Mostly Prior-Free Asset Allocation

Sylvain Chassang*

New York University

Forthcoming, *Journal of Risk*

Abstract

This paper develops a prior-free version of Markowitz (1952)'s efficient portfolio theory that allows the decision maker to express preferences over risk and reward, even though she is unable to express a prior over potentially non-stationary returns. The corresponding optimal allocation strategies are admissible, interior, and exhibit a form of momentum. Empirically, prior-free efficient allocation strategies successfully exploit time-varying risk premium present in historical returns.

KEYWORDS: prior-free asset allocation, drawdown control, non-stationary returns, time-varying risk premium, fear-of-missing-out, fear-of-loss, regret aversion, cost of robustness, approximate sample optimality.

1 Introduction

Financial markets are not stationary: they can change in durable ways. Sometimes change is anticipated. For instance, US treasury yields, which have been going down over the last 30 years, mechanically cannot keep going down much longer (see Figure 1). In this case, we know that the next 30 years must look different. Sometimes, change is only a possibility that decision makers are concerned with. For instance, an investor interested in investing

*I'm particularly indebted to Markus Brunnermeier and Harrison Hong for guidance and encouragement. The paper benefited from feedback by Yacine Aït-Shalia, Marie Brière, Valentin Haddad, Augustin Landier, Alexey Medvedev, Ulrich Müller, Wolfgang Pesendorfer, Chris Sims, David Sraer, Annette Vissen-Jorgensen, Pablo Winant, as well as seminar participants at Amundi, the Bank of England, Berkeley, Capital Fund Management, Toulouse, and Princeton. Alice Wang provided excellent research assistance.

in a smart-beta index-fund exploiting one of the familiar premium anomalies (e.g. value, momentum, low volatility, low beta ...) may be plausibly worried that those strategies will become crowded and fail to deliver advertised returns. In a non-stationary environment, past data provides limited guidance on future behavior which begs the following question: how to make practical risk management and portfolio allocation decisions in such a non-stationary world?

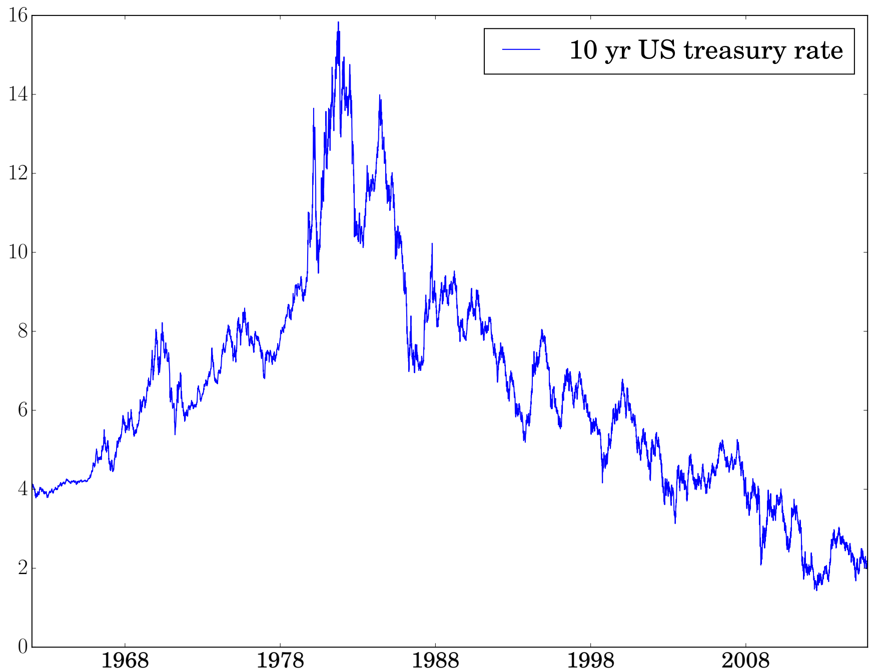

Figure 1: 10 year US treasury rate, 1962–2014.[1]

The benchmark framework for portfolio allocation, Markowitz (1952)'s efficient portfolio theory, is normatively attractive but requires the decision maker to specify priors over potential returns. This turns out to be practically difficult, even in static settings. Indeed, Black and Litterman (1992) show that when historical data is used to estimate a distribution of returns, plausible implementations of mean-variance optimal portfolios lead to sensitive corner allocations that are intuitively unappealing. In response, they suggest anchoring priors to a neutral prior under which owning a value-weighted portfolio of all assets is optimal. In

[1]Soure: FRED®, `https://research.stlouisfed.org/fred2/series/DGS10`.

dynamic environments with a time-varying and possibly non-stationary risk-premium, the difficulty of specifying priors is further increased.[2] The decision-maker must specify beliefs over the entire sequence of returns, which is tricky: in high dimensional state spaces, even full-support priors can have poor frequentist behavior (for instance failing to converge to true parameters, e.g. Sims, 1971, Diaconis and Freedman, 1986, Ghosh and Ramamoorthi, 2003). As a consequence, neither Markowitz (1952) nor Black and Litterman (1992) provide robust practical frameworks to guide asset allocation in dynamic environments where the process for returns may be non-stationary. Prior-free asset allocation seeks to provide such a framework by giving up on priors altogether.

The logic behind prior-free asset allocation matches Harry Markowitz' description of his actual rather than theoretical approach to portfolio construction (quoted in Zweig, 2007):

> "[...] I visualized my grief if the stock market went way up and I wasn't in it – or if it went way down and I was completely in it. My intention was to minimize my future regret. So I split my contributions 50/50 between bonds and equities."

This paper formalizes Markowitz' intuitive approach as an aversion to worst-case drawdowns (i.e. peak-to-trough losses) relative to reference safe and risky assets, in this case bonds and equities. It solves for the corresponding optimal dynamic asset allocation policy and argues that it provides a systematic framework for asset allocation in non-stationary, or novel environments. One simple takeaway is that Markowitz' 50/50 strategy is optimal in one-shot settings, but dominated in dynamic ones.

The model considers an agent who seeks to minimize the worst-case drawdowns of her portfolio relative to benchmark risky and risk-free assets (say the aggregate stock market and short-term US treasuries). As in other models of non-Bayesian decision making, such as the Gilboa and Schmeidler (1989) model of ambiguity aversion, the framework is game-theoretic. Nature is an adversary who seeks to maximize the agent's drawdowns relative to reference assets. In turn, the agent chooses the dynamic allocation policy that is least gameable by

[2] See Campbell (1984), Campbell and Viceira (1999) or Lettau and Ludvigson (2001, 2010) for evidence of time-variation in risk-premia.

nature. This yields a set of dynamic allocation strategies that achieve minimal worst-case drawdowns relative to all possible sequences of returns. Since these strategies are defined without reference to a prior over returns, the paper refers to these strategies as prior-free optimal.

The paper makes four points. The first is that prior-free optimal portfolios satisfy a form of robustness to non-stationarity which is not satisfied by more obvious approaches to asset allocation under a time-varying risk-premium. Intuitively, asset allocation strategies that experience large drawdowns with respect to either the safe or the risky asset misjudge average geometric returns over a large time window. More generally large drawdowns can be interpreted as sample violations of optimality conditions. Because a prior-free optimal allocation strategy guarantees small worst-case drawdowns there is no large time window over which it makes ex post suboptimal allocation choices. In contrast, any allocation strategy that is Bayesian-optimal for a full-support prior over finite hidden Markov models (Baum and Petrie, 1966) is gameable by nature: there exists a sequence of returns for which it experiences large drawdowns compared to one of the two reference assets.

The second point is that prior-free asset allocation lets decision makers express preferences over risk and reward in the same way that modern portfolio theory does. Indeed, prior-free optimal strategies define an entire frontier of minimal drawdowns. At one extreme, being fully invested in the market guarantees zero drawdowns against the market at the cost of high worst-case drawdowns against the safe asset. Inversely, being fully invested in the safe asset guarantees zero drawdowns against the safe asset at the cost of high potential drawdowns against the risky asset. The frontier of points in-between lets the decision maker express tradeoffs between fear-of-losing (drawdowns against the safe asset) and fear-of-missing-out (drawdowns against the risky asset). Points on this frontier map to dynamic allocation strategies that move smoothly from aggressive to cautious.

The third point is that prior-free optimal strategies are amenable to numerical computation. The agent's worst-case drawdown minimization problem admits a Bellman representation in which returns are not exogenously drawn from a prior, but rather endogenously picked by nature. The corresponding value function and strategy can be expressed as a function of a four dimensional summary statistic of past history. This representation allows to establish some theoretical results of interest: prior-free optimal portfolios are largely interior, and they satisfy a form of momentum. A multi-asset version of the prior-free framework admits an equally tractable representation after appropriate relaxation.

Finally, the paper provides a brief numerical and empirical exploration of this prior-free approach to portfolio optimization. Under worst-case analysis, prior-free portfolios improve significantly over popular portfolio construction rules, such as regularly rebalanced portfolios (RRP) and constant proportion portfolio insurance (CPPI, Black and Perold (1992)), both of which sit quite far away from the minimal drawdown frontier. More generally, this provides a systematic framework in which to evaluate technical trading rules: any fully specified allocation strategy (this includes RRP, CPPI, time-series momentum, volatility control, moving average rules) can be mapped against the worst-case drawdown frontier. The benefits of any strategy of interest (e.g. the in-sample performance of a volatility-control strategy) can then be weighed against the potential worst-case drawdowns it may experience.

In principle, the high degree of robustness required from prior-free optimal strategies may come at a cost. Indeed, if the true process for returns were i.i.d., a fixed regularly rebalanced portfolio may deliver better performance than a prior-free optimal portfolio whose allocation changes with realized market returns. This is not the case in the empirical sample of returns. Prior-free optimal strategies perform well in the historical time-series of returns, which suggests that they are able to capture time-variation in risk-premium present in the data. This is confirmed by a Henriksson and Merton (1981) test. Prior-free portfolios achieve asymmetric β exposures to the market in good and bad years (.7 vs .4). Importantly, there is little scope for data-snooping bias (Lo and MacKinlay, 1990) when backtesting prior-free optimal strategies. Prior-free optimal portfolios have a single free parameter, the potential magnitude of moves of nature, and it is set in advance of any exposure to data.

The paper connects to an applied literature in portfolio management that seeks to usefully operationalize Markowitz (1952)'s approach. Black and Litterman (1992) also place priors at the center of their analysis. They show that naïve implementations of Markowitz (1952) are extremely sensitive to prior assumptions over returns, or equivalently, to the sample of data used to calibrate parameters. To address the issue, they suggest anchoring priors to a default prior that rationalizes owning the market portfolio. Constant proportion portfolio insurance (CPPI), developed in Perold (1986), Black and Jones (1987) and Black and Perold (1992), also avoids priors and proposes a simple class of investment rules that provide risky upside exposure, while providing prior-free downside risk protection. The approach proceeds by using a cushion of safe assets, and leveraging funds above this cushion. However, CPPI can experience large drawdowns with respect to both the safe asset and the risky asset. Grossman and Zhou (1993) tackle the issue of drawdown control in a Bayesian setting where a fund manager wishes to exploit an asset with known fixed expected returns, but is subject to drawdown constraints versus a safe asset.

The worst-case approach emphasized in this paper is related to models of ambiguity aversion axiomatized by Gilboa and Schmeidler (1989) and to multiplier preferences popularized in Macroeconomics and Finance by Hansen and Sargent (2001, 2008). Cai et al. (2000) and Pflug and Wozabal (2007) apply the ambiguity averse framework to static portfolio construction, where it leads to more conservative allocations. Glasserman and Xu (2013, 2014) extend the approach to dynamic environments with trading costs and argue that it leads to better out-of-sample performance. Note that models based on multiplier preferences still rely on an anchoring prior that nature can perturb at a cost. This theoretical literature has an applied counterpart (see for instance Ceria and Stubbs, 2006, Asl and Etula, 2012) that seeks to better take into account model uncertainty when making portfolio allocation decisions.

Because drawdowns use reference assets to benchmark performance, the preferences explored in this paper are related to regret-averse and reference-dependent preferences that have received attention in the statistical (Wald, 1950, Savage, 1951, Milnor, 1954, Stoye, 2008) and behavioral literatures (Tversky and Kahneman, 1991, Kőszegi and Rabin, 2006). It is closely related to the question of online regret minimization originally studied in Blackwell (1956), Hannan (1957) (see Cesa-Bianchi and Lugosi, 2006, for a recent reference). The portfolio allocation problem studied here is also connected to Cover (1991) and DeMarzo et al. (2009), both of which derive prior-free lower bounds on portfolio performance.

More broadly, this paper contributes to a growing agenda in economics which seeks to rethink economic design questions from a prior-free perspective. Segal (2003), Bergemann and Schlag (2008), Hartline and Roughgarden (2008), Madarász and Prat (2016) and Brooks (2014) study auctions, pricing, and screening. Chassang (2013), Carroll (2014) and Antic (2014) study incentive provision. The current paper adds to this agenda in two ways. First, it provides a prior-free version of Markowitz (1952) which allows decision makers to express meaningful preferences over risk and reward while allowing for arbitrary non-stationarity in returns. Second, it provides an empirical evaluation of prior-free approaches in a practical context. It shows that the cost of robustness need not be large, and that in fact, prior-free optimization may improve over existing benchmarks in the realized historical sample. In addition, prior-free approaches reduce concerns of data-snooping bias.

The paper is structured as follows. Section 2 defines the framework and the prior-free asset allocation problem. Section 3 quantifies robustness to non-stationarity and shows that it is not achieved by a natural class of Bayesian optimal policies. Section 4 provides a general Bellman characterization of prior-free optimal allocation policies. Section 5 establishes qualitative properties assuming that trading costs are equal to zero. Section 6 extends the framework to multiple assets and reinterprets low drawdowns as sample versions of optimality conditions. Section 7 provides a brief empirical evaluation of prior-free asset allocation strategies. Section 8 concludes. Appendix A extends the empirical analysis and discusses decision-theoretic foundations, as well as possible Bayesian refinements, of the prior-free approach. Proofs are contained in Appendix B unless mentioned otherwise.

2 Framework

2.1 Setup

Returns. An investor with finite horizon $N \in \mathbb{N}$ allocates resources across two assets: a safe asset with returns at time $t \in \mathbb{N}$ denoted by r_t^0, as well as a risky asset — say the market — with returns at time t denoted by r_t^1 (see Section 6 for an extension to multiple risky assets). The set of possible returns $r_t = (r_t^0, r_t^1)$ at time t is denoted by $M \subset \mathbb{R}^2$ and referred to as moves of nature. For simplicity, and anticipating computational implementation, it is assumed that set M is finite and satisfies the following minimal richness assumption.

Assumption 1. *There exists $r \in M$ such that $r^0 = r^1$. There exist $(r, \widehat{r}) \in M^2$ such that $r^0 > r^1$, and $\widehat{r}^0 < \widehat{r}^1$. Set M contains at least three non-diagonal returns r such that $r^0 \neq r^1$.*

In computational applications, set M will take the form $M = \{r^0\} \times \{r^0 + n\Delta | n = -k, \cdots, +k\}$ for $k \in \mathbb{N}$. Let us denote by $\overline{r} = \max\{|r^0|, |r^1| \text{ for } r \in M\}$ an upper bound to the magnitude of returns in M.

Allocations. The set of possible allocations $A = \Delta(\{0, 1\}) \subset \mathbb{R}$ is compact and convex. An allocation $a_t \in A \subset \mathbb{R}^2$ at the beginning of period t yields a return $r_t^a = \langle a_t, r_t \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the usual dot product.

Given returns r_t for period t, and invested wealth w_t at t, wealth and asset shares at the beginning of period $t + 1$ (denoted $t + 1^-$) are given by

$$w_{t+1^-} = w_t(1 + r_t^a); \quad a_{t+1^-}^0 = \frac{a_t^0(1 + r_t^0)}{1 + r_t^a}; \quad a_{t+1^-}^1 = \frac{a_t^1(1 + r_t^1)}{1 + r_t^a}. \tag{1}$$

It is possible to reallocate assets at the beginning of each period, but reallocation is costly. Specifically, moving from a_{t+1^-} to a_{t+1} costs a proportion $c(a_{t+1^-}, a_{t+1}) \geq 0$ of the existing asset base. Denote by $\overline{c} \equiv \max_{a,\widehat{a}} c(a, \widehat{a})$ the highest trading cost. In numerical applications, trading costs will take the form $c(a, \widehat{a}) = c_1|a^0 - \widehat{a}^0|$, with $c_1 = .002$ (i.e. 20 bps). Invested wealth after reallocation is $w_{t+1} = [1 - c(a_{t+1^-}, a_{t+1})]w_{t+1^-}$.

Allocation strategies. An allocation strategy α maps each history of returns $h_t = (r_s)_{s\in\{0,\cdots,t-1\}} \in \mathcal{H}$ to an allocation a_t. We denote by $\mathcal{A}$ the set of possible allocation strategies. Taking into account trading costs, the returns associated to allocation strategy α in period t, denoted by r_t^α take the form

$$1 + r_t^\alpha = (1 + \langle \alpha(h_t), r_t \rangle)(1 - c(a_{t^-}, \alpha(h_t)). \tag{2}$$

For a state space Z sufficiently large, any allocation strategy can be described as an automaton depending on state $z \in Z$ with transition rule ϕ:

$$\alpha : z \in Z \mapsto \alpha(z) \in A \tag{3}$$

$$\phi : (z, r) \in Z \times M \mapsto \phi(z, r) \in Z. \tag{4}$$

An initial allocation a_0 and a sequence of returns $\mathbf{r} = (r_t)_{t\in\mathbb{N}}$ induces the sequence of allocations $(a_t)_{t\geq 1}$ defined by

$$\forall t \geq 1, \quad z_t = \phi(z_{t-1}, r_{t-1}) \quad \text{and} \quad a_t = \alpha(z_t).$$

This paper seeks to formalize the following normative question: what are good dynamic asset allocation strategies for a decision maker who worries that returns may be arbitrarily non-stationary?

