Heterogeneous beliefs and short selling taxes: A note

Michael Hatcher,* Sep 2024

Abstract

Short selling is widespread in financial markets but regulators can ban short positions. The intermediate policy of *taxing* short sellers has been studied in an asset pricing model with evolutionary competition of *two* belief types (Anufriev and Tuinstra, 2013). We extend this approach to an *arbitrary number* of belief types H, giving $3^H - 2^H$ cases to check each period in the worst-case scenario. We provide analytic expressions for asset prices along with conditions on beliefs (optimism) that determine which types take long, short or zero asset positions at the market-clearing price. We use these results to construct a fast solution algorithm (quadratic in H) which can solve models with hundreds or thousands of types in a matter of seconds. A numerical example with a short-selling tax and many heterogeneous beliefs in evolutionary competition shows that price dynamics can differ substantially relative to the benchmark of few types.

1 Introduction

Short selling is widespread in financial markets but is widely regulated by policymakers. When investors take a 'short' position, they borrow and immediately sell a financial asset before repurchasing and returning the asset to the lender, closing their position. Whereas a long position can be thought of as a bet that asset prices will increase, short-selling allows investors to bet on falling asset prices. It has been argued that such betting may increase financial market volatility, such as price downturns. A common policy response among market regulators has been to restrict or ban short selling; for example, many countries introduced short-selling bans following sharp declines in asset prices in the 2007-9 Financial Crisis. Similar short-selling bans were reintroduced in some European economies during the 2011-12 sovereign debt crisis and the Covid-19 outbreak (Siciliano and Ventoruzzo, 2020). Given the link to policy and outcomes in financial markets, it is important that researchers be able to solve asset pricing models with short-selling regulations in an efficient manner.

*Department of Economics, University of Southampton, SO17 1BJ, m.c.hatcher@soton.ac.uk.
I am grateful for financial support during the early stages of this project from the Economic and Social Research Council (ESRC), via the Rebuilding Macroeconomics Network (Grant Ref: ES/R00787X/1). I thank Tim Hellmann and an anonymous referee for providing useful feedback.

In this note, we take on this challenge. While some previous works consider a *full ban* on short selling, Anufriev and Tuinstra (2013) study the novel, intermediate measure of *taxing* (rather than banning) short selling in a model of two beliefs in evolutionary competition.[1] The problem of solving asset pricing models with short-selling constraints and *many* belief types is studied in Hatcher (2024), but only for the case of a full short-selling ban. Here, we extend both the analytical results and solution algorithm to the much harder problem of a short-selling *tax*. Our results enable a short-selling tax to be studied in models with a large number of competing beliefs, as in many real-world asset markets, since our algorithm can solve models with hundreds or thousands of belief types in a matter of seconds.

The difficulty with a short-selling tax is that the optimal asset demands are highly *non-linear* in the price: investors may decide to (i) buy the risky asset; (ii) take a negative position and pay the short-selling tax; or (iii) take a zero position if it is not beneficial after tax (in expected terms) to short sell. Accordingly, there are many additional cases into which investors can sort relative to the case of a full ban on short-selling, and finding a solution is computationally *expensive* – especially in a dynamic setting. Our results cut this computational cost and simulation times dramatically, thus making way for analyses of short-selling regulation in asset markets with many competing heterogeneous beliefs.

We build on the benchmark Brock and Hommes (1998) asset pricing model in which heterogeneous beliefs (or predictors) compete for investors via an evolutionary mechanism. A many-types version of the model is studied by Brock et al. (2005) when short-selling is unrestricted, and by Hatcher (2024) under a short-selling ban that precludes negative positions. In the latter case, some types are short-selling constrained if *belief dispersion* is large enough. Here, we consider the somewhat harder case of a short-selling *tax*, which is of policy relevance given the high costs and disruption caused by full bans on shorting,[2] and leads to sorting that depends on belief dispersion and the size of the tax penalty. We illustrate the utility of our results with a numerical example that demonstrates the speed and accuracy of our algorithm, and shows how heterogeneity of many beliefs types in evolutionary competition affects price dynamics when investors are taxed if they hold a short position.

