Online Appendix:

Smart Stochastic Discount Factors

SOFONIAS ALEMU KORSAYE, ALBERTO QUAINI and FABIO TROJANI*

This Online Appendix is organized as follows. Section A contains the proofs for all mathematical statements presented in the main text. Section B extends the theoretical framework developed in the main text by discussing: (a) the foundation of S–SDFs in viable markets with convex costs; (b) S–SDFs in markets with ambiguity, where the probability measure used to describe asset payoffs differs from the unobserved marginal investor's beliefs; (c) the link between S–SDFs and asset pricing theories that minimize pricing errors; (d) the role of conditional pricing restrictions and their implications. Section C presents additional empirical evidence on minimum dispersion APT S–SDFs. Finally, Section D collects the figures illustrating the empirical findings from Section C.

*Sofonias Alemu Korsaye (email: Sofonias.Korsaye@jhu.edu) is with Carey Business School, Johns Hopkins University; Alberto Quaini (email: Quaini@ese.eur.nl) is with Erasmus School of Economics, Erasmus University Rotterdam; Fabio Trojani (email: Fabio.Trojani@unige.ch) is with Geneva Finance Research Institute & Swiss Finance Institute, University of Geneva.

A Proofs

Proof of Proposition 1. By [Rockafellar, 1970, Thm. 13.2],

$$\sigma_h(\boldsymbol{\theta}) := \sup_{\boldsymbol{\eta}\in\mathbb{R}^N} \{\boldsymbol{\theta}'\boldsymbol{\eta} : h(\boldsymbol{\eta}) \leq \tau\}$$

is proper, closed and sublinear.[1] This implies that the pricing functional π is sublinear and that the set of payoffs, $\mathcal{Z}$ is a convex cone. Hence, by [Chen, 2001, Thm. 1, 5 and Cor. 1], absence of free lunches is equivalent to the existence of an almost surely strictly positive SDF, M, such that $\mathbb{E}[MZ] \leq \pi(Z)$ for any $Z \in \mathcal{Z}$.[2]

By definition of cost functional π and payoff space $\mathcal{Z}$, it equivalently follows for any $\boldsymbol{\theta} \in \mathbb{R}^N$ that $\boldsymbol{\theta}'\boldsymbol{P} + \sigma_h(\boldsymbol{\theta}) \geq \mathbb{E}[M\boldsymbol{\theta}'\boldsymbol{X}]$, i.e., $\boldsymbol{\theta}'\left(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}\right) \leq \sigma_h(\boldsymbol{\theta})$. By [Bauschke and Combettes, 2011, Prop. 13.10 (i)] this inequality holds if and only if

$$(\sigma_h)^*(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \leq 0 \; . \tag{OA-1}$$

By [Rockafellar, 1970, Thm. 13.2], the convex conjugate $(\sigma_h)^*$ is given by the characteristic function δ_h, i.e.,

$$(\sigma_h)^*(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) = \begin{cases} 0 & h(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \leq \tau \\ +\infty & \text{else} \end{cases} .$$

Thus, inequality (OA-1) is equivalent to condition $h(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \leq \tau$. □

Proof of Proposition 2. We first prove that $\Pi(h, \tau) = -\Delta(h, \tau)$. To do so, we rewrite

[1] These properties follow from the fact that σ_h can be expressed as the support function of set $C(h, \tau) := \{\boldsymbol{\eta} \in \mathbb{R}^N : h(\boldsymbol{\eta}) \leq \tau\}$, which is not empty by assumption.

[2] Technically, in order to use [Chen, 2001, Thm. 5], we require that there exists a traded payoff that is strictly positive almost surely. Such assumption is easily satisfied, e.g., when there exists a risk-free asset with positive risk-free payoff.

problem $\Pi(h,\tau)$ according to the notation given in Bauschke et al. [2017] and obtain its Fenchel-Rockafellar's dual problem. Let $A : L^q \to \mathbb{R}^N$ be the linear function $A(M) := \mathbb{E}[M\boldsymbol{X}]$, and let $g : \mathbb{R}^N \to (-\infty,+\infty]$ denote the function $g(\boldsymbol{\beta}) := \delta_R(\boldsymbol{\beta})$ with $R := \{\boldsymbol{v} \in \mathbb{R}^N : h(\boldsymbol{v} - \boldsymbol{P}) \le \tau\}$. With the notation $I_f := \mathbb{E}[f(\cdot)]$ for any function $f : \mathbb{R} \to [-\infty,+\infty]$, it follows that $I_{\phi_+}(M) = I_\phi(M) + \delta_{L^q_+}(M)$.[3] Thus we can write:

$$\Pi(h,\tau) = \inf_{M\in L^q} \{I_{\phi_+}(M) + g(A(M))\} \; .$$

As payoffs are in L^p with $1/p + 1/q = 1$, A is continuous, while the properties of ϕ_+ imply that I_{ϕ_+} is a closed convex function. Since $C(h,\tau) := \{\boldsymbol{\eta} \in \mathbb{R}^N : h(\boldsymbol{\eta}) \le \tau\}$ is convex, closed and non-empty, g is convex, closed and proper. In view of these properties, we can obtain the dual problem of $\Pi(h,\tau)$ via [Bauschke et al., 2017, Thm. 15.23]. In order to apply the mentioned theorem, [Bauschke et al., 2017, Prop. 15.24] shows that we only need to check that:[4]

$$\mathrm{ri}(\mathrm{dom}(g)) \cap A[\mathrm{qri}\;\; \mathrm{dom}(I_{\phi_+})] \neq \emptyset \; . \qquad \text{(OA-2)}$$

As shown in [Borwein and Lewis, 1991, Cor. 2.6], our requirement $(0,+\infty) \subset \mathrm{dom}\,\phi$ implies that $A(\mathrm{dom}\, I_\phi \cap L^q_{++}) \subset A[\mathrm{qri}(\mathrm{dom}\, I_{\phi_+})]$. Hence, since $\mathrm{dom}(g) = R$, showing that $A(\mathrm{dom}\, I_\phi \cap L^q_{++}) \cap \mathrm{ri}(R) \neq \emptyset$ is enough to prove (OA-2). Under the assumption that the economy $(\mathcal{Z}_{\sigma_{\tilde{h}}}, \pi_{\sigma_{\tilde{h}}}, \mathbb{P})$ admits no free lunches, with $C(\tilde{h},\tau) \subset \mathrm{ri}(C(h,\tau))$, by Proposition 1 in the main text there exists $\tilde{M} \in \mathrm{dom}\, I_\phi \cap L^q_{++}$, hence $A(\tilde{M}) \in A(\mathrm{dom}\, I_\phi \cap L^q_{++})$. Moreover, $\mathbb{E}[\tilde{M}\boldsymbol{X}] - \boldsymbol{P} \in C(\tilde{h},\tau)$, i.e., $A(\tilde{M}) - \boldsymbol{P} \in C(\tilde{h},\tau) \subset \mathrm{ri}(R)$ holds and condition (OA-2) is satisfied.

Thus, by [Bauschke et al., 2017, Thm. 15.23 and 15.24] we obtain $\Pi(h,\tau) = -\Delta(h,\tau)$

[3] Remember that ϕ_+ is the restriction of ϕ to $\mathbb{R}_+$, i.e., $\phi_+(x) = \phi(x)$ if $x \ge 0$ and $\phi_+(x) = +\infty$ if $x < 0$.