2.2 Bayesian optimal asset allocation

A standard model of dynamic asset allocation might take the following form. A decision maker with investment horizon N and log-utility over final wealth is able to place a prior $\mu \in \Delta(M^{N+1})$ over possible returns. Her optimal asset allocation strategy α then solves

$$\max_{\alpha\in\mathcal{A}} \mathbb{E}_\mu \left[\sum_{t=0}^{N} \log(1 + r_t^\alpha) \right]. \tag{5}$$

Unfortunately, this positive description of behavior provides little normative guidance to investors. This becomes particularly clear after mapping the problem of choosing an asset allocation strategy in the axiomatic framework of subjective utility theory (Savage, 1972). A realized sequence of returns $\mathbf{r} = (r_t)_{t\geq 0}$ is an event. A dynamic asset allocation strategy is a Savage act mapping events to financial outcomes for the decision maker. Provided that a decision-maker has well-behaved preferences over acts, subjective expected utility theory tells us that the decision-maker's behavior can be represented as maximizing an expected utility function.

Subjective utility theory obviously isn't a normative framework. Priors are inferred from preferences over acts; optimal acts are not obtained from priors. Still, subjective utility theory is routinely used for normative purposes. Black and Litterman (1992) implicitly highlight some of the difficulties that normative uses of subjective expected utility generate. They specify a Gaussian prior over returns, and use historical estimates to set mean and covariance parameters. Presuming mean-variance preferences, they show that such beliefs imply extreme corner allocations that are intuitively unappealing. They then propose using priors that would justify holding a value-weighted portfolio. Black and Litterman (1992)'s ex post assessment that extreme allocations are unappealing, and their response — modifying priors until they yield a more palatable allocation — demonstrate that priors are an output, inferred from preferences over actions, not a primitive of the decision problem.

The normative decision rule that would start by eliciting beliefs and then maximizing utility is even more tricky to implement in the dynamic setting considered in this paper. When the process for returns is potentially non-stationary, picking well behaved priors turns out to be difficult. The literature on frequentist properties of Bayesian estimates (Diaconis and Freedman, 1986, Ghosh and Ramamoorthi, 2003) shows that generic priors over large dimensional objects (here sequences of returns) may fail to satisfy consistency properties that common frequentist estimators robustly satisfy. Section 3 makes this point concretely.

2.3 Mostly prior-free asset allocation

This paper is written from a normative perspective. It specifies preferences over allocation strategies, argues that they are intuitively appealing, and studies the strategies that maximize them. Markowitz' description of his own investment behavior, suggests the three following observations:

- decision-makers fear net losses;
- decision-makers fear missing out on potential gains;
- decision-makers do not have sophisticated beliefs over patterns of returns.

Definitions 1, 2, 3 formalize preferences over allocation strategies that capture these premises.

Definition 1 (relative drawdowns)**.** *Given an allocation strategy* $\alpha : \mathcal{H} \rightarrow \Delta(\{0,1\})$ *and a sequence of returns* $\mathbf{r} \equiv (r_t)_{t\in\{0,\cdots,N\}}$, *drawdowns* $\mathcal{D}^0_N$ *and* $\mathcal{D}^1_N$ *relative to the safe and risky asset are defined as*

$$\mathcal{D}^0_N(\alpha,\mathbf{r}) = \max_{\substack{T\in\{0,\cdots,N\}\\ T'\in\{0,\cdots,T+1\}}} \sum_{t=T'}^{T} \log(1+r^0_t) - \log(1+r^\alpha_t) \tag{6}$$

$$\mathcal{D}^1_N(\alpha,\mathbf{r}) = \max_{\substack{T\in\{0,\cdots,N\}\\ T'\in\{0,\cdots,T+1\}}} \sum_{t=T'}^{T} \log(1+r^1_t) - \log(1+r^\alpha_t). \tag{7}$$

Given realized returns $\mathbf{r} \equiv (r_t)_{t\in\{0,\cdots,N\}}$, the relative drawdowns of strategy α correspond to strategy α's maximum relative losses against the safe and risky assets over arbitrary subperiods $[T',T] \subset [0,N]$. Figure 2 shows how to compute the drawdowns of a 50/50 fixed-weights strategy over the market and the risk-free rate during the 2007–2012 period.[3] Note that the time window $[T',T]$ over which each drawdown occurs are different for the risk-free and risky assets.

Definition 2 (worst-case drawdowns)**.** *Given a strategy* α, *worst-case drawdowns are defined by*

[3]Returns are obtained from Kenneth French's data library available at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

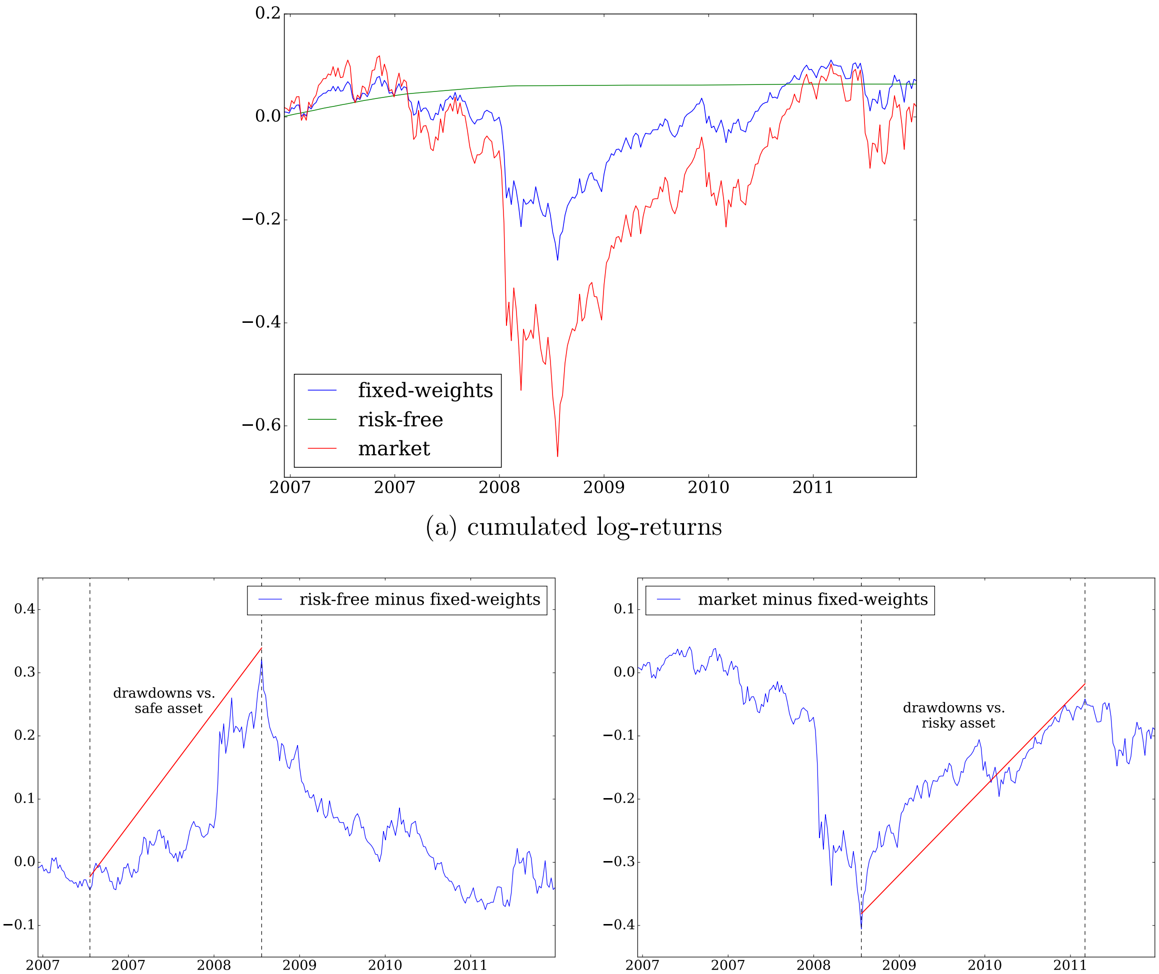

(a) cumulated log-returns

(b) drawdowns, risk-free minus fixed-weights

(c) drawdowns, market minus fixed-weights

Figure 2: drawdowns relative to the safe and risky assets for a 50/50 fixed-weights strategy, 2007–2012.

$$\overline{\mathcal{D}}^0_N(\alpha) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{D}^0_N(\alpha, \mathbf{r}) \tag{8}$$

$$\overline{\mathcal{D}}^1_N(\alpha) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{D}^1_N(\alpha, \mathbf{r}). \tag{9}$$

Potential net losses are captured by strategy α's worst-case drawdown $\overline{\mathcal{D}}^0_N(\alpha)$ against the safe asset. Potential foregone gains are captured by strategy α's worst-case drawdown $\overline{\mathcal{D}}^1_N(\alpha)$ against the risky asset.

The decision maker's fear-of-loss and fear-of-missing out are expressed by her willingness to trade-off drawdowns against the safe and the risky asset. Allocation strategies that attain optimal trade-offs are referred to as prior-free optimal.

Definition 3 (prior-free efficient portfolios)**.** *A portfolio allocation strategy* α *is prior-free efficient if there exists* $\lambda \in \Delta(\{0,1\})$ *such that* α *solves*

$$\min_{\widehat{\alpha}\in\mathcal{A}} \max_{i\in\{0,1\}} \lambda_i \overline{\mathcal{D}}^i_N(\widehat{\alpha}). \tag{P_λ}$$

Given $\lambda \in \Delta(\{0,1\})$*, denote by* α_λ *a solution to* (P_λ)*. The corresponding drawdowns are denoted by* $\overline{\mathcal{D}}^{i,*}_N(\lambda) \equiv \overline{\mathcal{D}}^i_N(\alpha_\lambda)$*. The minimal drawdown frontier* Γ *is described by*

$$\Gamma \equiv \left\{ \left(\overline{\mathcal{D}}^{i,*}_N(\lambda)\right)_{i\in\{0,1\}} \middle| \lambda \in \Delta(\{0,1\}) \right\}.$$

Define the associated function $\gamma(\mathcal{D}^0) = \inf\{\mathcal{D}^1|(\widehat{\mathcal{D}}^0, \mathcal{D}^1) \in \Gamma \text{ for } \widehat{\mathcal{D}}^0 \leq \mathcal{D}^0\}$.

Lemma 1. *Frontier mapping* γ *is continuous and strictly decreasing.*

Frontier γ lets the investor make continuous trade-offs between fear-of-loss and fear-of-missing-out in a simple and straightforward manner. Given a tolerable worst case drawdown $\overline{\mathcal{D}}_0$ against the safe asset, it returns the best possible drawdown guarantee $\overline{\mathcal{D}}_1$ against the risky asset.

Two extreme points. Two points of the frontier are easily characterized. At one extreme, it is possible to ensure no drawdowns against the safe asset by being entirely invested in the safe asset. This results in the largest possible drawdowns against the risky asset. Inversely, it is possible to ensure no drawdowns against the risky asset by being entirely invested in the risky asset. This results in the largest possible drawdowns against the safe asset.

The remainder of this paper is interested in the set of points in between these two extremes. It argues that the corresponding prior-free asset allocation strategies achieve the following desiderata: (i) they provide robust performance guarantees for arbitrarily non-stationary processes for returns; (ii) they let the decision-maker express meaningful risk-preferences over complex acts in a simple manner; (iii) they perform well in the data.

3 Robustness to Non-Stationarity

This section motivates the use of prior-free optimal strategies by: (i) highlighting their robustness to non-stationarity, (ii) highlighting the difficulty of finding priors that lead to robust Bayesian-optimal strategies. Section 6 further motivates the use of drawdown-minimizing strategies by reinterpreting drawdowns as sample versions of standard optimality conditions.

3.1 Drawdown control and robustness

Intuitively, strategies that guarantee low drawdowns perform well in environments with time varying risk-premia since they guarantee performance close to that of the best-performing asset over any subperiod $[T', T] \subset [0, N]$. This is captured by the following performance bound.

Proposition 1 (a performance bound)**.** *Consider a strategy α, a sequence of returns $\mathbf{r}$. For all time periods $T_1 < T_2 \leq N$, and for all $i \in \{0, 1\}$,*

$$\sum_{t=T_1}^{T_2} \log(1 + r_t^\alpha) \geq \left[\sum_{t=T_1}^{T_2} \log(1 + r_t^i)\right] - \overline{\mathcal{D}}_N^i(\alpha).$$

For all time sequences $0 = T_1 < T_2 < \cdots < T_n = N + 1$,

$$\sum_{t=0}^{N} \log(1 + r_t^\alpha) \geq \sum_{k=1}^{n-1} \max_{i \in \{0,1\}} \left\{ \left[\sum_{t=T_k}^{T_{k+1}-1} \log(1 + r_t^i)\right] - \overline{\mathcal{D}}_N^i(\alpha) \right\}.$$

In other terms, drawdown guarantees imply lower bounds on the performance of strategy

α. Up to a penalty $\overline{\mathcal{D}}^i(\alpha)$, it performs at least as well as asset i over any subperiod $[T_1, T_2]$.

Conversely, a strategy that experiences a large drawdown (say of order N) vis à vis either asset is making a binary allocation error (whether to be invested in the safe or risky asset) over a long period of time. This motivates the following definition.

Definition 4. *A sequence of asset allocation strategies $(\alpha_N)_{N\in\mathbb{N}}$ (indexed on increasing time horizon N) is said to be robust to non-stationarity if and only if*

$$\forall i \in \{0,1\}, \quad \lim_{N\to\infty} \frac{\overline{\mathcal{D}}^i_N(\alpha_N)}{N} = 0.$$

We now show that for a natural class of full-support priors, Bayesian optimal strategies are not robust to non-stationarity.

3.2 Fragility of finite hidden Markov models

Hidden Markov models are a popular and flexible way to model time-varying processes. However, Proposition 2 (below) shows that priors over hidden Markov models lead to strategies that can be defeated by an adversarial nature.

A K-state hidden Markov chain over returns with states in $Z = \{1, \cdots, K\}$ is described by $m = (\phi, \xi) \in (\Delta(Z))^Z \times (\Delta(M))^Z$, where $\phi : Z \to \Delta(Z)$ is a Markov chain describing transitions between unobserved states $z \in Z$ (with initial state normalized to 1), and ξ maps states $z \in Z$ into distributions over observed returns $r \in M$. A hidden Markov chain induces a stochastic process over *unobserved* states $(z_t)_{t\geq 0}$ and *observed* returns $(r_t)_{t\geq 0}$ defined by $z_0 = 1$ and

$$\forall t \geq 0, \quad z_{t+1} \sim \phi(z_t) \quad \text{and} \quad r_t \sim \xi(z_t).$$

Note that the set $\mathcal{M}_K$ of hidden Markov chains m with less than K states is finite dimensional and compact. This implies that one can easily define full-support priors μ over $\mathcal{M}_K$.

Any Bayesian prior $\mu \in \Delta(\mathcal{M}_K)$ over K-state hidden Markov chains is associated with Bayesian optimal policies α_μ^{B} solving

$$\max_{\alpha \in \mathcal{A}} \mathbb{E}_\mu \left[\sum_{t=0}^{N} \log(1 + r_t^\alpha) \right].$$

Such a Bayesian-optimal strategy reflects the investor's updating over the likelihood of different underlying Markov chains, as well as the state these chains may be in. Note that the investor's posterior belief is itself a Markov chain with infinite (in fact continuous) state space $\Delta(\mathcal{M}_K)$. As a result, it is able to capture many transient patterns of returns, and it is a plausible guess that the corresponding allocation policy could be robust in the sense of Definition 4. Proposition 2 shows that this isn't the case.

Proposition 2. *Take K, c, M as given. For any full support prior $\mu \in \Delta(\mathcal{M}_K)$ there exists $\nu > 0$ such that for all $N \in \mathbb{N}$ and any Bayesian optimal strategy α_μ^B,*

$$\max_{\mathbf{r} \in M^{N+1}} \max_{i \in \{0,1\}} \mathcal{D}_N^i(\alpha_\mu^B, \mathbf{r}) > \nu N.$$

In words, any Bayesian-optimal allocation policy derived from a full support prior over finite hidden Markov chains is susceptible to drawdowns of order N. As a result, it is not robust to non-stationarity.

3.3 The possibility of robustness

To be useful as a selection criterion, robustness to non-stationarity needs to be non-empty. Proposition 3 shows that indeed, there exist allocation strategies that guarantee sublinear drawdowns for all possible realized sequences of returns.

Proposition 3 (robustness)**.** *For all $\overline{c} < 1$ and $\overline{r}$, there exists $h > 0$ and a strategy α such that*

$$\forall N \in \mathbb{N}, \quad \max_{\mathbf{r} \in M^{N+1}} \max_{i \in \{0,1\}} \mathcal{D}_N^i(\alpha, \mathbf{r}) \leq h\sqrt{N}. \tag{10}$$

Together, Propositions 2 and 3 establish that it is possible to find strategies that are robust to non-stationarity, but they cannot be obtained by modeling the underlying returns as an unknown finite hidden Markov model.

An immediate corollary of Proposition 3 is that prior-free optimal strategies are robust to non-stationarity. Indeed they achieve the smallest possible drawdowns.

Corollary 1. *For any* $\lambda \in (0,1)^2$, *the sequence of prior-free optimal strategies* $(\alpha_{\lambda,N})_{N\in\mathbb{N}}$ *(indexed on time horizon* N*) is robust to non-stationarity.*

It is important to note that robustness to non-stationarity is only an asymptotic property which is achieved by many strategies. In this respect, focusing on prior-free optimal strategies, which achieve exact minimal drawdowns, has several benefits:

- it provides the best possible control on drawdowns, optimizing constants, which could matter for empirical evaluations with moderate investment horizon N;
- by providing a uniquely optimal strategy, it limits the scope for specification search, alleviating concerns of overfitting prevalent in the asset-pricing literature (Lo and MacKinlay, 1990, Novy-Marx, 2014);
- it provides a benchmark by which to evaluate the robustness of asset allocation strategies that are attractive for other reasons (e.g. in-sample performance);
- it provides a systematic framework for optimal dynamic allocation which can incorporate relevant economic features of the problem, such as trading costs, or restrictions on the process of returns (e.g. bounds on P/E ratios . . .).[4]

4 Computing Prior-Free Optimal Strategies

This section shows how to express the problem of computing prior-free optimal asset allocation strategies as a manageable dynamic programming problem. The first step is to identify a convenient state space. Consider an allocation strategy α.

