
SCIENTIFIC REPORTS

Quantitative Finance

01 Jun 2026–19 Jun 2026

Organizing Committee

Daniel Bartl

National University of Singapore

Agostino Capponi

Columbia University

Ying Chen

National University of Singapore

Min Dai

The Hong Kong Polytechnic University

Julian Sester

National University of Singapore

Nizar Touzi

New York University

Marko Weber

National University of Singapore

CONTENTS PAGE

		Page
Anna Aksamit The University of Sydney, Australia	Partially Ordered Peacocks	03
Sara Biagini LUISS Rome, Italy	Short-lived Gases, Carbon Markets, and Climate Risk Mitigation	04
Umut Cetin London School of Economics, UK	Market Segmentation and Arbitrage	09
Patrick Cheridito ETH Zurich, Switzerland	Deep Learning for Continuous-time Stochastic Control with Jumps	13
Rama Cont Oxford University, UK	Dynamic Hedging under Model Uncertainty	14
Paolo Guasoni Dublin City University, Ireland	Hedging: Holding Stocks, Trading Bonds	15
Anastasis Kratsios McMaster University, Canada	Neural Chaos: Minimax-Optimal Approximation of Predictable Processes without Iterated Integrals or Fixed Bases	17
Kasper Larsen Rutgers University, USA	Strategic Trading when Agents Forecast the Forecasts of Others: Existence and Convergence	18
Kasper Larsen Rutgers University, USA	Existence of an Equilibrium with Limited Stock Market Participation and Power Utilities	21
Mathieu Lauriere NYU Shanghai, China	Dual Approaches to Stochastic Control vis APDEs and the Pathwise Hopf Formula	24
Xiaolu Tan Chinese University of Hong Kong, Hong Kong SAR	Partial Method for Mean-field Control Problems with Partial Observation	25
Hao Xing Boston University, USA	Flexible Information Acquisition in the Kyle Model	27

PARTIALLY ORDERED PEACOCKS

ANNA AKSAMIT

Martingale optimal transport problem of transporting one measure to another under a martingale constraint emerges in the robust approach to pricing and hedging. The aim is to find optimal bounds on the prices of exotic derivatives in terms of the (market) prices of call options. No explicit assumptions are made about the dynamics of the price process of the underlying asset and information about the distribution of asset prices is derived from the call prices.

Kellerer's theorem shows that any peacock — that is, a family of probability measures increasing in convex order — can be embedded in a Markovian martingale whose marginals match those measures. This is a cornerstone result closely linked to martingale optimal transport. In our work, we look at what happens when the measures are only partially ordered in convex order — partial peacock — and study when an appropriate embedding is possible. This concept may be applied to verify compatibility of vanilla and barrier product prices.

In general, such embeddings do not always exist, and even simple cases can exhibit non-trivial interactions. By focusing on the subclass of min-stable embeddings, where the embedding extends to the lower semilattice completion of the measures, we will provide necessary and sufficient conditions for the existence of these min-stable embeddings. We will also present a convex-concave decomposition of measure pairs, as well as geometric characterisations of such embeddings.

THE UNIVERSITY OF SYDNEY, AUSTRALIA

Email address: `anna.aksamit@sydney.edu.au`

BEYOND CO₂ EQUIVALENCE: CARBON MARKETS AND CLIMATE RISK

SARA BIAGINI (BASED ON A JOINT WORK WITH ENRICO BIFFIS AND KAVEH SALEHZADEH-NOBARI)

Classification AMS 2020: 91B76, 91G80, 60H20, 93E20.

Keywords: carbon markets; greenhouse gases; stochastic emissions; radiative forcing; temperature stabilization; emission permits; carbon accounting.

Carbon markets and carbon-accounting standards typically aggregate greenhouse-gas emissions in CO₂-equivalent units. The conversion is usually based on Global Warming Potential (GWP), which assigns a fixed exchange rate between gases over a prescribed horizon. This is convenient, but it can be misleading when gases differ substantially in atmospheric persistence and radiative forcing. Methane, for example, has a strong but relatively short-lived effect on temperature, whereas part of the warming generated by CO₂ is effectively irreversible. Thus, finding the optimal abatement mix across greenhouse gases is a non-trivial issue, which becomes more pronounced when gas-specific abatement costs and the associated social costs are taken into account.

We develop a stochastic, continuous-time model in which emissions are kept in gas-specific mass units. Firms choose abatement and trade gas-specific permits over a finite horizon $[0, T]$, while a regulator sets emission targets so as to satisfy a temperature constraint at a possibly longer horizon $\tilde{T} \geq T$ and to minimize social costs. The model links stochastic emissions to radiative forcing through gas-specific pulse-response functions and then maps aggregate forcing into temperature using an energy-balance model. Rather than imposing CO₂ equivalence ex ante, this framework provides a direct, and economically sound, way to compare gases that incorporates abatement costs.

1. EMISSIONS, RADIATIVE FORCING AND TEMPERATURE

There are K sectors and each sector j emits for simplicity one greenhouse gas. Firm ij is the i -th firm in sector j and N_j denotes the sector size. The cumulative controlled emissions of firm ij are:

$$dE_{ij}^a(t) = (\mu_{ij} - a_{ij}(t))dt + \sigma_{ij}(t)dW_{ij}(t), \quad t \in [0, T],$$

where a_{ij} denotes abatement. The Brownian motions may contain idiosyncratic, sector-specific and common components, so that emission shocks can be correlated both within and across sectors. Each sector is regulated independently, namely there are K markets of emissions permits, and permits are not fungible across sectors. Firms in sector j trade permits at rate θ_{ij} , and terminal net emissions are:

$$X_{ij}^{a,\theta}(T) = E_{ij}^a(T) - \int_0^T \theta_{ij}(t)dt.$$

For a given permit price process P_j , firm ij solves:

$$(1.1) \quad \inf_{a_{ij}, \theta_{ij}} \mathbb{E} \left[\int_0^T \left(h_{ij} a_{ij}(t) + \frac{a_{ij}(t)^2}{2\eta_{ij}} + \theta_{ij}(t) P_j(t) \right) dt \right]$$

subject to

$$\mathbb{E} \left[(X_{ij}^{a, \theta}(T) - b_{ij})^2 \right] \leq \varepsilon^2.$$

The parameter b_{ij} is the regulatory target and ε captures tolerance for random deviations from it. We impose the condition

$$(1.2) \quad \mathbb{E} \left[\int_0^T \frac{\sigma_{ij}^2(t)}{T-t} dt \right] < \infty,$$

which guarantees an absolutely continuous representation of the optimal terminal permit position (see [3], [2]).

The climate related part starts from the contribution of each firm to radiative forcing. If ψ_j denotes the pulse-response measure associated with gas j , we set:

$$R_{ij}^a(t) = \int_0^{T \wedge t} \psi_j([s-t, 0]) dE_{ij}^a(s).$$

A sufficient condition for R_{ij}^a to be a square-integrable semimartingale is:

$$(1.3) \quad \psi_j([v, 0]) = \int_v^0 f_j(u) du, \quad f_j \in L_{\text{loc}}^\infty((-\infty, 0]).$$

Short-lived gases like CH_4 can be represented by exponential kernels: for instance $f_j(u) = d_j e^{\beta_j u}$, while long-lived gases like CO_2 are naturally described by kernels f_j which are mixtures of exponentials and include a non-decaying term.

