are calculated from differences ratherthan percentage changes

Table 2. In-sample performance comparisons using the full data series with robustness in factor dynamics (v_t) only.

		Obj × 10^{-6}				Mean/std			
	θ	Gross	t-stat	Net	t-stat	Gross	t-stat	Net	t-stat
Nonrobust		0.70		0.62		0.60		0.56	
Robust	10^{10}	0.70	0.40	0.62	0.53	0.60	0.80	0.56	0.89
	10^9	0.71	0.38	0.62	0.52	0.60	0.81	0.57	0.89
	10^8	0.72	0.35	0.65	0.48	0.61	0.98	0.57	1.05
	10^7	0.72	0.07	0.66	0.19	0.65	1.15	0.62	1.19
	10^6	0.51	−0.50	0.48	−0.37	0.72	0.97	0.68	0.96
	10^5	0.23	−1.02	0.22	−0.87	0.75	0.84	0.71	0.77

Note. For each θ, the t-stats compare performance of the robust rule at that θ with the nonrobust case, based on grouping the data into 40 batches.

because the assets are futures contracts—each contract has zero initial value, and total portfolio value can become negative. The columns labeled "Obj" report the objective function value

$$\text{Mean(daily \$ profit)} - \frac{\gamma}{2}\text{Variance(daily \$ profit)}.$$

The difference between gross and net performance is the effect of transaction costs.

To provide a rough indication of the statistical significance of our comparisons, we group the data into consecutive batches and calculate standard errors across batches. For each level of the robustness parameter θ, we calculate an approximate t-statistic (using the first estimator in Theorem 1 of Muñoz and Glynn 1997) for the difference in performance between the robust and nonrobust strategies. With sufficient stationarity and mixing in the underlying data, these statistics are indeed asymptotically t-distributed, although the conditions required are not guaranteed to hold in practice. We report results based on 40 batches and have obtained very similar results with 30 batches and with 50 batches.

Not surprisingly, ignoring transactions costs leads to good gross performance and terrible net performance in the first row of Table 1. The myopic rule produces less extreme differences, but overall poor results. These portfolios help illustrate the value of dynamic control rules. The nonrobust rule is optimized to the net objective function, so there is no reason to expect to see any benefit to robustness by this criterion. In Table 1, we see some deterioration in the net objective as we increase robustness (decrease θ); in Table 2, the net objective is relatively insensitive over a wide range of θ values. Interestingly, when we compare performance based on Sharpe ratios, for which none of the rules has been optimized, adding robustness appears to improve performance in both cases, although the differences are not significant as measured by our t-statistics. The in-sample improvement in the Sharpe ratio for the robust portfolios is primarily due to a reduction in the standard deviation in the denominator. At high robustness levels, the net excess return of the robust portfolio can be lower than that of the nonrobust portfolio, but with smaller standard deviation.

6.3. Out-of-Sample Tests

To compare out-of-sample performance, we reestimate the model parameters each week from 01/01/1996 through 04/09/2010 using the previous six months of data. Each time the parameters are estimated, the investment control rule remains fixed for one week until the parameters are next updated. Thus, at each point in time, the investment policy is based solely on prior market data. Updating the parameter estimates based on a rolling six-month window is also more reflective of how such a model would be used in practice.

Table 3 (with robustness to both u_t and v_t) and Table 4 (with robustness to v_t only) summarize the results. Over a wide range of θ values, the robust control rules show improved net performance as measured by either the objective function value or the Sharpe ratio. In effect, the robust rules acknowledge the uncertainty in the estimated model, and thus trade less aggressively than the nonrobust rule, and this improves out-of-sample performance. Allowing robustness to both the model of returns and the model of factor dynamics (Table 3) results in somewhat better results overall than focusing robustness on the factor dynamics.

As with the in-sample tests, the improvement mainly comes from the reduction in risk. In Table 5, t-statistics for the difference in net returns between the robust and nonrobust portfolios are estimated using batch means with 40 batches. None of the robust portfolios has a significantly better net return than that of the nonrobust portfolio.

Our subsequent analysis focuses on the less favorable case in which robustness is limited to the factor dynamics.

To illustrate the effect of robustness, Figure 1 shows the evolution of the positions in gold and crude oil under various strategies. Ignoring transaction costs leads to wild swings on a much wider scale, so we omit this case from the graph. Positions under the robust rules (shown at

Table 3. Out-of-sample performance comparisons using a rolling six-month estimation window with robustness in returns (u_t) and factor dynamics (v_t).

