



Research Paper

Mostly prior-free asset allocation

Sylvain Chassang

New York University Department of Economics, 19 West 4th Street, New York, NY 10011, USA;
email: chassang@nyu.edu

(Received September 12, 2017; accepted October 30, 2017)


ABSTRACT

This paper develops a prior-free version of Harry Markowitz's efficient portfolio theory, which allows the decision maker to express their preferences with regard to risk and reward, even though they are unable to express a prior over potentially nonstationary returns. The corresponding optimal allocation strategies are admissible and interior, and they exhibit a form of momentum. Empirically, prior-free efficient allocation strategies successfully exploit the time-varying risk premiums present in historical returns.

Keywords: prior-free asset allocation; minmax drawdown control; nonstationary returns; fear-of-missing-out; fear-of-loss; regret aversion.

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 for the last thirty years, mechanically cannot keep going down much longer (see Figure 1). In this case, we know that the next thirty years must look different. Sometimes, change is only a possibility with which decision makers are concerned. For instance, an investor interested in investing in a smart-beta index fund exploiting one of the familiar premium anomalies (eg, value, momentum, low volatility, low beta) may be

FIGURE 1 The ten-year US Treasury rate, 1962–2014.

Source: FRED. URL: <http://bit.ly/1MViNEQ>.

plausibly worried that those strategies will become crowded and fail to deliver advertised returns. In a nonstationary environment, past data provides limited guidance on future behavior, which begs the question: how do we make practical risk management and portfolio allocation decisions in such a nonstationary world?

The benchmark framework for portfolio allocation, the efficient portfolio theory of Markowitz (1952), 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 dynamic environments with a time-varying and possibly nonstationary risk premium, the difficulty of specifying priors is further increased.¹ 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 exhibit poor frequentist behavior (for instance, failing to converge

¹ See Campbell (1984), Campbell and Viceira (1999) or Lettau and Ludvigson (2001, 2009) for evidence of time variation in risk premiums.

to true parameters; see, for example, Sims (1971), Diaconis and Freedman (1986) and 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 nonstationary. Prior-free asset allocation seeks to provide such a framework by giving up on priors all together.

The logic behind prior-free asset allocation matches Harry Markowitz's 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's intuitive approach as an aversion to worst-case drawdowns (ie, 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 nonstationary or novel environments. One simple takeaway is that Markowitz's 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 their 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 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, this paper refers to these strategies as prior-free optimal.

Our paper makes four points. The first is that prior-free optimal portfolios satisfy a form of robustness to nonstationarity that 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 with one of the two reference assets.

The second point is that prior-free asset allocation lets decision makers express preferences with regard to 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 trade-offs 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 us 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, this paper provides a brief numerical and empirical exploration of this prior-free approach to portfolio optimization. Under worst-case analysis, prior-free portfolios significantly improve on popular portfolio construction rules, such as regularly rebalanced portfolios (RRPs) 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 and moving-average rules) can be mapped against the worst-case drawdown frontier. The benefits of any strategy of interest (eg, 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 independent and identically distributed (iid), a fixed regularly rebalanced portfolio may deliver a 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 premiums 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 (0.7 versus 0.4). 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 this is set in advance of any exposure to data.

The paper connects to an applied literature in portfolio management that seeks to usefully operationalize the approach of Markowitz (1952). Black and Litterman (1992) also place priors at the center of their analysis. They show that naive 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. 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 a 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 example, Asl and Etula 2012; Ceria and Stubbs 2006) 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 (Milnor 1954; Savage 1951; Stoye 2008; Wald 1950) and behavioral literatures (Kőszegi and Rabin 2006; Tversky and Kahneman 1991). It is closely related to the question of online regret minimization originally studied in Blackwell (1956) and Hannan (1957) (for a recent reference, see Cesa-Bianchi and Lugosi (2006)). 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 that 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 with regard to risk and reward while allowing for arbitrary nonstationarity 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 on existing benchmarks in the realized historical sample. In addition, prior-free approaches reduce concerns of data-snooping bias.

This paper is structured as follows. Section 2 defines the framework and the prior-free asset allocation problem. Section 3 quantifies robustness to nonstationarity 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 (available online) 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 (also available online) 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 2.1 *There exists $r \in M$ such that $r^0 = r^1$. There exist $(r, \hat{r}) \in M^2$ such that $r^0 > r^1$, and $\hat{r}^0 < \hat{r}^1$. Set M contains at least three nondiagonal returns r such that $r^0 \neq r^1$.*

In computational applications, set M will take the form

$$M = \{r^0\} \times \{r^0 + n\Delta \mid n = -k, \dots, +k\} \quad \text{for } k \in \mathbb{N}.$$

Let us denote by $\bar{r} = \max\{|r^0|, |r^1| \text{ for } r \in M\}$ 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

$$w_{t+1^-} = w_t(1 + r_t^a), \quad a_{t+1^-}^0 = \frac{a_t^0(1 + r_t^0)}{1 + r_t^a}, \quad a_{t+1^-}^1 = \frac{a_t^1(1 + r_t^1)}{1 + r_t^a}. \quad (2.1)$$

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 $\bar{c} \equiv \max_{a, \hat{a}} c(a, \hat{a})$ the highest trading cost. In numerical applications, trading costs will take the form $c(a, \hat{a}) = c_1|a^0 - \hat{a}^0|$, with $c_1 = 0.002$ (ie, 20 basis points (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 $h_t = (r_s)_{s \in \{0, \dots, t-1\}} \in \mathcal{H}$ to an allocation a_t . We denote by \mathcal{A} the set of possible allocation strategies.

Taking into account trading costs, the returns associated with allocation strategy α in period t , denoted by r_t^α , take the form

$$1 + r_t^\alpha = (1 + \langle \alpha(h_t), r_t \rangle)(1 - c(a_{t-1}, \alpha(h_t))). \quad (2.2)$$

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

$$\alpha: z \in Z \mapsto \alpha(z) \in A, \quad (2.3)$$

$$\phi: (z, r) \in Z \times M \mapsto \phi(z, r) \in Z. \quad (2.4)$$

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

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

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 nonstationary?

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. Their optimal asset allocation strategy α then solves

$$\max_{\alpha \in \mathcal{A}} \mathbb{E}_\mu \left[\sum_{t=0}^N \log(1 + r_t^\alpha) \right]. \quad (2.5)$$

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 is not 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. The ex post assessment of Black and Litterman (1992) 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, and not a primitive of the decision problem.

The normative decision rule, which 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 nonstationary, 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's description of his own investment behavior suggests the following three observations:

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

Definitions 2.2, 2.3 and 2.4 formalize preferences over allocation strategies that capture these premises.