We assume that contributions of individual gases aggregate linearly, and $R^a = \sum_{ij} R_{ij}^a$ denotes aggregate forcing. Adapting the carbon two-box model by Pierrehumbert [5] to a simplified one-layer energy-balance, the temperature impact I^a satisfies:

$$I^a(t) = -\frac{\lambda}{C} \int_0^t I^a(s) ds + \frac{1}{C} R^a(t),$$

with $C > 0$ the heat capacity and $\lambda > 0$ the climate-feedback parameter of the layer. Setting $\alpha = \lambda/C$, under condition (1.3) the temperature impact from human activities is given by:

$$I^a(t) = \frac{1}{C} \int_0^t e^{\alpha(s-t)} dR^a(s).$$

Let us note, for later use, that the stochastic Fubini Theorem implies:

$$(1.4) \quad dR_{ij}^a(t) = dt \int_0^{T \wedge t} f_j(s-t) dE_{ij}^a(s).$$

2. EQUILIBRIUM IN EMISSION PERMIT MARKETS

The individual optimization problem (1.1) is convex and, by monotonicity in the abatement, the regulatory constraint binds at the optimum. The first-order conditions imply that marginal abatement cost is equal to the permit price. Hence

$$(2.1) \quad \hat{a}_{ij}(t) = \eta_{ij}(\hat{P}_j(t) - h_{ij}).$$

The optimal trading strategy is not unique, because the objective is not strictly convex in θ_{ij} , but an admissible optimal representation exists under the quadratic-variation assumption (1.2). The Lagrange multiplier is common across firms within each sector and satisfies

$$\hat{y}_j = \frac{\sqrt{\mathbb{E}[\hat{P}_j(T)^2]}}{2\varepsilon}.$$

Let $\bar{\eta}_j, \bar{\mu}_j, \bar{c}_j = N_j^{-1} \sum_i \eta_{ij} h_{ij}$ and $\bar{b}_j$ denote the corresponding sectoral averages, and define:

$$\bar{W}_j(t) = \frac{1}{N_j} \sum_{i=1}^{N_j} \int_0^t \sigma_{ij}(s) dW_{ij}(s).$$

Under a sufficiently stringent average cap and tolerance level, there exists a unique non-zero equilibrium permit price, which is a martingale with dynamics:

$$d\hat{P}_j(t) = \frac{2\hat{y}_j}{1 + 2\hat{y}_j\bar{\eta}_j(T-t)} d\bar{W}_j(t),$$

with

$$\hat{P}_j(0) = \frac{2\hat{y}_j((\bar{\mu}_j + \bar{c}_j)T - \bar{b}_j)}{1 + 2\hat{y}_j\bar{\eta}_jT}.$$

The multiplier $\hat{y}_j$ is determined by the compliance tolerance ε . In particular, tighter caps and smaller tolerance levels increase the shadow cost of regulation and, through the equilibrium price, induce stronger abatement. Firms with relatively high abatement costs are net buyers of permits, while firms with relatively low costs are net sellers.

The limiting case $\varepsilon \downarrow 0$ corresponds to perfect compliance. Then:

$$\hat{P}_j^L(0) = \frac{(\bar{\mu}_j + \bar{c}_j)T - \bar{b}_j}{\bar{\eta}_jT}.$$

Thus, equilibrium permit prices depend on the corresponding abatement costs and regulatory targets.

3. OPTIMAL EMISSION TARGETS

The regulator chooses the vector of emission targets so as to minimize aggregate equilibrium abatement costs subject to a temperature target at date $\tilde{T}$. If I^0 denotes the temperature impact generated by business-as-usual emissions (zero abatement), then the laissez-faire expected temperature impact is:

$$(3.1) \quad U_0 := \mathbb{E}[I^0(\tilde{T})] = \sum_{j=1}^K N_j \bar{\mu}_j A_j(T, \tilde{T}), \quad \Delta_U = U_0 - U.$$

in which $A_j(T, \tilde{T})$ is the expectation of the temperature effect in $\tilde{T}$ of a unit flow of gas j emitted over $[0, T]$, namely:

$$(3.2) \quad A_j(T, \tilde{T}) = \frac{1}{C} \int_0^{\tilde{T}} e^{\alpha(s-\tilde{T})} \int_0^{T \wedge s} f_j(v-s) dv ds.$$

This easily follows from (1.4). We will refer to $A_j(T, \tilde{T})$ as an annuity factor. It plays the role of a gas-specific discount factor from emissions into temperature. It depends on both the emission horizon and the temperature horizon and incorporates the atmospheric

persistence of the gas directly. As the horizon $\tilde{T}$ increases, the relative contribution of long-lived gases becomes more important.

If all gas markets are at equilibrium, for a given set of emissions targets and tolerance, one has:

$$(3.3) \quad \mathbb{E}[\hat{I}(\tilde{T})] = \sum_{j=1}^K N_j (\bar{\mu}_j - \bar{a}_j^*(0)) A_j(T, \tilde{T}),$$

so that the temperature constraint can be written as

$$(3.4) \quad \sum_{j=1}^K N_j A_j(T, \tilde{T}) \bar{a}_j^*(0) = \Delta_U.$$

The regulator therefore allocates the mitigation effort across gases by combining the relevant annuity factors with equilibrium abatement costs. This is the main difference with an exogenous GWP conversion rule. The regulator then minimizes of the social cost goal function, defined as the sum of the individual firms' value functions (namely, at equilibrium for a given emission reduction target b_{ij}), over targets. The solution(s) provides the optimal mix across gases.

For a single gas, the expectation of the optimal abatement is $\bar{a}^*(0) = \Delta_U / [N A(T, \tilde{T})]$, and the permit price is obtained by adding the corresponding marginal abatement-cost term. With two gases and the same number of firms in each sector, $N_1 = N_2 = N$,

$$(3.5) \quad A_1(T, \tilde{T}) \bar{a}_1^*(0) + A_2(T, \tilde{T}) \bar{a}_2^*(0) = \frac{\Delta_U}{N}.$$

If gas 1 is long-lived and gas 2 is short-lived, the equilibrium permit price for gas 2 satisfies

$$(3.6) \quad P_2^*(0) = p_2(0) - \frac{A_1(T, \tilde{T}) \bar{a}_1^*(0)}{A_2(T, \tilde{T}) \bar{\eta}_2},$$

where $p_2(0)$ is the price that would prevail if gas 2 were the only gas in the economy. The ratio of annuity factors plays a role similar to a conversion factor, but equilibrium prices also depend on abatement costs and on how the temperature target is shared across sectors.

The calibrated examples focus on CO₂ and methane. Methane abatement is relatively more effective at short horizons; as $\tilde{T}$ increases, its contribution decays, while the persistent component of CO₂ becomes dominant. The optimal target mix therefore shifts towards stronger CO₂ abatement.

This also identifies a weakness of carbon markets based on fixed CO₂-equivalent conversion factors. Cross-trading at GWP weights can induce firms to substitute cheaper abatement of a short-lived gas for more costly abatement of a persistent one. The CO₂e footprint may then satisfy a short-term target while the long-run temperature path deteriorates. In contrast to [4], and in the spirit of [1], the model therefore supports gas-specific reporting and permit markets when the policy objective is temperature stabilization.

REFERENCES

- [1] M. R. Allen, G. P. Peters, K. P. Shine et al. Indicate separate contributions of long-lived and short-lived greenhouse gases in emission targets. *npj Climate and Atmospheric Science*, 5:5, 2022.
- [2] P. Bank, H. M. Soner, M. Voß. Hedging with temporary price impact. *Mathematics and Financial Economics*, 11(2):215-239, 2017.
- [3] S. Biagini, G. Žitković. Representation of random variables as Lebesgue integrals. *Bernoulli*, 30(3):1878-189, 2024.
- [4] J. Morris, S. Paltsev and J. Reilly. Marginal abatement costs and marginal welfare costs for greenhouse gas emissions reductions: Results from the EPPA model. *Environmental Modeling & Assessment*, 17:325–336, 2012.
- [5] R. Pierrehumbert, Short-lived climate pollution, *Annual Review of Earth and Planetary Sciences*, 42:341–379, 2014.
- [6] K. Tanaka and B. C. O'Neill. The Paris Agreement zero-emissions goal is not always consistent with the 1.5°C and 2°C temperature targets. *Nature Climate Change*, 8:319–324, 2018.

HIGH-FREQUENCY TRADING WITH PUBLIC AND PRIVATE INFORMATION: LIQUIDITY AND PRICE IMPACT IN A KYLE–BACK MODEL

UMUT ÇETİN AND EDUARDO FERIOLI GOMES

Classification AMS 2020: 91G15; 91B44; 60H30; 93E11.

Keywords: Kyle–Back equilibrium; insider trading; high-frequency trading; market liquidity; price impact; public information; signal aggregation; market fragmentation; Gaussian filtering.

1. INTRODUCTION

Empirical studies of algorithmic order flow identify a persistent group of *directional* high-frequency traders who take positions in anticipation of future price moves. Such traders place aggressive orders that are strongly correlated with public news and with short-horizon returns, yet they realise a substantial part of their profits over longer horizons, indicating that they trade on the long-run fundamental rather than on the news reaction alone [3, 4]. Algorithm-level data from Euronext Amsterdam confirm that liquidity-providing members split into stable behavioural clusters — *directional*, *opportunistic*, and *market making* — with only about a third acting as market makers [5]. This motivates a strategic-trading model whose cast is a single directional trader facing competitive, liquidity-providing market makers, in an environment where public information flows continuously.