[4]See Appendix A for a discussion of various ways to place restrictions, including probabilistic ones, on the set of possible returns.

For $i \in \{0,1\}$, and $T \in \{0, \cdots, N\}$, define regrets

$$\mathcal{R}_T^i(\alpha, \mathbf{r}) \equiv \max_{T' \in \{0,\cdots,T+1\}} \sum_{t=T'}^{T} \log(1+r_t^i) - \log(1+r_t^\alpha). \tag{11}$$

Regret $\mathcal{R}_T^i$ differs from drawdown $\mathcal{D}_T^i$ in that the endpoint T of the period over which under-performance is measured is fixed. In fact, we have that $\mathcal{D}_T^i = \max_{T' \leq T} \mathcal{R}_{T'}^i$. Lemma 2 shows that $\mathcal{R}_T^i$ can be used to compute worst-case drawdowns, and is described by a simple dynamic process.

Lemma 2. *For all $\alpha \in \mathcal{A}$,*

(i) $\forall i \in \{0,1\}, \quad \max_{\mathbf{r} \in M^{N+1}} \mathcal{D}_N^i(\alpha, \mathbf{r}) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{R}_N^i(\alpha, \mathbf{r})$*;*

(ii) $\forall T < N$, $\mathcal{R}_{T+1}^i = \left[\mathcal{R}_T^i + \log(1+r_{T+1}^i) - \log(1+r_{T+1}^\alpha)\right]^+$*.*

Point (i) implies that to compute drawdown-minimizing strategies, it is sufficient to compute regret-minimizing strategies (this result uses the fact that there exists a return r such that $r^0 = r^1$). Point (ii) clarifies why this observation is valuable: regrets at time $T+1$ can be computed as a function of regrets at time T and returns at time $T+1$. In contrast, drawdowns at time $T+1$ depend on drawdowns at time T, returns at time $T+1$, but also on returns at previous periods.

Denote by $\mathcal{R}_T \equiv (\mathcal{R}_T^i)_{i \in \{0,1\}}$ the vector of regrets, and define state $z_t = (t, a_{t^-}, \mathcal{R}_{t-1})$. For any $\lambda \in \Delta(\{0,1\})$, value function W_λ over states z is recursively defined as follows.

$$W_\lambda(z_T) = \begin{cases} \max_{i \in \{0,1\}} \lambda_i \mathcal{R}_T^i & \text{if} \quad T = N+1 \\ \min_{a_T \in A} \max_{r \in M} W_\lambda(z_{T+1}) & \text{if} \quad T \leq N \end{cases} \tag{12}$$

where $z_{T+1} = \left(T+1, a_{T+1^-}, \left[\mathcal{R}_{T-1}^i + \log(1+r_T^i) - \log(1+r_T^\alpha)\right]^+\right)_{i \in \{0,1\}}$.

This provides a straightforward way to compute prior-free optimal allocation strategies.

Proposition 4 (Bellman formulation)**.** *Let $z_0 = (0, a_0, 0, 0)$. The following hold.*

(i) $\min_\alpha \max_{i\in\{0,1\}} \lambda_i \overline{\mathcal{D}}^i_N(\alpha) = W_\lambda(z_0)$.

Drawdown minimizing policy α^*_λ *depends only on states* $(z_t)_{t\geq 0}$ *and is defined by*

$$\forall z_t, \quad \alpha^*_\lambda(z_t) \in \arg\min_{a_t\in A} \max_{r_t\in M} W_\lambda(z_{t+1}).$$

(ii) *The Pareto frontier of worst-case drawdowns is described by*

$$\Gamma = \left\{ \left(\frac{1}{\lambda_i} W_\lambda(z_0) \right)_{i\in\{0,1\}} \textit{for } \lambda \in \Delta(\{0,1\}) \right\}. \tag{13}$$

Figure 3 represents the Pareto frontier of minimal drawdowns for moves of nature $M = \{0\} \times \{-.02, -.01, 0, .01, .02\}$ and time horizon $N = 260$ (i.e. 5 years, each period corresponding to a week), computed using the algorithm laid out in Proposition 4. Each direction $\lambda \in \{(.5,.5), (.55,.45), (.6,.4)\}$ maps to the prior-free optimal allocation strategy α_λ such that $\lambda_0 \overline{\mathcal{D}}^N_0(\alpha) = \lambda_1 \overline{\mathcal{D}}^N_1(\alpha)$.

The frontier is convex and sits well to the South-West of the line-segment between extreme points corresponding to $\lambda = (0,1)$ and $\lambda = (1,0)$. This suggests that it is possible to find attractive trade-offs between fear-of-loss and fear-of-missing-out.

Computing the worst-case drawdowns of arbitrary strategies. This paper focuses on strategies minimizing worst-case drawdowns, however, worst-case drawdowns need not be the only criterion on which allocation strategies are evaluated. In that case, the Pareto frontier Γ characterized by Proposition 4 remains useful as a benchmark to evaluate the robustness of alternative strategies that are attractive according to other criteria (e.g. in sample performance).

To do this, it is necessary to compute the worst-case drawdowns of arbitrary alternative strategies. Corollary 2 (below) shows that the Bellman approach remains useful in this case. Consider an asset allocation strategy defined by an automaton (α, ϕ) over some state space

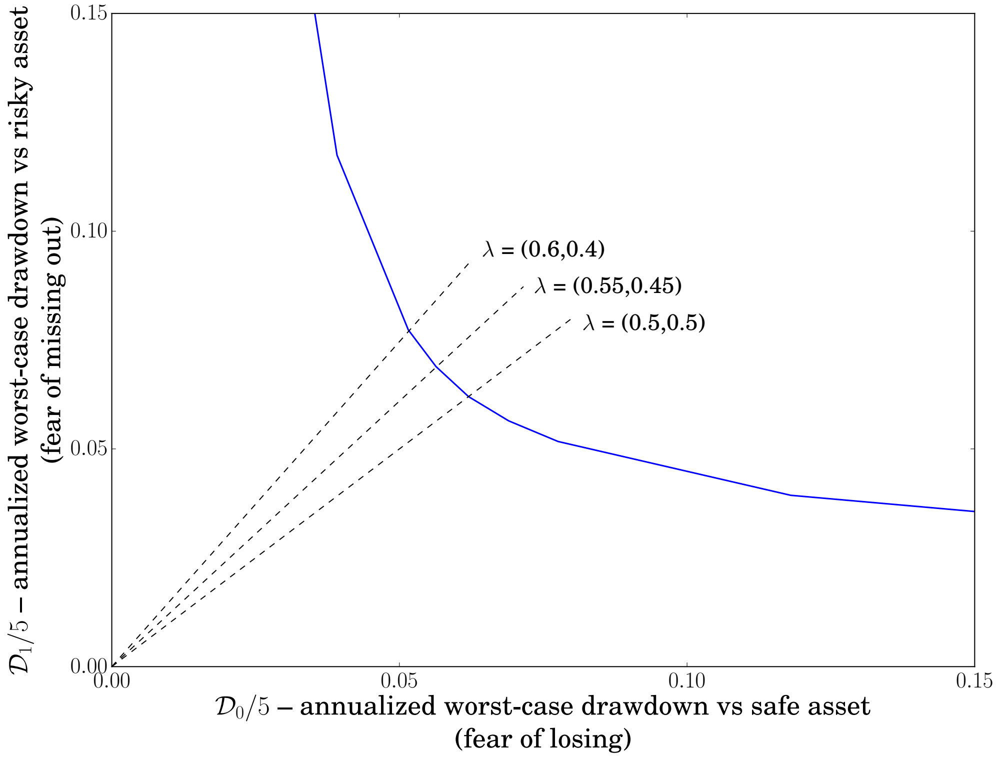

Figure 3: the drawdown frontier, $\lambda \in \{(.5, .5), (.55, .45), (.6, .4)\}$, $N = 260$ (weeks), $M = \{0\} \times \{-.02, -.01, 0, .01, .02\}$.

Z. For each $i \in \{0, 1\}$, define state $x_t^i = (t, a_{t^-}, z_t, \mathcal{R}_{t-1}^i)$ and introduce value function V_α^i over states x^i recursively defined as follows.

$$V_\alpha^i(x_T^i) = \begin{cases} \mathcal{R}_{T-1}^i & \text{if} \quad T = N + 1 \\ \max_{r \in M} V_\alpha(x_{T+1}^i) & \text{if} \quad T \leq N. \end{cases} \tag{14}$$

Corollary 2. *For any $i \in \{0, 1\}$, let $x_0^i = (0, a_0, z_0, 0)$. We have that $\overline{\mathcal{D}}_N^i(\alpha) = V_\alpha^i(x_0^i)$.*

Of course, this Bellman representation is only useful if state space Z has small dimensionality. When strategy α depends on a large state-space (which can be the case for strategies depending on truncated moving averages), worst-case drawdowns should be evaluated using Monte Carlo or genetic algorithms approaches designed for high-dimensional numerical optimization (Golberg, 1989, Glasserman, 2003).

5 Qualitative Properties

Proposition 4 provides a computational method to characterize drawdown minimizing policies for arbitrary trading costs and arbitrary moves of nature M. It is instructive to derive qualitative properties of prior-free optimal allocation strategies under the simplifying assumption that *trading costs c are equal to 0* and that $M = \{0\} \times \{-\overline{r}, 0, \overline{r}\}$.

We first note that prior-free optimal strategies are admissible. Recall that α_λ denotes the solution to the original max-min problem (P_λ).

Proposition 5 (admissibility)**.** *For every $\lambda \in \Delta(\{0, 1\})$, there exists a prior $\mu_\lambda \in \Delta(M^{N+1})$ such that*

$$\alpha_\lambda \in \arg\max_{\alpha \in \mathcal{A}} \mathbb{E}_{\mu_\lambda} \left[\sum_{t=0}^{N} \log(1 + r_t^\alpha) \right].$$

In other terms, there always exists a prior over returns for which a prior-free optimal strategy is also Bayesian optimal. Of course, as was emphasized in Sections 2 and 3, the difficulty is coming up with such a prior. The thesis defended by this paper is that expressing preferences over the properties of allocation strategies directly (here low drawdowns) is both a decision theoretically correct, and practical way to approach dynamic asset allocation.

What prior over moves of nature rationalizes prior-free optimal strategies can be further understood by taking a game-theoretic perspective. Note that since trading costs are equal to 0, allocation a_t is no longer a state variable. Abusing notation, $z_t \equiv (t, \mathcal{R}_t^i)_{i \in \{0,1\}}$ is now a sufficient state. For any z_t, a, r, define payoff function

$$U_\lambda(z_t, a, r) \equiv W_\lambda(z_{t+1})$$

where $z_{t+1} = \left(t + 1, \left[\mathcal{R}_{t-1}^i + \log(1 + r^i) - \log(1 + r^a) \right]^+ \right)_{i \in \{0,1\}}$.

Lemma 3. *(i) Payoff function $U_\lambda(z_t, a, r)$ is convex in a.*

(ii) Optimal allocation $\alpha_\lambda(z_t)$ is a Nash equilibrium strategy in the zero-sum game against nature with actions (a, r) and payoffs $-U_\lambda(z_t, a, r)$ to the investor.

(iii) *The drawdown frontier* $\gamma : \mathcal{D}_0 \mapsto \gamma(\mathcal{D}_0)$ *is convex in* $\mathcal{D}_0$.

The investor plays a stochastic zero-sum game against nature and prior-free optimal allocation strategies are Nash equilibria of this game. This game-theoretic interpretation is helpful in characterizing optimal policies.

5.1 A game-theoretic characterization

Explicit characterization for $T = N$. Solving for the optimal policy in period $T = N$ helps delineate the mechanics of optimal drawdown control. Note that since returns to the safe asset r^0 are equal to 0, nature's only choice is to pick returns to the risky asset $r^1 \in \{-\bar{r}, 0, \bar{r}\}$.

At $T = N$, for regrets $\mathcal{R}^0, \mathcal{R}^1$, the investor picks the allocation $a^* = (a^{*,0}, a^{*,1})$ solving

$$\min_{a^1 \in [0,1]} \max_{r^1 \in \{-\bar{r}, 0, \bar{r}\}} \left\{ \left[\mathcal{R}^0 - \log(1 + a^1 r^1)\right]^+, \left[\mathcal{R}^1 + \log(1 + r^1) - \log(1 + a^1 r^1)\right]^+ \right\}.$$

Lemma 4. *(i) Whenever* $\mathcal{R}^0 - \mathcal{R}^1 \in (\log(1-\bar{r}), \log(1+\bar{r}))$*, the optimal allocation* a^* *is strictly interior, i.e.* $a^{1,*} \in (0, 1)$.

(ii) If $\mathcal{R}^0 - \mathcal{R}^1 \geq \log(1 + \bar{r})$*, the optimal allocation is* $a^{1,*} = 0$*. If* $\mathcal{R}^0 - \mathcal{R}^1 \leq \log(1 - \bar{r})$*, the optimal allocation is* $a^{1,*} = 1$.

Proof. Point (ii) is immediate. Whenever $\mathcal{R}^0 - \mathcal{R}^1 \leq \log(1-\bar{r})$, for any allocation $a^1 \in (0, 1)$ and $r \in M$, $\mathcal{R}^1 + \log(1 + r^1) - \log(1 + a^1 r^1) > \mathcal{R}^0 - \log(1 + a^1 r^1)$, which implies that the optimal allocation a^1 minimizes $\max_r \mathcal{R}^1 + \log(1+r^1) - \log(1+a^1 r^1)$, i.e., $a^{1,*} = 1$. A similar reasoning holds when $\mathcal{R}^0 - \mathcal{R}^1 \geq \log(1 + \bar{r})$.

Point (i) exploits the fact that optimal allocation a^* is a Nash equilibrium in the zero-sum game against nature with payoffs $-U(z_t, a, r)$. Hence, it is sufficient to show that neither $a^{1,*} = 0$ nor $a^{1,*} = 1$ can be part of a Nash equilibrium. Indeed, if $a^{1,*} = 0$, then, since $\mathcal{R}^0 - \mathcal{R}^1 \in (\log(1 - \bar{r}), \log(1 + \bar{r}))$, nature's strict best response is to set $r = \bar{r}$. This yields worst-case regrets $\mathcal{R}^1 + \log(1 + \bar{r}) - \log(1 + a^1\bar{r})$, inducing best response $a^{1,*} > 0$ from the investor. A similar reasoning shows that $a^{1,*} = 1$ cannot be part of an equilibrium either. This implies that the only equilibrium is in mixed strategies. □

An immediate corollary is that whenever $\mathcal{R}^0 - \mathcal{R}^1 \notin (\log(1-\overline{r}), \log(1+\overline{r}))$, $W(z_N) = \max\{\mathcal{R}_0, \mathcal{R}_1\}$. When $\mathcal{R}^0 - \mathcal{R}^1 \in (\log(1-\overline{r}), \log(1+\overline{r}))$, since $a^{1,*} \in (0,1)$, it is strictly optimal for nature to pick returns in $\{-\overline{r}, \overline{r}\}$, which implies that $W(z_N) > \max\{\mathcal{R}^0, \mathcal{R}^1\}$. Furthermore, nature must be indifferent between picking $-\overline{r}$ and $\overline{r}$ (otherwise, the investor would not use an interior strategy). This implies that an optimal allocation a^* must satisfy

$$
\begin{aligned}
&\mathcal{R}^0 - \log(1 - a^{1,*}\overline{r}) = \mathcal{R}^1 + \log(1+\overline{r}) - \log(1 + a^{1,*}\overline{r}) \\
\iff & a^{1,*} = \frac{1}{\overline{r}} \times \frac{(1+\overline{r})\exp(\mathcal{R}^1 - \mathcal{R}^0) - 1}{(1+\overline{r})\exp(\mathcal{R}^1 - \mathcal{R}^0) + 1}.
\end{aligned}
\tag{15}
$$

This yields value function

$$
W_N(\mathcal{R}^0, \mathcal{R}^1) = \begin{cases} \max\{\mathcal{R}^0, \mathcal{R}^1\} & \text{if} \quad \mathcal{R}^0 - \mathcal{R}^1 \notin (\log(1-\overline{r}), \log(1+\overline{r})) \\ \log\left[(1+\overline{r})\exp\mathcal{R}^1 + \exp\mathcal{R}^0\right] - \log 2 & \text{otherwise.} \end{cases}
$$

Characterization for $T < N$. Lemma 4 extends as follows. Define $\hbar_0 = -\log(1-\overline{r})$, and $\hbar_1 = \log(1+\overline{r})$.

Proposition 6. *For all $t \in \{0, \cdots, N\}$, $\lambda \in \Delta(\{0,1\})$, $\alpha_\lambda(\mathcal{R}^0, \mathcal{R}^1, t)$ is continuous in $(\mathcal{R}^0, \mathcal{R}^1)$. Furthermore, for all $i \in \{0,1\}$,*

(i) if $\lambda_i \mathcal{R}^i \geq \lambda_{-i}\left[\mathcal{R}^{-i} + (N-t)\hbar_i\right]$, then $\alpha^i_\lambda(\mathcal{R}^i, \mathcal{R}^{-i}, t) = 1$;

(ii) if $\lambda_i \mathcal{R}^i - \lambda_{-i}\mathcal{R}^{-i} \in (-(N-t)\hbar_{-i}, (N-t)\hbar_i)$, then $\alpha^i_\lambda(\mathcal{R}^i, \mathcal{R}^{-i}, t) \in (0,1)$ and

$$
U_\lambda(z_t, \alpha(z_t), \overline{r}) = U_\lambda(z_t, \alpha(z_t), -\overline{r}). \tag{16}
$$

Lemma 4 and Proposition 6 have substantial implications.

5.2 Qualitative implications

Interior allocations. As Black and Litterman (1992) emphasize, naïve implementations of Markowitz (1952)'s approach to portfolio allocation frequently generate extreme corner allocations. One possible fix is to place ad hoc constraints on allocations. Alternatively, Black and Litterman (1992) suggest anchoring priors to benchmark priors constrained to justify holding the market portfolio. An immediate corollary of Proposition 6 is that the prior-free approach naturally leads to non-corner solutions, without the help of ad hoc side-constraints.