[4] See [Bauschke et al., 2017, Def. 6.9] for the definition of relative interior, ri, and quasi relative interior, qri.

where

$$\Delta(h,\tau) = -\min_{\boldsymbol{\theta}\in\mathbb{R}^N} \{I^*_{\phi_+}(-{}^tA(\boldsymbol{\theta})) + g^*(\boldsymbol{\theta})\} , \tag{OA-3}$$

where $I^*_{\phi_+} : L^p \to [0,+\infty]$ is the conjugate function of I_{ϕ_+} and ${}^tA : \mathbb{R}^N \to L^p$ is the adjoint map of A, given by ${}^tA(\boldsymbol{\theta}) = \boldsymbol{X}'\boldsymbol{\theta}$.[5] As ϕ_+ is convex and closed, we can apply [Rockafellar, 1968, Thm. 2] to obtain $I^*_{\phi_+} = I_{\phi^*_+}$. Moreover, for every $\boldsymbol{\theta} \in \mathbb{R}^N$,

$$\begin{aligned} g^*(-\boldsymbol{\theta}) &= \delta^*_R(-\boldsymbol{\theta}) \\ &= \sup_{\boldsymbol{v}\in\mathbb{R}^N} \left\{-\boldsymbol{\theta}'\boldsymbol{v} \ : \ h(\boldsymbol{v}-\boldsymbol{P}) \le \tau\right\} \\ &= -\boldsymbol{P}'\boldsymbol{\theta} + \sigma_h(-\boldsymbol{\theta}) \ . \end{aligned}$$

Thus (OA-3) reads:

$$\Delta(h,\tau) = -\min_{\boldsymbol{\theta}\in\mathbb{R}^N} \left\{\mathbb{E}[\phi^*_+(-\boldsymbol{X}'\boldsymbol{\theta})] + \boldsymbol{P}'\boldsymbol{\theta} + \sigma_h(\boldsymbol{\theta})\right\} \ ,$$

thereby proving strong duality between $\Pi(h,\tau)$ and $-\Delta(h,\tau)$.

We now prove the relation between the optimal solutions of $\Pi(h,\tau)$ and $\Delta(h,\tau)$ given in Proposition 2 of the main text. By [Bauschke et al., 2017, Prop. 19.4], if $\boldsymbol{\theta}_0$ is a dual solution such that $I_{\phi^*_+}$ is Gateaux differentiable in $-\boldsymbol{X}'\boldsymbol{\theta}_0$, then the primal problem either has no solution or it has a unique solution given by $\nabla I_{\phi^*_+}(-\boldsymbol{X}'\boldsymbol{\theta}_0)$. Hereafter we show that indeed $I_{\phi^*_+}$ is Gateaux differentiable and that the derivative is given by $M_0 = (\phi^*_+)'(-\boldsymbol{X}'\boldsymbol{\theta}_0)$.

The assumption that $-\boldsymbol{X}'\boldsymbol{\theta}_0 < \lim_{y\to\infty}\phi(y)/y$ a.s., by [Borwein and Lewis, 1991, Lem. 4.2], implies that $-\boldsymbol{X}'\boldsymbol{\theta}_0$ belongs a.s. to the interior of the domain of ϕ^*_+. This and the strict convexity of ϕ_+ on its domain imply that ϕ^*_+ is differentiable a.s. in $-\boldsymbol{X}'\boldsymbol{\theta}_0$, see [Borwein and Lewis, 1991, Thm. 4.6]. Let us denote such derivative by $M_0(\omega) := (\phi^*_+)'(-\boldsymbol{X}'(\omega)\boldsymbol{\theta}_0)$.

[5] The adjoint map of A, ${}^tA$, is characterized by the identity $\mathbb{E}[{}^tA(\boldsymbol{\theta})M] = \boldsymbol{\theta}'\mathbb{E}[M\boldsymbol{X}]$, for each $M \in L^q$ and each portfolio weights $\boldsymbol{\theta} \in \mathbb{R}^N$.

Hereafter in order to prove that $I_{\phi_+^*}$ is Gateaux differentiable in $-\boldsymbol{X}'\boldsymbol{\theta}_0$, we show that M_0 is the unique element of $\partial I_{\phi_+^*}(-\boldsymbol{X}'\boldsymbol{\theta}_0)$. Now let us consider any element $\bar{M} \in \partial I_{\phi_+^*}(-\boldsymbol{X}'\boldsymbol{\theta}_0)$, then by [Rockafellar, 1970, Thm. 23.5]:

$$I_{\phi_+^*}(-\boldsymbol{X}'\boldsymbol{\theta}_0) + I_{\phi_+}(\bar{M}) = \langle -\boldsymbol{X}'\boldsymbol{\theta}_0, \bar{M}\rangle \ .$$

Explicitly, this gives $\int \left[\phi_+^*(-\boldsymbol{X}'(\omega)\boldsymbol{\theta}_0) + \phi_+(\bar{M}(\omega)) + \boldsymbol{X}'(\omega)\boldsymbol{\theta}_0\bar{M}(\omega)\right] d\mathbb{P}(\omega) = 0$. Here, the integrand, since by Fenchel's inequality is nonnegative, is zero a.s.. Applying again [Rockafellar, 1970, Thm. 23.5] to this integrand, we obtain $\bar{M}(\omega) \in \partial\phi_+^*({}^tA(-\boldsymbol{\theta}_0)(\omega))$ a.s.. However, since ϕ_+^* is differentiable a.s. in $-\boldsymbol{X}'\boldsymbol{\theta}_0$, $\partial\phi_+^*({}^tA(-\boldsymbol{\theta}_0)(\omega))$ has a unique element. Hence, we have that $\bar{M} = M_0$ a.s., which proves that $\partial I_{\phi_+^*}(-\boldsymbol{X}'\boldsymbol{\theta}_0)$ has a unique element and that $I_{\phi_+^*}$ is Gateaux differentiable in $-\boldsymbol{X}'\boldsymbol{\theta}_0$ with derivative given by $(\phi_+^*)'(-\boldsymbol{X}'\boldsymbol{\theta}_0)$. This concludes the proof. $\square$

Proof of Proposition 3. (i) Suppose that M_0 solves $\Pi(h,\tau)$, and let $\nu(\tau) := \mathbb{E}[\phi(M_0)]$. By strict convexity of ϕ and [Borwein and Lewis, 1991, Prop. 2.11], M_0 is the unique solution of $\Pi(h,\tau)$. Therefore, M_0 is the unique element of

$$\{0 \leq M \in L^q \ : \ \mathbb{E}[\phi(M)] \leq \nu(\tau),\ h(\mathbb{E}[M\boldsymbol{X} - \boldsymbol{P}]) \leq \tau\}.$$

Thus, M_0 solves $\mathcal{P}(\nu(\tau))$. (ii) Suppose that M_0 solves $\mathcal{P}(h,\nu)$, and let $\tau(\nu) := h(\mathbb{E}[M_0\boldsymbol{X} - \boldsymbol{P}])$. By strict convexity of ϕ, for all $0 \leq M \in L^q$ such that $h(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \leq \tau(\nu)$ we have $\mathbb{E}[\phi(M)] \geq \nu$. If follows that M_0 solves $\Pi(h,\tau(\nu))$. $\square$

Define the APT norms h_1 and h_∞ for every $\boldsymbol{\eta} \in \mathbb{R}^N$ and for $\lambda \in [0,1]$ by:

$$h_{1,\boldsymbol{\Sigma}_{\boldsymbol{\zeta}}^{-1/2}}(\boldsymbol{\eta}) := (1-\lambda)\left\|\boldsymbol{\eta}\right\|_{1,\boldsymbol{\Sigma}_{\boldsymbol{\zeta}}^{-1/2}} + \lambda\left\|\boldsymbol{\eta}\right\|_{2,\boldsymbol{\Sigma}_{\boldsymbol{\zeta}}^{-1/2}} \ ,$$

and

$$h_{\infty,\boldsymbol{\Sigma}_{\zeta}^{-1/2}}(\boldsymbol{\eta}) := (1-\lambda)\sqrt{N}\left\|\boldsymbol{\eta}\right\|_{\infty,\boldsymbol{\Sigma}_{\zeta}^{-1/2}} + \lambda\left\|\boldsymbol{\eta}\right\|_{2,\boldsymbol{\Sigma}_{\zeta}^{-1/2}} .$$

Lemma OA-1. *Penalty function $\sigma_{h_i}(\boldsymbol{\theta}) = \tau h_{i*,\boldsymbol{\Sigma}_{\zeta}^{1/2}}(\boldsymbol{\theta})$ $(i = 1, \infty)$ is given in closed-from as follows.*

(i) If $\lambda = 1$: $h_{1*}(\boldsymbol{\theta}) = h_{\infty*}(\boldsymbol{\theta}) = \left\|\boldsymbol{\theta}\right\|_2$.

(ii) If $\lambda = 0$: $h_{1*}(\boldsymbol{\theta}) = \left\|\boldsymbol{\theta}\right\|_{\infty}$ and $h_{\infty*}(\boldsymbol{\theta}) = \left\|\boldsymbol{\theta}\right\|_1/\sqrt{N}$.