We study a continuous-time Kyle–Back model in which the informed trader and the market makers both learn about the fundamental value from Gaussian signals that sharpen gradually over the trading interval. In contrast to Kyle [9] and Back [1], *all* participants observe a public signal, and the informed trader observes a *second* signal in addition. The framework admits two economically distinct readings of the same equations. In the first, the second signal is genuine private information and the trader is a classical insider who also benefits from public news. In the second, there is no private information: two venues quote the same asset, each market maker sees only its own venue’s public signal and order flow, and the “informed” trader is an agent with cross-venue access who aggregates both venues’ public signals.

2. THE MODEL

The fundamental value is $V = \sigma_V \eta + \mu$, with $\eta \sim N(0, 1)$ and $\sigma_V > 0$, revealed at $t = 1$. All agents observe a public signal X^M , a continuous Gaussian bridge converging to η ,

$$(2.1) \quad dX_t^M = \sigma_M(t) dB_t^M + \sigma_M^2(t) \frac{\eta - X_t^M}{1 - \Sigma_M(t)} dt, \quad \Sigma_M(t) = \int_0^t \sigma_M^2(s) ds, \quad \Sigma_M(1) = 1,$$

so that $\frac{1}{1 - \Sigma_M(t)}$ is the precision of the public signal. The informed trader also observes a signal X^I of the same type, with $X_0^I \sim N(0, c^2)$ and precision index Σ_I , $\Sigma_I(1) = 1$. Noise traders submit a Brownian demand B ; the market makers observe the public signal and

the total order flow $Y = \theta + B$, where θ is the trader's strategy, and, competing à la Bertrand, quote $S_t = \mathbb{E}[V | \mathcal{F}_t^M]$. Following the literature, admissible (Markovian, linear) pricing rules take the form $S_t = H(t, X_t, X_t^M)$ with $H(t, x, u) = \beta_0(t) + \sigma_V x + \beta_2(t)u$, where the market makers' *processed signal* X solves

$$(2.2) \quad dX_t = w(t) dY_t + (r_0(t) + r_1(t)X_t + r_2(t)X_t^M) dt,$$

for deterministic $w > 0$ and (r_i) . The risk-neutral trader maximises $\mathbb{E}^z[(V - S_{1-})\theta_{1-} + \int_0^{1-} \theta_s dS_s]$ over admissible strategies. An equilibrium is a pricing rule and a strategy such that prices are rational given the strategy and the strategy is optimal given the rule.

The trader's effective information is the fusion of the two signals. Writing $Z_t := \mathbb{E}^z[\eta | \mathcal{F}_t^I]$, one has the following.

Theorem 2.1 (Signal aggregation). $Z_t = \pi_0(t)X_t^I + \pi_1(t)X_t^M$ with $\pi_0 = \frac{1-\Sigma_Z}{1-\Sigma_I}$, $\pi_1 = \frac{1-\Sigma_Z}{1-\Sigma_M}$, where the aggregate precision index Σ_Z satisfies

$$(2.3) \quad \frac{\Sigma_Z(t)}{1-\Sigma_Z(t)} = \frac{\Sigma_I(t)}{1-\Sigma_I(t)} + \frac{\Sigma_M(t)}{1-\Sigma_M(t)}, \quad t \in [0, 1).$$

In particular $\Sigma_Z \geq \max\{\Sigma_I, \Sigma_M\}$, and, given $\mathcal{F}_t^I$, η is Gaussian with mean Z_t and variance $1 - \Sigma_Z(t)$.

Precisions add: the trader always outperforms both signals, so the private signal need not be more informative than the public one — only the *aggregate* must be. Under the second reading, (2.3) is cross-venue signal aggregation with no privileged information.

3. MAIN RESULTS

Optimality and the equilibrium system. For an admissible rule define the nonnegative functional $\phi(t, x, u, z) = \frac{(H(t, x, u) - \sigma_V z - \mu)^2}{2\sigma_V w(t)}$. A martingale/jump argument shows that all randomness in the trader's problem enters the value function linearly through the order flow, which yields a sharp characterisation of optimality.

Theorem 3.1 (Insider optimality). *Let (H, w, r) be an admissible pricing rule. Any admissible strategy that is of finite variation (no martingale component) and satisfies $\lim_{t \rightarrow 1} \phi(t, X_t, X_t^M, Z_t) = 0$, $\mathbb{P}^z$ -a.s., is optimal, and the associated expected terminal wealth is finite and given in closed form.*

Combining rational pricing (the market makers filter V from X^M and Y ; the trader's drift is invisible to them) with Theorem 3.1 reduces the equilibrium to a coupled system for the order-flow weight w and the conditional variance v of η :

$$(3.1) \quad \begin{aligned} \frac{w'(t)}{w(t)} &= -2(v(t) + 1 - \Sigma_Z(t)) \frac{\sigma_M^2(t)}{(1 - \Sigma_M(t))^2}, \\ \sigma_Z^2(t) - v'(t) &= w^2(t) + \frac{\sigma_M^2(t)}{(1 - \Sigma_M(t))^2} (v(t) + 1 - \Sigma_Z(t))^2, \quad v(0) = c^2, \end{aligned}$$

where $\sigma_Z^2 = \Sigma'_Z$. The initial weight $w(0)$ is not a free parameter: it is pinned down by the terminal condition $v(1) = 0$.

Theorem 3.2 (Existence). *There is a pair (w, v) solving (3.1) with $w^{-1/2}(0) = \int_0^1 w^{3/2}(s) ds < \infty$. Moreover w and $v + 1 - \Sigma_Z$ are decreasing, $v \leq \Sigma_Z - \Sigma_M$,*

and both w and $v + 1 - \Sigma_Z$ are independent of Σ_Z . One has $\lim_{t \rightarrow 1} w(t) = 0$ and $w(t)/w(0) \geq (1 - \Sigma_M(t))^2$.

Theorem 3.3 (Equilibrium). *With w, v from Theorem 3.2, $\beta_0 \equiv \mu$ and $\beta_2 = \sigma_V \frac{v+1-\Sigma_Z}{1-\Sigma_M}$, the pricing rule*

$$H^*(t, x, u) = \mu + \sigma_V x + \beta_2(t)u, \quad r_0^* \equiv 0, \quad r_1^* = \frac{w'}{2w}, \quad r_2^* = \frac{\beta_2}{\sigma_V} \left(\left(1 - \frac{\beta_2}{\sigma_V}\right) \frac{\sigma_M^2}{1-\Sigma_M} - \frac{\beta_2'}{\beta_2} \right),$$

together with the trader's strategy $d\theta_t^* = \alpha_t^* dt$, $\theta_0^* = 0$,

$$\alpha_t^* = w(t) \frac{Z_t - X_t - \frac{\beta_2(t)}{\sigma_V} X_t^M}{v(t)},$$

form an equilibrium. The equilibrium order flow $Y^* = B + \theta^*$ is a $(\mathbb{P}, \mathcal{F}^M)$ -Brownian motion, and $H^*(t, X_t, X_t^M) \rightarrow V$ as $t \rightarrow 1$, $\mathbb{P}$ -a.s.

Liquidity, price impact, and the trader's strategy. In equilibrium the price follows $dS_t = \sigma_V w(t) dY_t + \beta_2(t) dX_t^M$, so Kyle's lambda is $\lambda(t) = \sigma_V w(t)$. Since w is decreasing to 0 (Theorem 3.2), *liquidity rises over the horizon and price impact vanishes at the close* — in sharp contrast with the constant λ of the classical model. As w depends only on the public flow Σ_M , so does liquidity.

Proposition 3.4. $\lambda = \sigma_V w$ is a convex, decreasing function of the public-signal precision $\frac{1}{1-\Sigma_M}$: sharper public information makes the market more liquid, at an increasing rate.

Unlike the infinite-variation bridge strategy of the classical Kyle–Back equilibrium, here it is the *market makers* who drive the price to V , by placing vanishing weight on order flow near the close so as to deny the trader profitable manipulation; correspondingly the trader's strategy is of finite variation: when $\Sigma_I \equiv 1$, $\int_0^1 \mathbb{E}^z[w(t)|\alpha_t^*|] dt < \infty$.