Corollary 3. *For all $\lambda \in \Delta(\{0,1\})$, there exists $\nu > 0$ such that for all N and all sequences of returns $\mathbf{r}$,*

$$\frac{1}{N+1}\sum_{t=0}^{N}\mathbf{1}_{\alpha_{\lambda,N}(z_t)\in(0,1)^2} \geq 1-\nu\frac{1}{\sqrt{N+1}}.$$

In other words, for any realization of returns, prior-free optimal allocations are interior for a share of periods asymptotically equal to 1.

Momentum. An influential literature documents the profitability of momentum strategies which buy recent overperforming stocks, while selling recent underperforming stocks (Jegadeesh and Titman, 1993, 2001, Barberis et al., 1998, Hong and Stein, 1999, Hong et al., 2000, Moskowitz et al., 2012, Asness et al., 2013), and proposes behavioral explanations for this apparent departure from the efficient market hypothesis.

Even though prior-free optimal strategies are not calibrated using historical data, they also exhibit momentum. This reflects the fact that they attempt to optimize asset allocation in arbitrarily non-stationary environments. Returns need not go back to the mean, and momentum emerges as a response to the fact that one asset may well keep performing better than an other. Formally the following result holds.

Corollary 4. *Let $T_1 = \lfloor \rho_1 N \rfloor < \lfloor \rho_2 N \rfloor = T_2$, with ρ_1, ρ_2 fixed.*[5] *For any $\epsilon > 0$, consider a probability measure μ over $(r_t^1)_{t\in\{T_1,T_1+1,\cdots,T_2\}}$ such that $\forall t \in \{T_1,\cdots,T_2\}$, $\mathbb{E}_\mu\left[\frac{r_t^1}{1+r_t^1}\big|h_t\right] > \epsilon$*

[5] As usual, $\lfloor x \rfloor$ denotes the integer part of x, defined as $\lfloor x \rfloor = \max\{n \in \mathbb{N} | n \leq x\}$.

then,

$$\mu \ a.s., \quad \lim_{N\to+\infty} \frac{1}{T_2 - T_1} \sum_{t=T_1}^{T_2} \alpha^1_{\lambda,N}(h_t) = 1.$$

The condition $\mathbb{E}_\mu \left[\frac{r_t^1}{1+r_t^1} \middle| h_t \right] > \epsilon$ — expected returns weighed by marginal utility are sufficiently high — implies that a history h_t, allocating all wealth to the risky asset yields strictly higher expected log-returns than any other allocation in $\Delta(\{0,1\})$.

Corollary 4 states that if the returns from the risky asset are drawn from a distribution with sufficiently positive mean over the time interval $[T_1, T_2]$, then the allocation to the risky asset must go to one. A converse holds if the returns to the risky asset are drawn from a distribution with sufficiently negative mean. Note that to be compatible with Corollary 3, Corollary 4 requires the allocation to converge to a corner allocation from the interior.

History dependence. Another notable property of prior-free optimal allocation strategies is that even though log preferences do not exhibit a wealth effect, past investment experience affects the investors' continuation behavior. Investors investing in the same period who have experienced different histories of returns will choose different allocations. This is because drawdown minimization is a reference-dependent objective, with the reference point being dependent on the investor's personal history.

This can be illustrated with a simple example. Consider a two-period investment problem, with time denoted by $t \in \{0, 1\}$. A young investor is born in period 1, and an old investor is born in period 0. Both investors share the same preference parameter $\lambda = (.5, .5)$ and the magnitude of potential returns is $\overline{r} = .02$. Expression (15) implies that the young investor will allocate 49.5% of her wealth to the risky asset. The allocation of the old investor depends on her experience at time $t = 0$. We know by Proposition 6 that she must have chosen an interior allocation in period $t = 0$. If the risky asset yielded returns $r_0^1 = -\overline{r}$, she experienced drawdowns against the safe asset, and by expression (15) must place a weight strictly less than 49.5% on the risky asset. Inversely, if the risky asset yielded returns $r_0^1 = \overline{r}$, she experienced drawdowns against the risky asset, and by expression (15) must place a weight strictly higher than 49.5% on the risky asset.

This property echoes Malmendier and Nagel (2011)'s finding that investors exhibit heterogenous risk preferences as a function of their personal histories. Specifically, they show that poor realized returns make investors significantly more risk-averse, in a way that's not quantitatively explained by wealth effects.

6 Multi-Asset Allocation

So far the analysis has focused on allocating resources to a single risky asset. This section extends the prior-free approach to several risky assets. Along the way, it suggests a reinterpretation of drawdowns as optimality conditions.

Framework. Consider an environment with several risky assets $i \in \mathcal{I} = \{1, \cdots, I\}$ and a single risk-free asset denoted by 0. For simplicity, trading costs are set to zero. Let $a_t^{\mathcal{I}} = (a_t^i)_{i \in \mathcal{I}}$ and a_t^0 respectively denote allocations to the risky and risk-free assets at time t. Allocations $a^{\mathcal{I}}$ to risky assets must belong to the product set $A = \prod_{i \in \mathcal{I}} A^i$ with $A^i = [\underline{a}^i, \overline{a}^i]$. In addition total allocation weights must sum to one, so that $a^0 = 1 - \sum_{i \in \mathcal{I}} a^i$. Note that short-selling is implicitly allowed. The overall allocation is denoted by $a = (a^0, a^{\mathcal{I}})$. For simplicity, returns $r^0 \geq 0$ to the risk-free asset are constant over time. Risky returns $r^{\mathcal{I}} = (r^i)_{i \in \mathcal{I}}$ belong to a set $M^{\mathcal{I}}$ of moves of nature taking the form $M^{\mathcal{I}} = \prod_{i \in \mathcal{I}} M^i \subset (-\overline{r}, \overline{r})^{\mathcal{I}}$, with $\overline{r} \in (0, 1)$. Let $r = (r^0, r^{\mathcal{I}})$.

Consider now the problem of a Bayesian investor maximizing her subjective expected utility. In each period t, the investor chooses the allocation $a_t^{\mathcal{I}}$ that solves

$$\max_{a^{\mathcal{I}} \in A} \mathbb{E}[\log(1 + \langle r_t, a \rangle) | \mathcal{F}_t] \tag{17}$$

where $\mathcal{F}_t$ denotes the investor's information set at t. Take as given $\Delta > 0$. For $\sigma \in \{-, +\}$ and any allocation a, denote by $\phi^+(a, i)$ and $\phi^-(a, i)$ allocations identical to a except that the

i^{th} coordinate is shifted up or down by an amount Δ. Under the paper's notation, $\phi^\sigma(a,i)^j$ denotes the weight assigned to asset j by allocation $\phi^\sigma(a,i)$. Formally, we have that

$$\phi^+(a,i)^i = a^{i,+} \equiv \min\{a^i + \Delta, \overline{a}^i\} \quad \text{and} \quad \forall j \in \mathcal{I} \setminus i,\ \phi^+(a,i)^j = a^j;$$
$$\phi^-(a,i)^i = a^{i,-} \equiv \max\{a^i - \Delta, \underline{a}^i\} \quad \text{and} \quad \forall j \in \mathcal{I} \setminus i,\ \phi^-(a,i)^j = a^j.$$

By definition, for any $T' \leq T \leq N$, solutions $(a_t^*)_{t\in\{T',\cdots,T\}}$ to (17) cannot be improved by shifting the allocation in any direction. As a result, for all $i \in \mathcal{I}$, under the investor's prior,

$$\sum_{t=T'}^{T} \mathbb{E}[\log(1 + \langle r_t, \phi^+(a_t^*, i)\rangle)|\mathcal{F}_t] - \mathbb{E}[\log(1 + \langle r_t, a_t^*\rangle)|\mathcal{F}_t] \leq 0 \tag{18}$$
$$\sum_{t=T'}^{T} \mathbb{E}[\log(1 + \langle r_t, \phi^-(a_t^*, i)\rangle)|\mathcal{F}_t] - \mathbb{E}[\log(1 + \langle r_t, a_t^*\rangle)|\mathcal{F}_t] \leq 0. \tag{19}$$

Using finite sample versions of the central limit theorem,[6] this implies that when returns are drawn from the investor's prior, then, with probability approaching 1 for N large, for all $T' \leq T \leq N$,

$$\sum_{t=T'}^{T} \log(1 + \langle r_t, \phi^+(a_t^*, i)\rangle) - \log(1 + \langle r_t, a_t^*\rangle) \leq O(\sqrt{N})$$
$$\text{and,} \quad \sum_{t=T'}^{T} \log(1 + \langle r_t, \phi^-(a_t^*, i)\rangle) - \log(1 + \langle r_t, a_t^*\rangle) \leq O(\sqrt{N}).$$

Define drawdowns $\mathcal{D}_N^{i,+}$ and $\mathcal{D}_N^{i,-}$ as

$$\mathcal{D}_N^{i,+} = \max_{T' \leq T \leq N} \sum_{t=T'}^{T} \log(1 + \langle r_t, \phi^+(a_t^*, i)\rangle) - \log(1 + \langle r_t, a_t^*\rangle),$$
$$\mathcal{D}_N^{i,-} = \max_{T' \leq T \leq N} \sum_{t=T'}^{T} \log(1 + \langle r_t, \phi^-(a_t^*, i)\rangle) - \log(1 + \langle r_t, a_t^*\rangle).$$

[6] Specifically, the Hoeffding-Azuma inequality. See Cesa-Bianchi and Lugosi (2006), Lemma A.7 for a reference.

These drawdowns capture maximal losses relative to strategies that systematically increase or decrease their exposure to a specific asset. Keeping these drawdowns low is a sample expression of optimality conditions (18) and (19).

Note that the optimality conditions being tested depend on the step of the deviation Δ. Indeed, the drawdowns sudied in Sections 2 to 5 correspond to setting $\Delta = 1$, $\underline{a}^i = 0$ and $\overline{a}^i = 1$. In that case, the drawdowns $\mathcal{D}_N^0$ and $\mathcal{D}_N^1$ of Sections 2 to 5 satisfy $\mathcal{D}_N^0 = \mathcal{D}_N^{1,-}$ and $\mathcal{D}_N^1 = \mathcal{D}_N^{1,+}$. Smaller steps Δ, correspond to more local deviations, and potentially allow for finer optimization. For any asset allocation strategy $\alpha \in \mathcal{A}$, mapping public histories to allocations, define maximum drawdowns as follows:

$$\forall i \in \mathcal{I}, \sigma \in \{+,-\}, \quad \overline{\mathcal{D}}_N^{i,\sigma}(\alpha) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{D}_N^{i,\sigma}(\alpha, \mathbf{r})$$

Take as given a deviation step $\Delta > 0$, and let Λ be the set of weights $\lambda = (\lambda^+, \lambda^-)$, such that $\lambda^+ + \lambda^- = 1$.[7]

Definition 5 (prior-free allocation strategies)**.** *A dynamic asset allocation strategy $\alpha \in \mathcal{A}$ is prior-free optimal if there exists $\lambda \in \Lambda$ such that α solves*

$$\min_{\alpha \in \mathcal{A}} \max_{i \in \mathcal{I}} \max_{\sigma \in \{-,+\}} \lambda^\sigma \overline{\mathcal{D}}_N^{i,\sigma}(\alpha). \tag{$P_\lambda^{\mathcal{I}}$}$$

Computation and key properties. The remainder of this section clarifies difficulties in finding numerical solutions to $P_\lambda^{\mathcal{I}}$ and identifies an approximately optimal class of strategies amenable to numerical computation and theoretical analysis.

For all $T \leq N$, let

$$\mathcal{R}_T^{i,\sigma}(\alpha, r) \equiv \max_{T' \leq T} \sum_{t=T'}^{T} \log(1 + \langle r_t, \phi^\sigma(a_t^*, i)\rangle) - \log(1 + \langle r_t, a_t^* \rangle).$$

An argument identical to that of Lemma 2 implies that for all $\alpha \in \mathcal{A}$,

[7] Weights λ^σ could also be indexed on i.

$$\overline{\mathcal{D}}_N^{i,\sigma}(\alpha) = \overline{\mathcal{R}}_N^{i,\sigma}(\alpha)$$

$$\mathcal{R}_{T+1}^{i,\sigma} = \max\{0, \mathcal{R}_T^{i,\sigma} + \log(1 + \langle r_{T+1}, \phi^\sigma(\alpha_{T+1}, i)\rangle) - \log(1 + \langle r_{T+1}, \alpha_{T+1}\rangle)\}.$$

Since there are no trading costs, current allocations are not a state variable in Problem $P_\lambda^{\mathcal{I}}$. An argument identical to that of Proposition 4 implies that optimal asset allocation policies α are a function of state $z_t = (t, \mathcal{R}_t^{i,+}, \mathcal{R}_t^{i,-})_{i\in\mathcal{I}}$. Unfortunately, the problem of picking an optimal allocation $\alpha_{\mathcal{I}}$ over risk-assets cannot be separated in I independent problems. The returns r^{-i} of assets other than i affect the optimal allocation to asset i. This implies that problem $P_\lambda^{\mathcal{I}}$ becomes numerically intractable as the number I of risky assets becomes large.

Fortunately, a relaxed problem admits computationally tractable solutions that are approximate solutions to $P_\lambda^{\mathcal{I}}$. For any i, let

$$M^{\neg i} = \left\{ \sum_{j\in\mathcal{I}\setminus i} a^j(r^j - r^0) \;\middle|\; a^{\mathcal{I}} \in A,\, r^{\mathcal{I}} \in M \right\}.$$

$M^{\neg i}$ is the set of possible returns differentials due to allocations to assets other than i.

Recall the notation $a^{i,+} = \min\{a^i + \Delta, \overline{a}^i\}$ and $a^{i,-} = \max\{a^i - \Delta, \overline{a}^i\}$. For all $a^i \in [\underline{a}^i, \overline{a}^i]$, $r^i \in M^i$, $r^{\neg i} \in M^{\neg i}$, define

$$g^\sigma(a^i, r^i, r^{\neg i}) \equiv \log(1 + a^{i,\sigma} r^i + (1 - a^{i,\sigma})r^0 - r^{\neg i}) - \log(1 + a^i r^i + (1 - a^i)r^0 - r^{\neg i}).$$

For any $i \in \mathcal{I}$,

$$\mathcal{R}_T^{i,\sigma}(\alpha, r) = \max_{T'\le T} \sum_{t=T'}^{T} g^\sigma(\alpha_t^i, r_t^i, r_t^{\neg i}) \quad \text{with} \quad r_t^{\neg i} = \sum_{j\in\mathcal{I}\setminus i} \alpha_t^j (r_t^j - r^0).$$

In Problem $P_\lambda^{\mathcal{I}}$, returns $r^{\neg i}$ are not freely chosen by nature. They are jointly determined by the allocation $a^{\mathcal{I}}$ and returns $r^{\mathcal{I}}$. The relaxed problem increases nature's degrees of freedom by allowing it to independently pick $(r^i)_{i\in\mathcal{I}}$ and $(r^{\neg i})_{i\in\mathcal{I}}$. Denote by $\mathbf{r}^i$ and $\mathbf{r}^{\neg i}$

sequences $(r_t^i)_{t\in\{0,\cdots,N\}}$, $(r_t^{\neg i})_{t\in\{0,\cdots,N\}}$. For any strategy $\alpha^i \in \mathcal{A}^i$, mapping histories to the weight assigned to asset i, let

$$\overline{G}_N^{i,\sigma}(\alpha^i) \equiv \max_{\mathbf{r}^i,\mathbf{r}^{\neg i}} \max_{T' \leq T \leq N} \sum_{t=T'}^{T} g^\sigma(\alpha_t^i, r_t^i, r_t^{\neg i}).$$

Let $\alpha_N^{*,i}$ denote a solution to

$$\min_{\alpha^i \in \mathcal{A}^i} \max_{\sigma \in \{-,+\}} \lambda^\sigma \overline{G}_N^{i,\sigma}(\alpha). \tag{$\widehat{P}_\lambda^i$}$$

This problem is associated with states $z_t^i = (t, \mathcal{R}_t^{i,+}, \mathcal{R}^{i,-})$ and value function

$$W_{\lambda_i}^i(z_T^i) = \begin{cases} \max_{\sigma\in\{+,-\}} \lambda^\sigma \mathcal{R}_T^{i,\sigma} & \text{if} \quad T = N+1 \\ \min_{a_T^i \in A^i} \max_{r^i \in M^i} \max_{r^{\neg i} \in M^{\neg i}} W_\lambda^i(z_{T+1}^i) & \text{if} \quad T \leq N. \end{cases} \tag{20}$$

Proposition 7. *The following hold:*

(i) $\forall \alpha \in \mathcal{A}, \overline{\mathcal{D}}_N^{i,\sigma}(\alpha) \leq \overline{G}_N^{i,\sigma}(\alpha^i)$*;*

(ii) $\forall i \in \mathcal{I}$, $\forall \sigma \in \{+,-\}$, $\lambda^\sigma \overline{G}_N^{i,\sigma}(\alpha_N^{*,i}) = O(\sqrt{N})$*;*

(iii) $\alpha_N^{*,i}$ *depends only on states* z_t^i *and* $\alpha_N^{*,i}(z_t^i)$ *solves*

$$\min_{a^i \in A^i} \max_{r^i \in M^i} \max_{r^{\neg i} \in M^{\neg i}} W_{\lambda_i}^i(z_{t+1}^i).$$

In words, solutions to problem $\widehat{P}_\lambda^i$ are easily computable, and provide an approximate solution to problem $\widehat{P}_\lambda^i$. The fact that drawdowns are sublinear in N implies the following extension of Corollary 4.

Let $T_1 = \lfloor \rho_1 N \rfloor < \lfloor \rho_2 N \rfloor = T_2$, with ρ_1, ρ_2 fixed. Assume that for $t \in \{T_1, \cdots, T_2\}$, returns r_t are i.i.d. with a distribution μ. Furthermore, assume that for all $i \in \mathcal{I}$, there exists $\sigma_i \in \{+1, -1\}$ such that

$$\forall a^{\mathcal{I}} \in A, \quad \sigma_i \mathbb{E}_\mu \left[\frac{r^i - r^0}{1 + \langle a, r \rangle} \right] > 0 \tag{21}$$

where $a = (a_0, a^{\mathcal{I}})$. This implies that problem $\max_{a^{\mathcal{I}} \in A} \mathbb{E}_\mu[\log(1+\langle a, r\rangle)]$ has a unique corner solution $a^{*,\mathcal{I}} \in \prod_{i\in\mathcal{I}}\{\underline{a}^i, \overline{a}^i\}$.