DEFINITION 2.2 (Relative drawdowns) Given an allocation strategy $\alpha: \mathcal{H} \rightarrow \Delta(\{0, 1\})$ and a sequence of returns $\mathbf{r} \equiv (r_t)_{t \in \{0, \dots, N\}}$, drawdowns \mathcal{D}_N^0 and \mathcal{D}_N^1 relative to the safe and risky asset are defined as

$$\mathcal{D}_N^0(\alpha, \mathbf{r}) = \max_{\substack{T \in \{0, \dots, N\} \\ T' \in \{0, \dots, T+1\}}} \sum_{t=T'}^T \log(1 + r_t^0) - \log(1 + r_t^\alpha), \quad (2.6)$$

$$\mathcal{D}_N^1(\alpha, \mathbf{r}) = \max_{\substack{T \in \{0, \dots, N\} \\ T' \in \{0, \dots, T+1\}}} \sum_{t=T'}^T \log(1 + r_t^1) - \log(1 + r_t^\alpha). \quad (2.7)$$

Given realized returns $\mathbf{r} \equiv (r_t)_{t \in \{0, \dots, N\}}$, 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-weight strategy over the market and the risk-free rate during the 2007–12 period.² Note that the time windows $[T', T]$ over which each drawdown occurs are different for the risk-free and risky assets.

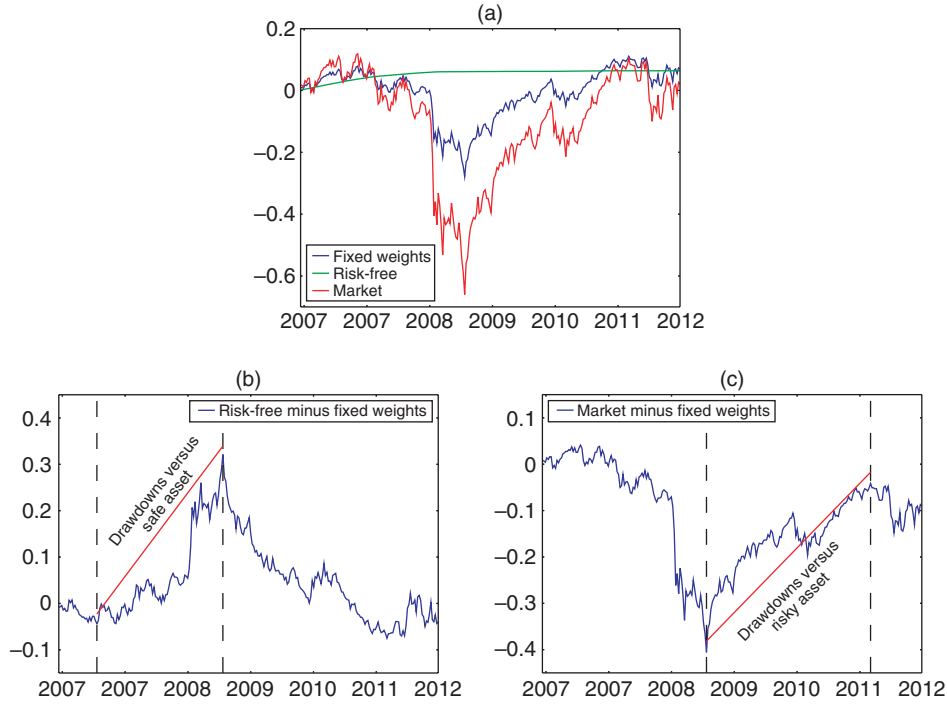
DEFINITION 2.3 (Worst-case drawdowns) Given a strategy α , worst-case drawdowns are defined by

$$\bar{\mathcal{D}}_N^0(\alpha) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{D}_N^0(\alpha, \mathbf{r}), \quad (2.8)$$

$$\bar{\mathcal{D}}_N^1(\alpha) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{D}_N^1(\alpha, \mathbf{r}). \quad (2.9)$$

² Returns are obtained from Kenneth French's data library, available at <http://bit.ly/1jwasZk>.

FIGURE 2 Drawdowns relative to the safe and risky assets for a 50/50 fixed-weight strategy, 2007–12.



(a) Cumulated log returns. (b) Drawdowns, risk-free minus fixed weights. (c) Drawdowns, market minus fixed weights.

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

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

DEFINITION 2.4 (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

$$\min_{\hat{\alpha} \in \mathcal{A}} \max_{i \in \{0, 1\}} \lambda_i \bar{\mathcal{D}}_N^i(\hat{\alpha}). \quad (P_\lambda)$$

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

by

$$\Gamma \equiv \{(\tilde{\mathcal{D}}_N^{i,*}(\lambda))_{i \in \{0,1\}} \mid \lambda \in \Delta(\{0,1\})\}.$$

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

LEMMA 2.5 *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 $\tilde{\mathcal{D}}_0$ against the safe asset, it returns the best possible drawdown guarantee $\tilde{\mathcal{D}}_1$ 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 nonstationary processes for returns; (ii) they let the decision maker express meaningful risk preferences over complex acts in a simple manner; and (iii) they perform well in the data.

3 ROBUSTNESS TO NONSTATIONARITY

This section motivates the use of prior-free optimal strategies by (i) highlighting their robustness to nonstationarity, and (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 premiums, 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 3.1 (A performance bound) *Consider a strategy α and a sequence of returns \mathbf{r} . For all time periods $T_1 < T_2 \leq N$, and for all $i \in \{0, 1\}$,*

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

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

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

In other words, drawdown guarantees imply lower bounds on the performance of strategy α . Up to a penalty $\bar{\mathcal{D}}^i(\alpha)$, 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 3.2 A sequence of asset allocation strategies $(\alpha_N)_{N \in \mathbb{N}}$ (indexed on increasing time horizon N) is said to be robust to nonstationarity if and only if

$$\lim_{N \rightarrow \infty} \frac{\bar{\mathcal{D}}_N^i(\alpha_N)}{N} = 0 \quad \text{for all } i \in \{0, 1\}.$$

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

3.2 Fragility of finite hidden Markov models

Hidden Markov models are a popular and flexible way to model time-varying processes. However, Proposition 3.3 (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 $Z = \{1, \dots, K\}$ is described by $m = (\phi, \xi) \in (\Delta(Z))^Z \times (\Delta(M))^Z$, where $\phi: Z \rightarrow \Delta(Z)$ 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

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

Note that the set \mathcal{M}_K of hidden Markov chains m with fewer 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 3.2. Proposition 3.3 shows that this is not the case.

PROPOSITION 3.3 *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 other 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 nonstationarity.