(iii) If $\lambda \in (0,1)$:

$$h_{1*}(\boldsymbol{\theta}) = \min_{\boldsymbol{z}\in\mathbb{R}^N}\left\{\max\left\{\frac{\left\|\boldsymbol{z}\right\|_{\infty}}{1-\lambda}, \frac{\left\|\boldsymbol{\theta}-\boldsymbol{z}\right\|_2}{\lambda}\right\}\right\} ,$$

$$h_{\infty*}(\boldsymbol{\theta}) = \min_{\boldsymbol{z}\in\mathbb{R}^N}\left\{\max\left\{\frac{\left\|\boldsymbol{z}\right\|_1}{\sqrt{N}(1-\lambda)}, \frac{\left\|\boldsymbol{\theta}-\boldsymbol{z}\right\|_2}{\lambda}\right\}\right\} .$$

Proof of Lemma OA-1. Since the cases $\lambda = 0, 1$ are obvious, let $\lambda \in (0,1)$. Denoting by $\left\|\cdot\right\|_A$, $\left\|\cdot\right\|_B$ two norms, following identity holds for any $\boldsymbol{\eta} \in \mathbb{R}^N$:

$$\left\|\boldsymbol{\eta}\right\| := \lambda\left\|\boldsymbol{\eta}\right\|_A + (1-\lambda)\left\|\boldsymbol{\eta}\right\|_B = \left\|\begin{pmatrix} \lambda\left\|\boldsymbol{\eta}\right\|_A \\ (1-\lambda)\left\|\boldsymbol{\eta}\right\|_B \end{pmatrix}\right\|_1 .$$

From [Combettes et al., 2019, Thm. 2.5], we obtain for any $\boldsymbol{\theta} \in \mathbb{R}^N$:

$$\begin{aligned}\left\|\boldsymbol{\theta}\right\|_* &= \inf_{(\boldsymbol{\theta}_1,\boldsymbol{\theta}_2)\in\mathbb{R}^{2N}}\left\{\left\|\begin{pmatrix} \frac{1}{\lambda}\left\|\boldsymbol{\theta}_1\right\|_{A*} \\ \frac{1}{1-\lambda}\left\|\boldsymbol{\theta}_2\right\|_{B*} \end{pmatrix}\right\|_{1*} : \boldsymbol{\theta}_1 + \boldsymbol{\theta}_2 = \boldsymbol{\theta}\right\} \\ &= \min_{\boldsymbol{z}\in\mathbb{R}^N}\left\{\max\left\{\frac{1}{\lambda}\left\|\boldsymbol{z}\right\|_{A*}, \frac{1}{1-\lambda}\left\|\boldsymbol{\theta}-\boldsymbol{z}\right\|_{B*}\right\}\right\} ,\end{aligned}$$

since $\left\|\cdot\right\|_{1*} = \left\|\cdot\right\|_{\infty}$. The statement of the lemma finally follows by further recalling that $\left\|\cdot\right\|_{2*} = \left\|\cdot\right\|_2$ and $\left\|\cdot\right\|_{\infty*} = \left\|\cdot\right\|_1$. □

B Theory of S–SDFs and applications

B.1 Foundation of S–SDFs in markets with convex costs

We have shown in Section 3 of the main text that S–SDFs are characterized by the absence of free lunches in an economy with sublinear costs. In this section, we show that S–SDFs arise more generally also in an economy with convex costs, under a no-arbitrage condition stronger than the absence of free lunches, i.e., market viability, as defined in, e.g., Harrison and Kreps [1979].[6] To this end, let portfolio positions involve costs measured by the closed convex function $\sigma : \mathbb{R}^N \to [0, \infty]$. As in the main text, we define the set of traded payoffs as the set of portfolio payoffs involving finite costs:

$$\mathcal{Z} := \{Z = \boldsymbol{X}'\boldsymbol{\theta} \,:\, \sigma(\boldsymbol{\theta}) < +\infty\} \ .$$

Accordingly, we define a pricing functional π on $\mathcal{Z}$ by:

$$\pi(Z) := \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \{\boldsymbol{P}'\boldsymbol{\theta} + \sigma(\boldsymbol{\theta}) \,:\, Z = \boldsymbol{X}'\boldsymbol{\theta}\} \ .$$

Since the cost function σ is closed and convex, π is a convex pricing functional and $\mathcal{Z}$ is a closed and convex set of traded payoffs. With these definitions, the next proposition provides the closed-form supporting market in which the existence of a S–SDF is characterized by market viability.

Proposition OA-1. *Market $(\mathcal{Z}, \pi)$ is viable if and only if there exists an almost surely*

[6] A market $(\mathcal{Z}, \pi, \mathbb{P})$ is viable if there exists an agent (represented by preference $\gtrsim$) and $Z^* \in \mathcal{Z}$ such that $\pi(Z^*) \leq 0$ and $Z^* \gtrsim Z$ for all $Z \in \mathcal{Z}$ with $\pi(Z) \leq 0$. The preference relation $\gtrsim$ is assumed to be convex, continuous and strictly increasing.

strictly positive SDF $M \in L^q$, such that

$$\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P} \in C_\sigma \ ,$$

with the closed convex set

$$C_\sigma := \{\boldsymbol{\eta} \in \mathbb{R}^N \ : \ \boldsymbol{\eta}'\boldsymbol{\theta} \leq \sigma(\boldsymbol{\theta}) \ \ \textit{for all } \boldsymbol{\theta} \in \mathbb{R}^N\} \ .$$

Proof of Proposition OA-1. Convexity of σ implies convexity of pricing functional π and the set of traded payoffs $\mathcal{Z}$. By [Chen, 2001, Thm. 1 and Cor. 1], market viability is equivalent to the existence of an a.s. strictly positive SDF, M, such that $\mathbb{E}[MZ] \leq \pi(Z)$ for any $Z \in \mathcal{Z}$. Now, by definition of pricing functional π and traded payoff space $\mathcal{Z}$, it equivalently follows for any $\boldsymbol{\theta} \in \mathbb{R}^N$ that $\boldsymbol{\theta}'\boldsymbol{P} + \sigma(\boldsymbol{\theta}) \geq \mathbb{E}[M\boldsymbol{\theta}'\boldsymbol{X}]$, i.e.,

$$\boldsymbol{\theta}' \left(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}\right) \leq \sigma(\boldsymbol{\theta}) \ . \tag{OA-4}$$

Equivalently, by definition of C_σ, we have $\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P} \in C_\sigma$. Finally, since the cost function σ is closed and convex, C_σ is a closed and convex set. □

Given a closed convex set $C = \{\boldsymbol{\eta} \in \mathbb{R}^N \ : \ h(\boldsymbol{\eta}) \leq \tau\}$ bounding the pricing errors of a S–SDF with closed convex pricing error metric h and threshold $\tau \geq 0$, the convex arbitrage-free economy supporting such constraint set C is based on a cost function given by:

$$\sigma = h^* + \tau \ .$$

While in Proposition 1 in the main text, the necessary and sufficient condition for the existence of a strictly positive S–SDF in an economy with sublinear cost function σ is the absence of free-lunches, when σ is not sublinear the no-arbitrage condition equivalent to the existence of a strictly positive S–SDF is market viability.