Invariance to the quality of the trader's signal. Let $u := v + 1 - \Sigma_Z = \text{Var}(\eta \mid \mathcal{F}^M)$, which is fixed by Σ_M alone, and let $\kappa := v/u \in (0, 1]$ be the trader's *information share* ($\kappa = 1$ is perfect information). Because w and u are independent of Σ_Z (Theorem 3.2), the price process, market liquidity, and the trader's ex-ante expected profit are *invariant* to the quality of the trader's signal. Degrading the signal re-splits the fixed u into v and $1 - \Sigma_Z$ and leaves prices untouched; only the trading intensity responds, $\mathbb{E}^z|\alpha_t^*| = \sqrt{2/\pi} w/\sqrt{v} = \kappa^{-1/2} \times (\text{perfect-information intensity})$. Quantity substitutes for quality, and signal quality is undetectable from the order flow.

4. DISCUSSION AND RELATED WORK

The model sits in the Kyle–Back tradition [9, 1, 2], in which a risk-neutral insider trades against competitive market makers and prices converge to the fundamental; the continuous-time analysis and filtering techniques build on Çetin and Danilova [6]. In that literature Kyle's λ is constant and the insider's information-revealing strategy is of infinite variation. Our two departures are that public information flows continuously to all agents, and that revelation is achieved by the market makers' choice of weights rather than by the insider; as a consequence λ is time-varying and decreasing, liquidity is endogenously rising, and the equilibrium strategy is of finite variation.

The closest work in the high-frequency literature is Foucault, Hombert and Roşu [7], who also embed news in a Kyle-type model and distinguish a *fast* from a *slow* speculator

according to whether the dealer prices off the contemporaneous news innovation. Their equilibrium features a *constant* λ and a comparative-static “more informative news $\Rightarrow$ more liquid”; ours delivers a dynamic mechanism in which liquidity rises within the trading day as the market makers withdraw weight from order flow. The cross-venue reading connects to the “toxic arbitrage” of Foucault, Kozhan and Tham [8], who study a liquidity provider blind to a second venue facing an agent who sees across venues. There the relevant edge is *latency* — a stale-quote race that widens spreads — whereas here the edge is *informational*: because each venue’s own public signal converges to the fundamental, the informational gap closes endogenously and liquidity *improves*.

REFERENCES

- [1] K. Back. Insider trading in continuous time. *Review of Financial Studies*, 5(3):387–409, 1992.
- [2] K. Back and H. Pedersen. Long-lived information and intraday patterns. *Journal of Financial Markets*, 1(3–4):385–402, 1998.
- [3] J. Brogaard, T. Hendershott, and R. Riordan. High-frequency trading and price discovery. *Review of Financial Studies*, 27(8):2267–2306, 2014.
- [4] A. Carrion. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets*, 16(4):680–711, 2013.
- [5] Á. Cartea, S. N. Cohen, R. Graumans, S. Labyad, L. Sánchez-Betancourt, and L. van Veldhuijzen. Statistical predictions of trading strategies in electronic markets. *Journal of Financial Econometrics*, 23(2), 2025.
- [6] U. Çetin and A. Danilova. *Dynamic Markov Bridges and Market Microstructure: Theory and Applications*. Probability Theory and Stochastic Modelling, vol. 90, Springer, 2018.
- [7] T. Foucault, J. Hombert, and I. Roşu. News trading and speed. *Journal of Finance*, 71(1):335–382, 2016.
- [8] T. Foucault, R. Kozhan, and W. W. Tham. Toxic arbitrage. *Review of Financial Studies*, 30(4):1053–1094, 2017.
- [9] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

DEPARTMENT OF STATISTICS, LONDON SCHOOL OF ECONOMICS, UK
Email address: u.cetin@lse.ac.uk

MATHEMATICS AND STATISTICS INSTITUTE, UNIVERSIDADE FEDERAL FLUMINENSE, BRASIL
Email address: eduardoferioli@id.uff.br

DEEP LEARNING FOR CONTINUOUS-TIME STOCHASTIC CONTROL WITH JUMPS

PATRICK CHERIDITO

Classification AMS 2020:

Keywords: Machine Learning (cs.LG), Systems and Control (eess.SY), Optimization and Control (math.OC), Portfolio Management (q-fin.PM)

In this paper, we introduce a model-based deep-learning approach to solve finite-horizon continuous-time stochastic control problems with jumps. We iteratively train two neural networks: one to represent the optimal policy and the other to approximate the value function. Leveraging a continuous-time version of the dynamic programming principle, we derive two different training objectives based on the Hamilton–Jacobi–Bellman equation, ensuring that the networks capture the underlying stochastic dynamics. Empirical evaluations on different problems illustrate the accuracy and scalability of our approach, demonstrating its effectiveness in solving complex high-dimensional stochastic control tasks. Code is available at <https://github.com/jdupret97/Deep-Learning-for-CT-Stochastic-Control-with-Jumps>

arXiv:2505.15602v3

DEPARTMENT OF MATHEMATICS, ETH ZURICH, SWITZERLAND
Email address: patrickc@ethz.ch

DYNAMIC HEDGING UNDER MODEL UNCERTAINTY

RAMA CONT

It is commonly assumed that a detailed and accurate description of the joint dynamics of risk factors and asset prices is a pre-requisite for the design of hedging and risk management strategies for derivatives portfolios. This has motivated the development of increasingly complex stochastic models with many risk factors and parameters, which are challenging to estimate and implement, but whose benefit in terms of risk management is not clear.

We revisit the dynamic hedging problem as a feedback control problem under model uncertainty. We argue that the widely used approach of continuously recalibrating an auxiliary pricing model is an effective way of treating model uncertainty and link this approach with Model Predictive Control (MPC) and reinforcement learning (RL). We show that any auxiliary pricing model capable of calibrating the cross section of market prices of hedging instruments and satisfying an identifiability condition may be used in this way to compute adequate robust dynamic hedging strategies without explicit knowledge of market dynamics. This procedure captures exposures to latent risk factors through parameter recalibration and leads to a self-financing strategy with predictable loss. We provide explicit formulas for hedge ratios in dynamically recalibrated models, and provide a detailed decomposition of the gain of the dynamically hedged position. We argue that this approach is an effective recipe for dynamic hedging under model uncertainty and pleads in favor of computationally exible models with good inversion properties. Our results provide insights for model design and clarify the role of parameter calibration in risk management.

UNIVERSITY OF OXFORD, UK

Email address: `rama.cont@maths.ox.ac.uk`

HOLDING STOCKS, TRADING BONDS

PAOLO GUASONI AND MARKO HANS WEBER

Classification AMS 2020: 91B51, 91G80.

Keywords: equilibrium, heterogeneous preferences, trading volume, term structure, continuous time.

The paper studies why rational investors actively trade long-term bonds even when information and beliefs are common. The economy contains long-lived agents with constant absolute risk aversion α_i , impatience β_i , initial wealth, and additive nontraded income. A stock in unit net supply pays an arithmetic dividend stream whose expected growth rate is mean reverting,

$$dD_t = \mu_t dt + \sigma_D dW_t^D, \quad d\mu_t = a(\bar{\mu} - \mu_t) dt + \sigma_\mu dW_t^\mu.$$

Agents trade the stock, borrow and lend at an endogenous short rate, and trade a consol in zero net supply. Personal income may load on dividend shocks, growth shocks, and idiosyncratic shocks. CARA preferences permit agents with different risk aversion, impatience, and income exposures to coexist indefinitely. More importantly, the constant-growth benchmark is a no-trade economy: after the initial allocation, the same heterogeneity produces no rebalancing.

Let $1/\bar{\alpha} = \sum_i 1/\alpha_i$, and let $\bar{\beta}$ be the corresponding risk-tolerance-weighted impatience rate. The equilibrium is available in closed form. For suitable deterministic functions f and p ,

$$(0.1) \quad r_t = \bar{r} + \bar{\alpha}(\mu_t - \bar{\mu}), \quad C_t = f(\mu_t), \quad P_t = D_t f(\mu_t) + p(\mu_t),$$

where C and P denote the consol and stock prices. Thus stochastic growth generates a Vasicek short rate endogenously. The market prices of traded Brownian risks are constant, while investment opportunities vary through the short rate and the prices of long-term claims.