3.3 The possibility of robustness

To be useful as a selection criterion, robustness to nonstationarity needs to be nonempty. Proposition 3.4 shows that there exist allocation strategies which guarantee sublinear drawdowns for all possible realized sequences of returns.

PROPOSITION 3.4 (Robustness) *For all $\bar{c} < 1$ and \bar{r} , there exists $h > 0$ and a strategy α such that*

$$\max_{\mathbf{r} \in M^{N+1}} \max_{i \in \{0,1\}} \mathcal{D}_N^i(\alpha, \mathbf{r}) \leq h\sqrt{N} \quad \text{for all } N \in \mathbb{N}. \quad (3.1)$$

Together, Propositions 3.3 and 3.4 establish that it is possible to find strategies that are robust to nonstationarity, but they cannot be obtained by modeling the underlying returns as an unknown finite hidden Markov model.

An immediate corollary of Proposition 3.4 is that prior-free optimal strategies are robust to nonstationarity. Indeed, they achieve the smallest possible drawdowns.

COROLLARY 3.5 *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 nonstationarity.*

It is important to note that robustness to nonstationarity is only an asymptotic property that 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 the 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 (eg, in-sample performance).
- It provides a systematic framework for optimal dynamic allocation that can incorporate relevant economic features of the problem, such as trading costs or restrictions on the process of returns (bounds on P/E ratios).³

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 .

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

$$\mathcal{R}_T^i(\alpha, \mathbf{r}) \equiv \max_{T' \in \{0, \dots, T+1\}} \sum_{t=T'}^T \log(1 + r_t^i) - \log(1 + r_t^\alpha). \quad (4.1)$$

Regret \mathcal{R}_T^i differs from drawdown \mathcal{D}_T^i in that the endpoint T of the period over which underperformance is measured is fixed. In fact, we have that $\mathcal{D}_T^i = \max_{T' \leq T} \mathcal{R}_{T'}^i$. Lemma 4.1 shows that \mathcal{R}_T^i can be used to compute worst-case drawdowns. It is described by a simple dynamic process.

LEMMA 4.1 *For all $\alpha \in \mathcal{A}$,*

- (i) $\max_{\mathbf{r} \in M^{N+1}} \mathcal{D}_N^i(\alpha, \mathbf{r}) = \max_{\mathbf{r} \in M^{N+1}} \mathcal{R}_N^i(\alpha, \mathbf{r})$ for all $i \in \{0, 1\}$;
- (ii) $\mathcal{R}_{T+1}^i = [\mathcal{R}_T^i + \log(1 + r_{T+1}^i) - \log(1 + r_{T+1}^\alpha)]^+$ for all $T < N$.

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

³ See Appendix A (available online) for a discussion of various ways to place restrictions, including probabilistic ones, on the set of possible returns.

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 $T + 1$ and returns in 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:

$$W_\lambda(z_T) = \begin{cases} \max_{i \in \{0,1\}} \lambda_i \mathcal{R}_T^i & \text{if } T = N + 1, \\ \min_{a_T \in A} \max_{r \in M} W_\lambda(z_{T+1}) & \text{if } T \leq N, \end{cases} \quad (4.2)$$

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

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

PROPOSITION 4.2 (Bellman formulation) *Let $z_0 = (0, a_0, 0, 0)$. The following hold.*

- (i) $\min_\alpha \max_{i \in \{0,1\}} \lambda_i \bar{\mathcal{D}}_N^i(\alpha) = W_\lambda(z_0)$. Drawdown minimizing policy α_λ^* depends only on states $(z_t)_{t \geq 0}$ and is defined by

$$\alpha_\lambda^*(z_t) \in \arg \min_{a_t \in A} \max_{r_t \in M} W_\lambda(z_{t+1}) \quad \text{for all } z_t.$$

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

$$\Gamma = \left\{ \left(\frac{1}{\lambda_i} W_\lambda(z_0) \right)_{i \in \{0,1\}} \text{ for } \lambda \in \Delta(\{0, 1\}) \right\}. \quad (4.3)$$

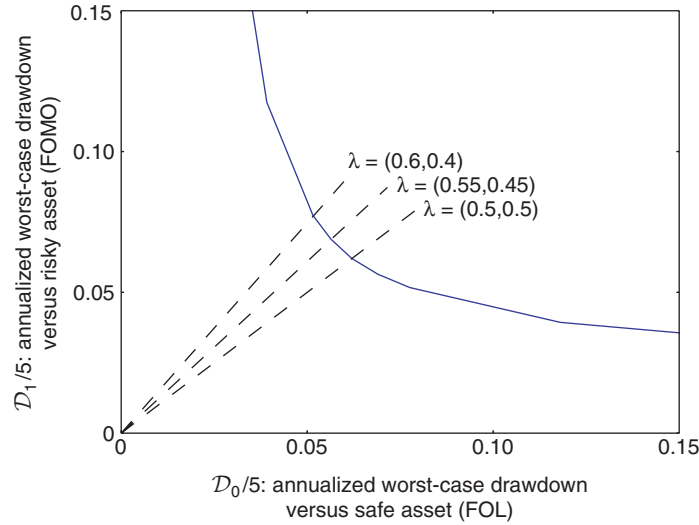
Figure 3 represents the Pareto frontier of minimal drawdowns for moves of nature $M = \{0\} \times \{-0.02, -0.01, 0, 0.01, 0.02\}$ and time horizon $N = 260$ (ie, five years, each period corresponding to a week), computed using the algorithm laid out in Proposition 4.2. Each direction $\lambda \in \{(0.5, 0.5), (0.55, 0.45), (0.6, 0.4)\}$ maps to the prior-free optimal allocation strategy α_λ such that $\lambda_0 \bar{\mathcal{D}}_0^N(\alpha) = \lambda_1 \bar{\mathcal{D}}_1^N(\alpha)$.

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 4.2 remains

FIGURE 3 The drawdown frontier, $\lambda \in \{(0.5, 0.5), (0.55, 0.45), (0.6, 0.4)\}$, $N = 260$ (weeks), $M = \{0\} \times \{-0.02, -0.01, 0, 0.01, 0.02\}$.



FOMO: fear of missing out. FOL: fear of losing.

useful as a benchmark to evaluate the robustness of alternative strategies that are attractive according to other criteria (eg, in sample performance).