B.2 Foundation in markets with ambiguity

In this section, we show how S–SDFs can arise also in frictionless markets, when the unobserved marginal investor belief and the probability belief $\mathbb{P}$ used to describe asset payoffs are not equivalent, offering a convenient framework to formulate robust asset pricing predictions in a variety of settings where an SDF may not exist under belief $\mathbb{P}$.[7] To clarify this, let $\tilde{\mathbb{P}}$ be the marginal investor belief and let frictionless market, denoted by the triplet $(\tilde{Z}_0, \tilde{\pi}_0, \tilde{\mathbb{P}})$, admit no free lunches, where

$$\tilde{Z}_0 := \{Z = \boldsymbol{X}'\boldsymbol{\theta} : \boldsymbol{\theta} \in \mathbb{R}^N\} \quad ; \quad \tilde{\pi}_0(Z) := \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \{\boldsymbol{P}'\boldsymbol{\theta} : Z = \boldsymbol{X}'\boldsymbol{\theta}\} \ .$$

From Proposition 1 in the main text, there exists a SDF $\tilde{M}$ such that:

$$\tilde{\mathbb{E}}[\tilde{M}\boldsymbol{X}] = \boldsymbol{P} \ , \tag{OA-5}$$

where $\tilde{\mathbb{E}}[\cdot]$ denotes expectations under probability $\tilde{\mathbb{P}}$. Importantly, an exact price representation of the form (OA-5) can hold for a strictly positive SDF M with respect to a probability belief $\mathbb{P}$ if and only if $\mathbb{P}$ and $\tilde{\mathbb{P}}$ are equivalent, in which case:

$$\boldsymbol{P} = \tilde{\mathbb{E}}[\tilde{M}\boldsymbol{X}] = \mathbb{E}[M\boldsymbol{X}] \ , \tag{OA-6}$$

with the Radon-Nykodim derivative $\psi := \frac{d\tilde{\mathbb{P}}}{d\mathbb{P}}$ and $M := \tilde{M}\psi$. Conversely, no strictly positive SDF satisfying equation (OA-6) under probability belief $\mathbb{P}$ can exist when some zero probability assessments under $\mathbb{P}$ and $\tilde{\mathbb{P}}$ differ. However, from Proposition 1 in the main

[7] For instance, in all continuous-time theoretical asset pricing settings, non equivalent beliefs $\tilde{\mathbb{P}}$ and $\mathbb{P}$ arise whenever the associated return volatility processes are not identical.

text we also know more generally that for any closed convex pricing error metric h, such that $h(\mathbf{0}) = 0$, market $(\mathcal{Z}_h, \pi_h, \mathbb{P})$ admits no free lunches if and only if a strictly positive S–SDF M exists such that:

$$h(\mathbb{E}[M\boldsymbol{X}] - \tilde{\mathbb{E}}[\tilde{M}\boldsymbol{X}]) \leq \tau \ . \tag{OA-7}$$

Here, set $C(h, \tau) = \{\boldsymbol{\eta} \in \mathbb{R}^N : h(\boldsymbol{\eta}) \leq \tau\}$ can be interpreted as an implicit constraint on the degree of ambiguity between beliefs $\tilde{\mathbb{P}}$ and $\mathbb{P}$. Under such constraint, SDF $\tilde{M}$ and S–SDF M may imply distinct asset valuations under the corresponding beliefs, however with differences restricted to belong to set $C(h, \tau)$. Such differences arise from the non-equivalence of $\mathbb{P}$ and $\tilde{\mathbb{P}}$, which renders the set of arbitrage opportunities perceived under each of these beliefs different.

Corollary OA-1 (S–SDFs and ambiguity). *Markets $(\tilde{Z}_0, \tilde{\pi}_0, \tilde{\mathbb{P}})$ and $(\mathcal{Z}_h, \pi_h, \mathbb{P})$ admit no free lunches if and only if there exist strictly positive SDF $\tilde{M}$ and strictly positive S—SDF M such that constraints (OA-7) hold, with pricing functional π_h explicitly given for any $Z \in \mathcal{Z}_h$ by:*

$$\pi_C(Z) = \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \left\{ \tilde{\mathbb{E}}[\tilde{M}\boldsymbol{X}'\boldsymbol{\theta}] + \sup_{\boldsymbol{\eta} \in \mathbb{R}^N} \left\{ \boldsymbol{\eta}'\boldsymbol{\theta} : h(\boldsymbol{\eta}) \leq \tau \right\} \ : \ \mathbb{P}(Z = \boldsymbol{X}'\boldsymbol{\theta}) = 1 \right\} \ .$$

Proof of Corollary OA-1. By Proposition 1 in the main text, absence of free lunches in market $(\tilde{Z}_0, \tilde{\pi}_0, \tilde{\mathbb{P}})$ is equivalent to the existence of a strictly positive SDF $\tilde{M}$ such that $\tilde{\mathbb{E}}[\tilde{M}\boldsymbol{X}] = \boldsymbol{P}$. Moreover, the absence of free lunches in market $(Z_h, \pi_h, \mathbb{P})$ is equivalent to the existence of a strictly positive SDF M such that $h(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \leq \tau$.

□

π_h in Corollary OA-1 is interpretable as the min-max replication cost of a payoff Z, after considering all ambiguous asset valuations under belief $\mathbb{P}$, which result from admissible price adjustments $\boldsymbol{\eta}$ relative to the marginal investor valuation $\boldsymbol{P} = \tilde{\mathbb{E}}[\tilde{M}\boldsymbol{X}]$. Since the no free lunches condition in market $(\mathcal{Z}_h, \pi_h, \mathbb{P})$ is weaker than the no free lunches condition in market $(\mathcal{Z}_0, \pi_0, \mathbb{P})$, because by definition $\mathcal{Z}_0 \supset \mathcal{Z}_h$ and $\pi_h(Z) \geq \pi_0(Z)$ for all $Z \in \mathcal{Z}_h$, the pricing predictions of S–SDFs in ambiguous frictionless settings are also more robust than those of SDFs in unambiguous frictionless settings.

B.3 Theories of Good-Deal-Bounds and robust identification of investor beliefs

Consider first the Good-Deal Bounds theory in Cochrane and Saa-Requejo [2000], which specifies a maximal price $\overline{P}_2$ and a minimal price $\underline{P}_2$ for a newly issued security with payoff X_2, given an exact pricing condition for a set of traded payoffs $\boldsymbol{X}_1$ with prices $\boldsymbol{P}_1$ and under an upper bound on the maximum attainable Sharpe Ratio. This bound is introduced to rule out unreasonably good deals in the extended market including the new security. Let M_S be the minimum dispersion SDF that minimizes the SDF dispersion for a given convex dispersion function ϕ and which prices exactly the set of sure securities.[8] Therefore, M_0 solves the optimization problem:

$$\inf_{M \in \mathcal{M}} \mathbb{E}[\phi(M)] \ , \quad \text{where } \mathcal{M} := \{M \in L^q \ : \ M \geq 0, \ \mathbb{E}[M\boldsymbol{X}_1] = \boldsymbol{P}_1\} \ .$$

Let $\nu_0 := \mathbb{E}[\phi(M_0)]$ and denote by $P_2 := \mathbb{E}[M_0 X_2]$ the reference price for the newly issued payoff under SDF M_0. The minimal and maximal price for the newly issued security are

[8] In Cochrane and Saa-Requejo [2000]'s good-deal bounds theory the notion of dispersion minimized is the SDF variance.

obtained for an upper dispersion threshold $\nu \geq \nu_0$ as follows:

$$\underline{P}_2 := P_2 + \inf_{M \in \mathcal{M}} \{\mathbb{E}[MX_2] - P_2 \ : \ \mathbb{E}[\phi(M)] \leq \nu\} \ , \tag{OA-8}$$

$$\overline{P}_2 := P_2 - \inf_{M \in \mathcal{M}} \{-(\mathbb{E}[MX_2] - P_2) \ : \ \mathbb{E}[\phi(M)] \leq \nu\} \ . \tag{OA-9}$$

By construction, these prices satisfy the inequalities $\underline{P}_2 \leq P_2 \leq \overline{P}_2$. Moreover, they are supported by corresponding S–SDFs $\underline{M}_0$ and $\overline{M}_0$, which solve each one of the minimizations on the RHS of identities (OA-8)-(OA-9):

$$\underline{M}_0 := \arg\min_{M \in L^q} \{\underline{h}(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \ : \ M \geq 0, \ \mathbb{E}[\phi(M)] \leq \nu\} \ , \tag{OA-10}$$

$$\overline{M}_0 := \arg\min_{M \in L^q} \left\{\overline{h}(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P})) \ : \ M \geq 0, \ \mathbb{E}[\phi(M)] \leq \nu\right\} \ , \tag{OA-11}$$

with corresponding pricing error function given for $\boldsymbol{\eta} = (\boldsymbol{\eta}_1', \eta_2)')$ by $\underline{h}(\boldsymbol{\eta}) = \delta_{\{\mathbf{0}\}}(\boldsymbol{\eta}_1) + \eta_2$ and $\overline{h}(\boldsymbol{\eta}) = \delta_{\{\mathbf{0}\}}(\boldsymbol{\eta}_1) - \eta_2$. It thus follows from Proposition 3 in the main text that $\underline{P}_2$ and $\overline{P}_2$ can be equivalently obtained by computing the price of the new security's payoff X_2 under these two minimum dispersion S–SDFs.