		Obj $\times 10^{-6}$				Mean/std			
	θ	Gross	t-stat	Net	t-stat	Gross	t-stat	Net	t-stat
No TC		−59.76		−8,257.60		0.53		−5.74	
Myopic		−11.80		−11.90		−0.57		−0.59	
Nonrobust		−36.78		−42.92		0.35		0.04	
Robust	10^9	−8.17	2.24	−10.50	2.32	0.39	0.93	0.15	3.07
	10^8	0.29	2.06	0.02	2.18	0.40	0.21	0.26	1.76
	10^7	0.10	1.99	0.08	2.13	0.46	0.38	0.37	1.73
	10^6	0.01	1.99	0.01	2.12	0.52	0.66	0.49	2.21
	10^5	0.00	1.99	0.00	2.12	0.53	0.69	0.52	2.32
	10^4	0.00	1.99	0.00	2.12	0.53	0.69	0.53	2.34

Note. For each θ, the t-stats compare performance of the robust rule at that θ with the nonrobust case, based on grouping the data into 40 batches.

Table 4. Out-of-sample performance comparisons using a rolling six-month estimation window with robustness in factor dynamics (v_t) only.

	θ	Obj $\times 10^{-6}$ Gross	t-stat	Net	t-stat	Mean/std Gross	t-stat	Net	t-stat
Nonrobust		−36.78		−42.92		0.35		0.04	
Robust	10^9	−31.16	1.69	−36.38	1.72	0.37	1.21	0.08	2.67
	10^8	−18.70	1.91	−22.16	1.95	0.40	1.06	0.16	2.65
	10^7	−5.10	2.05	−6.70	2.12	0.44	0.58	0.25	1.99
	10^6	−0.28	2.03	−0.85	2.13	0.42	−0.39	0.27	0.81
	10^5	0.27	2.01	0.11	2.13	0.37	−0.73	0.26	0.38
	10^4	0.18	2.00	0.14	2.13	0.41	−0.26	0.33	0.86

Note. For each θ, the t-stats compare performance of the robust rule at that θ with the nonrobust case, based on grouping the data into 40 batches.

Table 5. t-statistics of the difference of net returns between robust and nonrobust portfolios for out-of-sample tests, based on grouping the data into 40 batches.

θ	10^9	10^8	10^7	10^6	10^5	10^4
Robust in v only	0.96	0.69	0.33	0.03	−0.09	−0.12
Robust in v and u	0.24	−0.07	−0.14	−0.15	−0.15	−0.15

$\theta = 10^7$ and $\theta = 10^4$ with robustness to v_t only) fluctuate less than those chosen by the nonrobust rule. At the same time, by anticipating the evolution of the factors, the robust rules are quicker to respond than the myopic portfolio. The figures and numerical results suggest that $\theta = 10^7$ provides a reasonable level of robustness and $\theta = 10^4$ is overly conservative. The third panel scales the nonrobust positions to facilitate comparison. We discuss scaling strategies in §6.4.

Figure 2 compares net returns over the full time period and provides further insight into differences across strategies. Ignoring transaction costs results in disastrously poor performance, so this case is omitted from the figure. The performance of the myopic portfolio degrades over the time. Interestingly, much of the benefit of the robust rule, compared with the nonrobust rule, appears to be because of a small number of days. The nonrobust rule can outperform

Figure 1. Positions in gold in out-of-sample tests under various control rules.