Previous works studied short-selling constraints and few competing belief types (Anufriev and Tuinstra, 2013; in't Veld, 2016; Dercole and Radi, 2020). Recently, Hatcher (2024) extends the analysis to an *arbitrary number* of belief types – for a *ban* on shorting as in Dercole and Radi (2020) – and provides conditions on beliefs that determine the sets of unconstrained and short-selling constrained types at each date, plus a fast solution algorithm. In this note we follow a similar approach but take on the more difficult challenge of a short-selling *tax* and many belief types. The latter requires new analytical results and a novel solution algorithm for numerical simulations, and we show that price dynamics with many competing beliefs can differ substantially relative to the benchmark of few types.

[1] A short-selling constraint appears to have first been studied, in a static model, by Miller (1977).

[2] For example, Beber and Pagano (2013) study short-selling bans during the 2007-9 Financial Crisis and conclude that such bans significantly disrupted market liquidity and slowed down price discovery.

2 Model

Consider a finite set of myopic, risk-averse investor types $\mathcal{H} = \{h_1, ..., h_H\}$. At each date $t \in \mathbb{N}_+$, each type $h \in \mathcal{H}$ chooses a portfolio of a risky asset $z_{t,h}$ and a riskless bond with return $r > 0$ in order to maximize a mean-variance utility function over future wealth with risk-aversion parameter $a > 0$. The risky asset has current price p_t, future price p_{t+1}, and pays stochastic dividends d_{t+1}, which are exogenous and IID. Investors form subjective expectations of the future price and dividends of the risky asset as described below. The underlying model follows Brock and Hommes (1998), except that the risky asset is in positive net supply $\overline{Z} > 0$ and investors who short sell must pay a cost $T \in (0, \infty)$ per share as in Anufriev and Tuinstra (2013), such that date t transaction costs of a short seller are $T|z_{t,h}|$.

2.1 Asset demand

We denote the subjective expectation of type h at date t by $\tilde{E}_{t,h}[.]$, and the subjective variance by $\tilde{V}_{t,h}[.]$. The portfolio choice of type $h \in \mathcal{H}$ at date t solves the problem:

$$\max_{z_{t,h}} \tilde{E}_{t,h}[w_{t+1,h}] - \frac{a}{2}\tilde{V}_{t,h}[w_{t+1,h}] \tag{1}$$

where future wealth is $w_{t+1,h} = (1+r)w_{t,h} + (p_{t+1} + d_{t+1} - (1+r)p_t)z_{t,h} - (1+r)T|z_{t,h}|\mathbb{1}_{\{z_{t,h}<0\}}$, $\mathbb{1}_{\{z_{t,h}<0\}}$ equals 1 if $z_{t,h} < 0$ and 0 otherwise, and $\tilde{V}_{t,h}[w_{t+1,h}] = \sigma^2 z_{t,h}^2$, with $\sigma^2 > 0$.

Given the short-selling tax T, the date t demand of each investor type $h \in \mathcal{H}$ is

$$z_{t,h} = \begin{cases} \frac{\tilde{E}_{t,h}[p_{t+1}]+\overline{d}-(1+r)p_t}{a\sigma^2} & \text{if } p_t \leq p_t^h \\ 0 & \text{if } p_t \in (p_t^h, p_t^h + T] \\ \frac{\tilde{E}_{t,h}[p_{t+1}]+\overline{d}-(1+r)(p_t-T)}{a\sigma^2} & \text{if } p_t > p_t^h + T \end{cases} \tag{2}$$

where $p_t^h := \frac{\tilde{E}_{t,h}[p_{t+1}]+\overline{d}}{1+r}$ and $\tilde{E}_{t,h}[d_{t+1}] = \overline{d}$ is used.[3]

If the price p_t is small enough, type h's expected excess return is positive and they take a long position which decreases with the price; this is the standard demand function in Brock and Hommes (1998), where short-selling costs are absent. However, if the price is high enough to make the expected excess return of type h *negative*, they will choose either to short by taking a negative position and paying the tax (if the after-tax expected return is negative) or they will take a zero position in the risky asset (if shorting would be optimal *absent* a short-selling tax but is *not* optimal given the tax T).[4] From Equation (2) we see that types with sufficiently *pessimistic* price expectation, $\tilde{E}_{t,h}[p_{t+1}]$, will want to short sell and pay the tax, while more optimistic types will either buy or take a zero position.