To do this, it is necessary to compute the worst-case drawdowns of arbitrary alternative strategies. Corollary 4.3 (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 Z . For each $i \in \{0, 1\}$, define state $x_t^i = (t, a_t, z_t, \mathcal{R}_{t-1}^i)$ and introduce the value function V_α^i over states x^i recursively defined as follows:

$$V_\alpha^i(x_T^i) = \begin{cases} \mathcal{R}_{T-1}^i & \text{if } T = N + 1, \\ \max_{r \in M} V_\alpha(x_{T+1}^i) & \text{if } T \leq N. \end{cases} \quad (4.4)$$

COROLLARY 4.3 For any $i \in \{0, 1\}$, let $x_0^i = (0, a_0, z_0, 0)$. We have that $\bar{\mathcal{D}}_N^i(\alpha) = V_\alpha^i(x_0^i)$.

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 (Glasserman 2003; Golberg 1989).

5 QUALITATIVE PROPERTIES

Proposition 4.2 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 $M = \{0\} \times \{-\bar{r}, 0, \bar{r}\}$.

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

PROPOSITION 5.1 (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 words, 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 consistent with the revealed preference approach, and a 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 the payoff function

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

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

LEMMA 5.2

- (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 equilibriums 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

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

LEMMA 5.3

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

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

$$\begin{aligned} \mathcal{R}^0 - \log(1 - a^{1,*} \bar{r}) &= \mathcal{R}^1 + \log(1 + \bar{r}) - \log(1 + a^{1,*} \bar{r}) \\ \iff a^{1,*} &= \frac{1}{\bar{r}} \times \frac{(1 + \bar{r}) \exp(\mathcal{R}^1 - \mathcal{R}^0) - 1}{(1 + \bar{r}) \exp(\mathcal{R}^1 - \mathcal{R}^0) + 1}. \end{aligned} \quad (5.1)$$

This yields the value function

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

Characterization for $T < N$

Lemma 5.3 extends as follows. Define $\hbar_0 = -\log(1 - \bar{r})$ and $\hbar_1 = \log(1 + \bar{r})$.

PROPOSITION 5.4 *For all $t \in \{0, \dots, N\}$, $\lambda \in \Delta(\{0, 1\})$, $\alpha_\lambda(\mathcal{R}^0, \mathcal{R}^1, t)$ is continuous in $(\mathcal{R}^0, \mathcal{R}^1)$. Further, for all $i \in \{0, 1\}$,*

- (i) *if $\lambda_i \mathcal{R}^i \geq \lambda_{-i}[\mathcal{R}^{-i} + (N - t)\hbar_i]$, then $\alpha_\lambda^i(\mathcal{R}^i, \mathcal{R}^{-i}, t) = 1$;*
- (ii) *if $\lambda_i \mathcal{R}^i - \lambda_{-i} \mathcal{R}^{-i} \in (-(N - t)\hbar_{-i}, (N - t)\hbar_i)$, then $\alpha_\lambda^i(\mathcal{R}^i, \mathcal{R}^{-i}, t) \in (0, 1)$,*
and

$$U_\lambda(z_t, \alpha(z_t), \bar{r}) = U_\lambda(z_t, \alpha(z_t), -\bar{r}). \quad (5.2)$$

Lemma 5.3 and Proposition 5.4 have substantial implications.

5.2 Qualitative implications

Interior allocations

As Black and Litterman (1992) emphasize, naive implementations of the Markowitz (1952) 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 5.4 is that the prior-free approach naturally leads to noncorner solutions, without the help of ad hoc side constraints.

COROLLARY 5.5 *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 that buy recent overperforming stocks, while selling recent underperforming stocks (Asness *et al* 2013; Barberis *et al* 1998; Hong and Stein 1999; Hong *et al* 2000; Jegadeesh and Titman 1993, 2001; Moskowitz *et al* 2012), 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, too, exhibit momentum. This reflects the fact that they attempt to optimize asset allocation in arbitrarily nonstationary 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 another. Formally, the following result holds.

COROLLARY 5.6 *Let $T_1 = \lfloor \rho_1 N \rfloor < \lfloor \rho_2 N \rfloor = T_2$, with ρ_1, ρ_2 fixed.⁴ For any $\epsilon > 0$, consider a probability measure μ over $(r_t^1)_{t \in \{T_1, T_1+1, \dots, T_2\}}$ such that, for all $t \in \{T_1, \dots, T_2\}$,*

$$\mathbb{E}_\mu \left[\frac{r_t^1}{1 + r_t^1} \mid h_t \right] > \epsilon.$$

Then,

$$\mu\text{-almost surely, } \lim_{N \rightarrow +\infty} \frac{1}{T_2 - T_1} \sum_{t=T_1}^{T_2} \alpha_{\lambda, N}^1(h_t) = 1.$$

The condition $\mathbb{E}_\mu[r_t^1/(1 + r_t^1) \mid h_t] > \epsilon$ – 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 5.6 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 5.5, Corollary 5.6 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 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 an investor's personal history.

⁴ As usual, $\lfloor x \rfloor$ denotes the integer part of x , defined as $\lfloor x \rfloor = \max\{n \in \mathbb{N} \mid n \leq x\}$.

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 $\lambda = (0.5, 0.5)$ and the magnitude of potential returns is $\bar{r} = 0.02$. Expression (5.1) implies that the young investor will allocate 49.5% of their wealth to the risky asset. The allocation of the old investor depends on their experience at time $t = 0$. We know by Proposition 5.4 that the latter must have chosen an interior allocation in period $t = 0$. If the risky asset yielded returns $r_0^1 = -\bar{r}$, they experienced drawdowns against the safe asset, and by (5.1) must place a weight strictly less than 49.5% on the risky asset. Inversely, if the risky asset yielded returns $r_0^1 = \bar{r}$, they experienced drawdowns against the risky asset, and by (5.1) must place a weight strictly higher than 49.5% on the risky asset.

This property echoes the finding of Malmendier and Nagel (2011) that investors exhibit heterogeneous 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 is 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 $i \in \mathcal{I} = \{1, \dots, I\}$ 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 $A^i = [\underline{a}^i, \bar{a}^i]$. In addition, total allocation weights must sum to 1, 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 $M^{\mathcal{I}} = \prod_{i \in \mathcal{I}} M^i \subset (-\bar{r}, \bar{r})^{\mathcal{I}}$, with $\bar{r} \in (0, 1)$. Let $r = (r^0, r^{\mathcal{I}})$.