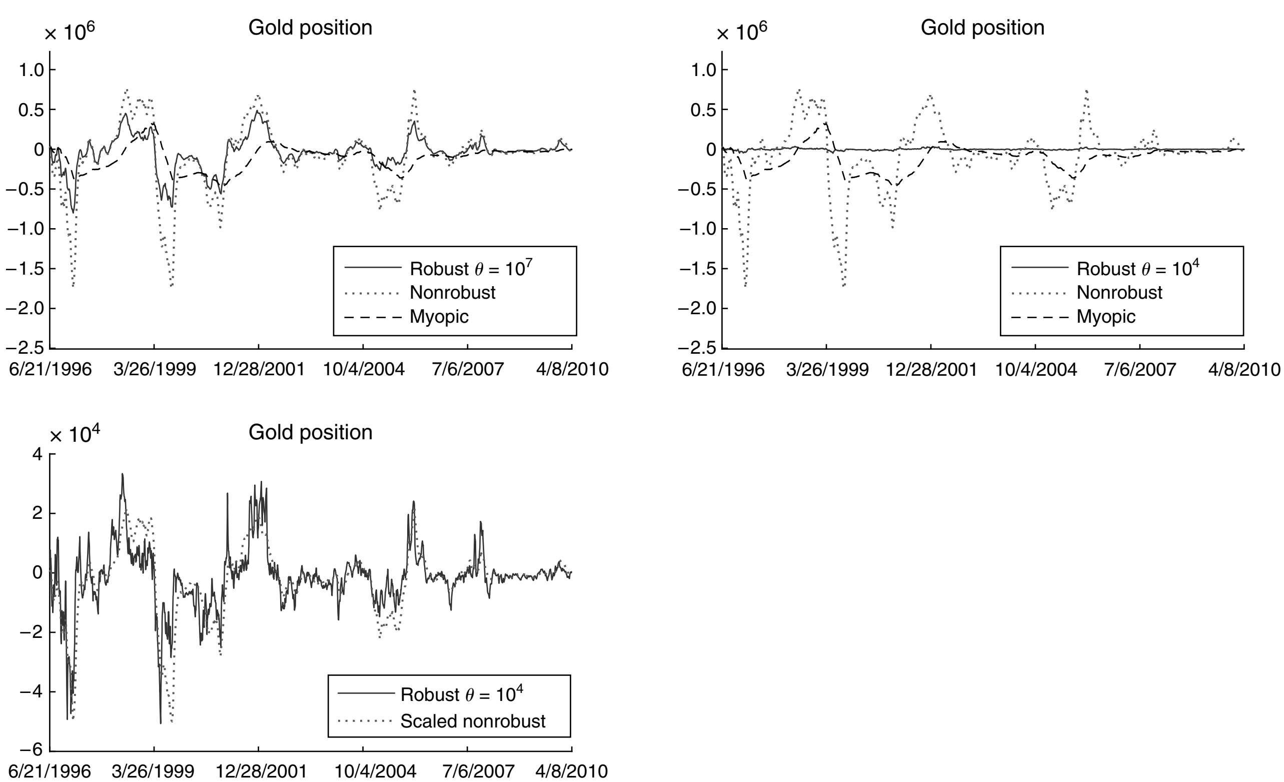

Note. The lower-left figure shows the positions for risk-scaled nonrobust portfolio and corresponding robust portfolio.

Figure 2. Net portfolio returns in out-of-sample tests.

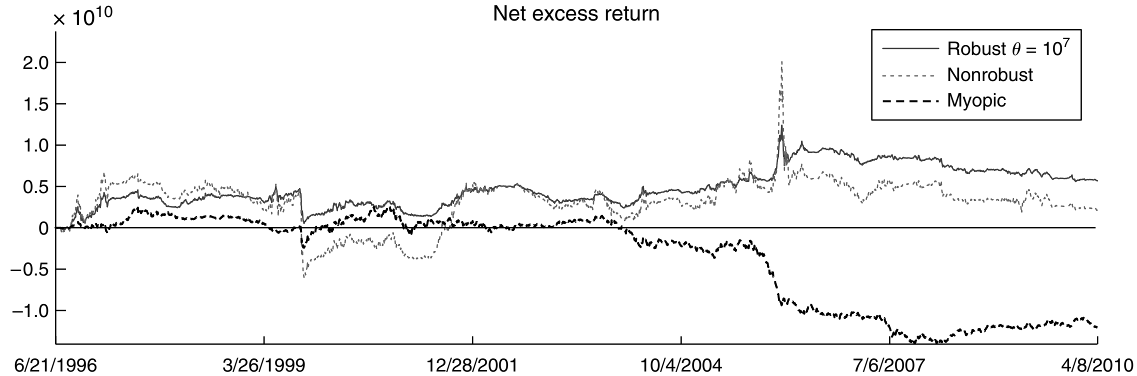

the robust rule over long periods of time; adding robustness reduces the impact of a small number of bad bets by trading less aggressively on the signals from the factors. Consistent with what we see in Figure 2, the improved Sharpe ratio under the robust rules results mainly from a smaller denominator rather than a larger numerator. We have also observed in QQ-plots (not included) that the tails of the out-of-sample distributions of daily returns of the robust portfolio are lighter than those of the nonrobust portfolio.

The largest losses in Figure 2 occur near September 27, 1999, and February 2, 2006, so we examine events around these days in greater detail. Leading up to this date, the loading matrix (B) and the mean-reversion matrix (C) were relatively slow moving. As shown in Figure 3, both portfolios had short positions in gold, although more aggressively under the nonrobust rule. On September 26, 15 European central banks signed an agreement to limit gold sales (Weber 2003); the price of gold rose 6% the next day and 11% the day after. This spike results in large losses for the short positions in our test portfolios, but the loss is tempered under the robust rule.

To ensure that our results are not overly influenced by a single day, we repeat the comparison removing days September 27–28, 1999, from the data. Table 6 shows that the robust portfolio still outperforms the nonrobust portfolio.

Whereas the large price change on September 27, 1999 was limited to gold, changes around February 2, 2006 were spread across multiple commodities, and the portfolio losses resulted from large positions rather than large price changes. The largest positions for both the robust and nonrobust portfolios on that date are in aluminum, zinc, gold, and sugar. The prices for these commodities are shown in Figure 4.

The position sizes for these commodities are shown in Figure 5. The steady price increases in the first half of Figure 4 lead to growing positions, particularly for the nonrobust portfolio. The positions change smoothly; this is consistent with the representation in (29)—more precisely, the infinite-horizon version without the superscripts (t, T)—of the portfolio as a weighted average of the previous position and a target position, together with the observation that the factors are moving smoothly as a consequence of the pattern of price changes. The large positions produce large losses on February 2. The price drop in sugar, for example, is barely perceptible, yet it produces

Figure 3. Parameter estimates near September 29, 1999.