[3] Dividends are $d_t = \overline{d} + \epsilon_t$, where $\overline{d} > 0$ and ϵ_t is IID, mean zero and has fixed variance. We assume all types know the dividend process, such that $\tilde{E}_{t,h}[d_{t+1}] = \overline{d}$ for all t, h as in Anufriev and Tuinstra (2013).

[4] In the special case $T \to \infty$, short selling becomes prohibitively costly, i.e. the case of a full ban.

2.2 Beliefs and population shares

To keep contact with the related literature, we assume that price beliefs and population shares are determined as in Brock and Hommes (1998); see Assumptions 1–2 below.

Assumption 1 *All price beliefs are of the form:*

$$\tilde{E}_{t,h}\left[p_{t+1}\right] = E_t[p^*_{t+1}] + f_h(x_{t-1}, ..., x_{t-L}) \tag{3}$$

where p^*_t is the fundamental price, E_t is the conditional expectations operator, $x_t := p_t - p^*_t$ is the deviation of price from the fundamental price, and $f_h : \mathbb{R}^L \to \mathbb{R}$ is a deterministic function that can differ across investor types h.

The price beliefs in Assumption 1 are boundedly-rational and do not depend on the current price p_t or any future values. Given our assumption of IID and mean-zero dividend shocks, the fundamental price is constant at $p^*_t = \overline{p} := \frac{\overline{d} - a\sigma^2\overline{Z}}{r}$, and hence $E_t[p^*_{t+1}] = \overline{p}$. The fundamental price $\overline{p}$ is lower than in Brock and Hommes (1998) because the risky asset is in positive net supply, $\overline{Z} > 0$.[5] By (3), beliefs in *deviations* from the fundamental price are

$$f_{t,h} := \tilde{E}_{t,h}\left[x_{t+1}\right] = f_h(x_{t-1}, ..., x_{t-L}) \tag{4}$$

where $\tilde{E}_{t,h}\left[x_{t+1}\right] := \tilde{E}_{t,h}\left[p_{t+1}\right] - E_t[p^*_{t+1}]$.

Thus, expressed in price deviations $x_t := p_t - \overline{p}$ and beliefs $f_{t,h}$, the demands in (2) are

$$z_{t,h} = \begin{cases} \frac{f_{t,h} + a\sigma^2\overline{Z} - (1+r)x_t}{a\sigma^2} & \text{if } x_t \leq x^h_t \\ 0 & \text{if } x_t \in (x^h_t, x^h_t + T] \\ \frac{f_{t,h} + a\sigma^2\overline{Z} - (1+r)(x_t - T)}{a\sigma^2} & \text{if } x_t > x^h_t + T \end{cases} \tag{5}$$

where $x^h_t := \frac{f_{t,h} + a\sigma^2\overline{Z}}{1+r}$ and $x^h_t + T$ are 'kink points' in the demand schedule of type h.

Aggregate demand for the risky asset is $\sum_{h\in\mathcal{H}} n_{t,h} z_{t,h}$, where $n_{t,h}$ is the population share of type h at date t and $\sum_{h\in\mathcal{H}} n_{t,h} = 1$. Following Brock and Hommes (1997, 1998), we assume the population shares are given by a discrete choice logistic model (Assumption 2), such that the population shares $n_{t,h}$ are endogenously determined and are time-varying but do *not* depend on the contemporaneous price x_t (or p_t) or any future values.

Assumption 2 *Population shares are updated using a discrete choice logistic model:*

$$n_{t+1,h} = \frac{exp(\beta U_{t,h})}{\sum_{h\in\mathcal{H}} exp(\beta U_{t,h})} \tag{6}$$

where $\beta \in [0, \infty)$ is the intensity of choice and $U_{t,h} \in \mathbb{R}$ is fitness of predictor h at date t.

[5] As in Brock and Hommes (1998) the fundamental price is the (hypothetical) price solution when all in investors are fundamentalists with common rational expectations and speculative bubbles are absent.