Corollary 5. *If condition* (21) *holds, then prior-free optimal strategy* $\alpha_N^{*,i}$ *satisfies*

$$\forall i \in \mathcal{I},\ \mu\text{-}\ a.s., \quad \lim_{N\to+\infty} \frac{1}{T_2 - T_1} \sum_{t=T_1}^{T_2} \alpha_N^{*,i}(h_t) = a^{*,i}.$$

In other terms, the prior-free optimal strategy must approach the Bayesian optimal allocation when it takes extreme values. It is worth noting that Corollary 5 holds regardless of the history of returns occuring before time T_1. Even after long histories, prior-free optimal allocation strategies do not become doctrinaire. They adapt to new circumstances.

Corollary 5 also clarifies that although Problems $P_\lambda^{\mathcal{I}}$ and $\widehat{P}_\lambda^i$ let nature pick returns to each asset independently, the resulting strategies respond to correlation between assets. For instance, if one asset is redundant because it is highly correlated to another asset with higher returns, then, prior-free asset allocation strategies (relaxed or not) will assign minimal weight to this asset.

7 Practical Evaluation

This section evaluates the behavior of prior-free optimal asset allocation against two benchmark risk-management strategies (applied to the risk-free asset, and a single risky asset):

- the first is the simple $1/n$ benchmark, whose out-of-sample robustness is emphasized in DeMiguel et al. (2009);
- the second is constant proportions portfolio insurance (CPPI Black and Perold, 1992).

The $1/n$ solution is implemented as a quarterly-rebalanced portfolio, targeting a 50/50 fixed-weight allocation between the safe and risky assets.[8] CPPI is an especially relevant

[8]Here, the 1 over n strategy serves as the simplest possible benchmark, similar to a popular 60/40 allocation. The robustness of the 1 over n approach documented by DeMiguel et al. (2009) is important when the number of assets becomes large, so that the correlation matrix between assets can become near singular. The current paper makes no empirical claim regarding such large asset allocation problems.

benchmark since its goal is also to provide prior-free performance guarantees. The version of CPPI tested in this section takes the following form: the investor tracks her counterfactual wealth $\widehat{w}_t$ if she had invested only in the safe asset; a share of her actual wealth w_t equal to 75% of her counterfactual wealth $\widehat{w}_t$ is invested in the safe asset. The remaining cushion $w_t - .75\widehat{w}_t$ is leveraged once and invested in the risky asset. If the price process is continuous and rebalancings occur sufficiently frequently, CPPI guarantees the investor 75% of her wealth if she had invested in the safe asset, while also providing exposure to the risky asset.

7.1 Worst-case performance

Figure 4 plots the worst-case drawdowns of both the fixed-weights portfolio and the CPPI portfolio against the prior-free efficient frontier in the case where a period corresponds to a week, $N = 260$, and $M = \{0\} \times \{-.02, -.01, 0, .01, .02\}$ and trading cost c is equal to 20 basis points. Mechanically, both the fixed-weights portfolio and CPPI must sit to the

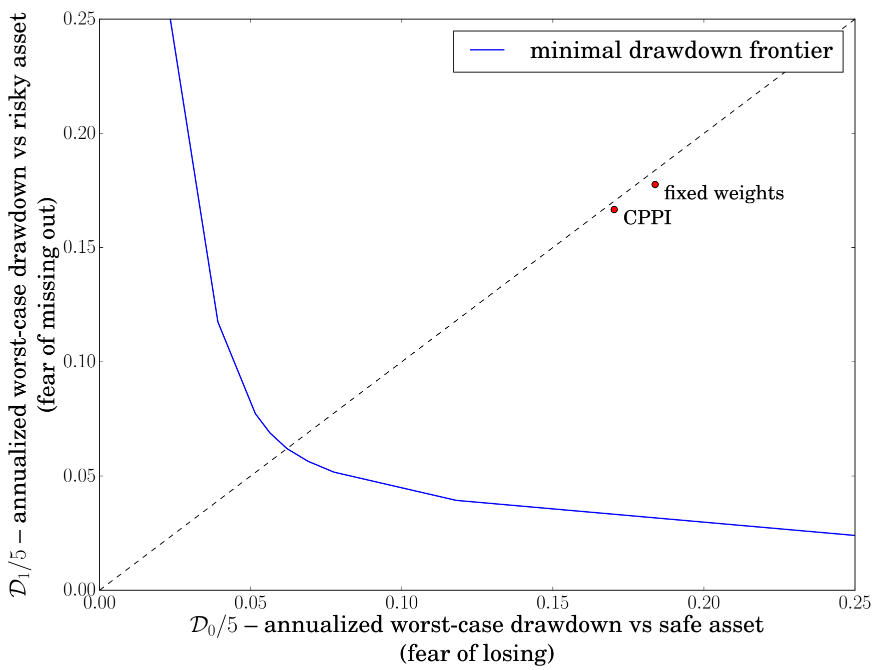

Figure 4: worst-case drawdowns for prior-free efficient, fixed-weights and CPPI strategies.

North-East of the efficient frontier. In fact, they sit quite far away from the efficient frontier.

The reason why the fixed-weights portfolio can experience large drawdowns is clear. It keeps the allocation close to 50/50 even if the risky asset keeps yielding positive (or negative) returns.

A more surprising finding is that CPPI can also experience large drawdowns even though, it is designed to provide performance guarantees. As Figure 5 illustrates, CPPI experiences drawdowns against the safe asset if the risky asset experiences large gains followed by equally large losses. Indeed, after large gains, CPPI will keep a large exposure to the risky asset until those gains are lost. Inversely, CPPI experiences large drawdowns versus the risky asset if large losses are followed by equally large gains. Indeed, if cumulated losses over a large number of periods make CPPI approach the 75% mark, CPPI will then limit its holding of the risky-asset for a commensurate number of periods, resulting in large drawdowns versus the risky asset during the rebound.

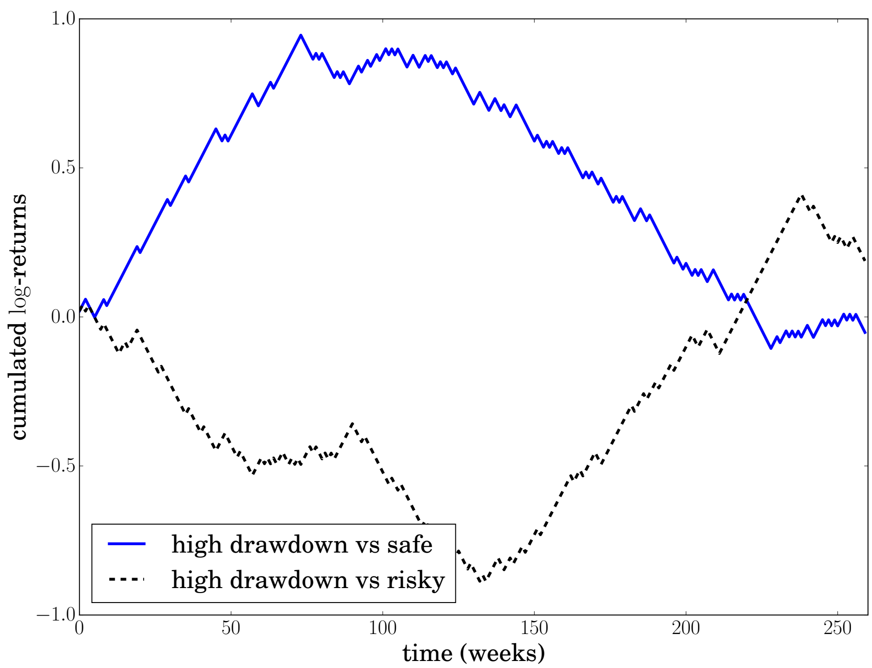

Figure 5: returns generating high drawdowns for the CPPI strategy.

7.2 Historical performance

If market returns were i.i.d., prior-free optimal strategies could not possibly improve on the performance of a fixed-weight allocation. However, if risk-premia exhibit significant variation, prior-free optimal strategies may out-perform fixed-weight strategies. Whether this is the case is an empirical question.

Table 1 reports findings from running the fixed-weights, CPPI, and prior-free optimal strategies defined above, on the sample of market and risk-free returns from January 1^{st} 1927 to December 31^{st} 2014, obtained from Kenneth French's data library. Each strategy is implemented over rolling periods of 5 years, starting January 1^{st} of each of the 88 years in the sample. Trading costs are set to 20 basis points.

Denoting by $\widehat{\mathbb{E}}$ expectations under the empirical sample of daily returns, the following statistics are reported (counting 252 trading days in a year):

- net annualized performance $\mathsf{net\ perf} = 252 \times \widehat{\mathbb{E}}\,[r_t^\alpha - r_t^0]$;
- annualized Sharpe ratios

$$\mathsf{Sharpe} \equiv \frac{\widehat{\mathbb{E}}\,[r^\alpha - r^0]}{\sqrt{\widehat{\mathbb{E}}\,[(r^\alpha - r^0)^2]}}\sqrt{252};$$

- worst-case 5 year relative drawdowns $\mathcal{D}_0$, $\mathcal{D}_1$ versus the safe and risky asset over the entire period;
- net-performance to drawdown ratio

$$\mathsf{net\text{-}perf\ to\ drawdown} \equiv \frac{\mathsf{net\ perf}}{\mathcal{D}_0}. \tag{22}$$

- parameters α and β from CAPM regression

$$r^\alpha - r^0 \sim \alpha + \beta(r^1 - r^0) + \varepsilon. \tag{23}$$

  estimated using annualized returns (N=88), and reporting robust standard errors.

The net performance-to-drawdown ratio defined in (22) summarizes each strategy's ability to capture upside while reducing drawdowns.

	fixed-weights	CPPI	prior-free
net perf	3.7%	4.7%	6.2%
Sharpe	.45	.48	.56
$\mathcal{D}_0$	.54	.37	.29
$\mathcal{D}_1$	.44	.53	.37
net-perf to drawdown	.07	.12	.21
α	.000	.004	.014***
	(.001)	(.005)	(.006)
β	.49***	.56***	.62***
	(.003)	(.024)	(.027)

*, ** and *** respectively denote effects significant at the .1, .05 and .01 level, standard errors are given in parentheses.

Table 1: In-sample performance of fixed-weights, CPPI, and prior-free optimal strategies, 1927–2014, $N = 88$.

The main finding is that instead of reducing in-sample performance, prior-free optimal strategies improve the Sharpe, the performance, and especially the performance-to-drawdown ratio of the underlying portfolios. Prior-free asset allocation strategies successfully capture time-varying risk-premium.

Figure 6 reports the cumulative log-returns of being long the prior-free portfolio and short the fixed-weights portfolio. Dotted lines separate the sample period in periods of thirty years: 1927–1957, 1957–1987, 1987–2014. The prior-free optimal portfolio over-performs in each subsample, but especially so in the 1927–1957 sample where large swings in returns make drawdown control especially valuable. The long-short strategy's Sharpe ratio over these three subsamples is respectively .54 (1927–1957), .20 (1957–1987), and .30 (1987–2014).

Figure 7 provides further insight into the circumstances in which the prior-free optimal strategy improves on the fixed-weight portfolio. It plots the quantiles of the distribution of returns under the prior-free portfolio against quantiles of the distribution of returns of the fixed-weight portfolio. The prior-free optimal strategy improves both the left and the right tail of returns, but this comes at a cost for yearly returns in the $[-.05, .05]$ range. This makes intuitive sense: in a range-bound market, the prior-free optimal strategy shifts its allocation following small up and down movements. These adjustments guarantee limited drawdowns

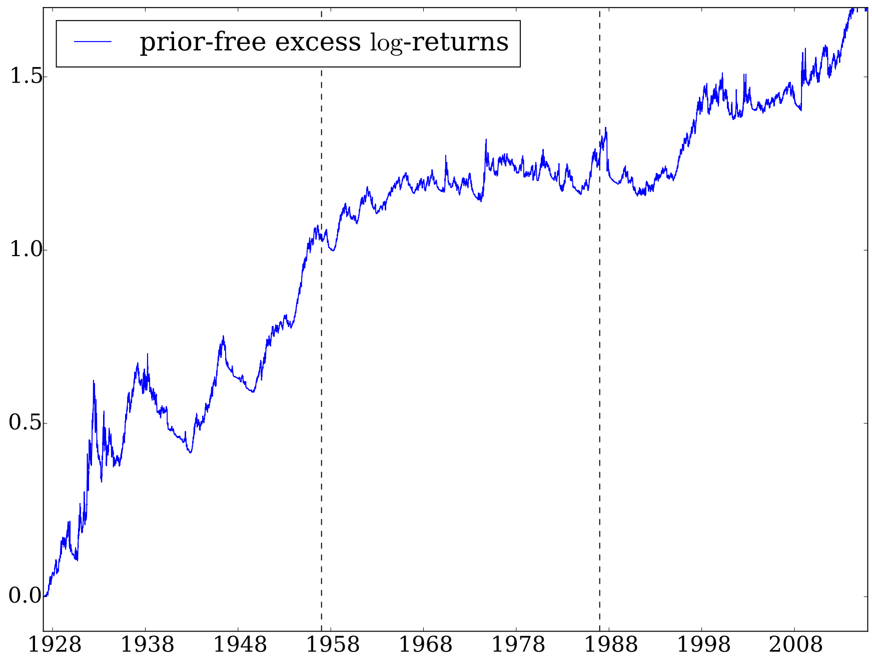

Figure 6: excess log-returns of prior-free optimal portfolio over fixed-weights strategy, 1927–2014.

in case a bull or bear market should emerge. However, if the market remains range-bound, this results in unnecessary transaction costs.

Appendix A reports further empirical findings. First, a Henriksson and Merton (1981) market-timing regression shows that prior-free optimal asset allocation strategies achieve asymmetric β exposure to the market in good and bad years (.7 vs. .4). Second, the prior-free allocation strategy improves on a strategy that goes long the market and hedges large losses using put options. Third, the main empirical findings are not sensitive to the choice of parameters used in setting up drawdown-control problem P_λ.

8 Conclusion

This paper provides a prior-free framework for asset allocation in arbitrarily non-stationary environments. The framework allows decision makers to express risk-preferences by trading off fear-of-loss (potential drawdowns against the safe asset) and fear-of-missing-out (potential

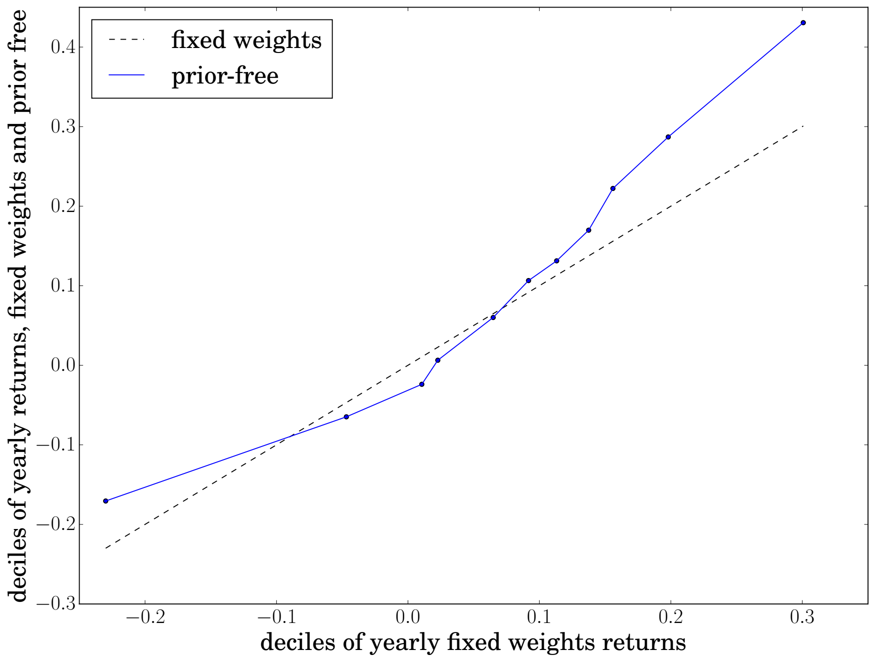

Figure 7: quantile-quantile plot of yearly returns for prior-free optimal strategy and fixed-weights strategy, 1927–2014.

drawdowns against the risky asset). Prior-free optimal allocation strategies are amenable to numerical computation, they are largely interior, and they satisfy a form of momentum. Finally, they are history dependent.

Practically, prior-free optimal strategies offer worst-case drawdown guarantees that improve significantly over those offered by fixed-weight strategies or CPPI. In addition, prior-free optimal strategies perform well in the sample of historical returns, showing that the cost of robustness need not be prohibitive. This is encouraging evidence for a growing agenda that seeks to rethink economic design without probabilistically sophisticated decision makers.

Apendix A presents some extensions. It reports additional information on the behavior of prior-free optimal strategies in data, including robustness checks. In addition, it provides further decision-theoretic perspective on the approach, as well as a discussion of how to extend the prior-free framework to include incomplete probabilistic insights, i.e. restrictions on the likelihood of aggregate events.

Appendix

A Extensions

A.1 Empirical analysis

Market timing and asymmetric risk exposure. Prior-free optimal strategies are intrinsically market timing strategies seeking to achieve asymmetric exposure to market gains and losses. Merton (1981) and Henriksson and Merton (1981) provide a framework to evaluate market timing strategies built on the observation that a good market-timing strategy essentially provides a cheap option, delivering asymmetric exposure to gains and losses.