The main result separates equity and fixed-income positions. Agent i holds a constant number of stock shares,

$$(0.2) \quad \hat{\theta}_t^{P,i} = \frac{\bar{\alpha}}{\alpha_i} + B_D^i - B_\mu^i \frac{\sigma_\mu \tilde{\rho}}{\sigma_D},$$

where B_D^i and B_μ^i measure the gaps between the agent's own income exposures and the agent's risk-tolerance-weighted shares of aggregate dividend and growth exposures. The stock allocates direct dividend risk and offsets the dividend-risk spillover caused by correlation between dividend and growth shocks. By contrast, the consol position depends on μ_t and is generically a diffusion. Growth shocks move interest rates and change the present value of heterogeneous consumption tilts and income hedges; agents therefore rebalance long-term bonds to restore the desired allocation. Deterministic coupons make bonds the targeted hedge: they transmit discount-rate and growth-state risk without adding direct contemporaneous dividend cash-flow risk. When $\sigma_\mu = 0$, both asset positions remain constant after date zero.

The result qualifies the no-trade conclusions of Judd, Kubler, and Schmedders [3] and Berrada, Hugonnier, and Rindisbacher [1]. The former relies on a dynamically complete stationary Lucas economy with common time preference, while the latter characterizes restrictions under which the equilibrium sharing rule remains compatible with zero trade. Stochastic growth and heterogeneous intertemporal or income exposures make the desired fixed-income exposure state dependent. The mechanism is also distinct from asymmetric-information or disagreement models: relevant aggregate risks are spanned, and

trade reflects equilibrium hedging. Relative to partial-equilibrium analyses of long-term bond demand [2], prices and trading policies are determined jointly in general equilibrium. The role of the consol also complements dynamic separation results in Schmedders [4].

A diffusion target has infinite total variation, so frictionless turnover is not directly observable. Let Θ^i be agent i 's consol target and let $\Theta^{i,\varepsilon}$ be the minimal finite-variation strategy that remains within a symmetric band of half-width ε . On every finite horizon $[0, T]$,

$$(0.3) \quad \varepsilon \text{ TV}(\Theta^{i,\varepsilon}; [0, T]) \longrightarrow \frac{1}{2} \langle \Theta^i \rangle_T \quad \text{in probability.}$$

Hence quadratic variation, rather than literal diffusion variation, determines leading-order executable volume. For one-sided aggregate volume in the zero-net-supply bond market, the leading term is $(4\varepsilon)^{-1} \sum_i \langle \Theta^i \rangle_T$.

The multivariate extension contains several stocks and m persistent growth factors. A single consol generally cannot span all fixed-income hedging directions. The paper therefore introduces m sinking funds with deterministic coupons $e^{-\kappa_d t}$. For $\kappa_d > 0$, the coupon-weighted average remaining maturity is $1/\kappa_d$, making these securities continuous-time analogues of constant-maturity bond funds. If $f_{\kappa_d}(\boldsymbol{\mu})$ is the price of fund d , the full-rank condition

$$\det D_{\boldsymbol{\mu}}(f_{\kappa_1}, \dots, f_{\kappa_m})(\boldsymbol{\mu}) \neq 0$$

ensures that the funds span all orthogonal growth shocks. Stock holdings remain constant, while sinking-fund holdings adjust dynamically. Zero-coupon yields are affine in the growth state; their factor dimension is the Krylov dimension generated by the short-rate loading and the mean-reversion matrix. Persistent directions produce level-like movements, faster directions produce slope-like movements, and directions outside that Krylov space are unspanned by risk-free yields. The structural conclusion is that active bond trading can arise from rational risk sharing itself, in an economy deliberately constructed to make trade difficult.

REFERENCES

- [1] T. Berrada, J. Hugonnier, and M. Rindisbacher. Heterogeneous preferences and equilibrium trading volume. *Journal of Financial Economics*, 83:719–750, 2007.
- [2] J. Y. Campbell and L. M. Viceira. Who should buy long-term bonds? *American Economic Review*, 91:99–127, 2001.
- [3] K. L. Judd, F. Kubler, and K. Schmedders. Asset trading volume with dynamically complete markets and heterogeneous agents. *Journal of Finance*, 58:2203–2229, 2003.
- [4] K. Schmedders. Two-fund separation in dynamic general equilibrium. *Theoretical Economics*, 2:135–161, 2007.

DUBLIN CITY UNIVERSITY, SCHOOL OF MATHEMATICAL SCIENCES, DUBLIN 9, IRELAND; AND UNIVERSITÀ DI BOLOGNA, DEPARTMENT OF STATISTICAL SCIENCES, 40126 BOLOGNA, ITALY

Email address: paolo.guasoni@dcu.ie

NATIONAL UNIVERSITY OF SINGAPORE, DEPARTMENT OF MATHEMATICS, SINGAPORE 119076

Email address: matmhw@nus.edu.sg

NEURAL CHAOS: MINIMAX-OPTIMAL APPROXIMATION OF PREDICTABLE PROCESSES WITHOUT ITERATED INTEGRALS OR FIXED BASES

ANASTASIS KRATSIOS

Predictable square-integrable processes are central in continuous-time stochastic control, reinforcement learning, and financial applications. Wiener–chaos-type expansions provide powerful representations, but their computational use is limited by large chaos dictionaries and high-order iterated stochastic integrals. We address this bottleneck with NeuralChaos, a finitely parameterized class of H_T -valued processes that uses only finitely many evaluations of the driving Brownian motion while preserving predictability and square-integrability by construction.

We prove that NeuralChaos achieves minimax-optimal nonlinear approximation rates, equivalently best N -term sparse approximation rates. For smooth and compressible predictable processes, these rates are sublinear, and we show that compressibility is generic: a process drawn from H_T according to a natural sub-Gaussian sampling model is compressible with high probability.

In contrast, we show that finite-dimensional Markovian neural-SDE-type parametrizations are not representative of H_T : their time-discretized predictable surrogates form a meagre and Gaussian-null subset of H_T . Numerical experiments on hedging and calibration problems for stochastic-volatility models illustrate the practical effectiveness of the NeuralChaos architecture.

MCMaster UNIVERSITY, CANADA
Email address: `kratsioa@mcmaster.ca`

Strategic trading when agents forecast the forecasts of others: existence and convergence

Kasper Larsen
Rutgers University

Joint work with Jin Hyuk Choi.

Kyle's model of informed trading [8] describes how a privately informed trader optimally exploits long-lived information while competitive market makers learn from aggregate order flow. Holden and Subrahmanyam [7] extend the model to several traders who all know the terminal value. Foster and Viswanathan [6] consider the substantially more intricate case in which each informed trader knows only one component of the terminal value. Each trader must then forecast not only the asset payoff, but also what the other informed traders infer from past order flow. This creates a tension between waiting, which preserves private information, and trading aggressively before competitors learn it. Foster and Viswanathan derive the linear equilibrium equations and study them numerically, but a general existence proof was not available. The purpose of the joint work summarized here is to prove existence for every finite number of trading rounds and to show that these discrete-time equilibria converge to the continuous-time equilibrium of Back, Cao, and Willard [2].