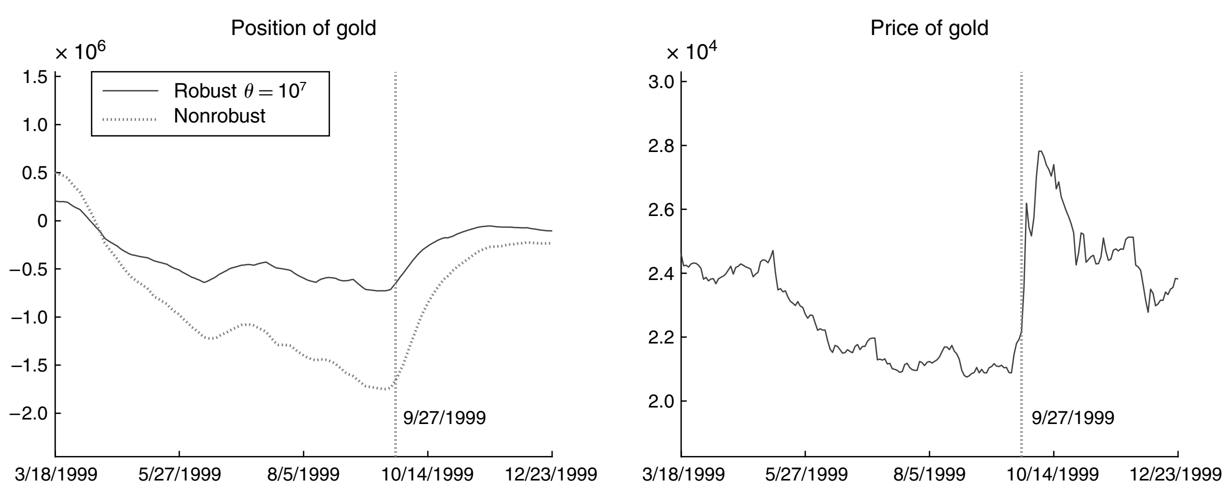

Table 6. Out-of-sample results with two extreme days (9/27–28/1999) removed from the data.

	θ	Obj $\times 10^{-6}$ Gross	t-stat	Net	t-stat	Mean/std Gross	t-stat	Net	t-stat
Nonrobust		−32.11		−37.98		0.49		0.18	
Robust	10^7	−3.89	1.89	−5.45	1.97	0.56	1.05	0.37	2.58
	10^6	−0.02	1.86	−0.58	1.98	0.48	−0.03	0.33	1.24

the largest losses of any of the commodities because of the large position accumulated. The robust portfolio suffers smaller losses because it is less aggressive in building up large positions in response to the increasing factor levels. Interestingly, the two portfolios hold fairly similar positions in zinc and gold, despite the large difference in their sugar positions. The nonrobust portfolio positions continue to grow quickly following the price drop. We attribute this, informally, to the nonrobust portfolio ascribing greater persistence to the factors.

Table 7 lists relative entropy values for in-sample tests with standard errors reported in parentheses. Using the results in Theorem 1, 10^4 sample paths with the same length as the history for in-sample tests are simulated using the estimated model, and the relative entropy is estimated using (6). For each M_t, the conditional relative entropy $E_t[m_{t+1}\log m_{t+1}]$ is calculated using the closed-form expression (6) in the online e-companion (available as supplemental material at http://dx.doi.org/10.1287/opre.2013.1180). To achieve the similar level of relative entropy for the robustness only in v with $\theta = 10^9$, one needs to set $\theta = 10^7$ when robustness in both v and u are considered. Both objective function and Sharpe ratio with $\theta = 10^7$ in Table 3 are better than those with $\theta = 10^9$ in Table 4.

Figure 4. Prices for aluminum, gold, zinc, and sugar before and after February 2, 2006.

Note. These are the commodities in which the portfolios hold the largest positions on that date.

Figure 5. Four largest positions for the nonrobust (left) and robust (right) portfolios around February 2, 2006.