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

$$\max_{a^{\mathcal{I}} \in A} \mathbb{E}[\log(1 + \langle r_t, a \rangle) \mid \mathcal{F}_t], \quad (6.1)$$

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

$$\begin{aligned}\phi^+(a, i)^i &= a^{i,+} \equiv \min\{a^i + \Delta, \bar{a}^i\} \text{ and, for all } j \in \mathcal{I} \setminus i, \phi^+(a, i)^j = a^j, \\ \phi^-(a, i)^i &= a^{i,-} \equiv \max\{a^i - \Delta, \underline{a}^i\} \text{ and, for all } j \in \mathcal{I} \setminus i, \phi^-(a, i)^j = a^j.\end{aligned}$$

By definition, for any $T' \leq T \leq N$, solutions $(a_t^*)_{t \in \{T', \dots, T\}}$ to (6.1) cannot be improved by shifting the allocation in any direction. As a result, for all $i \in \mathcal{I}$, under the investor's prior,

$$\sum_{t=T'}^T \mathbb{E}[\log(1 + \langle r_t, \phi^+(a_t^*, i) \rangle) | \mathcal{F}_t] - \mathbb{E}[\log(1 + \langle r_t, a_t^* \rangle) | \mathcal{F}_t] \leq 0, \quad (6.2)$$

$$\sum_{t=T'}^T \mathbb{E}[\log(1 + \langle r_t, \phi^-(a_t^*, i) \rangle) | \mathcal{F}_t] - \mathbb{E}[\log(1 + \langle r_t, a_t^* \rangle) | \mathcal{F}_t] \leq 0. \quad (6.3)$$

Using finite sample versions of the central limit theorem, this implies that when returns are drawn from the investor's prior, then, with probability approaching 1 for N large, for all $T' \leq T \leq N$,⁵

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

and

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

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

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

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 the optimality conditions (6.2) and (6.3).

⁵ Specifically, the Hoeffding–Azuma inequality. See Cesa-Bianchi and Lugosi (2006, Lemma A.7) for a reference.

Note that the optimality conditions being tested depend on the step of the deviation Δ . Indeed, the drawdowns studied in Sections 2 to 5 correspond to setting $\Delta = 1$, $\underline{a}^i = 0$ and $\bar{a}^i = 1$. In that case, the drawdowns \mathcal{D}_N^0 and \mathcal{D}_N^1 of Sections 2–5 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:

$$\bar{\mathcal{D}}_N^{i,\sigma}(\alpha) = \max_{r \in M^{N+1}} \mathcal{D}_N^{i,\sigma}(\alpha, r) \quad \text{for all } i \in \mathcal{I}, \sigma \in \{+, -\}.$$

Take as given a deviation step $\Delta > 0$, and let Λ be the set of weights $\lambda = (\lambda^+, \lambda^-)$ such that $\lambda^+ + \lambda^- = 1$.⁶

DEFINITION 6.1 (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

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

Computation and key properties

The remainder of this section clarifies difficulties in finding numerical solutions to (P_λ^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 4.1 implies that, for all $\alpha \in \mathcal{A}$,

$$\begin{aligned} \bar{\mathcal{D}}_N^{i,\sigma}(\alpha) &= \bar{\mathcal{R}}_N^{i,\sigma}(\alpha), \\ \mathcal{R}_{T+1}^{i,\sigma} &= \max\{0, \mathcal{R}_T^{i,\sigma} + \log(1 + \langle r_{T+1}, \phi^\sigma(\alpha_{T+1}, i) \rangle) \\ &\quad - \log(1 + \langle r_{T+1}, \alpha_{T+1} \rangle)\}. \end{aligned}$$

Since there are no trading costs, current allocations are not a state variable in (P_λ^I) . An argument identical to that of Proposition 4.2 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 α_t 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 (P_λ^I) becomes numerically intractable as the number I of risky assets becomes large.

⁶ Weights λ^σ could also be indexed on i .

Fortunately, a relaxed problem admits computationally tractable solutions that are approximate solutions to (P_λ^I) . For any i , let

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

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

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

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

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

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

In (P_λ^I) , returns r^{-i} are not freely chosen by nature. They are jointly determined by the allocation a^I and returns r^I . The relaxed problem increases nature's degrees of freedom by allowing it to independently pick $(r^i)_{i \in \mathcal{I}}$ and $(r^{-i})_{i \in \mathcal{I}}$. Denote by \mathbf{r}^i and \mathbf{r}^{-i} sequences $(r_t^i)_{t \in \{0, \dots, N\}}$, $(r_t^{-i})_{t \in \{0, \dots, N\}}$. For any strategy $\alpha^i \in \mathcal{A}^i$, mapping histories to the weight assigned to asset i , let

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

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

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

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

$$W_{\lambda_i}^i(z_T^i) = \begin{cases} \max_{\sigma \in \{+, -\}} \lambda^\sigma \mathcal{R}_T^{i,\sigma} & \text{if } T = N + 1, \\ \min_{a_T^i \in A^i} \max_{r^i \in M^i} \max_{r^{-i} \in M^{-i}} W_{\lambda_i}^i(z_{T+1}^i) & \text{if } T \leq N. \end{cases} \quad (6.4)$$

PROPOSITION 6.2 *The following hold:*

- (i) $\bar{\mathcal{D}}_N^{i,\sigma}(\alpha) \leq \bar{G}_N^{i,\sigma}(\alpha^i)$ for all $\alpha \in \mathcal{A}$;
- (ii) $\lambda^\sigma \bar{G}_N^{i,\sigma}(\alpha_N^{*,i}) = O(\sqrt{N})$ for all $i \in \mathcal{I}$ and $\sigma \in \{+, -\}$;

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

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

In other words, solutions to (\hat{P}_λ^i) are easily computable and provide an approximate solution to (\hat{P}_λ^i) . The fact that drawdowns are sublinear in N implies the following extension of Corollary 5.6.

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

$$\sigma_i \mathbb{E}_\mu \left[\frac{r^i - r^0}{1 + \langle a, r \rangle} \right] > 0 \quad \text{for all } a^i \in A, \quad (6.5)$$

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

COROLLARY 6.3 *If (6.5) holds, then prior-free optimal strategy $\alpha_N^{*,i}$ satisfies, for all $i \in \mathcal{I}$,*

$$\mu\text{-almost surely,} \quad \lim_{N \rightarrow +\infty} \frac{1}{T_2 - T_1} \sum_{t=T_1}^{T_2} \alpha_N^{*,i}(h_t) = a^{*,i}.$$

In other words, the prior-free optimal strategy must approach the Bayesian optimal allocation when it takes extreme values. It is worth noting that Corollary 6.3 holds regardless of the history of returns occurring before time T_1 . Even after long histories, prior-free optimal allocation strategies do not become doctrinaire. They adapt to new circumstances.

Corollary 6.3 also clarifies that although (P_λ^i) and (\hat{P}_λ^i) let nature pick returns for 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 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.⁷ CPPI is an especially relevant benchmark since its goal is to provide prior-free performance guarantees. The version of CPPI tested in this section takes the following form: the investor tracks their counterfactual wealth \hat{w}_t as if they invested only in the safe asset; a share of their actual wealth w_t equal to 75% of their counterfactual wealth \hat{w}_t is invested in the safe asset. The remaining cushion $w_t - 0.75\hat{w}_t$ is leveraged once and invested in the risky asset. If the price process is continuous and rebalancings occur frequently enough, CPPI guarantees the investor 75% of their wealth if they invested in the safe asset, while also providing exposure to the risky asset.