Assumption 2 says that the population share $n_{t+1,h} \in (0,1)$ of predictor h at date $t+1$ depends on the relative performance of predictor h against all other predictors $h_1, \ldots, h_H$, as judged by their past observed levels of fitness $U_{t,h}$. The intensity of choice parameter β determines how fast agents switch to better-performing predictors. In the special case $\beta = 0$, no switching occurs and population shares are fixed at $n_{t.h} = 1/H$ for all t and h. In financial market settings, the fitness measure $U_{t,h}$ is typically net trading profits.

2.3 Market clearing

The asset market is in equilibrium when the aggregate demand equals outside supply:

$$\sum\nolimits_{h\in\mathcal{H}} n_{t,h} z_{t,h} = \overline{Z} \qquad \text{subject to (5),(6).} \tag{7}$$

We now show how the price x_t and demands $z_{t,h}$ are determined – via the market clearing condition (7) – as a function of price beliefs and the short-selling tax, T. The resulting analytic results have a nice economic interpretation and are central to our algorithm.

3 Solving the model

Given positive outside supply $\overline{Z} > 0$, there exists a unique price x_t satisfying Equation (7) (see Anufriev and Tuinstra, 2013, Proposition 2.1). We first present our main analytical result in Proposition 1; we then use this result to build an algorithm for numerical simulations.

Proposition 1 *Let x_t be the market-clearing price at date $t \in \mathbb{N}_+$. Let $\mathcal{B}_t \subseteq \mathcal{H}$ be the non-empty set of buyers at date t, let $\mathcal{S}_{1,t} \subseteq \mathcal{H} \setminus \mathcal{B}_t$ ($\mathcal{S}_{2,t} = \mathcal{H} \setminus (\mathcal{B}_t \cup \mathcal{S}_{1,t})$) be the sets of zero-position types (short-sellers) and $\tilde{T}_t := (1+r)T \sum_{h\in\mathcal{H}\setminus\mathcal{S}_{1,t}} n_{t,h}$. Then the following holds:*

1. If $\sum_{h\in\mathcal{H}} n_{t,h}\,(f_{t,h} - \min_{h\in\mathcal{H}}\{f_{t,h}\}) \leq a\sigma^2\overline{Z}$, all types are buyers ($\mathcal{B}^*_t = \mathcal{H}$, $\mathcal{S}^*_{1,t} = \mathcal{S}^*_{2,t} = \emptyset$), demands are $z_{t,h} = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)x_t) \geq 0\ \forall h \in \mathcal{H}$, and the price is

$$x_t = \frac{\sum_{h\in\mathcal{H}} n_{t,h} f_{t,h}}{1+r} := x^*_t. \tag{8}$$

2. If $\sum_{h\in\mathcal{H}} n_{t,h}\,(f_{t,h} - \min_{h\in\mathcal{H}}\{f_{t,h}\}) > a\sigma^2\overline{Z}$, one or more types are non-buyers at date t (i.e. $\mathcal{B}^*_t \subset \mathcal{H}$, $\mathcal{S}^*_{1,t} \cup \mathcal{S}^*_{2,t} = \mathcal{H} \setminus \mathcal{B}^*_t \neq \emptyset$) and we have the following:

(i) If $\exists\ \mathcal{B}^*_t, \mathcal{S}^*_{1,t} = \mathcal{H} \setminus \mathcal{B}^*_t$ s.t. $\max\{d_{\mathcal{B}^*_t}, d_{\mathcal{S}^*_{1,t}}\} \leq a\sigma^2\overline{Z} < \sum_{h\in\mathcal{B}^*_t} n_{t,h}(f_{t,h} - \max_{h\in\mathcal{S}^*_{1,t}}\{f_{t,h}\})$, then $z_{t,h} = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)x_t) \geq 0\ \forall h \in \mathcal{B}^*_t$, $z_{t,h} = 0\ \forall h \in \mathcal{S}^*_{1,t}$, price is