Consider now Chen et al. [2020] theory for a robust identification of investor beliefs. This theory relies on the determination of a S–SDF $\underline{M}_0$ solving problem (OA-10) under robust notions of S–SDF dispersion (Example 3, cases 1. and 3.(a) of Appendix Section B of the main text) and while imposing the normalization constraint $\mathbb{E}[M] = 1$, which can be naturally incorporated in constrained set $\mathcal{M}$ as an exact pricing constraint. Such normalization gives rise to the interpretation of $\underline{M}_0$ as a density reproducing a recovered robust investor belief with respect to the reference probability used to describe risky asset payoffs. Such interpretation if further motivated by the fact that in this theory "sure payoffs" $\boldsymbol{X}_1$ are not asset payoffs, but instead asset payoffs that have been already stochastically discounted by a parametric SDF $M^\star$ implied by economic theory. Therefore, in this setting $\mathbb{E}[\underline{M}_0 X_2]$ has the interpretation of a conservative expected payoff under the recovered investor belief, instead of a lowest arbitrage-free price for payoff X_2.

B.4 S–SDFs and conditional pricing restrictions

Our S–SDF framework is grounded in pricing restrictions formulated in terms of unconditional moments, taking the form

$$h(\mathbb{E}[M\boldsymbol{X}] - \boldsymbol{P}) \leq \tau \ . \tag{OA-12}$$

By Proposition 1 in the main text, these pricing restrictions arise in arbitrage-free markets where portfolio formation may entail costs. Specifically, purchasing portfolio payoffs $Z = \boldsymbol{\theta}'\boldsymbol{X}$ constructed with weights $\boldsymbol{\theta} \in \mathbb{R}^N$ incur formation costs captured by a sublinear cost function $\sigma(\boldsymbol{\theta})$, that is closely related to the pricing error function h. In particular, the cost function is one-to-one related to the set of admissible pricing errors $C := \{\boldsymbol{\eta} \in \mathbb{R}^N : h(\boldsymbol{\eta}) \leq \tau\}$; see Remark 1(i) in the main text.

Implicitly, our framework is set in a one-period economy where, at time $t = 0$, investors can acquire traded payoffs $\boldsymbol{\theta}'\boldsymbol{X}$ maturing at $t + 1 = 1$. The price of such a portfolio is given by the functional

$$\pi_\sigma(Z) := \min_{\boldsymbol{\theta} \in \mathbb{R}^N} \{\boldsymbol{P}'\boldsymbol{\theta} + \sigma(\boldsymbol{\theta}) : Z = \boldsymbol{X}'\boldsymbol{\theta}\} \ ,$$

which accounts for the quoted prices and the cost of forming portfolios. That is, the market effectively prices payoffs based on a modified cost structure, reflecting transaction frictions and other market imperfections.

Extending this framework to a multi-period setting introduces additional complexities that fall beyond the scope of this paper. Nevertheless, an informal discussion on how conditional pricing restrictions can be integrated into the framework provides insights into the dynamic structure of S–SDFs.

To formalize this, consider an information filtration $\{\mathcal{I}_t : t \in \mathbb{Z}\}$, where $\mathcal{I}_t$ represents the information set available to market participants at time t. The evolution of the market can be characterized by the following stationary and ergodic random sequences adapted to this filtration: (i) the basis asset payoffs and their quoted prices, $\{(\boldsymbol{X}_{t+1}, \boldsymbol{P}_t) : t \in \mathbb{Z}\}$, (ii) the threshold parameters $\{\tau_t \geq 0 : t \in \mathbb{Z}\}$, and (iii) a sequence of S–SDFs $\{M_t : t \in \mathbb{Z}\}$ satisfying the conditional pricing restriction

$$h(\mathbb{E}[M_{t+1}\boldsymbol{X}_{t+1}|\mathcal{I}_t] - \boldsymbol{P}_t) \leq \tau_t \ , \tag{OA-13}$$

almost surely.

Remark 1. These conditional restrictions imply unconditional pricing constraints of the form (OA-12), analogously to the classical exact pricing case. Specifically, by leveraging the convexity of h and applying Jensen's inequality along with the law of iterated expectations, it follows that

$$\begin{aligned} h\left(\mathbb{E}[M_{t+1}\boldsymbol{X}_{t+1} - \boldsymbol{P}_t]\right) &= h\left(\mathbb{E}\left[\mathbb{E}[M_{t+1}\boldsymbol{X}_{t+1}|\mathcal{I}_t] - \boldsymbol{P}_t\right]\right) \\ &\leq \mathbb{E}\left[h\left(\mathbb{E}_t[Mt+1\boldsymbol{X}_{t+1}] - \boldsymbol{P}_t\right)\right] \leq \mathbb{E}[\tau_t] \ , \end{aligned}$$

under appropriate regularity conditions.

The economic rationale underpinning conditional restrictions (OA-13) is a multi-period market characterized by a sequence of time-varying cost functions $\{\sigma_t : \mathbb{R}^N \to [0, +\infty] : t \in \mathbb{Z}\}$, which capture the evolution of transaction costs and other market frictions. At each time t, the cost function σ_t determines the admissible set of pricing deviations through the constraint

$$C_t := \{\boldsymbol{\eta} \in \mathbb{R}^N : h(\boldsymbol{\eta}) \leq \tau_t\} \ . \tag{OA-14}$$

An example arises when these pricing frictions are linked to time-varying bid-ask spreads. Specifically, if market liquidity fluctuates over time, the threshold τ_t can be modeled as an increasing function of prevailing transaction costs, capturing how wider bid-ask spreads relax the pricing constraints and allow for larger deviations from exact valuation.

Working with conditional market restrictions poses several challenges already in the classical setting of exact pricing, including (i) specifying a flexible statistical model for the conditional cross-moments between asset payoffs and the SDF, and (ii) identifying an appropriate set of conditioning variables that adequately summarize the information set $\mathcal{I}_t$ available to agents at time t. In the S–SDF framework, two additional considerations arise. The first is the need to model the time-variation in the threshold parameters τ_t. Since these thresholds can be linked to transaction costs, an empirical specification should account for the potentially dynamic structure of market frictions. The second is the proper construction of the conditioning variables so that the geometry of the conditional pricing restrictions is preserved when they are transformed into unconditional ones.

As an illustration, consider the simple case of short–sale constraints, where $\sigma_t(\boldsymbol{\theta}) = \delta_{\mathbb{R}^N_+}(\boldsymbol{\theta})$ and the associated pricing restrictions are

$$\mathbb{E}[M_{t+1}\, \boldsymbol{X}_{t+1} \mid \mathcal{I}_t] - \boldsymbol{P}_t \in \mathbb{R}^N_- \ ;$$

see Luttmer [1996]. Suppose further that we can transform these restrictions into unconditional pricing constraints using K instruments $\boldsymbol{Z}_t$ that summarize the information set $\mathcal{I}_t$. In this case, $\boldsymbol{Z}_t$ must be nonnegative to ensure that the unconditional pricing restrictions remain in the nonnegative orthant:

$$\mathbb{E}[(M_{t+1}\,\boldsymbol{X}_{t+1} - \boldsymbol{P}_t) \otimes \boldsymbol{Z}_t] \in \mathbb{R}^N_- \ ,$$

where $\otimes$ denotes the Kronecker product. More generally, depending on the shape of the constraint set C_t given in (OA-14), instruments should be constructed so that

$$\mathbb{E}[(M_{t+1}\,\boldsymbol{X}_{t+1} - \boldsymbol{P}_t) \otimes \boldsymbol{Z}_t]$$

remains in the convex hull of the collection of conditional pricing constraints $\{C_t : t \in \mathbb{Z}\}$.

By explicitly incorporating transaction costs into the pricing framework, our approach generalizes the standard SDF framework and provides a tractable method for studying markets with imperfect liquidity. The conditional S–SDF framework thus offers a flexible tool for analyzing pricing distortions arising from evolving market frictions.