7.1 Worst-case performance

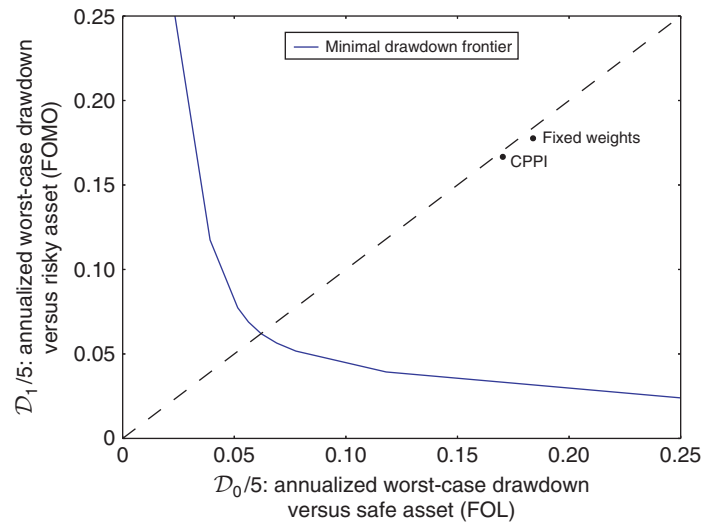
Figure 4 plots the worst-case drawdowns of both the fixed-weight portfolio and the CPPI portfolio against the prior-free efficient frontier in the case where a period corresponds to a week, $N = 260$, and $M = \{0\} \times \{-0.02, -0.01, 0, 0.01, 0.02\}$ and the trading cost c is equal to 20bps.

Mechanically, both the fixed-weight portfolio and CPPI must sit to the north-east of the efficient frontier. In fact, they sit quite far away from the efficient frontier.

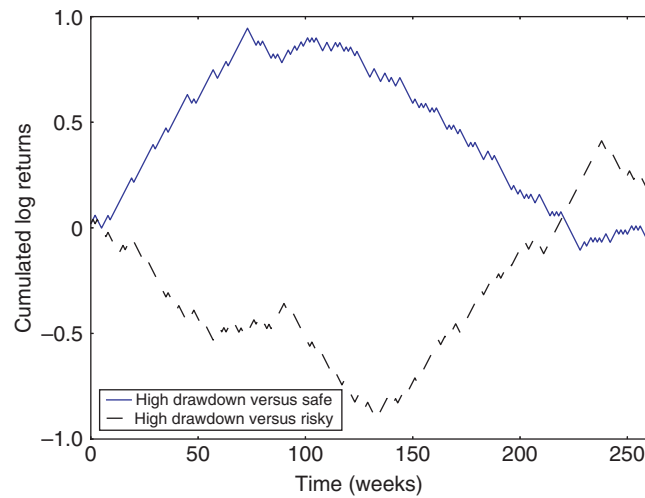
The reason why the fixed-weight portfolio can experience large drawdowns is clear. It keeps the allocation close to 50/50, even if the risky asset keeps yielding positive (or negative) returns.

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.

⁷ Here, the $1/n$ strategy serves as the simplest possible benchmark, similar to a popular 60/40 allocation. The robustness of the $1/n$ approach, as documented by DeMiguel *et al* (2009), is particularly important when the number of assets becomes large, since in that case the correlation matrix between assets becomes near singular. The current paper makes no empirical claim regarding such many-assets allocation problems.

FIGURE 4 Worst-case drawdowns for prior-free efficient, fixed-weight and CPPI strategies.

FOMO: fear of missing out. FOL: fear of losing.

FIGURE 5 Returns generating high drawdowns for the CPPI strategy.

7.2 Historical performance

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

Table 1 reports findings from running the fixed-weight, CPPI and prior-free optimal strategies defined above on a sample of market and risk-free returns from January 1, 1927 to December 31, 2014, obtained from Kenneth French's data library. Each strategy is implemented over rolling periods of five years, starting on January 1 for each of the eighty-eight years in the sample. Trading costs are set to 20bps.

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

- Net annualized performance: net performance = $252 \times \hat{\mathbb{E}}[r_t^\alpha - r_t^0]$.

- Annualized Sharpe ratios:

$$\text{Sharpe} \equiv \frac{\hat{\mathbb{E}}[r^\alpha - r^0]}{\sqrt{\hat{\mathbb{E}}[(r^\alpha - r^0)^2]}} \sqrt{252}.$$

- Worst-case five-year relative drawdowns $\mathcal{D}_0, \mathcal{D}_1$ versus the safe and risky asset over the entire period.
- Net performance-to-drawdown ratio:

$$\text{net performance-to-drawdown} \equiv \frac{\text{net performance}}{\mathcal{D}_0}. \quad (7.1)$$

- Parameters α and β from capital asset pricing model regression:

$$r^\alpha - r^0 \sim \alpha + \beta(r^1 - r^0) + \varepsilon, \quad (7.2)$$

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

The net performance-to-drawdown ratio defined in (7.1) summarizes each strategy's ability to capture upside while reducing drawdowns.

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

Figure 6 reports the cumulative log returns of being long the prior-free portfolio and short the fixed-weight portfolio. The dotted lines separate the sample period

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

	Fixed weights	CPPI	Prior free
Net performance	3.7%	4.7%	6.2%
Sharpe	0.45	0.48	0.56
\mathcal{D}_0	0.54	0.37	0.29
\mathcal{D}_1	0.44	0.53	0.37
Net performance to drawdown	0.07	0.12	0.21
α	0.000 (0.001)	0.004 (0.005)	0.014*** (0.006)
β	0.49*** (0.003)	0.56*** (0.024)	0.62*** (0.027)

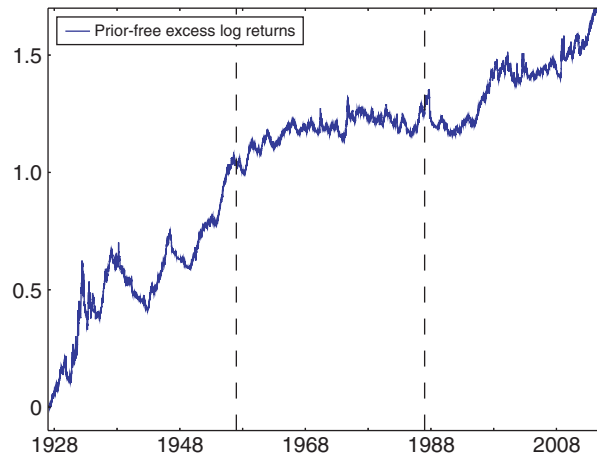
*, ** and *** respectively denote effects significant at the 0.1, 0.05 and 0.01 levels; standard errors are given in parentheses.

into periods of about thirty years: 1927–57, 1957–87 and 1987–2014. The prior-free optimal portfolio overperforms in each subsample, particularly in the 1927–57 sample, where large swings in returns make drawdown control especially valuable. The long–short strategy’s Sharpe ratio over these three subsamples is, respectively, 0.54 (1927–57), 0.20 (1957–87) and 0.30 (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 $[-0.05, 0.05]$ 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 in case a bull or bear market should emerge. However, if the market remains range bound, this results in unnecessary transaction costs.