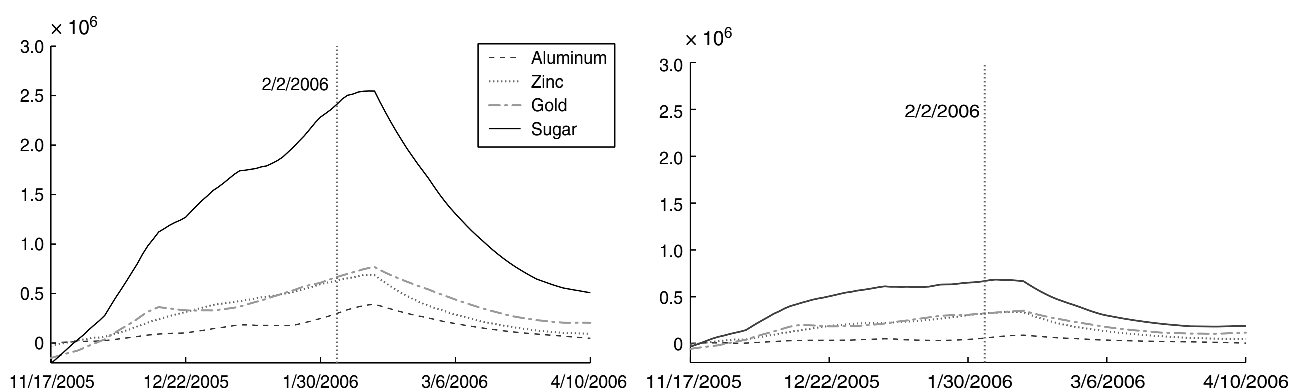

Table 7. Relative entropy for in-sample tests.

θ	Robustness in v	Robustness in v and u
10^{10}	0.020 (1.4×10^{-4})	0.017 (1.2×10^{-4})
10^{9}	0.30 (1.2×10^{-3})	0.17 (6.4×10^{-4})
10^{8}	1.1 (5.7×10^{-3})	0.27 (1.0×10^{-3})
10^{7}	2.4 (0.016)	0.31 (9.7×10^{-4})
10^{6}	5.0 (0.048)	0.40 (8.2×10^{-4})
10^{5}	12 (0.16)	0.44 (8.1×10^{-4})

Note. Standard errors are reported in parentheses.

The difference of these performance measurements suggests that the improvement is brought by considering the extra source of uncertainty from u.

For out-of-sample tests, there is no exact way to capture the relative entropy budget at each time, so we simply use the relative entropy for the same value of θ in in-sample tests.

6.4. Scaling and Trimming

In this section, we will compare the robust portfolios with simple heuristics to make the nonrobust portfolios less aggressive by adding constraints. We consider three alternatives.

Risk-scaled portfolio: For the out-of-sample test, at each time the model is updated, the position of the nonrobust portfolio is scaled by a factor. For given robustness level $\theta > 0$, the scaling factor is computed using the previous six-months' realized return so that the variance of the net return of the scaled portfolio equals that of the robust portfolio. The scaling factor is applied to those positions in the subsequent week. For the first six months, we still use their own performance for scaling.

Capital-scaled portfolio: First, define the total exposure to be the sum of absolute exposures, $\sum_i |x_{i,t}|p_{i,t}$, with $p_{i,t}$ being the price of the ith asset at time t. Whenever the total exposure of the nonrobust portfolio exceeds a predetermined threshold, it will be scaled down proportionately so that the total exposure of the scaled portfolio equals the threshold. We choose the maximum total exposure of a robust portfolio as the threshold.

Trimmed portfolio: Here we trim the nonrobust portfolio so that at any time t the position for each asset will be bounded by some upper and lower bounds. We set the bounds to be the maximum and minimum positions of the robust portfolio for each asset.

Tables 8, 9, and 10 report out-of-sample performance for these three portfolios. Columns labeled with "*RS*," "*CS*" and "*T*" refer to risk-scaled portfolio, capital-scaled portfolio, and trimmed portfolio, respectively. The $t-$statistics in parentheses, which compare performance of the robust portfolio with these three portfolios, indicate that none of these alternatives performs consistently as well as the robust method. Actually, the robust portfolio perform significantly better than these three portfolios as measured by the objective function when robustness is not too extreme, i.e., when θ is not too small.

Among the three constrained portfolios, the risk-scaled portfolios have relatively closer performance to the corresponding robust portfolios. Interestingly, there is a heuristic reason for this. Suppose that we scale down the nonrobust portfolio by a factor $s \in (0, 1)$, such that the positions of the resulting portfolio becomes $x_t^s = s x_t^\infty$, where x_t^∞ is the position of the nonrobust portfolio. Then

$$x_t^s = (I + 2\Lambda^{-1}A_{xx}^{(t,T)})x_{t-1}^s - 2\Lambda^{-1}A_{xx}^{(t,T)} \times s \times target,$$

where the matrix A is computed under the case $\theta = \infty$. Therefore, the scaled portfolio follows original nonrobust policy, but with scaled target. Gârleanu and Pedersen (2013, Proposition 3) show that under the specification $\Lambda = \lambda\Sigma_u$ for some $\lambda > 0$, the target portfolio of the nonrobust portfolio can be written as a discounted sum of expected myopic portfolios, $(\gamma\Sigma_u)^{-1}Bf_t$, at all future times. Therefore, scaling the target portfolio is very close to scaling up the risk-aversion parameter γ to $\gamma^s = \gamma/s$, although the discounting factor for myopic portfolios will change slightly when γ changes.