There are trading times n/N , $n = 1, \dots, N$, and an exogenous zero-interest bond. Noise traders submit independent Gaussian orders

$$\Delta w_n \sim \mathcal{N}(0, \sigma^2), \quad \sigma^2 = \frac{\sigma_w^2}{N}.$$

There are $M \geq 2$ risk-neutral informed traders. Their private signals $(\tilde{a}_1, \dots, \tilde{a}_M)$ are jointly Gaussian, independent of the noise orders, and satisfy

$$\mathbb{E}[\tilde{a}_i] = 0, \quad \mathbb{V}[\tilde{a}_i] = \sigma_a^2, \quad \text{cov}(\tilde{a}_i, \tilde{a}_j) = \rho \quad (i \neq j),$$

where

$$-\frac{\sigma_a^2}{M-1} < \rho < \sigma_a^2. \quad (1)$$

The terminal dividend is

$$\tilde{v} := \sum_{i=1}^M \tilde{a}_i. \quad (2)$$

Trader i initially observes only $\tilde{a}_i$. At trading time n , she also observes past aggregate order flow, but not the current noise order. If $\Delta\theta_{i,n}$ denotes her stock order, aggregate order flow is

$$\Delta y_n := \sum_{i=1}^M \Delta\theta_{i,n} + \Delta w_n, \quad y_0 = 0. \quad (3)$$

Competitive market makers impose the zero-profit pricing rule

$$p_n = \mathbb{E}[\tilde{v} \mid \sigma(y_1, \dots, y_n)], \quad p_0 = 0, \quad (4)$$

and trader i maximizes

$$\sup_{\theta_i \in \mathcal{A}_i} \mathbb{E} \left[\sum_{n=1}^N (\tilde{v} - p_n) \Delta\theta_{i,n} \mid \sigma(\tilde{a}_i) \right]. \quad (5)$$

The equilibrium is symmetric and linear. Its candidate processes have the form

$$\Delta\hat{\theta}_{i,n} = \beta_n(\tilde{a}_i - \hat{t}_{n-1}), \quad (6)$$

$$\Delta\hat{y}_n = \sum_{i=1}^M \Delta\hat{\theta}_{i,n} + \Delta w_n,$$

$$\Delta\hat{p}_n = \lambda_n \Delta\hat{y}_n, \quad \Delta\hat{t}_n = \zeta_n \Delta\hat{y}_n. \quad (7)$$

Here $\hat{t}$ is the public Gaussian estimate relevant for each trader's signal. To formulate Nash optimality, fix trader i and allow her to deviate. The other traders respond according to

$$\Delta\theta_{j,n} = \beta_n(\tilde{a}_j - \hat{t}_{n-1}), \quad j \neq i,$$

while market makers use $\Delta p_n = \lambda_n \Delta y_n$ and $\Delta t_n = \zeta_n \Delta y_n$. The deviating trader's relevant Markov states are

$$Y_{1,n} := \tilde{a}_i - \hat{t}_n, \quad Y_{2,n} := \hat{t}_n - t_n.$$

A quadratic value function in $(Y_{1,n}, Y_{2,n})$, together with Gaussian filtering, yields the optimizer

$$\Delta \theta_{i,n}^* = \beta_n (\tilde{a}_i - t_{n-1}) + \alpha_n (\hat{t}_{n-1} - t_{n-1}). \quad (8)$$

The second term records the strategic effect of trader i 's earlier deviations on the forecasts and future orders of the other traders.

The equilibrium coefficients can be encoded by a one-dimensional second-order recursion. For $y > 0$, set

$$y^+ := \frac{(M+1)y - 1 + \sqrt{4y + ((M+1)y - 1)^2}}{2}, \quad (9)$$

and define $c_0 := (M-1)(\rho - \sigma_a^2) < 0$. The main finite-horizon result is the following.

Theorem 1 (Existence of a discrete-time Nash equilibrium). *Assume $M \geq 2$, $N \in \mathbb{N}$, and (1). There is a constant $\Sigma_N > 0$ such that the backward recursion*

$$\Sigma_{N-1} = \Sigma_N^+, \quad (10)$$

$$\begin{aligned} & (\Sigma_n - \Sigma_{n+1}) \left((1 + (M+1)\Sigma_{n-1})\Sigma_n - (1 + \Sigma_{n-1})\Sigma_{n-1} \right)^2 \\ &= \Sigma_{n-1}\Sigma_{n+1}(\Sigma_{n-1} - \Sigma_n)(\Sigma_n + (M-1)\Sigma_{n+1})^2, \quad n = N-1, \dots, 1, \end{aligned} \quad (11)$$

with

$$\Sigma_n^+ > \Sigma_{n-1} > \Sigma_n, \quad n = N-1, \dots, 1,$$

produces the prescribed initial value

$$\Sigma_0 = \frac{M\sigma_a^2 + c_0}{-c_0} = \frac{\mathbb{V}[\tilde{v}]}{M(M-1)(\sigma_a^2 - \rho)}. \quad (12)$$

Let

$$\Sigma_{2,n} := (M-1)(\sigma_a^2 - \rho)\Sigma_n.$$

Then

$$\beta_n = \frac{\sigma \sqrt{\Sigma_{2,n-1} - \Sigma_{2,n}}}{\sqrt{M} \sqrt{\Sigma_{2,n-1}} \sqrt{\Sigma_{2,n}}}, \quad \zeta_n = \frac{\beta_n \Sigma_{2,n}}{\sigma^2}, \quad \lambda_n = M\zeta_n \quad (13)$$

produce a Nash equilibrium satisfying (4) and (8).

The proof is a discrete shooting argument. For a fixed terminal shot $\Sigma_N > 0$, equation (11) determines Σ_{n-1} from (Σ_n, Σ_{n+1}) . Although the equation is quartic in Σ_{n-1} , it has exactly one admissible root in (Σ_n, Σ_n^+) . The resulting backward solution depends continuously on the terminal shot. Varying Σ_N then produces the target value (12). This establishes the candidate coefficients. The remaining and technically most delicate step is verification of the informed traders' second-order conditions in the quadratic dynamic program.

The second main result connects the recursion to continuous time. Back, Cao, and Willard [2] construct their equilibrium from the decreasing solution of

$$\begin{cases} S''(t) = \frac{4(M-1)S(t) - 2}{MS(t)^2} S'(t)^2, & t \in (0, 1), \\ S(0) = \frac{\mathbb{V}[\tilde{v}]}{M(M-1)(\sigma_a^2 - \rho)}, & S(1) = 0. \end{cases} \quad (14)$$

The sign of S'' changes when $S = 1/(2(M-1))$, which gives the convex-to-concave information-revelation pattern conjectured by Foster and Viswanathan. The equation is singular at $t = 1$. Indeed, under the inverse substitution $S(f(x)) = x$, one obtains

$$f'(x) = C \exp\left(-\frac{2}{Mx}\right) x^{4/M-4}, \quad C < 0,$$

which also explains why analyticity at the terminal endpoint cannot be expected.

Theorem 2 (Convergence to the continuous-time equilibrium). *Let $\Sigma_0, \dots, \Sigma_N$ be the solution in Theorem 1, and let S solve (14). Then, for every $t \in [0, 1)$,*

$$\lim_{N \rightarrow \infty} \Sigma_{\lfloor tN \rfloor} = S(t). \quad (15)$$

Moreover, the discrete slopes converge locally away from the singular endpoint:

$$\lim_{N \rightarrow \infty} \frac{\Sigma_{\lfloor tN \rfloor + 1} - \Sigma_{\lfloor tN \rfloor}}{1/N} = S'(t), \quad t \in [0, 1).$$

Consequently, the discrete-time equilibrium coefficients converge to the coefficient functions of the Back–Cao–Willard equilibrium.

The convergence proof first establishes compactness of the interpolated recursions. A key estimate shows that, once $\Sigma_n < 1/(2M^2)$, the successive information reductions satisfy

$$\Sigma_{n-1} - \Sigma_n < \Sigma_n - \Sigma_{n+1},$$

and one obtains an N -uniform bound on $N(\Sigma_0 - \Sigma_1)$. Helly-type compactness then yields a subsequential limit. Expanding (11) with mesh $\Delta = 1/N$ identifies every such limit with (14); uniqueness of the ODE solution removes the need to pass to subsequences. The restriction $t < 1$ reflects the genuine singularity at the terminal boundary.

This convergence is not automatic in Kyle-type models. Caldentey and Stacchetti [3], for example, obtain a well-defined limit of discrete-time strategies that is not an equilibrium of the corresponding continuous-time model. Here the discrete shooting construction, the uniform estimates, and uniqueness for the limiting ODE together establish consistency between the Foster–Viswanathan and Back–Cao–Willard formulations.

References

- [1] K. Back, *Insider trading in continuous time*, Review of Financial Studies **5** (1992), 387–409.
- [2] K. Back, C. H. Cao, and G. A. Willard, *Imperfect competition among informed traders*, Journal of Finance **55** (2000), 2117–2155.
- [3] R. Caldentey and E. Stacchetti, *Insider trading with a random deadline*, Econometrica **78** (2010), 245–283.
- [4] J. H. Choi and K. Larsen, *Existence and convergence of discrete-time Kyle models with multiple insiders*, arXiv:2607.15057, 2026.
- [5] F. D. Foster and S. Viswanathan, *Strategic trading with asymmetrically informed traders and long-lived information*, Journal of Financial and Quantitative Analysis **29** (1994), 499–518.
- [6] F. D. Foster and S. Viswanathan, *Strategic trading when agents forecast the forecasts of others*, Journal of Finance **51** (1996), 1437–1478.
- [7] C. W. Holden and A. Subrahmanyam, *Long-lived private information and imperfect competition*, Journal of Finance **47** (1992), 247–270.
- [8] A. S. Kyle, *Continuous auctions and insider trading*, Econometrica **53** (1985), 1315–1335.