$$x_t = \frac{\sum_{h\in\mathcal{B}^*_t} n_{t,h} f_{t,h} - (1 - \sum_{h\in\mathcal{B}^*_t} n_{t,h}) a\sigma^2\overline{Z}}{(1+r)\sum_{h\in\mathcal{B}^*_t} n_{t,h}} := \tilde{x}_t > x^*_t \tag{9}$$

where $d_{\mathcal{B}^*_t} := \sum\limits_{h\in\mathcal{B}^*_t} n_{t,h}(f_{t,h} - \min\limits_{h\in\mathcal{B}^*_t}\{f_{t,h}\})$, $d_{\mathcal{S}^*_{1,t}} := \sum\limits_{h\in\mathcal{B}^*_t} n_{t,h}(f_{t,h} - \min\limits_{h\in\mathcal{S}^*_{1,t}}\{f_{t,h}\}) - (1+r)T \sum\limits_{h\in\mathcal{B}^*_t} n_{t,h}$.

(ii) *If* $\exists\, \mathcal{B}_t^*, \mathcal{S}_{2,t}^* = \mathcal{H} \backslash \mathcal{B}_t^*$ *s.t.* $\sum_{h\in\mathcal{H}} n_{t,h}(f_{t,h} - \min_{h\in\mathcal{B}_t^*}\{f_{t,h}\}) \leq a\sigma^2\overline{Z} - (1+r)T\sum_{h\in\mathcal{S}_{2,t}^*} n_{t,h} < \sum_{h\in\mathcal{H}} n_{t,h}(f_{t,h} - \max_{h\in\mathcal{S}_{2,t}^*}\{f_{t,h}\}) - (1+r)T$, *then* $z_{t,h} = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)x_t) \geq 0\ \forall h \in \mathcal{B}_t^*$, $z_{t,h} = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)(x_t - T)) < 0\ \forall h \in \mathcal{S}_{2,t}^*$, *and price is*

$$x_t = \frac{\sum_{h\in\mathcal{H}} n_{t,h} f_{t,h} + (1+r)T\sum_{h\in\mathcal{S}_{2,t}^*} n_{t,h}}{1+r} := \hat{x}_t > x_t^* \tag{10}$$

(iii) *Else,* $\exists\, \mathcal{B}_t^*, \mathcal{S}_{1,t}^*, \mathcal{S}_{2,t}^* \neq \emptyset$ *s.t.* $\max\{d_{1,t}, \tilde{d}_{1,t}\} \leq a\sigma^2\overline{Z} - (1+r)T\sum_{h\in\mathcal{S}_{2,t}^*} n_{t,h} < \min\{d_{2,t}, \tilde{d}_{2,t}\}$, $z_{t,h} = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)x_t) \geq 0\ \forall h \in \mathcal{B}_t^*$, $z_{t,h} = 0\ \forall h \in \mathcal{S}_{1,t}^*$, $z_{t,h} = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)(x_t - T)) < 0\ \forall h \in \mathcal{S}_{2,t}^*$ *and price is*

$$x_t = \frac{\sum\limits_{h\in\mathcal{H}\backslash\mathcal{S}_{1,t}^*} n_{t,h} f_{t,h} + (1+r)T \sum\limits_{h\in\mathcal{S}_{2,t}^*} n_{t,h} - (\sum\limits_{h\in\mathcal{S}_{1,t}^*} n_{t.h}) a\sigma^2\overline{Z}}{(1+r)\sum_{h\in\mathcal{H}\backslash\mathcal{S}_{1,t}^*} n_{t,h}} := \overline{x}_t > x_t^* \tag{11}$$

where $d_{1,t} = \sum_{h\in\mathcal{H}\backslash\mathcal{S}_{1,t}^*} n_{t,h}(f_{t,h} - \min_{h\in\mathcal{B}_t^*}\{f_{t,h}\})$, $d_{2,t} = \sum_{h\in\mathcal{H}\backslash\mathcal{S}_{1,t}^*} n_{t,h}(f_{t,h} - \max_{h\in\mathcal{S}_{1,t}^*}\{f_{t,h}\})$, $\tilde{d}_{1,t} = \sum_{h\in\mathcal{H}\backslash\mathcal{S}_{1,t}^*} n_{t,h}(f_{t,h} - \min_{h\in\mathcal{S}_{1,t}^*}\{f_{t,h}\}) - \tilde{T}_t$, $\tilde{d}_{2,t} = \sum_{h\in\mathcal{H}\backslash\mathcal{S}_{1,t}^*}(f_{t,h} - \max_{h\in\mathcal{S}_{2,t}^*}\{f_{t,h}\}) - \tilde{T}_t$.