On the other hand, from (16), (17), and online e-companion A (available as supplemental material at http://dx.doi.org/10.1287/opre.2013.1180), robustness in the mean of u with $\theta > 0$ is equivalent to increasing γ to γ_θ. Thus, scaling down the nonrobust portfolio by s is close to considering robustness in the mean of u with

$$\theta = \frac{s}{(1-s)\beta\gamma}. \tag{38}$$

The performance of the risk-scaled portfolios in Table 10 is close to that of the corresponding robust portfolios. This suggests that most of the improvement is explained by the robustness in u, since the gap between the risk-scaled portfolio and the robust portfolio can be considered as the

Table 8. Average of scaling parameters for out-of-sample tests, so that the realized variance of the scaled nonrobust portfolio equals that of the robust portfolio.

θ	10^9	10^8	10^7	10^6	10^5	10^4
Robust in v only	0.96	0.81	0.53	0.27	0.11	0.037
Robust in v and u	0.53	0.11	0.013	1.5×10^{-3}	1.5×10^{-4}	1.5×10^{-5}
γ/γ_θ	0.5	0.091	9.9×10^{-3}	1.0×10^{-3}	1.0×10^{-4}	1.0×10^{-5}

Table 9. Out-of-sample performance comparisons using a rolling six-month estimation window with robustness factor dynamics (v_t) only.

		Obj × 10^6								Mean/std							
	θ	Gross	RS	CS	T	Net	RS	CS	T	Gross	RS	CS	T	Net	RS	CS	T
Nonrobust		−36.78				−42.92				0.35				0.04			
Robust	10^9	−31.16	−31.59	−35.43	−34.67	−36.38	−36.97	−41.29	−40.16	0.37	0.35	0.37	0.38	0.08	0.06	0.06	0.08
			(−0.93)	(−2.13)	(−2.43)		(−0.96)	(−2.22)	(−2.79)		(−1.24)	(−1.05)	(−0.35)		(−1.48)	(−2.64)	(−1.68)
	10^8	−18.70	−19.70	−30.45	−27.31	−22.16	−23.35	−35.70	−32.27	0.40	0.35	0.42	0.43	0.16	0.10	0.12	0.13
			(−1.22)	(−3.24)	(−2.99)		(−1.19)	(−3.41)	(−3.17)		(−1.14)	(−0.88)	(0.17)		(−1.19)	(−2.47)	(−1.32)
	10^7	−5.10	−6.70	−21.07	−14.63	−6.70	−8.18	−25.21	−18.37	0.44	0.33	0.41	0.53	0.25	0.14	0.13	0.24
			(−1.35)	(−3.80)	(−3.03)		(−1.32)	(−4.19)	(−3.49)		(−0.63)	(−0.33)	(1.14)		(−0.49)	(−1.56)	(0.09)
	10^6	−0.28	−0.80	−7.92	−6.59	−0.85	−1.33	−11.20	−9.41	0.42	0.31	0.34	0.42	0.27	0.18	0.10	0.10
			(−0.79)	(−3.52)	(−2.71)		(−0.70)	(−4.19)	(−3.54)		(0.29)	(1.01)	(0.75)		(0.56)	(0.44)	(−0.30)
	10^5	0.27	0.15	−1.1	−2.78	0.11	0.05	−1.93	−4.54	0.37	0.30	0.39	0.17	0.26	0.23	0.21	−0.16
			(−0.39)	(−1.18)	(−2.05)		(−0.19)	(−1.77)	(−3.03)		(0.62)	(1.87)	(0.08)		(1.03)	(1.75)	(−0.80)
	10^4	0.18	0.14	0.36	−0.50	0.14	0.12	0.18	−1.48	0.41	0.31	0.44	0.23	0.33	0.29	0.34	−0.10
			(−0.41)	(0.45)	(−0.93)		(−0.17)	(0.12)	(−1.96)		(0.19)	(1.41)	(0.11)		(0.64)	(1.53)	(−0.77)

Notes. Scaled portfolio is derived from scaling the nonrobust portfolio so that the realized variance of its net return equals that of the corresponding robust portfolio. (RS) is for risk-scaled portfolio, and (CS) is for capital-scaled portfolio. (T) indicates trimmed portfolio is derived by trimming the position of the nonrobust portfolio so that the positions of trimmed portfolio are bounded by the positions of corresponding robust portfolio. The t-statistics in parentheses compare performance of the robust rule with scaling or trimming, based on grouping the data into 40 batches.

Table 10. Out-of-sample performance comparisons using a rolling six-month estimation window with robustness in returns (u_t) factor dynamics (v_t).