Appendix A (available online) reports further empirical findings. First, a Henriksen 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 (0.7 versus 0.4). 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 the drawdown-control problem (P_λ).

FIGURE 6 Excess log returns of prior-free optimal portfolio over fixed-weight strategy, 1927–2014.



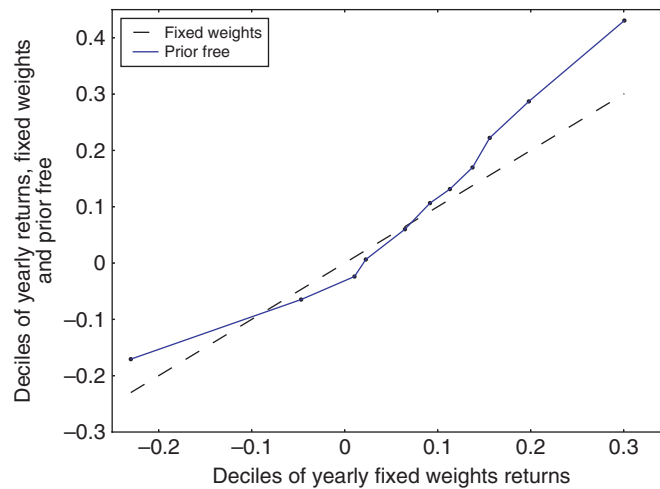
8 CONCLUSION

This paper provides a prior-free framework for asset allocation in arbitrarily nonstationary 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 drawdowns against the risky asset). Prior-free optimal allocation strategies are amenable to numerical computation, are largely interior and satisfy a form of momentum. Finally, they are history dependent.

Practically, prior-free optimal strategies offer worst-case drawdown guarantees that are a significant improvement on 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.

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

FIGURE 7 Q–Q plot of yearly returns for the prior-free optimal and fixed-weight strategies, 1927–2014.



DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

I am particularly indebted to Markus Brunnermeier and Harrison Hong for guidance and encouragement. This 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 and 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.

REFERENCES

- Antic, N. (2014). Contracting with unknown technologies. Unpublished Paper, Princeton University.
- Asl, F. M., and Etula, E. (2012). Advancing strategic asset allocation in a multi-factor world. *Journal of Portfolio Management* **39**, 59–66 (<https://doi.org/10.3905/jpm.2012.39.1.059>).
- Asness, C. S., Moskowitz, T. J., and Pedersen, L. H. (2013). Value and momentum everywhere. *Journal of Finance* **68**, 929–985 (<https://doi.org/10.1111/jofi.12021>).

- Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics* **49**, 307–343 ([https://doi.org/10.1016/S0304-405X\(98\)00027-0](https://doi.org/10.1016/S0304-405X(98)00027-0)).
- Baum, L. E., and Petrie, T. (1966). Statistical inference for probabilistic functions of finite state Markov chains. *Annals of Mathematical Statistics* **37**, 1554–1563 (<https://doi.org/10.1214/aoms/1177699147>).
- Bergemann, D., and Schlag, K. H. (2008). Pricing without priors. *Journal of the European Economic Association* **6**, 560–569 (<https://doi.org/10.1162/JEEA.2008.6.2-3.560>).
- Black, F., and Jones, R. W. (1987). Simplifying portfolio insurance. *Journal of Portfolio Management* **14**, 48–51 (<https://doi.org/10.3905/jpm.1987.409131>).
- Black, F., and Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal* **48**, 28–43 (<https://doi.org/10.2469/faj.v48.n5.28>).
- Black, F., and Perold, A. (1992). Theory of constant proportion portfolio insurance. *Journal of Economic Dynamics and Control* **16**, 403–426 ([https://doi.org/10.1016/0165-1889\(92\)90043-E](https://doi.org/10.1016/0165-1889(92)90043-E)).
- Blackwell, D. (1956). An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics* **6**, 1–8 (<https://doi.org/10.2140/pjm.1956.6.1>).
- Brooks, B. (2014). Surveying and selling: belief and surplus extraction in auctions. Technical Report, Princeton University.
- Cai, X., Teo, K.-L., Yang, X., and Zhou, X. Y. (2000). Portfolio optimization under a minimax rule. *Management Science* **46**, 957–972 (<https://doi.org/10.1287/mnsc.46.7.957.12039>).
- Campbell, J. Y. (1984). Asset duration and time varying risk premia. Unpublished Dissertation, Yale University.
- Campbell, J. Y., and Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics* **114**, 433–495 (<https://doi.org/10.1162/003355399556043>).
- Carroll, G. (2014). Robustness and linear contracts. *American Economic Review* **105**, 536–563 (<https://doi.org/10.1257/aer.20131159>).
- Ceria, S., and Stubbs, R. A. (2006). Incorporating estimation errors into portfolio selection: robust portfolio construction. *Journal of Asset Management* **7**, 109–127 (<https://doi.org/10.1057/palgrave.jam.2240207>).
- Cesa-Bianchi, N., and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press (<https://doi.org/10.1017/CBO9780511546921>).
- Chassang, S. (2013). Calibrated incentive contracts. *Econometrica* **81**, 1935–1971 (<https://doi.org/10.3982/ECTA9987>).
- Cover, T. M. (1991). Universal portfolios. *Mathematical Finance* **1**, 1–29 (<https://doi.org/10.1111/j.1467-9965.1991.tb00002.x>).
- DeMarzo, P., Kremer, I., and Mansour, Y. (2009). Hannan and Blackwell meet Black and Scholes: approachability and robust option pricing. Technical Report, Stanford University.
- DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: how inefficient is the $1/N$ portfolio strategy? *Review of Financial Studies* **22**, 1915–1953 (<https://doi.org/10.1093/rfs/hhm075>).
- Diaconis, P., and Freedman, D. (1986). On the consistency of Bayes estimates. *Annals of Statistics* **14**, 1–26 (<https://doi.org/10.1214/aos/1176349830>).
- Foster, D. P., and Vohra, R. (1999). Regret in the on-line decision problem. *Games and Economic Behavior* **29**, 7–35 (<https://doi.org/10.1006/game.1999.0740>).