Existence of an equilibrium with limited stock market participation and power utilities

Kasper Larsen
Rutgers University

Joint work with Paolo Guasoni and Giovanni Leoni, and with Heeyoung Kwon.

Limited stock market participation is a parsimonious way to model the fact that some households do not hold stocks, whether because of information costs, institutional restrictions, or other frictions. Basak and Cuoco [1] introduced a continuous-time pure-exchange equilibrium with two infinitely lived traders: one trader can invest in both the stock and the money market, whereas the other can invest only in the money market. Their explicit construction assumes logarithmic utility for the restricted trader. This assumption is particularly convenient because the restricted trader's incomplete-market consumption problem becomes explicitly solvable. The purpose of the work described here is to establish equilibrium existence when both traders have power utility, and then to examine how heterogeneous time preferences change their long-run survival.

The aggregate dividend rate follows a geometric Brownian motion

$$dD_t = D_t(\mu_D dt + \sigma_D dB_t), \quad D_0 > 0, \quad \sigma_D > 0. \quad (1)$$

The endogenous money-market account and stock price satisfy

$$\begin{aligned} dS_t^{(0)} &= r_t S_t^{(0)} dt, \quad S_0^{(0)} = 1, \\ dS_t &= (r_t S_t - D_t) dt + S_t \sigma_{S,t} (\kappa_t dt + dB_t). \end{aligned}$$

Trader 1 chooses stock holdings and consumption, while trader 2 is excluded from the stock market and chooses only consumption financed through the money market. Their preferences are

$$\frac{1}{1-\gamma} \mathbb{E} \left[\int_0^\infty e^{-\beta_i t} c_{i,t}^{1-\gamma} dt \right], \quad i = 1, 2, \quad \gamma \in (0, 1). \quad (2)$$

A Radner equilibrium consists of optimal controls and endogenous price processes such that

$$\hat{c}_{1,t} + \hat{c}_{2,t} = D_t, \quad \hat{\theta}_{1,t} = 1,$$

with money-market clearing following from Walras' law. The central state variable is trader 1's consumption share

$$Y_t := \frac{\hat{c}_{1,t}}{D_t} \in (0, 1), \quad \hat{c}_{1,t} = D_t Y_t, \quad \hat{c}_{2,t} = D_t (1 - Y_t).$$

The equilibrium construction reduces the original coupled optimization and market-clearing problem to a one-dimensional singular boundary-value problem.

Identical time preferences. First suppose that $\beta_1 = \beta_2 = \beta$ and define

$$A := \frac{2\beta + \sigma_D^2 - (1-\gamma)(2\mu_D - \gamma\sigma_D^2)}{\sigma_D^2}. \quad (3)$$

The integrability condition

$$\beta > \frac{1}{2}(1-\gamma)(2\mu_D - \gamma\sigma_D^2) \quad (4)$$

corresponds to $A > 1$. The key analytic problem is to find $h \in C^1([0, 1])$ satisfying

$$\begin{aligned} h'(y) &= \frac{1+\gamma}{y} (\gamma - h(y)) \\ &+ \frac{h(y)}{1-y} \left[\gamma + \left((A-\gamma) \exp \left\{ \int_y^1 \frac{1-h(q)}{1-q} dq \right\} - A \right) h(y) \right], \quad y \in (0, 1), \end{aligned} \quad (5)$$

with the two singular boundary conditions

$$h(0) = \gamma, \quad h(1) = 1. \quad (6)$$

Equation (5) is Riccati-type, singular at both endpoints, and path-dependent through the integral of h . Its solution cannot be obtained from standard local ODE theory alone. The proof introduces a shooting parameter, establishes comparison and monotonicity in that parameter, and identifies the unique shot connecting the two boundary values. One obtains the bounds

$$\gamma \leq h(y) \leq 1, \quad y \in [0, 1], \quad h'(1) > 0.$$

The function h determines the consumption-share diffusion through

$$\mu_Y(y) := \frac{\sigma_D^2(1-y)(1+\gamma+2\gamma y h(y)-2y(1+\gamma))}{2yh(y)^2}, \quad (7)$$

$$\sigma_Y(y) := \sigma_D \frac{1-y}{h(y)}, \quad (8)$$

so that

$$dY_t = \mu_Y(Y_t) dt + \sigma_Y(Y_t) dB_t. \quad (9)$$

The endpoint behavior of h gives the asymptotic expansions needed for Feller's boundary test. In particular, pathwise uniqueness and strong existence hold, the boundary 0 is an inaccessible entrance boundary, and the boundary 1 is inaccessible, natural, and attracting. Indeed,

$$\lim_{t \rightarrow \infty} Y_t = 1 \quad \text{almost surely.} \quad (10)$$

The equilibrium short rate and market price of risk are Markovian functions of Y :

$$r(y) = \beta + \gamma\mu_D - \frac{1}{2}\gamma(\gamma+1)\sigma_D^2 - \frac{\gamma(\gamma+1)\sigma_D^2(1-y)}{2yh(y)^2}, \quad (11)$$

$$\kappa(y) = \gamma\sigma_D \left(1 + \frac{1-y}{yh(y)}\right). \quad (12)$$

The remaining price and portfolio processes are then recovered from the agents' first-order conditions and budget identities. The verification step requires both duality and an infinite-horizon transversality argument. The following summarizes the existence result of Guasoni, Larsen, and Leoni [3].

Theorem 1 (Equilibrium existence). *Let $\gamma \in (0, 1)$, $\sigma_D > 0$, and suppose (4) holds. For every initial restricted-trader money-market endowment in a nonempty admissible interval, there is a unique corresponding $Y_0 \in (0, 1)$ and a Radner equilibrium with consumption rates $D_t Y_t$ and $D_t(1 - Y_t)$, where Y solves (9), and with $r_t = r(Y_t)$ and $\kappa_t = \kappa(Y_t)$ given by (11)–(12).*

The limit (10) means that the restricted trader becomes extinct in consumption-share terms, even though her level of consumption need not converge to zero when aggregate dividends grow sufficiently fast.

Heterogeneous time preferences and survival. A joint extension with Heeyoung Kwon [6] allows $\beta_1 \neq \beta_2$ while retaining the common risk-aversion coefficient γ . Set

$$\delta := \frac{2(\beta_2 - \beta_1)}{\sigma_D^2}, \quad A := \frac{2\beta_2 + \sigma_D^2 - (1-\gamma)(2\mu_D - \gamma\sigma_D^2)}{\sigma_D^2}. \quad (13)$$

Relative to the equal-discount problem, the governing ODE acquires the cubic term

$$\frac{\delta y}{1-y} h(y)^2 \left(1 - \frac{h(y)}{\gamma}\right). \quad (14)$$

This term is singular at $y = 1$ and destroys the Riccati structure. Nevertheless, the shooting and comparison method can be adapted. If

$$\delta \in (-\gamma, 0), \quad A > 1 + \delta - \frac{2\delta}{\gamma}, \quad (15)$$

then the cubic boundary-value problem has a global C^1 solution satisfying $\gamma \leq h \leq 1$, and a Radner equilibrium exists. In this equilibrium,

$$r(y) = \beta_2 + y(\beta_1 - \beta_2) + \gamma\mu_D - \frac{1}{2}\gamma(\gamma+1)\sigma_D^2 - \frac{\gamma(\gamma+1)\sigma_D^2(1-y)}{2yh(y)^2},$$

$$\kappa(y) = \gamma\sigma_D \left(1 + \frac{1-y}{yh(y)}\right),$$

and the equilibrium consumption rates again equal $D_t Y_t$ and $D_t(1 - Y_t)$.