		Obj × 10^6								Mean/std							
	θ	Gross	RS	CS	T	Net	RS	CS	T	Gross	RS	CS	T	Net	RS	CS	T
Nonrobust		−36.78				−42.92				0.35				0.04			
Robust	10^9	−8.17	−9.25	−26.98	−18.18	−10.5	−11.05	−31.85	−22.37	0.39	0.35	0.42	0.53	0.15	0.18	0.13	0.23
			(−1.83)	(−3.76)	(−3.74)		(−1.07)	(−4.11)	(−4.39)		(−0.75)	(−0.72)	(1.25)		(1.75)	(−2.41)	(0.01)
	10^8	0.29	0.16	−4.54	−4.11	0.02	0.07	−6.18	−6.24	0.40	0.35	0.34	0.19	0.26	0.31	0.11	0.15
			(−1.06)	(−2.59)	(−2.44)		(0.38)	(−3.25)	(−3.41)		(−0.07)	(0.97)	(−0.75)		(1.22)	(0.47)	(−1.72)
	10^7	0.10	0.08	0.31	−0.10	0.08	0.08	0.25	−0.62	0.46	0.35	0.44	0.17	0.37	0.34	0.38	−0.14
			(−0.59)	(0.98)	(−0.51)		(0.10)	(0.80)	(−1.31)		(−0.21)	(1.10)	(−0.72)		(0.49)	(1.47)	(−1.45)
	10^6	0.01	0.01	0.04	0.05	0.01	0.01	0.04	−0.01	0.52	0.34	0.43	0.22	0.49	0.34	0.42	−0.00
			(−0.73)	(1.37)	(0.18)		(−0.52)	(1.38)	(−0.39)		(−0.46)	(0.62)	(−0.60)		(−0.24)	(0.82)	(−1.14)
	10^5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.53	0.35	0.43	0.12	0.52	0.35	0.43	0.05
			(−0.71)	(1.41)	(0.05)		(−0.68)	(1.41)	(−0.15)		(−0.49)	(0.53)	(−0.77)		(−0.46)	(0.56)	(−0.95)
	10^4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.53	0.35	0.43	0.10	0.53	0.35	0.43	0.09
			(−0.70)	(1.41)	(−0.01)		(−0.69)	(1.41)	(−0.04)		(−0.49)	(0.52)	(−0.82)		(−0.49)	(0.53)	(−0.84)

Notes. Scaled portfolio is derived from scaling the nonrobust portfolio so that the realized variance of its net return equals that of the corresponding robust portfolio. (RS) is for risk-scaled portfolio, and (CS) is for capital-scaled portfolio. (T) indicates trimmed portfolio is derived by trimming the position of the nonrobust portfolio so that the positions of trimmed portfolio are bounded by the positions of corresponding robust portfolio.The t-statistics in parentheses compare performance of the robust rule with scaling or trimming, based on grouping the data into 40 batches.

extra benefit brought by considering robustness in v. This is consistent with the observation in Table 7, where the performance of the robust portfolio considering only the uncertainty in v is much less than the portfolio with robustness in both u and v at the same relative entropy level.

In Figure 1, the lower-left figure shows the position of gold for risk-scaled portfolio and the corresponding robust portfolio with $\theta = 10^4$. The robust portfolio is different from the risk-scaled portfolio, especially when it has some extreme positions.

Table 8 reports average scaling parameters over time. For the cases with robustness in both v and u, the scaling parameters are very close to γ/γ_θ, which supports our observation on the effect of scaling.

7. Concluding Remarks

We have developed robust portfolio control rules using a stochastic and dynamic notion of robustness to model error. Our analysis covers both finite- and infinite-horizon multiperiod problems. We work with a factor model of returns, in which factors evolve stochastically. The relationship between returns and factors and the evolution of the factors are subject to model error and are treated robustly. We incorporate transaction costs and develop simple optimal controls that remain tractable for multiple assets. Robustness significantly improves performance in out-of-sample tests on historical data.

Using this approach requires choosing a value for the parameter θ, which controls the degree of robustness or pessimism. In principle, one would want to select this parameter to reflect the reliability of a model based on the data available to support it. Conveniently, we find that our results are consistent over a wide range of θ values, so the exact choice of this parameter does not dominate our empirical results. Methods for selecting this parameter nevertheless remain an interesting topic for further investigation.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/opre.2013.1180.

Acknowledgments

The authors thank Nicolae Gârleanu and Lasse Pedersen for helpful discussions concerning their data and empirical results, and Ciamac Moallemi for helpful comments. They also thank the anonymous associate editor and referees for their very helpful comments. This research is supported by the National Science Foundation [Grant DMS 09-14539] and a fellowship from the Centre for Risk Studies at the University of Cambridge.

References

Anderson EW, Hansen LP, Sargent TJ (2003) A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *J. Eur. Econom. Assoc.* 1(1):68–123.

Asness C, Moskowitz T, Pedersen L (2013) Value and momentum everywhere. *J. Finance* 68(3):929–985.

Basak S, Chabakauri G (2010) Dynamic mean-variance asset allocation. *Rev. Financial Stud.* 23(8):2970–3016.

Ben-Tal A, Margalit T, Nemirovski A (2000) Robust modeling of multistage portfolio problems. *High Performance Optim.* 303–328.

Bertsimas D, Pachamanova D (2008) Robust multiperiod portfolio management in the presence of transaction costs. *Comput. Oper. Res.* 35(1):3–17.