Proof. See the Supplementary Appendix. ■

Proposition 1 gives the market-clearing price and demands for an arbitrary number of belief types $H = |\mathcal{H}|$ facing a short-selling tax T and *evolutionary competition*. Since the proposition applies at any date $t \in \mathbb{N}_+$, we can find a solution recursively for $t = 1, 2, ...$, starting from period 1. As compared to a short-selling ban (see Hatcher, 2024), there are *four* rather than two parts because *in addition to* all types long (Part 1) or some long and some with zero positions (Part 2(i)), all non-buyers could short sell and pay the tax T (if they are sufficiently pessimistic; see Part 2(ii)), or there may be a mix among the non-buyers, with the less pessimistic at positions of zero and the more pessimistic choosing to short sell and pay the tax (see Part 2(iii)). In all these cases, the market-clearing price depends on the beliefs of investor types with non-zero positions – the market *participants*.

An important difference relative to Proposition 2.1 in Anufriev and Tuinstra (2013) (which shows *existence* of a unique market-clearing price) is that Proposition 1 gives explicit conditions on beliefs that determine, for an arbitrarily large number of belief types, the sets of buyers $\mathcal{B}_t^*$, non-participants $\mathcal{S}_{1,t}^*$, and taxed short-sellers $\mathcal{S}_{2,t}^*$ (see parts 2(i)–(iii)) and therefore the market-clearing price and demands.[6] Note that the asset prices in (9)–(11) are strictly larger than the no-tax solution x_t^*; hence, if one or more types are (taxed) short sellers or have zero positions, the price is higher than in absence of short-selling costs. Hence, like a short-selling ban (Miller, 1977), a binding short-selling tax raises the asset price.

[6] Anufriev and Tuinstra (2013, p. 1529) do not provide an explicit solution to this problem, as they note after Proposition 2.1 in their paper: "Note however, that x_t is still implicitly defined by (10) since the right-hand side also depends upon x_t through the definition of the sets $P(x_t)$, $Z(x_t)$ and $N(x_t)$. Below we will derive the market equilibrium price x_t explicitly for some special cases" (they study a two-type example).

The part of Proposition 1 we are 'in' depends on investor beliefs. If *belief dispersion* is small enough relative to risk-adjusted supply $a\sigma^2\overline{Z}$, all types are 'buyers' and the price is given by the usual solution x_t^* when short-selling costs are absent (Proposition 1, Part 1). Hence, if beliefs are 'close to homogeneous' there is some consensus on the price and no type wants to short or to not participate at date t. Conversely, if belief dispersion exceeds this threshold, at least one type (and at most $H-1$ types) will short sell or take a zero position; see Part 2. Both belief dispersion and the short-selling tax T determine which 'branch' gives the market-clearing price and demands at date t; see the inequalities in parts 2(i)–(iii) that pin down the *marginal* long, short and zero-position type(s) based on their optimism.

There are many permutations of types which correspond to the sets $\mathcal{B}_t, \mathcal{S}_{1,t}, \mathcal{S}_{2,t}$ that partition the set of investor types, $\mathcal{H}$, in Proposition 1.[7] The unique partition for which the asset market clears corresponds to the sets $\mathcal{B}_t^*, \mathcal{S}_{1,t}^*, \mathcal{S}_{2,t}^*$ that satisfy the marginal conditions (on beliefs) just discussed and uniquely determines the asset price and demands. As noted by Anufriev and Tuinstra (2013), there are $3^H - 2^H$ different cases (or regions) in total when types are unordered, making a solution *computationally expensive*: for $H = 5$ types there are 665 cases to check in the worst case scenario, and for $H = 15$ more than 14.3 million! As a result, obtaining a fast solution is a non-trivial problem for a large number of investor types, as seems plausible in many real-world asset markets.