Table A.1 reports the results from running a Henrikssen-Merton regression of the form

$$r^{\alpha} - r^{0} \sim \beta_1(r^1 - r^0) + \beta_2(r^1 - r^0)^- + \varepsilon \tag{24}$$

where $(r^1 - r^0)^- \equiv \max\{0, -(r^1 - r^0)\}$. Coefficient β_1 represents the strategy's average exposure to the market when net market returns $r^1 - r^0$ are positive. Coefficient β_2 captures the reduction in market exposure when net market returns $r^1 - r^0$ are negative. Being fully invested in the market would yield $\beta_1 = 1$ and $\beta_2 = 0$. In contrast, a free at-the-money call option on the market would achieve perfect market timing with $\beta_1 = \beta_2 = 1$. Table A.1 shows that the prior-free optimal strategy achieves an exposure to the market of .7 when the market is going up, and an exposure to the market of $.4 = .7 - .3$ when the market is going down. As expected, the fixed-weight strategy achieves a roughly constant exposure to the market of roughly .5.

	fixed-weights	prior-free
β_1	0.49***	0.70***
	(.003)	(.022)
β_2	-0.02	0.30***
	(.006)	(.042)

Table A.1: Henrikssen-Merton regression, yearly returns, 1927–2014, $N = 88$.

It is instructive to price the returns of the prior-free optimal strategy against those of an option-based strategy achieving similar asymmetric exposure to the market and providing similar performance guarantees (worst case drawdowns under 30% over a 5 year period). For this exercise, the market is replaced with the S&P 500. The benchmark long/put strategy consists of being long the S&P 500 and buying one-year put options with strike price equal to 98% of the current market price. Returns are computed using option price data from the CBOE available for the time period 1996-2014. Because of this shorter time period, returns are aggregated semiannually rather than annually, meaning that standard errors should be interpreted more cautiously. Table A.2 reports coefficients from a Henrikssen-Merton regression (24) over the time-period 1996-2014. It confirms that the prior-free and long/put strategies offer comparable asymmetric exposures to positive and negative returns.

	prior-free	long/put
β_1	0.72***	0.66***
	(.045)	(.065)
β_2	.35***	0.19*
	(.071)	(.101)

Table A.2: Henrikssen-Merton regression, prior-free and long/put strategies, semiannual returns, 1996-2014, $N = 36$.

Table A.3 reports OLS estimates of regression

$$r^{\text{prior-free}} - r^0 = \alpha + \beta(r^{\text{long/put}} - r^0) + \varepsilon. \tag{25}$$

It confirms that the prior-free strategy is highly correlated to the long/put strategy, but offers asymmetric exposure to positive and negative returns at a significantly lower cost (i.e. a semiannual α of .0121).

α	β
0.0121***	0.857***
(.004)	(.059)

Table A.3: Pricing of prior-free returns, semiannual returns, 1996-2014, $N = 36$.

Sensitivity to parameters. Table A.4 reports the in-sample performance of prior-free allocation strategies other than the one studied in Section 7 ($\lambda = (.5, .5)$, $M = \{-.02, 0, .02\}$, $N = 260$). It considers

- variation in the weights: $\lambda = (.5, .5)$ vs. $(.45, .55)$ or $(.55, .45)$;
- variation in the possible moves of nature: $\pm 2\%$/week vs. $\pm 4\%$/week;

- variation in the time horizon: 5 years ($N = 260$) vs. 3 years ($N = 156$).

	fixed-weights	prior-free (pf), benchmark	pf, $\lambda = (.45, .55)$
net perf	3.7%	6.2%	6.5%
Sharpe	.45	.56	.54
$\mathcal{D}_0$	.54	.29	.30
$\mathcal{D}_1$	.44	.37	.36
net perf to drawdown	.07	.21	.22

	pf, $\lambda = (.55, .45)$	pf, $M = \{-.04, 0, .04\}$	pf, 3 year horizon
net perf	5.8%	5.8%	6.0%
Sharpe	.55	.55	.56
$\mathcal{D}_0$	.28	.31	.31
$\mathcal{D}_1$	.42	.37	.39
net perf to drawdown	.21	.19	.19

Table A.4: in-sample performance for various prior-free optimal strategies, 1927–2014, $N = 88$.

A.2 A decision theoretic perspective

It is useful to cast the approach to dynamic asset allocation developped in this paper in the framework of Savage (1972). As was already emphasized, an event is a realization of returns $\mathbf{r} = (r_t^0, r_t^1)_{t \in \{0, \cdots, N\}}$, and an act is an allocation strategy $\alpha \in \mathcal{A}$ mapping events $\mathbf{r}$ to realized portfolio returns $(r_t^\alpha)_{t \in \{0, \cdots, N\}}$. The premise of Savage (1972) is that it is possible to obtain from the decision maker a preference-ranking over acts. If such a ranking is well-behaved it is then possible to extract an implicit prior over returns $\mathbf{r}$.

Practically, the space of acts $\mathcal{A}$ — i.e. possible allocation strategies — is simply too big for it to be ranked by a decision maker. Instead, a plausible strategy consists of describing acts through a set of moments, and asking the decision-maker to express incomplete preferences over a small dimensional subset of moments providing a meaningful summary of each act. Formally, a moment function is a mapping $\phi : \alpha \in \mathcal{A} \mapsto \phi(\alpha) \in \mathbb{R}$. For a sufficiently large family of moment function $\Phi = (\phi_0, \phi_1, \cdots)$, an act α will be uniquely indentified by its moments $(\phi_0(\alpha), \phi_1(\alpha), \cdots)$. Incomplete preferences over a subset of moments induce incomplete preferences over acts.

This paper can be thought of as implementing such an approach with moments $\Phi = (\phi_0, \phi_1)$ with $\phi_0(\alpha) = \overline{\mathcal{D}}_N^0(\alpha)$ and $\phi_1(\alpha) = \overline{\mathcal{D}}_N^1(\alpha)$. A natural expansion would add in-sample performance as an additional moment, setting $\Phi = (\phi_0, \phi_1, \phi_2)$ with $\phi_2(\alpha) = \frac{1}{N+1}\sum_{t=0}^{N}\log(1+r_t^{\alpha,\mathsf{sample}})$ with $\mathbf{r}^{\mathsf{sample}}$ a relevant sample of returns. Further moments of interest may include sample drawdowns, sample Sharpe, worst-case drawdowns under different assumptions about nature, the distribution of the size and frequency of drawdowns, and so on.

From this perspective, the approach of this paper may be thought of as exploring an approximation of realistic preferences obtained by focusing on two moments of plausible first-order importance.

A.3 Bridging the prior-free and Bayesian approaches

The minimax framework used in this paper allows nature to pick any sequence of returns $\mathbf{r} \in M^{N+1}$. This ensures that prior-free optimal strategies are robust to arbitrary non-stationarity. As was discussed in Section 5, this implies that prior-free optimal strategies must satisfy some form of momentum. As a result, like momentum strategies, prior-free optimal strategies are unanchored to a fundamental value which leads to allocation errors if there exists a true process that exhibits some form of return to the mean (Stein, 2009, Lou and Polk, 2013). This concern indirectly reflects the fact that the decision maker is willing to express some fundamental restrictions on the process for returns. A natural approach would be to include these restrictions directly in the minimax optimization problem. A full fledged investigation is left for further work, but it can be cast along the following lines.

Two types of retrictions may be considered. Non-probabilistic restrictions along the lines already developed in this paper are straightforward. One would restrict the set $Q \subset M^{N+1}$ of possible returns and for every $\lambda \in \Delta(\{0,1\})$, solve the problem

$$\min_{\alpha \in \mathcal{A}} \max_{\mathbf{r} \in Q} \max_{i} \lambda_i \mathcal{D}_N^i(\alpha, \mathbf{r}). \tag{26}$$

For instance the decision maker may be willing to state hard bounds on the possible aggregated growth in asset prices, corresponding to a set Q of the form

$$Q = \left\{ \mathbf{r} \,\middle|\, \forall T \in \{0, \cdots, N\}, \sum_{t=0}^{T} \log(1 + r_t^1) \in \left[\underline{a}T + \underline{b},\, \overline{a}T + \overline{b}\right] \right\}$$

with parameters $(\underline{a}, \underline{b}, \overline{a}, \overline{b}) \in \mathbb{R}^4$ given by the decision-maker.

Alternatively, one may try to capture the problem of a partially sophisticated decision maker. The decision-maker is able to identify coarse sets of events $Q_k \subset M^{N+1}$ for $k \in$

$\{1, \cdots, K\}$, and places a probabilistic restriction $[\underline{\pi}_k, \overline{\pi}_k]$ on the likelihood that nature picks a sequence $\mathbf{r}$ from Q_k. The decision maker then picks the allocation strategy α solving

$$\min_{\alpha \in \mathcal{A}} \max_{\pi \in \Pi} \sum_{k=1}^{K} \pi_k \max_{\substack{\mathbf{r} \in Q_k \\ i \in \{0,1\}}} \mathcal{D}^i(\alpha, \mathbf{r}) \tag{27}$$

where $\pi = (\pi_k)_{k \in \{1, \cdots, K\}}$ and $\Pi = \Delta(\{1, \cdots, K\}) \cap \prod_{k=1}^{K} [\underline{\pi}_k, \overline{\pi}_k]$.

Problem (27) allows to capture the objectives of decision makers with various degrees of probabilistic sophistication: it bridges both the prior-free framework of this paper and the fully sophisticated Bayesian framework. The research agenda going forward is to find useful practical ways to do so.

B Proofs

B.1 Proofs for Section 2

Proof of Lemma 1: By construction, γ must be weakly decreasing. Furthermore, since both $\overline{\mathcal{D}}_N^{0,*}(\lambda)$ and $\overline{\mathcal{D}}_N^{1,*}(\lambda)$ are continuous in λ, it follows that Γ is connected and γ is also continuous. Let us show that γ is strictly decreasing. Pick $\mathcal{D}^0$ and $\epsilon > 0$. Let $\mathcal{D}^1 = \gamma(\mathcal{D}^0)$. We show that there exists a strategy α such that $\overline{\mathcal{D}}^0(\alpha) \leq \mathcal{D}^0 + \epsilon$ and $\overline{\mathcal{D}}^1(\alpha) < \mathcal{D}^1$.

By compactness of $\mathcal{A}$ and continuity of mappings $\overline{\mathcal{D}}^i$, there exists α^A that attains $\mathcal{D}^0$ and $\mathcal{D}^1$. Consider the allocation strategy α constructed as follows: in period 0, a share η of wealth is invested in asset 1 and is not rebalanced; a share $1-\eta$ is invested according to α^A. By continuity of $\overline{\mathcal{D}}^0$, for $\eta > 0$ small enough, $\overline{\mathcal{D}}^0(\alpha) < \mathcal{D}^0 + \epsilon$. Furthermore, for any $T' \leq T$,

$$\begin{aligned}
\sum_{T'}^{T} \log(1+r_t^1) - \log(1+r_t^\alpha) &\leq \log\left(\prod_{t=T'}^{T}(1+r_t^1)\right) - \log\left(\frac{\eta \prod_{t=0}^{T}(1+r_t^1) + (1-\eta)\prod_{t=0}^{T}(1+r_t^{\alpha^A})}{\eta \prod_{t=0}^{T'-1}(1+r_t^1) + (1-\eta)\prod_{t=0}^{T'-1}(1+r_t^{\alpha^A})}\right) \\
&\leq \log\left(\prod_{t=T'}^{T}(1+r_t^1)\right) - \log\left(\frac{a\prod_{t=T'}^{T}(1+r_t^1) + b\prod_{t=T'}^{T}(1+r_t^{\alpha^A})}{a+b}\right) \\
&\leq \frac{b}{a+b}\sum_{t=T'}^{T} \log(1+r_t^1) - \log(1+r_t^{\alpha^A})
\end{aligned}$$

where $a = \eta \prod_{t=0}^{T'-1}(1+r_t^1)$ and $b = (1-\eta)\prod_{t=0}^{T'-1}(1+r_t^{\alpha^A})$. Since a and b are bounded above and bounded away from 0, it follows that $\overline{\mathcal{D}}^1(\alpha) < \overline{\mathcal{D}}^1(\alpha^A)$. $\square$

B.2 Proofs for Section 3

Proof of Proposition 1: The proof is immediate from the definition of $\overline{\mathcal{D}}^i(\alpha)$. $\square$

Proof of Proposition 2: Pick $r_A, r_B, r_C \in M$ such that such that $r_X^0 < r_X^1$ for $X \in \{A, B\}$ and $r_C^0 > r_C^1$. Note that since M is bounded, there exists $\underline{\rho} \in (0,1)$ such that whenever $\mathsf{prob}(r \in \{r_A, r_B\})$ is close enough to 1, the allocation $a_0 = 0$, $a_1 = 1$ solves $\max_{a \in A} \mathbb{E}\left[\log(1 + r^a)\right]$.

Consider the hidden Markov chain $m_0 = (\phi_0, \xi_0)$ defined as follows:

$$\forall k \in \{1, \cdots, K\}, \quad \phi_0(k) = \delta_{(k \bmod K)+1}$$
$$\xi_0(k) = \begin{cases} r_A & \text{if} \quad k \neq K \\ r_B & \text{if} \quad k = K \end{cases}$$

where δ denotes the usual Dirac mass. Markov chain m_0 has a unique sequence of states $\widehat{\mathbf{z}}$, and a unique realization $\widehat{\mathbf{r}}$, which takes the form$(r_A, \cdots, r_A, r_B, r_A \cdots)$. Given a sequence $\mathbf{r}$ of returns, we define $\mathbf{r}^T \equiv (r_t)_{t \in \{0, \cdots, T\}}$ and $\mathbf{r}^{T':T} \equiv (r_t)_{t \in \{T', \cdots, T\}}$.

We identify the set of Markov chains $\mathcal{M}_K$ with the finite-dimensional compact set $(\Delta(Z))^Z \times (\Delta(M))^Z$, endowed with the sup norm. For $\eta > 0$, let $B_\eta(m_0)$ denote the open ball of Markov chains in $\mathcal{M}_K$ within distance η of m_0. For $\rho \in (0,1)$ define $\underline{M}_\rho$ the set of Markov chains in $\mathcal{M}_K$ such that $\mathsf{prob}(r \in \{r_A, r_B\} | z) \leq \rho$ for some state $z \in \{1, \cdots, K\}$. We use the following lemma, which we prove below.

Lemma B.1. *Pick $\rho > .5$. For all T and $m \in \underline{M}_\rho$,*

$$\log\left(\mathit{prob}_{m_0}(\widehat{\mathbf{r}}^T)\right) - \log\left(\mathit{prob}_m(\widehat{\mathbf{r}}^T)\right) > -\frac{T}{K} \log \rho. \tag{28}$$

For all $\eta < .5$ and $m \in B_\eta(m_0)$

$$\log\left(\mathit{prob}_{m_0}(\widehat{\mathbf{r}}^T)\right) - \log\left(\mathit{prob}_m(\widehat{\mathbf{r}}^T)\right) \leq 4\eta T. \tag{29}$$

We now exploit the fact that after N_0 periods under sequence $\widehat{\mathbf{r}}$, a Bayesian investor's beliefs are highly concentrated on Markov chains such that the optimal policy is to invest fully in asset 1, and will remain so for a large amount of time, even if the returns of asset 1 systematically underperform those of asset 0.

There exists $\rho \in (0,1)$ and $\epsilon > 0$ such that if $\mathsf{prob}_\mu(\underline{M}_\rho | h_T) \leq \epsilon$, then the optimal allocation in period T is $a_T^0 = 0$, $a_T^1 = 1$. Pick $\eta > 0$ such that $d_0 \equiv -\frac{1}{K}\log\rho + 2\log(1-\eta) > 0$. Let $D_\eta(m_0) \equiv \{(\phi, \xi) \in B_\eta(m_0) \mid \forall z, \xi(z)(r_C) \geq \eta/2\}$. From Lemma B.1 and Bayes law, it follows that

$$\begin{aligned}\mathsf{prob}_\mu\left(\underline{M}_\rho | \mathbf{r}^{N_0}\right) &\leq \frac{\mathsf{prob}_\mu\left(\underline{M}_\rho\right)\mathsf{prob}_\mu(\mathbf{r}^{N_0}|\underline{M}_\rho)}{\mathsf{prob}_\mu\left(D_\eta(m_0)\right)\mathsf{prob}_\mu(\mathbf{r}^{N_0}|D_\eta(m_0)) + \mathsf{prob}_\mu\left(\underline{M}_\rho\right)\mathsf{prob}_\mu(\mathbf{r}^{N_0}|\underline{M}_\rho)} \\ &\leq \frac{\mathsf{prob}_\mu\left(\underline{M}_\rho\right)}{\mathsf{prob}_\mu(D_\eta(m_0))}\exp\left(-d_0 N_0\right).\end{aligned}$$

Consider the history $\mathbf{r}^{N_0,N_1}$ consisting of $\widehat{\mathbf{r}}^{N_0}$ followed by a sequence $(r_C, \cdots, r_C)$ of N_1 realization r_C. Since r_C has probability bounded away from 0 under $m \in D_\eta(m_0)$, there exists $d_1 > 0$ such that

$$\mathsf{prob}_\mu\left(\underline{M}_\rho | \mathbf{r}^{N_0,N_1}\right) \leq \frac{\mathsf{prob}_\mu\left(\underline{M}_\rho\right)}{\mathsf{prob}_\mu(D_\eta(m_0))}\exp\left(-d_0 N_0 + d_1 N_1\right).$$

Hence, for N large, one can find N_0 and N_1 of order N such that under realization $\mathbf{r}^{N_0,N_1}$, the Bayesian optimal policy is to invest entirely in asset 1 over the time interval $\{N_0, \cdots, N_0 + N_1\}$. This results in a drawdown of at least $N_1[\log(1 + r_C^0) - \log(1 + r_C^1)]$ vis à vis asset 0. $\square$

Proof of Lemma B.1: Note that $\mathsf{prob}_{m_0}(\widehat{\mathbf{r}}^T) = 1$. For any $n \in \mathbb{N}$, and $m \in \underline{M}_\rho$,

$$\mathsf{prob}_m(\widehat{\mathbf{r}}^{nK+1:(n+1)K} | \widehat{\mathbf{r}}^{nK}) = \sum_{\mathbf{z}^{nK+1:(n+1)K}} \mathsf{prob}(\widehat{\mathbf{r}}^{nK+1:(n+1)K} | \mathbf{z}^{nK+1:(n+1)K})\mathsf{prob}(\mathbf{z}^{nK+1:(n+1)K} | \widehat{\mathbf{r}}^{nK}).$$

For all $\mathbf{z}^{nK+1:(n+1)K}$ such that sequence $\widehat{\mathbf{r}}^{nK+1:(n+1)K}$ has positive probability, either:

- there exists $z \in \{z_{nK+1}, \cdots, z_{(n+1)K}\}$ such that $\mathsf{prob}(r \in \{r_A, r_B\}) \leq \rho$,
- the same state z occurs twice in sequence $\mathbf{z}^{nK+1:(n+1)K}$ but is associated with two different realizations of returns r
- the same state z occurs twice in sequence $\mathbf{z}^{nK+1:(n+1)K-1}$ followed up with two distinct continuation states.