Bielecki TR, Pliska SR (1999) Risk-sensitive dynamic asset management. *Appl. Math. Optim.* 39(3):337–360.

Bielecki TR, Pliska SR, Sheu SJ (2005) Risk sensitive portfolio management with Cox–Ingersoll–Ross interest rates: The HJB equation. *SIAM J. Control Optim.* 44(5):1811–1843.

Campbell JY, Viceira LM (2002) *Strategic Asset Allocation: Portfolio Choice for Long-term Investors* (Oxford University Press, New York).

Chen Z, Epstein L (2003) Ambiguity, risk, and asset returns in continuous time. *Econometrica* 70(4):1403–1443.

Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* 58(3):595–612.

DeMiguel V, Garlappi L, Nogales FJ, Uppal R (2009) A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Sci.* 55(5):798–812.

Duffie D, Epstein LG (1992) Stochastic differential utility. *Econometrica: J. Econometric Soc.* 60(2):353–394.

El Ghaoui L, Oks M, Oustry F (2003) Worst-case value-at-risk and robust portfolio optimization: A conic programming approach. *Oper. Res.* 51(4):543–556.

Erb CB, Harvey CR (2006) The strategic and tactical value of commodity futures. *Financial Analysts J.* 62(2):69–97.

Fleming WH, Sheu SJ (2001) Risk-sensitive control and an optimal investment model. *Math. Finance* 10(2):197–213.

Gârleanu NB, Pedersen LH (2013) Dynamic trading with predictable returns and transaction costs. *J. Finance*. Forthcoming.

Goh J, Sim M (2010) Distributionally robust optimization and its tractable approximations. *Oper. Res.* 58(4):902–917.

Goldfarb D, Iyengar G (2003) Robust portfolio selection problems. *Math. Oper. Res.* 28(1):1–38.

Hansen LP, Sargent TJ (1995) Discounted linear exponential quadratic Gaussian control. *Automatic Control, IEEE Trans.* 40(5):968–971.

Hansen LP, Sargent TJ (2007) *Robustness* (Princeton University Press, Princeton, NJ).

Hansen LP, Sargent TJ, Turmuhambetova G, Williams N (2006) Robust control and model misspecification. *J. Econom. Theory* 128(1):45–90.

Iyengar GN (2005) Robust dynamic programming. *Math. Oper. Res.* 30(2):257–280.

Jacobson D (1973) Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *Automatic Control, IEEE Trans.* 18(2):124–131.

Lim AEB, Shanthikumar JG (2007) Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* 55(2):198–214.

Lim AEB, Shanthikumar JG, Watewai T (2011) Robust asset allocation with benchmarked objectives. *Math. Finance* 21(4):643–679.

Liu H (2010) Robust consumption and portfolio choice for time varying investment opportunities. *Ann. Finance* 6(4):435–454.

Maenhout PJ (2004) Robust portfolio rules and asset pricing. *Rev. Financial Stud.* 17(4):951–983.

Moallemi C, Saglam M (2012) Dynamic portfolio choice with linear rebalancing rules. Working paper, SSRN 2011605.

Moskowitz T, Ooi YH, Pedersen LH (2012) Time series momentum. *J. Financial Econom.* 104:228–250.

Muñoz DF, Glynn PW (1997) A batch means methodology for estimation of a nonlinear function of a steady-state mean. *Management Sci.* 43(8):1121–1135.

Natarajan K, Pachamanova D, Sim M (2008) Incorporating asymmetric distributional information in robust value-at-risk optimization. *Management Sci.* 54(3):573–585.

Natarajan K, Sim M, Uichanco J (2010) Tractable robust expected utility and risk models for portfolio optimization. *Math. Finance* 20(4):695–731.

Pesaran MH, Timmermann A (2012) Predictability of stock returns: Robustness and economic significance. *J. Finance* 50(4):1201–1228.

Petersen IR, James MR, Dupuis P (2000) Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Trans. Automatic Control* 45(3):398–412.

Ruszczyński A (2010) Risk-averse dynamic programming for Markov decision processes. *Math. Programming* 125(2):235–261.

Skiadas C (2003) Robust control and recursive utility. *Finance Stochastics* 7(4):475–489.

Uppal R, Wang T (2003) Model misspecification and underdiversification. *J. Finance* 58(6):2465–2486.

Weber EJ (2003) The misuse of central bank gold holdings. Tcha MJ, ed. *Gold and the Modern World Economy* (Routledge, London), 64–80.

Whittle P (1981) Risk-sensitive linear/quadratic/Gaussian control. *Adv. Appl. Probab.* 13(4):764–777.

Whittle P, Whittle PR (1990) *Risk-Sensitive Optimal Control* (Wiley, Chichester, UK).

Paul Glasserman is the Jack R. Anderson Professor of Business at the Columbia Business School, Columbia University.

Xingbo Xu is a Ph.D. candidate in the Industrial Engineering and Operations Research Department at Columbia University.