C Dispersion Metrics Comparison

Figure OA-4 in Section D of this Online Appendix illustrates how the choice of S–SDF dispersion measure influences the tradeoff between time series and cross-sectional explanatory powers. The two left panels show that, when using non-perturbed data, S–SDFs that minimize dispersion measures—such as the Kullback-Leibler or Hellinger divergences—exhibit a tradeoff dominated by the minimum variance S–SDFs.[9] Moreover, the minimum variance S–SDFs yield a more consistent tradeoff on perturbed data for both the low- and intermediate-dimensional datasets.

[9]Kullback-Leibler and Hellinger divergences are part of the Cressie-Read family of stochastic dispersion measures described in Appendix Section B of the main text (see Example 3 case 1. and case 3.(b) for $\gamma = 1/2$, respectively).

D Additional Figures

We report additional figures to complement the results in the main text. Figure OA-1 illustrates the empirical duality failure discussed in Section 3.3 of the main text. It replicates the results in Figure 1 of the main text, but for the intermediate dimensional dataset. Figure OA-2 illustrates the tradeoff between time series and cross-sectional fit for APT S–SDFs, non APT S–SDFs and PCA-based S–SDFs, but for the low dimensional dataset. Figure OA-3 illustrates the cross-sectional distribution of APT S–SDF betas and time series R^2 metrics as a function of threshold τ. Figure OA-4 illustrates the tradeoff between time series and cross-sectional APT S–SDF explanatory power, under different notions of S–SDF dispersion and data perturbations. Figure OA-5 illustrates the role of traded factors in minimum dispersion APT S–SDFs, where the cross-sectional fit is now measured by Hansen-Jagannathan distance. Finally, Figure OA-6 illustrates the tradeoff between time series and cross-sectional APT S–SDF explanatory power for the intermediate dimensional dataset, under different notions of data perturbations.

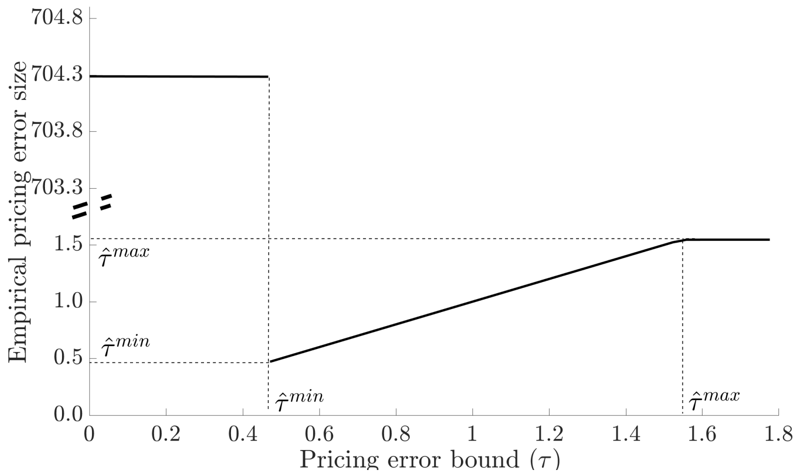

Figure OA-1: **Empirical duality failure.** We compute the minimum variance APT S–SDF M_0 with l_2-pricing function for varying thresholds $\tau \geq 0$. On the y-axis we report the estimate of the function value $h(\mathbb{E}[M_0\boldsymbol{X} - \boldsymbol{P}])$ for each τ. The point of discontinuity in the plot identifies the smallest threshold τ, for which a solution of the empirical primal S–SDF problem exists. The largest threshold $\hat{\tau}^{max}$ in the plot is computed as the sample version of the maximal threshold τ^{max} defined in the main text. Sorted portfolio returns are used to construct the S–SDFs without assuming any observable traded factor. All calculations are based on the intermediate dimensional dataset from June 1990 to June 2018.

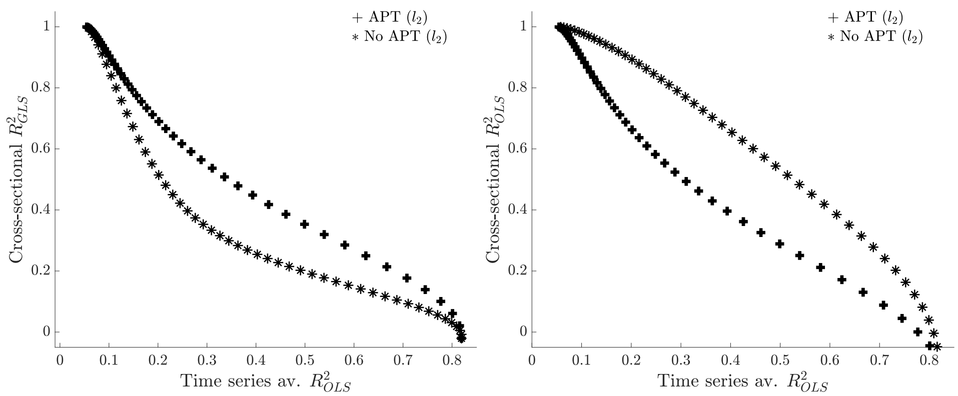

(1) **Panel (A)**

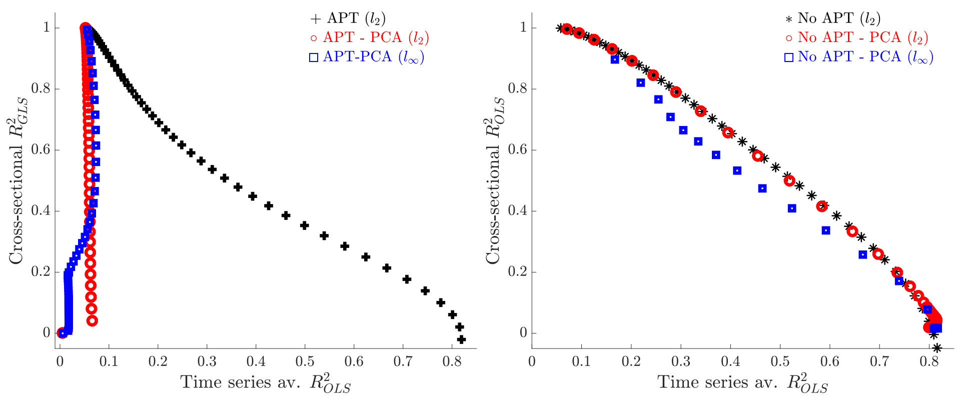

(2) **Panel (B)**

Figure OA-2: **Tradeoff between time series and cross-sectional S–SDF explanatory power of APT S–SDFs, non APT S–SDFs and principal components based S–SDFs.** For various minimum variance S–SDFs, the figure reports the tradeoff between average time series R^2 metric and cross-sectional GLS R^2 metric (denoted by R^2_{GLS}, left panels), or cross-sectional R^2 metric (denoted by R^2_{OLS}, right panels) defined in Section 4 of the main text. The following S–SDFs are compared in Panel (A). First, APT S–SDFs based on the l_2-metric. Second, non APT S–SDFs based on the l_2-metric, but with a diagonal weighting matrix $\boldsymbol{\Sigma_\eta} = \boldsymbol{I}_{N_D \times N_D}$. Both these S–SDFs take the first principal component of excess returns as the single traded factor. In each figure of Panel (B), results for two additional principal components based S–SDFs are reported, which impose the corresponding pricing constraints on the principal components of sorted portfolio excess returns, instead of the original excess returns. The first and second of these S–SDFs impose in the left (right) figure an APT l_2- and $l_\infty-$ pricing constraint (a non APT l_2- and $l_\infty-$ pricing constraint based on a diagonal weighting matrix $\boldsymbol{\Sigma_\eta} = \boldsymbol{I}_{N_D \times N_D}$), respectively. All results are based on the low dimensional dataset running from 1963 to 2018.