Under each of these events, for all $m \in \underline{M}_\rho$, $\mathsf{prob}_m(\widehat{\mathbf{r}}^{nK+1:(n+1)K} | \mathbf{z}^{nK+1:(n+1)K}) \leq \overline{\rho}$, which implies (28). Inequality (29) follows from the fact that for all $m \in B_\eta(m_0)$

$$\begin{aligned}&\mathsf{prob}_m(\widehat{\mathbf{r}}^T) \geq \mathsf{prob}_m(\widehat{\mathbf{r}}^T, \widehat{\mathbf{z}}^T) \geq (1-\eta)^{2T}\\ \Rightarrow &\log(\mathsf{prob}_m(\widehat{\mathbf{r}}^T)) \geq -4\eta T.\end{aligned}$$

□

Proof of Proposition 3: The proof is constructive an uses a slowly rebalanced strategy roughly adapted from Cover (1991). Let $\lfloor x \rfloor$ (resp. $\lceil x \rceil$) denote the highest (lowest) integer less than (higher than) x. Consider the following allocation strategy α whose decision to buy or sell assets depends on time alone:

$$\forall h_T, \quad \alpha(h_T) = \begin{cases} (.5, .5) & \text{if} \quad T \in \{k^2 \,|\, k \in \mathbb{N}\} \\ a_{T^-} & \text{otherwise.} \end{cases}$$

in words, strategy α is a buy-hold strategy which rebalances to symmetric allocation (.5, .5) every period $T \in \{k^2 \,|\, k \in \mathbb{N}\}$. By construction if $T = k^2$, there have been $k+1$ rebalancings. This implies that by period T, this strategy experienced trading costs in at most $\sqrt{T}+1$ periods.

Let us now bound the difference in log returns that can be accumulated over any time interval of the form $[k^2, (k+1)^2]$, excluding trading costs. We have that for all $i \in \{0, 1\}$,

$$\begin{aligned}\sum_{t=k^2}^{(k+1)^2} \log(1+r_t^i) - \log(1+r_t^\alpha) &= \log\left(\prod_{t=k^2}^{(k+1)^2}(1+r_t^i)\right) - \log\left(\frac{1}{2}\prod_{t=k^2}^{(k+1)^2}(1+r_t^0) + \frac{1}{2}\prod_{t=k^2}^{(k+1)^2}(1+r_t^1)\right)\\ &\leq \log 2.\end{aligned}$$

We now prove inequality (10). Define $\overline{c} \equiv \max_{a,a' \in A} c(a, a')$. We use the fact that there exists $h_0 > 0$ such that $\forall a \in A$, $i \in \{0, 1\}$, $r \in M$, $\log(1+r^i) - \log(1+r^a) \leq h_0$. We also use the fact that for any T,

$$\lceil \sqrt{T} \rceil^2 - \lfloor \sqrt{T} \rfloor^2 \leq 2\sqrt{T} + 1.$$

For any $T' \leq T$ and $i \in \{0,1\}$, the following hold

$$\begin{aligned}
\sum_{t=T'}^{T} \log(1+r_t^i) - \log(1+r_t^\alpha) &\leq \sum_{t=\lfloor\sqrt{T'}\rfloor^2}^{\lceil\sqrt{T}\rceil^2} \log(1+r_t^i) - \log(1+r_t^\alpha) + h_0(2\sqrt{T}+1) \\
&\leq \lceil\sqrt{T}\rceil \log 2 + \lceil\sqrt{T}\rceil \bar{c} + h_0(2\sqrt{T}+2) \\
&\leq h_1\sqrt{N}.
\end{aligned}$$

This concludes the proof. □

Proof of Corollary 1: Consider the strategy α defined in the proof of Proposition 3. There exists h such that $\forall i \in \{0,1\}$, $\overline{\mathcal{D}}_N^i(\alpha) \leq h\sqrt{N}$.

By definition, $\alpha_{\lambda,N}$ satisfies

$$\max_{i\in\{0,1\}} \lambda_i \overline{\mathcal{D}}_N^i(\alpha_{\lambda,N}) \leq \max_{i\in\{0,1\}} \lambda_i \overline{\mathcal{D}}_N^i(\alpha) \leq h\sqrt{N},$$

which implies that

$$\lim_{N\to\infty} \frac{\overline{\mathcal{D}}_N^i(\alpha_{\lambda,N})}{N} = 0.$$

□

B.3 Proofs for Section 4

Proof of Lemma 2: Let us begin with point (i). It is immediate that for all $\alpha, \mathbf{r}$, $\mathcal{D}_N^i(\alpha,\mathbf{r}) \geq \mathcal{R}_N^i(\alpha,\mathbf{r})$. Consider now $\mathbf{r}^* \in \arg\max_{\mathbf{r}\in M^{N+1}} \mathcal{D}_N^i(\alpha,\mathbf{r})$. By definition of $\mathcal{D}_N^i$, there exist T, T' such that

$$\mathcal{D}_N^i(\alpha,\mathbf{r}) = \sum_{t=T'}^{T} \log(1+r_t^i) - \log(1+r_t^\alpha).$$

Consider a sequence $\widehat{\mathbf{r}}$ such that

$$\widehat{r}_t = \begin{cases} r_t^* & \text{if} \quad t \leq T \\ (r^0, r^0) & \text{if} \quad t > T. \end{cases}$$

By construction we have that

$$\mathcal{D}_N^i(\alpha,\mathbf{r}^*) \leq \mathcal{R}_N^i(\alpha,\widehat{\mathbf{r}}).$$

Altogether, this implies that $\max_{\mathbf{r}} \mathcal{D}^i_N(\alpha, \mathbf{r}) = \max_{\mathbf{r}} \mathcal{R}^i_N(\alpha, \mathbf{r})$.

Point (ii) follows from the observation that

$$\begin{aligned}
\mathcal{R}^i_{T+1} &= \max_{T' \leq T+2} \sum_{t=T'}^{T+1} \log(1 + r^i_t) - \log(1 + r^\alpha_t) \\
&= \max \left\{ 0, \log(1 + r^i_{T+1}) - \log(1 + r^\alpha_{T+1}) + \max_{T' \leq T+1} \sum_{t=T'}^{T} \log(1 + r^i_t) - \log(1 + r^\alpha_t) \right\} \\
&= \left[\log(1 + r^i_{T+1}) - \log(1 + r^\alpha_{T+1}) + \mathcal{R}^i_T \right]^+ .
\end{aligned}$$

□

Proof of Proposition 4: Point (i) follows from the discussion preceding Proposition 4.

Point (ii) follows from the fact that if α solves $\min_{\alpha \in \mathcal{A}} \max_{\mathbf{r} \in M^{N+1}} \lambda_i \overline{\mathcal{D}}^i_N(\alpha)$, then it must be that $\lambda_0 \overline{\mathcal{D}}^0_N(\alpha) = \lambda_1 \overline{\mathcal{D}}^1_N(\alpha)$. Indeed, imagine we had $\lambda_0 \overline{\mathcal{D}}^0_N(\alpha) > \lambda_1 \overline{\mathcal{D}}^1_N(\alpha)$. For $\eta > 0$ consider the strategy α' such that in period 0, a share η of wealth is invested in asset 0 and is not rebalanced, while a share $1 - \eta$ is invested according to α. For $\eta > 0$ small enough, α' is such that $\lambda_0 \overline{\mathcal{D}}^0_N(\alpha) > \max\{\lambda_0 \overline{\mathcal{D}}^0_N(\alpha'), \lambda_1 \overline{\mathcal{D}}^1_N(\alpha')\}$, contradicting the optimality of α. Together, with point (i), this implies that

$$\lambda_0 \overline{\mathcal{D}}^0_N(\alpha) = \lambda_1 \overline{\mathcal{D}}^1_N(\alpha) = W_\lambda(z_0).$$

□

Proof of Corollary 2: Given any $T \leq N$, regret $\mathcal{R}^i_T$, and continuation sequence $\mathbf{r}^{T:N} \equiv (r_t)_{t \in \{T+1, \cdots, N\}}$, one can compute regret $\mathcal{R}^i_N$, denoted more explicitly by $\mathcal{R}^i_N(\alpha, \mathbf{r}^{T:N}, \mathcal{R}^i_T)$ using the recursion equation of Lemma 2 (ii). Let us show by induction that for any T

$$V^i_\alpha(x^i_T) = \max_{\mathbf{r}^{T:N}} \mathcal{R}^i_N(\alpha, \mathbf{r}^{T:N}, \mathcal{R}^i_T).$$

The result is obviously true for $T = N$. Assume that it is true for $T \leq N$, and let us show it holds for $T-1$:

$$\begin{aligned}
V_\alpha^i(x_{T-1}^i) &= \max_{r_{T-1}\in M} V_\alpha^i(x_T^i) \\
&= \max_{r_{T-1}\in M} \max_{\mathbf{r}^{T:N}} \mathcal{R}_N^i(\alpha, \mathbf{r}^{T:N}, \mathcal{R}_T^i). \\
&= \max_{\mathbf{r}^{T-1:N}} \mathcal{R}_N^i(\alpha, \mathbf{r}^{T-1:N}, \mathcal{R}_{T-1}^i).
\end{aligned}$$

Corollary 2 follows from applying this result to x_0^i. □

B.4 Proofs for Section 5

Proof of Proposition 5: The result is immediate. For any public history $h_t = (r_s)_{s\in\{0,\cdots,t-1\}}$ there exists a marginal distribution $\mu_{t|h_t}$ of r_t such that

$$\alpha_\lambda(h_t) \in \arg\max_{a\in A} \mathbb{E}_{\mu_{t|h_t}}\left[\log(1+r_t^a)\right].$$

The prior μ with conditional marginals $(\mu_{t|h_t})_{\substack{t\in\{0,\cdots,N\}\\ h_t\in\mathcal{H}}}$ admits α_λ as Bayesian optimal, expected utility maximizing, policy. □

Proof of Lemma 3: For any two allocations a^0 and a^1 in period T, denote by α^0 (resp. α^1) the allocation strategies with initial allocation a^0 (resp. a^1) and continuation allocation strategy α_λ. For any $\rho \in (0,1)$ define $\alpha^\rho = (1-\rho)\alpha^0 + \rho\alpha^1$ and $a^\rho = (1-\rho)a^0 + \rho a^1$. We have that

$$U(z_T, a, r) = \max_{\mathbf{r}\in\{r\}\times M^{N-T}} \max_{T'\in\{T,\cdots,N\}} \max_{i\in\{0,1\}} \lambda_i \left[\mathbf{1}_{T'=T}\mathcal{R}_T^i + \sum_{t=T'}^{N} \log(1+r_t^i) - \log(1+r_t^{\alpha^\rho})\right]$$

Concavity of log implies that for all $i \in \{0,1\}$ and all $\mathbf{r} \in \{r\}\times M^{N-T}$,

$$\begin{aligned}
\forall T' \leq N, \quad \sum_{t=T'}^{N} \log(1+r_t^i) - \log(1+r_t^{\alpha^\rho}) \leq& \rho \sum_{t=T'}^{N} \log(1+r_t^i) - \log(1+r_t^{\alpha^1}) \\
&+ (1-\rho)\sum_{t=T'}^{T} \log(1+r_t^i) - \log(1+r_t^{\alpha^0})
\end{aligned}$$

This implies that

$$\max_{\substack{\mathbf{r}\in\{r\}\times M^{N-T}\\ T'\in\{T,\cdots,N\}\\ i\in\{0,1\}}} \sum_{t=T'}^{N} \log(1+r_t^i) - \log(1+r_t^{\alpha^\rho}) \leq \max_{\substack{\mathbf{r}\in\{r\}\times M^{N-T}\\ T'\in\{T,\cdots,N\}\\ i\in\{0,1\}}} \Bigg[\rho \sum_{t=T'}^{N} \log(1+r_t^i) - \log(1+r_t^{\alpha^1}) + (1-\rho)\sum_{t=T'}^{T} \log(1+r_t^i) - \log(1+r_t^{\alpha^0})\Bigg]$$

It follows that $U(z_T, a, r)$ is indeed convex in a.

Point (ii) follows directly from the Minimax theorem (see Luenberger, 1968). Note that optimal allocation $\alpha_\lambda(z_T)$ solves

$$\min_{a\in\Delta(\{0,1\})} \max_{\mu\in\Delta(M)} \mathbb{E}_\mu U(z_t, a, r).$$

Denote by $\mu_\lambda(z_t)$ a solution to

$$\max_{\mu\in\Delta(M)} \min_{a\in\Delta(\{0,1\})} \mathbb{E}_\mu U(z_t, a, r).$$

Lemma 3 implies that $\mathbb{E}_\mu U(z_t, a, r)$ is convex in a and concave in μ. By the Minimax theorem, this implies that $(\alpha_\lambda(z_t), \mu_\lambda(z_t))$ is a mixed Nash equilibrium of the zero-sum game with payoffs to the investor equal to $-U(z_t, a, r)$.

Let us now establish point (iii). Take as given drawdowns $\mathcal{D}_A^0 \neq \mathcal{D}_B^0$ as well as the strategies α^A and α^B implementing the corresponding points on the frontier. Define $\alpha^\rho \equiv \rho\alpha^A + (1-\rho)\alpha^B$. By convexity of mapping $\overline{\mathcal{D}}^0(\cdot)$, we obtain that $\overline{\mathcal{D}}^0(\alpha^\rho) \leq \rho D_A^0 + (1-\rho) D_B^0$. Since γ is decreasing, this implies that

$$\begin{aligned}\gamma(\rho\mathcal{D}_A^0 + (1-\rho)\mathcal{D}_B^0) \leq& \gamma(\overline{\mathcal{D}}^0(\alpha^\rho)) \leq \overline{\mathcal{D}}^1(\alpha^\rho)\\ \leq& \rho\overline{\mathcal{D}}^0(\alpha^A) + (1-\rho)\overline{\mathcal{D}}^1(\alpha^B) = \rho\gamma(\mathcal{D}_A^0) + (1-\rho)\gamma(\mathcal{D}_A^1).\end{aligned}$$

This yields point (iii). □

As a preliminary to Proposition 6, let us prove the following lemmas.

Lemma B.2. *If* $\lambda_i\mathcal{R}^i - \lambda_{-i}\mathcal{R}^{-i} \notin (-\lambda_i(N-T)\hbar_i, \lambda_{-i}(N-T)\hbar_{-i})$*, then* $W_T(\mathcal{R}^0, \mathcal{R}^1) = \max_{i\in\{0,1\}} \mathcal{R}^i$ *and the optimal allocation is such that* $a_i = \mathbf{1}_{\lambda_i\mathcal{R}^i \geq \lambda_{-i}\mathcal{R}_{-i}}$.

Proof of Lemma B.2: Since nature can guarantee $\mathcal{R}^T = \mathcal{R}^{T+1}$ by selecting returns r

such that $r^0 = r^1$, it follows that for all T, $W_T(\mathcal{R}^0, \mathcal{R}^1) \geq \max_{i\in\{0,1\}} \lambda_i \mathcal{R}^i$. The constant allocation setting $a_i = 1$ in every period ensures that $W_T(\mathcal{R}^0, \mathcal{R}^1) = \max_{i\in\{0,1\}} \lambda_i \mathcal{R}^i$. Hence, constant allocation $a_i = 1$ is an optimal continuation strategy. It is in fact uniquely optimal since if $a_i \neq 1$ nature can ensure that $W_T(\mathcal{R}^0, \mathcal{R}^1) > \max_{i\in\{0,1\}} \lambda_i \mathcal{R}^i$ by selecting returns $r \in \arg\max_{r\in M} \max_{i\in\{0,1\}} \mathcal{R}^i + \log(1 + r^i) - \log(1 + r^a)$. □

Lemma B.3. *For all $T \leq N$, $i \in \{0,1\}$, whenever $\lambda_i \mathcal{R}^i - \lambda_{-i} \mathcal{R}^{-i} \in (-\lambda_i(N-T)\hbar_i, \lambda_{-i}(N-T)\hbar_{-i})$, then*

(i) $W_T(\mathcal{R}^0, \mathcal{R}^1) > W_{T+1}(\mathcal{R}^0, \mathcal{R}^1)$;

(ii) $W_T(\mathcal{R}^0, \mathcal{R}^1)$ *is strictly increasing in both* $\mathcal{R}^0$ *and* $\mathcal{R}^1$;

(iii) any Nash equilibrium (a, μ) of the zero-sum game with payoffs $-U(z_t, a, r)$ to the investor are such that a is in the interior of $\Delta(\{0,1\})$, and μ puts positive mass on both $(0, \bar{r})$ and $(0, -\bar{r})$.

Proof of Lemma B.3: The proof is by induction on $T \leq N$. The discussion of Section 5.1 establishes the result for $T = N$. We now show that if the result holds for $T + 1 \leq N$, it must hold for T.

Let us first establish that under the induction hypothesis at $T+1$, the optimal allocation at T is necessarily interior. The optimal allocation a^* at T is a Nash equilibrium of the zero-sum game with payoffs $-U(z_T, a, r)$. Let us show that we cannot have $a_1 \in \{0,1\}$. Indeed, if we had $a_1 = 1$, then nature's payoff take the form $W_{T+1}([\mathcal{R}_T^0 - \log(1+r)]^+, \mathcal{R}_T^1)$. Since (by the induction hypothesis) W_{T+1} is strictly increasing in $\mathcal{R}_{T+1}^0$, the unique best response by nature is to pick $r_1 = -\bar{r}$, inducing a best response $a_1 < 1$ from the investor. Hence $a_1 = 1$ is not part of an equilibrium. A similar reasoning holds if $a_1 = 0$. This implies that the optimal policy must set $a_1 \in (0,1)$.