- Ghosh, J. K., and Ramamoorthi, R. (2003). *Bayesian Nonparametrics*, Volume 1. Springer.
- Gilboa, I., and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* **18**, 141–153 ([https://doi.org/10.1016/0304-4068\(89\)90018-9](https://doi.org/10.1016/0304-4068(89)90018-9)).
- Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*, Volume 53. Springer Science & Business Media (<https://doi.org/10.1007/978-0-387-21617-1>).
- Glasserman, P., and Xu, X. (2013). Robust portfolio control with stochastic factor dynamics. *Operations Research* **61**, 874–893 (<https://doi.org/10.1287/opre.2013.1180>).
- Glasserman, P., and Xu, X. (2014). Robust risk measurement and model risk. *Quantitative Finance* **14**, 29–58 (<https://doi.org/10.1080/14697688.2013.822989>).
- Golberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley.
- Grossman, S. J., and Zhou, Z. (1993). Optimal investment strategies for controlling draw-downs. *Mathematical Finance* **3**, 241–276 (<https://doi.org/10.1111/j.1467-9965.1993.tb00044.x>).
- Hannan, J. (1957). Approximation to Bayes risk in repeated play. In *Contributions to the Theory of Games*, Dresher, M., Tucker, A., and Wolfe, P. (eds), Volume 3, pp. 97–139. Princeton University Press.
- Hansen, L. P., and Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review* **91**, 60–66 (<https://doi.org/10.1257/aer.91.2.60>).
- Hansen, L. P., and Sargent, T. J. (2008). *Robustness*. Princeton University Press (<https://doi.org/10.1515/9781400829385>).
- Hartline, J. D., and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *Symposium on Theory of Computing (STOC)*, pp. 75–84. ACM Special Interest Group on Algorithms and Computation Theory (<https://doi.org/10.1145/1374376.1374390>).
- Henriksson, R. D., and Merton, R. C. (1981). On market timing and investment performance II: statistical procedures for evaluating forecasting skills. *Journal of Business* **54**, 513–533 (<https://doi.org/10.1086/296144>).
- Hong, H., and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance* **54**, 2143–2184 (<https://doi.org/10.1111/0022-1082.00184>).
- Hong, H., Lim, T., and Stein, J. C. (2000). Bad news travels slowly: size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance* **55**, 265–295 (<https://doi.org/10.1111/0022-1082.00206>).
- Jegadeesh, N., and Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance* **48**, 65–91 (<https://doi.org/10.1111/j.1540-6261.1993.tb04702.x>).
- Jegadeesh, N., and Titman, S. (2001). Profitability of momentum strategies: an evaluation of alternative explanations. *Journal of Finance* **56**, 699–720 (<https://doi.org/10.1111/0022-1082.00342>).
- Kőszegi, B., and Rabin, M. (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics* **121**, 1133–1165 (<https://doi.org/10.1093/qje/121.4.1133>).
- Lettau, M., and Ludvigson, S. C. (2001). Resurrecting the (C) CAPM: a cross-sectional test when risk premia are time-varying. *Journal of Political Economy* **109**, 1238–1287 (<https://doi.org/10.1086/323282>).

- Lettau, M., and Ludvigson, S. C. (2009). Measuring and modeling variation in the risk-return trade-off. In *The Handbook of Financial Econometrics*, Volume 1, Chapter 11, pp. 617–690. Elsevier.
- Lo, A. W., and MacKinlay, A. C. (1990). Data-snooping biases in tests of financial asset pricing models. *Review of Financial Studies* **3**, 431–467 (<https://doi.org/10.1093/rfs/3.3.431>).
- Lou, D., and Polk, C. (2013). Comomentum: inferring arbitrage activity from return correlations. Working Paper, London School of Economics.
- Luenberger, D. G. (1968). *Optimization by Vector Space Methods*. Wiley.
- Madarász, K., and Prat, A. (2016). Screening with an approximate type space. *Review of Economic Studies* **84**(2), 790–815.
- Malmendier, U., and Nagel, S. (2011). Depression babies: do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics* **126**, 373–416 (<https://doi.org/10.1093/qje/qjq004>).
- Markowitz, H. (1952). Portfolio selection. *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* **54**, 363–406 (<https://doi.org/10.1086/296137>).
- Milnor, J. (1954). Games against nature. Working Paper, Rand Corporation.
- Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics* **104**, 228–250 (<https://doi.org/10.1016/j.jfineco.2011.11.003>).
- Novy-Marx, R. (2014). Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars. *Journal of Financial Economics* **112**, 137–146 (<https://doi.org/10.1016/j.jfineco.2014.02.002>).
- Perold, A. (1986). Constant proportion portfolio insurance. Unpublished Paper, Harvard Business School.
- Pflug, G., and Wozabal, D. (2007). Ambiguity in portfolio selection. *Quantitative Finance* **7**, 435–442 (<https://doi.org/10.1080/14697680701455410>).
- Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association* **46**, 55–67 (<https://doi.org/10.1080/01621459.1951.10500768>).
- Savage, L. J. (1972). *The Foundations of Statistics*. Courier Corporation.
- Segal, I. (2003). Optimal pricing mechanisms with unknown demand. *American Economic Review* **93**, 509–529 (<https://doi.org/10.1257/000282803322156963>).
- Sims, C. A. (1971). Distributed lag estimation when the parameter space is explicitly infinite-dimensional. *Annals of Mathematical Statistics* **42**, 1622–1636 (<https://doi.org/10.1214/aoms/1177693161>).
- Stein, J. C. (2009). Presidential address: sophisticated investors and market efficiency. *Journal of Finance* **64**, 1517–1548 (<https://doi.org/10.1111/j.1540-6261.2009.01472.x>).
- Stoye, J. (2008). Axioms for minimax regret choice correspondences. Unpublished Paper, New York University.
- Tversky, A., and Kahneman, D. (1991). Loss aversion in riskless choice: a reference-dependent model. *Quarterly Journal of Economics* **106**, 1039–1061 (<https://doi.org/10.2307/2937956>).
- Wald, A. (1950). Statistical decision functions. Working Paper, Wiley.
- Zweig, J. (2007). *Your Money and Your Brain: How the New Science of Neuroeconomics Can Help Make You Rich*. Simon & Schuster.