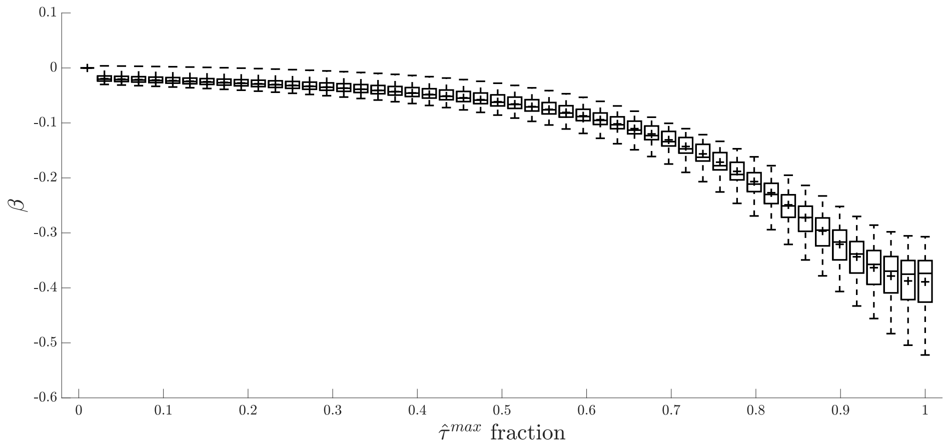

(1) **Panel (A)**

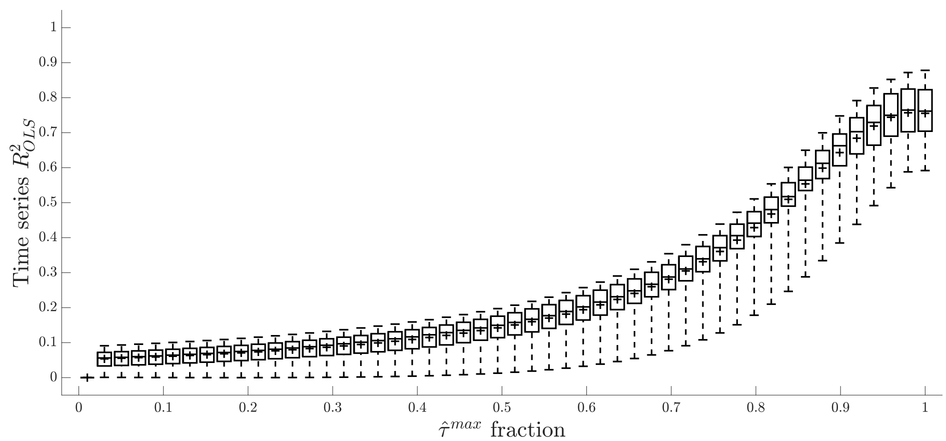

(2) **Panel (B)**

Figure OA-3: **Cross-sectional distributions of APT S–SDF betas and time series R^2 metrics.** The figure reports in **Panel (A)** (**Panel (B)**) the cross-sectional distribution of S–SDF betas (time series R^2 metrics defined in Section 4 of the main text) as a function of threshold τ, for minimum variance S–SDFs satisfying APT pricing constraint defined in Section 4 of the main text under an l_2-metric, and while taking the market excess return as the single traded factor. In both panels, plus signs (horizontal lines) in the boxes indicate cross-sectional averages (medians). All computations are based on the low dimensional dataset from 1963 to 2018.

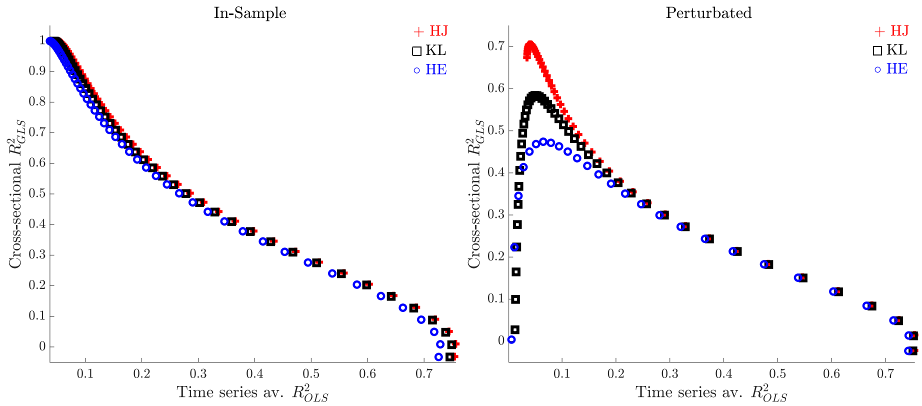

(1) **Panel (A)**

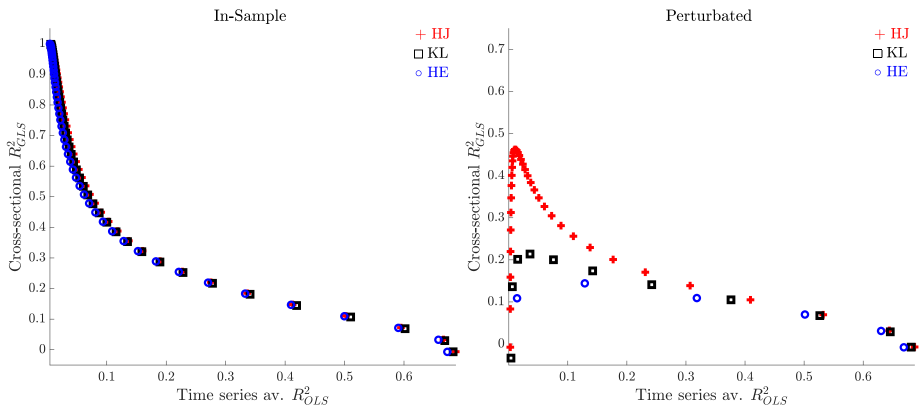

(2) **Panel (B)**

Figure OA-4: **Tradeoff between time series and cross-sectional APT S–SDF explanatory power, under different notions of S–SDF dispersion and data perturbations.** The figure reports the tradeoff between average time series R^2 metric and cross-sectional GLS R^2 metric defined in Section 4 of the main text for minimum dispersion APT S–SDFs based on the l_2-metric and on different notions of S–SDF dispersion from the Cressie and Read [1984] family in Appendix Section B of the main text: Variance (label HJ, case 3.(a)), Kullback-Leibler divergence (label KL, case 1.) and Hellinger divergence (label HE, case 3.(b)). **Panel (A)** reports results for the low dimensional dataset and **Panel (B)** for the intermediate dimensional dataset. Left panels report in-sample results, where the estimation and evaluation periods overlap on the 40 (50) year window 1965-2005 (1955-2005) for the low (intermediate) dimensional dataset. For the low (intermediate) dimensional dataset, right panels report results for a 40 (50) year estimation window 1965-2005 (1955-2015) and a translated 40 (50) year evaluation window 1975-2015 (1965-2015) of 10 years.