We turn to point (i). Define $H_T = \{\mathcal{R}_T \text{ s.t. } \lambda_1 \mathcal{R}_T^1 - \lambda_0 \mathcal{R}_T^0 \in (-\lambda_1(N-T)\hbar_1, \lambda_0(N-T)\hbar_0)\}$. Consider first the case where $\mathcal{R}_T \in H_T \setminus H_{T+1}$. By Lemma B.2, this implies that $W_{T+1}(\mathcal{R}^0, \mathcal{R}^1) = \max_{i\in\{0,1\}} \lambda_i \mathcal{R}^i$. Since the optimal allocation is interior, by picking returns $r_1 = -\bar{r}$ or $r_1 = \bar{r}$ nature can ensure that regrets are strictly greater than $\max_{i\in\{0,1\}} \mathcal{R}_T^i$. This implies point (i).

Consider the case where $\mathcal{R}_T \in H_{T+1}$ point (iii) of the induction hypothesis implies that there exists μ^* placing mass on both $(0, -\bar{r})$ and $(0, +\bar{r})$ such that $W_{T+1}(\mathcal{R}_T) =$

$\min_a \mathbb{E}_{\mu^*} W_{T+1}(\mathcal{R}_{T+1})$. Since μ places positive weight on both $(0, -\overline{r})$ and $(0, +\overline{r})$, it follows that for all $a \in \Delta(\{0,1\})$, with positive probability under μ, $\mathcal{R}_{T+1} \in H_{T+1}$. Hence, using the Minimax theorem, we have that

$$\begin{aligned} W_T(\mathcal{R}_T) &= \max_{\widehat{\mu} \in \Delta(M)} \min_{a \in \Delta(\{0,1\})} \mathbb{E}_{\widehat{\mu}} W_{T+1}(\mathcal{R}_{T+1}) \\ &\geq \min_{a \in \Delta(\{0,1\})} \mathbb{E}_{\mu^*} W_{T+1}(\mathcal{R}_{T+1}) \\ &> \min_{a \in \Delta(\{0,1\})} \mathbb{E}_{\mu^*} W_{T+2}(\mathcal{R}_{T+1}) = W_{T+1}(\mathcal{R}_T) \end{aligned}$$

where the last strict inequality follows from induction hypothesis (i) for $T+1$. This proves point (i) for T.

An implication on point (i) is that whenever $\mathcal{R}_T \in H_T$, nature's equilibrium strategy cannot put weight on returns $(0,0)$ since for this value of r, $U(\mathcal{R}_T, a, r) = W_{T+1}(\mathcal{R}_T)$. Together with the fact that a is interior, this implies point (iii).

We now turn to point (ii). Since the optimal allocation a is interior, nature's equilibrium strategy must be mixing between actions $(0, -\overline{r})$ and $(0, \overline{r})$. We have that

$$W_T(\mathcal{R}_T) = \min_{a \in \Delta(\{0,1\}} \max_{\mu \in \Delta(M)} \mathbb{E}_\mu \left[W_{T+1}(\mathcal{R} + T + 1) \right].$$

For all realizations of equilibrium distribution μ^*, $W_{T+1}(\mathcal{R}_{T+1})$ is weakly increasing in $\mathcal{R}^1$. Furthermore, it follows from the induction hypothesis that with strictly positive probability under μ^*, realized returns are such that $W_{T+1}(\mathcal{R}_{T+1})$ is strictly increasing in $\mathcal{R}^1$. This implies that $W_T(\mathcal{R}_T)$ is strictly increasing in $\mathcal{R}^1$. A similar reasoning implies that $W_T(\mathcal{R}_T)$ is strictly increasing in $\mathcal{R}^0$. This concludes the proof of Lemma B.3. □

Proof of Proposition 6: Proposition 6 follows directly from Lemmas B.2 and B.3. □

Proof of Corollary 3: The result follows from Propositions 3 and 6 (ii). Proposition 3 implies that there exist $h > 0$ such that for all N and all sequences of returns $\mathbf{r}$,

$$|\lambda_1 \mathcal{R}^1_N(\alpha_{\lambda,N}, \mathbf{r}) - \lambda_0 \mathcal{R}^0_N(\alpha_{\lambda,N}, \mathbf{r})| \leq h\sqrt{N}.$$

This and Proposition 6 implies that allocation $\alpha_{N,\lambda}(h_t)$ is a corner allocation for at most $O(\sqrt{N})$ periods. □

Proof of Corollary 4: The fact that $\mathbb{E}_\mu[\frac{r_t^1}{1+r_t^1}|h_t] > \epsilon$ implies that $\mathbb{E}_\mu[\log(1+r_t^1) - \log(1+a_1 r_t^1)] > \epsilon(1-a_1)$. Assume that there exists $\nu > 0$ such that

$$\mathsf{prob}\left(\liminf_N \frac{1}{T_2 - T_1}\sum_{t=T_1}^{T_2} \alpha^1_{\lambda,N}(h_t) < 1-\nu\right) > \nu.$$

This implies that

$$\begin{aligned}
\limsup \mathbb{E}_\mu & \left(\frac{1}{T_2 - T_1}\sum_{t=T_1}^{T_2} \log(1+r_t^1) - \log(1+r_t^{\alpha_{\lambda,N}})\right) \\
&\geq \limsup \epsilon \mathbb{E}_\mu \left[\frac{1}{T_2 - T_1}\sum_{t=T_1}^{T_2} 1 - \alpha^1_{\lambda,N}(h_t)\right] \\
&\geq \epsilon\nu^2,
\end{aligned}$$

which contradicts the fact that $\overline{\mathcal{D}}^1_N(\alpha_{\lambda,N}) = O(\sqrt{N})$. □

B.5 Proofs for Section 6

Proof of Proposition 7: Point (i) and (iii) are immediate. Point (ii) is proven by exhibiting an asset allocation strategy that guarantees regrets $\mathcal{R}^{i,\sigma}_N$ of order $O(\sqrt{N})$. The main step is to prove the approachability condition required in Blackwell (1956).

Consider some asset $i \in \mathcal{I}$. For any $\rho \in [0,1]$, let $\psi(\rho) \in [\underline{a}^i, \overline{a}^i]$ denote a solution to equation

$$\begin{aligned}
a_i &= \rho a^{i,+} + (1-\rho) a^{i,-} \\
&= \rho \min\{a_i + \Delta, \overline{a}^i\} + (1-\rho)\max\{a_i - \Delta, \underline{a}^i\}
\end{aligned}$$

Elementary algebra shows that: existence is immediate; unicity holds whenever $\rho \neq \frac{1}{2}$; ψ is increasing in ρ.

Regets $\mathcal{R}^{i,\sigma}_T$ satisfy

$$\begin{aligned}
\mathcal{R}^{i,\sigma}_{T+1} =& [\mathcal{R}^{i,\sigma}_T + g^\sigma(a^i_{T+1}, r^i_{T+1}, r^{\neg i}_{T+1})]^+ \\
=& [\mathcal{R}^{i,\sigma}_T + \log(1 + a^{i,\sigma}_{T+1} r^i_{T+1} + (1 - a^{i,\sigma}_{T+1}) r^0 - r^{\neg i}_{T+1}) \\
& - \log(1 + a^i_{T+1} r^i_{T+1} + (1 - a^i_{T+1}) r^0 - r^{\neg i}_{T+1})]^+.
\end{aligned}$$

By concavity of the log, it follows from simple algebra that setting

$$a^i_{T+1} = \psi\left(\frac{\mathcal{R}^{i,+}_T}{\mathcal{R}^{i,-}_T + \mathcal{R}^{i,+}_T}\right) \tag{30}$$

ensures that the following approachability condition holds

$$\forall r^i_{T+1}, \forall r^{\neg i}_{T+1}, \quad \mathcal{R}^{i,+}_T \times g^+(a^i_{T+1}, r^i_{T+1}, r^{\neg i}_{T+1}) + \mathcal{R}^{i,-}_T \times g^-(a^i_{T+1}, r^i_{T+1}, r^{\neg i}_{T+1}) \leq 0.$$

It follows from Blackwell (1956) (see also Foster and Vohra (1999) for a modern treatment) that the policy defined by (30) ensures regrets of order $\sqrt{N}$. □

Proof of Corollary 5: Condition (21) implies that for all $i \in \mathcal{I}$, there exists $\nu > 0$ and $\sigma \in \{+1, -1\}$ such that

$$\nu\sigma(a^{\sigma,i} - a^i) \leq \mathbb{E}_\mu[g^\sigma(a^i, r^i, r^{\neg i})]. \tag{31}$$

Without loss of generality, we can consider the case where $\sigma_i = +1$. This implies that with probability approaching 1 for N large, the allocations $(a^i_t)_{t\in\{T_1,\cdots,T_2\}}$ generated by prior-free strategy $\alpha^{*,i}$ satisfy

$$\frac{\nu}{2}\sum_{t=T_1}^{T_2} a^{+,i}_t - a^i_t \leq \sum_{t=T_1}^{T_2} g^+(a^i_t, r^i_t, r^{\neg i}_t) = O(\sqrt{N}). \tag{32}$$

Since $a^{+,i} = \max\{\overline{a}^i, a^i + \Delta\}$, (32) implies that

$$\sum_{t=T_1}^{T_2} a^{+,i}_t - a^i_t = \sum_{t=T_1}^{T_2} (\overline{a}^i - a^i_t)\mathbf{1}_{\overline{a}^i < a^i+\Delta} + \Delta\sum_{t=T_1}^{T_2} \mathbf{1}_{\overline{a}^i > a^i+\Delta} = O(\sqrt{N}).$$

This implies that $\sum_{t=T_1}^{T_2} \mathbf{1}_{\overline{a}^i > a^i+\Delta} = O(\sqrt{N})$ and therefore that $\sum_{t=T_1}^{T_2} \overline{a}^i - a^i_t = O(\sqrt{N})$. Since $T_2 - T_1$ is of order N, this implies that

$$\mu\text{-a.s.}, \quad \frac{1}{T_2 - T_1}\sum_{t=T_1}^{T_2} a^i_t \to \overline{a}^i.$$

Since a similar argument holds for $\sigma_i = -1$, this implies Corollary 5. □

References

Antic, N. (2014): "Contracting with Unknown Technologies," *unpublished paper, Princeton University.*

Asl, F. M. and E. Etula (2012): "Advancing strategic asset allocation in a multi-factor world," *The Journal of Portfolio Management*, 39, 59–66.

Asness, C. S., T. J. Moskowitz, and L. H. Pedersen (2013): "Value and momentum everywhere," *The Journal of Finance*, 68, 929–985.

Barberis, N., A. Shleifer, and R. Vishny (1998): "A model of investor sentiment," *Journal of financial economics*, 49, 307–343.

Baum, L. E. and T. Petrie (1966): "Statistical inference for probabilistic functions of finite state Markov chains," *The annals of mathematical statistics*, 1554–1563.

Bergemann, D. and K. H. Schlag (2008): "Pricing without priors," *Journal of the European Economic Association*, 6, 560–569.

Black, F. and R. W. Jones (1987): "Simplifying portfolio insurance," *The Journal of Portfolio Management*, 14, 48–51.

Black, F. and R. Litterman (1992): "Global portfolio optimization," *Financial Analysts Journal*, 28–43.

Black, F. and A. Perold (1992): "Theory of constant proportion portfolio insurance," *Journal of Economic Dynamics and Control*, 16, 403–426.

Blackwell, D. (1956): "An Analog of the Minimax Theorem for Vector Payoffs," *Pacific Journal of Mathematics*, 6, 1–8.

Brooks, B. (2014): "Surveying and selling: Belief and surplus extraction in auctions," Tech. rep., Princeton University working paper.

Cai, X., K.-L. Teo, X. Yang, and X. Y. Zhou (2000): "Portfolio optimization under a minimax rule," *Management Science*, 46, 957–972.

Campbell, J. Y. (1984): "Asset Duration and Time Varying Risk Premia," *unpublished dissertation, Yale University.*

Campbell, J. Y. and L. M. Viceira (1999): "Consumption and Portfolio Decisions when Expected Returns are Time Varying," *The Quarterly Journal of Economics*, 114, 433–495.

Carroll, G. (2014): "Robustness and Linear Contracts," *the American Economic Review*.

Ceria, S. and R. A. Stubbs (2006): "Incorporating estimation errors into portfolio selection: Robust portfolio construction," *Journal of Asset Management*, 7, 109–127.

Cesa-Bianchi, N. and G. Lugosi (2006): *Prediction, Learning, and Games*, Cambridge University Press.

Chassang, S. (2013): "Calibrated Incentive Contracts," *Econometrica*.

Cover, T. M. (1991): "Universal portfolios," *Mathematical finance*, 1, 1–29.

DeMarzo, P., I. Kremer, and Y. Mansour (2009): "Hannan and Blackwell meet Black and Scholes: Approachability and Robust Option Pricing," Tech. rep., Working paper, Stanford University.

DeMiguel, V., L. Garlappi, and R. Uppal (2009): "Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?" *Review of Financial Studies*, 22, 1915–1953.

Diaconis, P. and D. Freedman (1986): "On the consistency of Bayes estimates," *The Annals of Statistics*, 1–26.

Foster, D. P. and R. Vohra (1999): "Regret in the on-line decision problem," *Games and Economic Behavior*, 29, 7–35.

Ghosh, J. K. and R. Ramamoorthi (2003): *Bayesian nonparametrics*, vol. 1, Springer.

Gilboa, I. and D. Schmeidler (1989): "Maxmin expected utility with non-unique prior," *Journal of mathematical economics*, 18, 141–153.

Glasserman, P. (2003): *Monte Carlo methods in financial engineering*, vol. 53, Springer Science & Business Media.

Glasserman, P. and X. Xu (2013): "Robust portfolio control with stochastic factor dynamics," *Operations Research*, 61, 874–893.

——— (2014): "Robust risk measurement and model risk," *Quantitative Finance*, 14, 29–58.

Golberg, D. E. (1989): "Genetic algorithms in search, optimization, and machine learning," *Addion Wesley*, 1989.

Grossman, S. J. and Z. Zhou (1993): "Optimal investment strategies for controlling drawdowns," *Mathematical Finance*, 3, 241–276.

Hannan, J. (1957): "Approximation to Bayes Risk in Repeated Play," in *Contributions to the Theory of Games*, ed. by M. Dresher, A. Tucker, and P. Wolfe, vol. 3, 97–139.

Hansen, L. P. and T. J. Sargent (2001): "Robust control and model uncertainty," *American Economic Review*, 91, 60–66.

——— (2008): *Robustness*, Princeton university press.

Hartline, J. D. and T. Roughgarden (2008): "Optimal Mechanism Design and Money Burning," in *Symposium on Theory Of Computing (STOC)*, 75–84.

Henriksson, R. D. and R. C. Merton (1981): "On market timing and investment performance. II. Statistical procedures for evaluating forecasting skills," *Journal of business*, 513–533.

Hong, H., T. Lim, and J. C. Stein (2000): "Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies," *The Journal of Finance*, 55, 265–295.

Hong, H. and J. C. Stein (1999): "A unified theory of underreaction, momentum trading, and overreaction in asset markets," *The Journal of Finance*, 54, 2143–2184.

Jegadeesh, N. and S. Titman (1993): "Returns to buying winners and selling losers: Implications for stock market efficiency," *The Journal of Finance*, 48, 65–91.

——— (2001): "Profitability of momentum strategies: An evaluation of alternative explanations," *The Journal of Finance*, 56, 699–720.

Kőszegi, B. and M. Rabin (2006): "A model of reference-dependent preferences," *The Quarterly Journal of Economics*, 1133–1165.

Lettau, M. and S. Ludvigson (2001): "Resurrecting the (C) CAPM: A cross-sectional test when risk premia are time-varying," *Journal of Political Economy*, 109, 1238–1287.

Lettau, M. and S. C. Ludvigson (2010): "Measuring and Modeling Variation in the Risk-Return Trade-off," *Financial Econometrics*, 617.

Lo, A. W. and A. C. MacKinlay (1990): "Data-snooping biases in tests of financial asset pricing models," *Review of financial studies*, 3, 431–467.

Lou, D. and C. Polk (2013): "Comomentum: Inferring Arbitrage Activity from Return Correlations," .

Luenberger, D. G. (1968): *Optimization by vector space methods*, John Wiley & Sons.

Madarász, K. and A. Prat (2016): "Screening with an Approximate Type Space," *Review of Economic Studies*.

Malmendier, U. and S. Nagel (2011): "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics*, 126.

Markowitz, H. (1952): "Portfolio selection," *The Journal of Finance*, 7, 77–91.

Merton, R. C. (1981): "On market timing and investment performance. I. An equilibrium theory of value for market forecasts," *Journal of business*, 363–406.

Milnor, J. (1954): "Games against nature," .

Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen (2012): "Time series momentum," *Journal of Financial Economics*, 104, 228–250.

Novy-Marx, R. (2014): "Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars," *Journal of Financial Economics*, 112, 137–146.

Perold, A. (1986): "Constant proportion portfolio insurance," .

Pflug, G. and D. Wozabal (2007): "Ambiguity in portfolio selection," *Quantitative Finance*, 7, 435–442.

Savage, L. J. (1951): "The theory of statistical decision," *Journal of the American Statistical association*, 46, 55–67.

——— (1972): *The foundations of statistics*, Courier Corporation.

Segal, I. (2003): "Optimal pricing mechanisms with unknown demand," *The American economic review*, 93, 509–529.

Sims, C. A. (1971): "Distributed lag estimation when the parameter space is explicitly infinite-dimensional," *The Annals of Mathematical Statistics*, 42, 1622–1636.

Stein, J. C. (2009): "Presidential address: Sophisticated investors and market efficiency," *The Journal of Finance*, 64, 1517–1548.

Stoye, J. (2008): "Axioms for minimax regret choice correspondences," *New York University, unpublished*.

Tversky, A. and D. Kahneman (1991): "Loss aversion in riskless choice: A reference-dependent model," *The quarterly journal of economics*, 1039–1061.

Wald, A. (1950): "Statistical decision functions." .

Zweig, J. (2007): *Your money and your brain: how the new science of neuroeconomics can help make you rich*, Simon and Schuster.