Following Kogan, Ross, Wang, and Westerfield [5] and Yan [9], trader 1 is called extinct if $Y_t \rightarrow 0$, and trader 2 is called extinct if $1 - Y_t \rightarrow 0$. Otherwise the trader survives. Under equal discounting, the restricted trader is extinct by (10). In contrast, the heterogeneous-discount equilibrium can support coexistence:

Theorem 2 (Long-run survival). *Under (15), if in addition*

$$\delta \in (-\gamma, -\gamma^2), \quad (16)$$

then both traders survive in the long run.

The proof uses the scale and speed measures of the diffusion Y . Under condition (16), the scale function diverges at both endpoints and the diffusion transformed to natural scale has a finite speed measure. Consequently, neither consumption share converges to zero. Economically, $\delta < 0$ means $\beta_1 > \beta_2$: the unrestricted trader is more impatient than the restricted trader. Sufficient patience can therefore offset the restricted trader's inability to hold the stock.

The analysis complements several extensions of Basak and Cuoco. Hugonnier [4] studies more general portfolio constraints and rational bubbles. Prieto [7] permits power utility for the unrestricted trader but retains log utility for the restricted trader. Chabakauri [2] studies heterogeneous preferences, beliefs, and portfolio constraints numerically. Weston [8] proves global existence with exponential preferences and general dividend streams through a system of quadratic backward stochastic differential equations. The present approach instead gives a global analytic construction for power utilities by solving a singular, path-dependent ODE and then classifying the boundaries of the endogenous consumption-share diffusion. Whether the resulting power-utility equilibrium stock price contains a bubble, as in the logarithmic Basak–Cuoco model, remains open.

References

- [1] S. Basak and D. Cuoco, *An equilibrium model with restricted stock market participation*, Review of Financial Studies **11** (1998), 309–341.
- [2] G. Chabakauri, *Asset pricing with heterogeneous preferences, beliefs, and portfolio constraints*, Journal of Monetary Economics **75** (2015), 21–34.
- [3] P. Guasoni, K. Larsen, and G. Leoni, *Existence of an equilibrium with limited stock market participation and power utilities*, Journal of Differential Equations **448** (2025), Article 113694.
- [4] J. Hugonnier, *Rational asset pricing bubbles and portfolio constraints*, Journal of Economic Theory **147** (2012), 2260–2302.
- [5] L. Kogan, S. A. Ross, J. Wang, and M. M. Westerfield, *The price impact and survival of irrational traders*, Journal of Finance **61** (2006), 195–229.
- [6] H. Kwon and K. Larsen, *Long-run survival in limited stock market participation models with power utilities*, arXiv:2512.14680, 2025.
- [7] R. Prieto, *Dynamic equilibrium with heterogeneous agents and risk constraints*, working paper, 2013.
- [8] K. Weston, *Existence of an equilibrium with limited participation*, Finance and Stochastics **28** (2024), 329–361.
- [9] H. Yan, *Natural selection in financial markets: Does it work?*, Management Science **54** (2008), 1935–1950.

DUAL APPROACHES TO STOCHASTIC CONTROL VIA SPDES AND THE PATHWISE HOPF FORMULA

MATHIEU LAURIÈRE

Keywords: Numerical method, stochastic control, duality theory, deep learning, generalized Hopf formula, deterministic control

We develop dual approaches for continuous-time stochastic control problems, enabling the computation of robust dual bounds in high-dimensional state and control spaces. Building on the dual formulation proposed in [L. C. G. Rogers, *SIAM Journal on Control and Optimization*, 46(2007), pp. 1116–1132], we first formulate the inner optimization problem as a stochastic partial differential equation (SPDE); the expectation of its solution yields the dual bound. Curse-of-dimensionality-free methods are proposed based on the Pontryagin maximum principle and the generalized Hopf formula. In the process, we prove the generalized Hopf formula, first introduced as a conjecture in [T. Y. Chow, J. Darbon, S. Osher, and W. Yin, *Journal of Computational Physics* 387 (2019), pp. 376–409], under mild conditions. Numerical experiments demonstrate that our dual approaches effectively complement primal methods, including the deep BSDE method for solving high-dimensional PDEs and the deep actor-critic method in reinforcement learning.

NYU SHANGHAI, CHINA

Email address: mathieu.lauriere@nyu.edu

PARTICLE METHOD FOR MEAN-FIELD CONTROL PROBLEMS WITH PARTIAL OBSERVATION

XIAOLU TAN

Classification AMS 2020: 65C35, 49M25, 93C41

Keywords: Optimal control, partial observation, mean-field, particle method

We present a general framework for the numerical approximation and theoretical analysis for solving continuous-time stochastic optimal control problems under partial observation, a class of problems where the decision-maker must act based on incomplete information about the underlying system state. The partially observed control problem is first reformulated as a special mean-field control problem where the control is required to be adapted to the common noise filtration. We then introduce a natural controlled particle system for the numerical approximation.

In a first step, we study in [1] a general mean-field control problem with common noise adapted controls, and establish the limit theory. Namely, we introduce a centralized finite-population control problems and prove its convergence to their mean-field limit with common-noise-adapted controls. Crucially, it establishes an optimal convergence rate using a BSDE characterization when only the drift is controlled, and further demonstrates C^1 and then C^2 -regularity of the value function in a Markovian setting, leading to a weak convergence rate of order $O(1/N)$. In a second step, we develop in [2] a practical, deep learning-based numerical method by leveraging a particle approximation to transform the infinite-dimensional filtering problem into a finite-dimensional, permutation-invariant control system. As a result, we obtain the convergence of the fully discretized particle scheme to the original continuous-time problem.

The core challenge of the partially observed control problem is that the problem's dynamic programming formulation is inherently infinite-dimensional, residing on the space of conditional distributions (filtering states) of the hidden process. The particle-based approximation methods replace the infinite-dimensional filtering state with a finite, weighted particle system. A key insight is the permutation-invariant structure of the resulting high-dimensional particle system, which is naturally exploited using symmetric neural network architectures like DeepSets. We consider two deep learning algorithms built upon this particle approximation: a direct optimization method that learns a symmetric feedback control policy by minimizing the expected cost, and a Deep BSDE method that solves the associated backward stochastic differential equation when the diffusion coefficient is uncontrolled. The framework is also extended to the more complex setting of partially observed mean-field control (MFC) problems, a class that has seen limited numerical development.

The practical utility and accuracy of the proposed methods are demonstrated through several numerical experiments. A linear-quadratic benchmark with a known analytical solution validates the approach, showing that the direct method achieves a relative

error of 0.04% in the optimal value, outperforming a recurrent neural network (RNN)-based baseline. The framework is further tested on a nonlinear partially observed MFC problem, where the results from the direct and Deep BSDE solvers cross-validate each other, confirming numerical consistency. Finally, two financial applications—portfolio liquidation and asset allocation—illustrate the method’s effectiveness and versatility in real-world decision-making contexts.

REFERENCES

- [1] Bruno Bouchard, Xiaolu Tan Limit theory for mean-field control problems with common noise adapted controls *ArXiv Preprint*, 2509.14734, 2025.
- [2] Mathieu Laurière, Xiaolu Tan, Jiefei Yang Particle Methods with Deep Learning for Stochastic Control under Partial Observation *ArXiv Preprint*, 2606.09055, 2026.

DEPARTMENT OF MATHEMATICS, THE CHINESE UNIVERSITY OF HONG KONG, SHATIN, N.T., HONG KONG
Email address: xiaolu.tan@cuhk.edu.hk

FLEXIBLE INFORMATION ACQUISITION IN THE KYLE MODEL

HAO XING

Keywords: Kyle model, information acquisition, mutual information, entropy regularized optimal transport

We study an information acquisition problem in which an informed trader acquires costly information prior to trading in the Kyle equilibrium. The cost of information acquisition is represented by an entropy cost. Regardless of the prior distribution of the asset payoff, continuous signals are optimal. Moreover, any continuously distributed signal, together with an associated logit-type posterior distribution of the payoff, yields the same ex-ante value for the informed trader, the same distribution of posterior expected payoff, and the same unconditional distribution of the informed traders trading strategy. Consequently, a normally distributed signal can be adopted without loss of generality. We further show that when the information acquisition cost increases or the volatility of noise trades decreases, the variance of the posterior expected payoff declines, the profit potential from trading diminishes, meanwhile the posterior expected payoff increasingly resembles a normal distribution, and the information leakage cost from trading decreases.

BOSTON UNIVERSITY, USA

Email address: haoxing@bu.edu