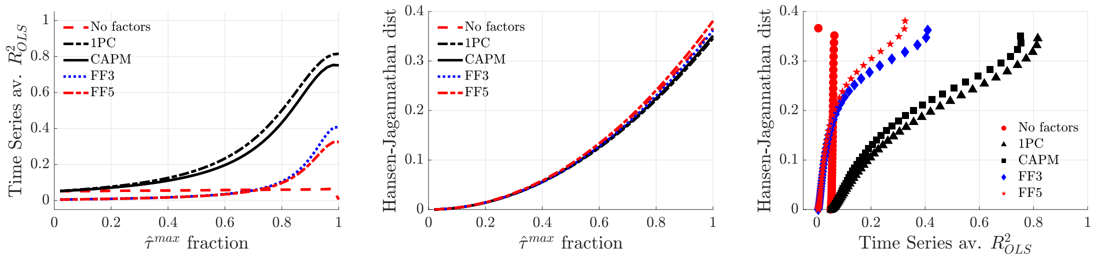

(1) **Panel (A)**

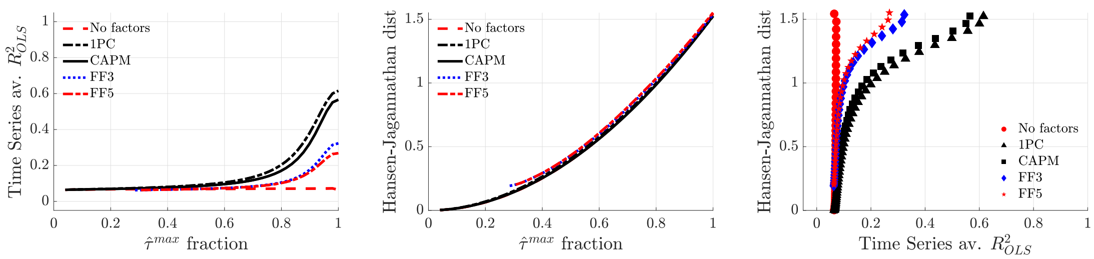

(2) **Panel (B)**

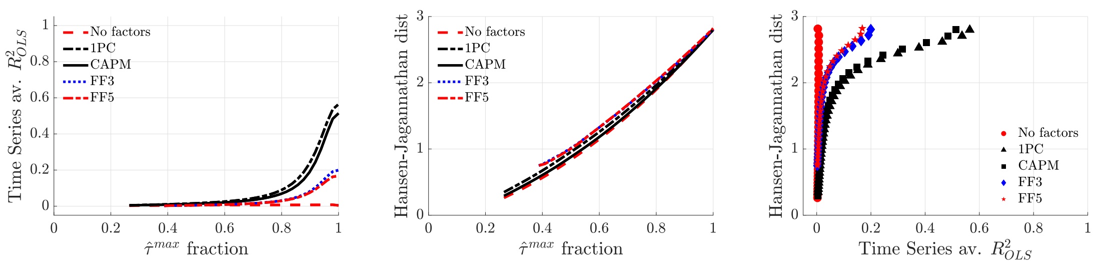

(3) **Panel (C)**

Figure OA-5: **Role of traded factors in minimum dispersion APT S–SDFs (Hansen-Jagannathan metric).** The figure illustrates the average time-series R^2 metric (left panel) and the cross-sectional pricing ability, measured by the Hansen–Jagannathan distance, for various minimum variance APT S–SDFs, constructed using different sets of traded risk factors, in the low dimensional (**Panel (A)**), intermediate dimensional (**Panel (B)**), and high dimensional (**Panel (C)**) datasets. Left (middle) panel reports the time-series (cross-sectional) fit as a function of the pricing error threshold. hold. The right panel shows the tradeoff between goodness of time-series and cross-sectional fit. We consider the following choices for the traded factors: (1) No factors; (2) 1PC, corresponding to the first principal component of excess returns as the single traded factor (2) CAPM, corresponding to the market return as the single traded factor; (3) FF3, corresponding to the three Fama-French factor returns; (4) FF5, corresponding to the five Fama-French factor returns.

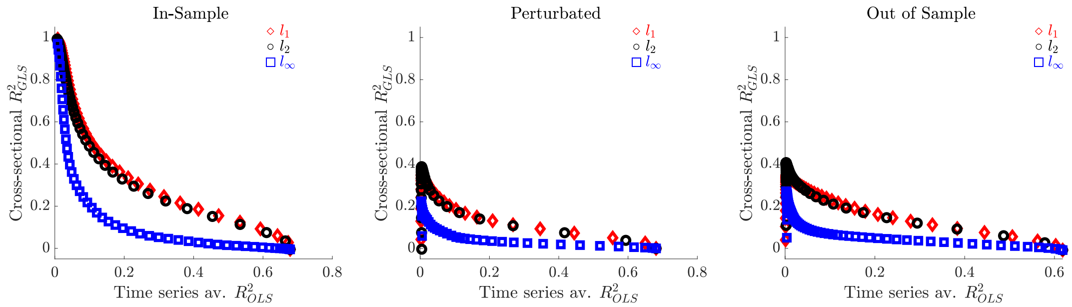

(1) **Panel (A)**

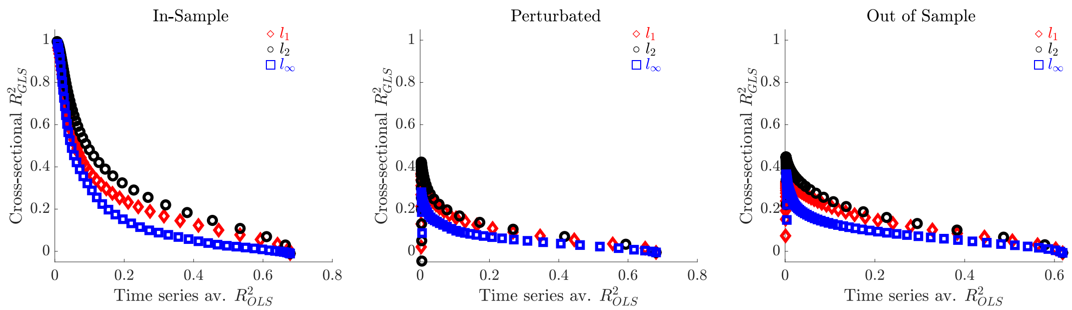

(2) **Panel (B)**

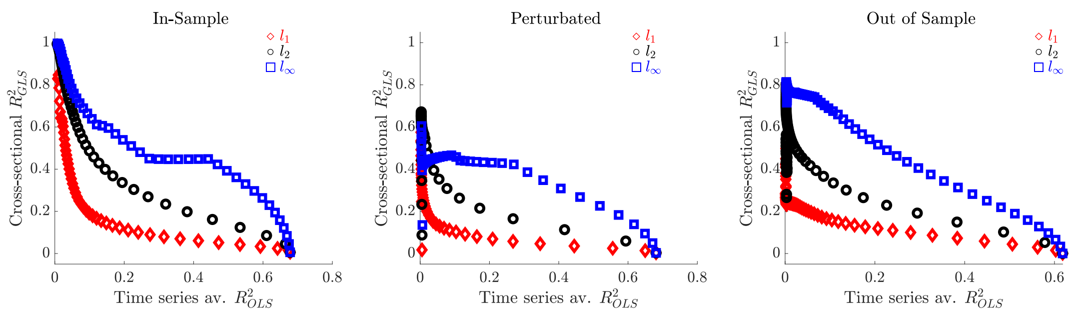

(3) **Panel (C)**

Figure OA-6: **Tradeoff between time series and cross-sectional APT S–SDF explanatory power (Intermediate-dimensional data)** The figure presents the average time-series R^2 metric and the cross-sectional GLS R^2 metric, calculated using the l_1-norm in **Panel (A)**, the l_2-norm in **Panel (B)**, and the (scaled) l_∞-norm in **Panel (C)**. Left panels show in-sample results where the S–SDFs are both estimated and evaluated over the same 40-year period (1975–2015). Middle panels present results where S–SDFs are estimated on a 40-year period (1965–2005) and evaluated on the 40-year period (1975–2015). Right panels provide out-of-sample results, with S–SDFs parameters estimated on rolling 30-year windows and returns of the subsequent 6 months. Each figure compares minimum variance APT S–SDFs estimated with the l_1-, l_2-, and (scaled) l_∞-norm pricing error metrics and varying pricing error bounds. Market return is used as the single risk factor.

References

Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.

Heinz H Bauschke, Patrick L Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces.* Springer, 2017.

Jonathan M Borwein and Adrian S Lewis. Duality relationships for entropy-like minimization problems. *SIAM Journal on Control and Optimization*, 29(2):325–338, 1991.

Xiaohong Chen, Lars Peter Hansen, and Peter G Hansen. Robust identification of investor beliefs. *Proceedings of the National Academy of Sciences*, 117(52):33130–33140, 2020.

Zhiwu Chen. Viable costs and equilibrium prices in frictional securities markets. *Annals of Economics and Finance*, 2(2):297–323, 2001.

John H Cochrane and Jesus Saa-Requejo. Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *Journal of Political Economy*, 108(1):79–119, 2000.

Patrick L Combettes, Andrew M McDonald, Charles A Micchelli, and Massimiliano Pontil. Learning with optimal interpolation norms. *Numerical Algorithms*, 81(2):695–717, 2019.

Noel Cressie and Timothy RC Read. Multinomial goodness-of-fit tests. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 46(3):440–464, 1984.

J Michael Harrison and David M Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic theory*, 20(3):381–408, 1979.

Erzo GJ Luttmer. Asset pricing in economies with frictions. *Econometrica*, 64(6):1439–1467, 1996.

Ralph T Rockafellar. Integrals which are convex functionals. *Pacific Journal of Mathematics*, 24(3):525–539, 1968.

Ralph T Rockafellar. *Convex analysis*. Princeton University Press, 1970.