We now show how the analytic results in Proposition 1 can be used – in conjunction with an *optimism ranking* – to reduce computational burden substantially, making simulations of models with hundreds or thousands of belief types tractable on a standard desktop or laptop computer. We start with an example that helps motivate our computational approach.

Example 1 *Suppose that there are $H = 2$ types, h_1 and h_2. Then by Proposition 1, there are $3^2 - 2^2 = 5$ different regions (or cases) and the equilibrium price is given by*

$$x_t = \begin{cases} \frac{\sum_{h\in\{h_1,h_2\}} n_{t,h} f_{t,h}}{1+r} & \text{if } -\frac{a\sigma^2\overline{Z}}{n_{t,h_2}} \le f_{t,h_1} - f_{t,h_2} \le \frac{a\sigma^2\overline{Z}}{n_{t,h_1}} \\ \frac{n_{t,h_1} f_{t,h_1} - (1-n_{t,h_1}) a\sigma^2\overline{Z}}{(1+r) n_{t,h_1}} & \text{if } \frac{a\sigma^2\overline{Z}}{n_{t,h_1}} < f_{t,h_1} - f_{t,h_2} \le \frac{a\sigma^2\overline{Z}}{n_{t,h_1}} + (1+r)T \\ \frac{n_{t,h_2} f_{t,h_2} - (1-n_{t,h_2}) a\sigma^2\overline{Z}}{(1+r) n_{t,h_2}} & \text{if } -\frac{a\sigma^2\overline{Z}}{n_{t,h_2}} - (1+r)T \le f_{t,h_1} - f_{t,h_2} < -\frac{a\sigma^2\overline{Z}}{n_{t,h_2}} \\ \frac{\sum_{h\in\{h_1,h_2\}} n_{t,h} f_{t,h} + (1+r)T n_{t,h_1}}{1+r} & \text{if } f_{t,h_1} - f_{t,h_2} < -\frac{a\sigma^2\overline{Z}}{n_{t,h_2}} - (1+r)T \\ \frac{\sum_{h\in\{h_1,h_2\}} n_{t,h} f_{t,h} + (1+r)T n_{t,h_2}}{1+r} & \text{if } f_{t,h_1} - f_{t,h_2} > \frac{a\sigma^2\overline{Z}}{n_{t,h_1}} + (1+r)T \end{cases} \tag{12}$$

which matches the result in Anufriev and Tuinstra (2013, Proposition 2.2).

Suppose that beliefs follow the two-type Brock and Hommes (1998) model: type h_1 is a fundamentalist with $\tilde{E}_{t,h_1}[p_{t+1}] = \overline{p}$, where $\overline{p} = (\overline{d} - a\sigma^2\overline{Z})/r$ is the fundamental price, and

[7] We use the term 'partition' loosely since $\mathcal{S}_{1,t}$ and $\mathcal{S}_{2,t}$ may be empty sets. In Part 1 of Proposition 1, both $\mathcal{S}_{1,t}$ and $\mathcal{S}_{2,t}$ are the empty set; in Part 2(i), $\mathcal{S}_{2,t}$ is the empty set; in Part 2(ii) $\mathcal{S}_{1,t}$ is the empty set.

h_2 is a 1-lag chartist: $\tilde{E}_{t,h_2}[p_{t+1}] = \overline{p} + \overline{g}x_{t-1}$, where $\overline{g} > 0$. Note that these beliefs imply that $f_{t,h_1} = 0$ and $f_{t,h_2} = \overline{g}x_{t-1}$; see (3)–(4) for reference. Assuming $x_{t-1} > 0$, the chartist h_2 is more optimistic at date t and must buy the asset. We can therefore narrow down to

$$x_t = \begin{cases} \frac{n_{t,h_2}\overline{g}x_{t-1}}{1+r} & \text{if } n_{t,h_2}\overline{g}x_{t-1} \leq a\sigma^2\overline{Z} \\ \frac{n_{t,h_2}\overline{g}x_{t-1}-(1-n_{t,h_2})a\sigma^2\overline{Z}}{(1+r)n_{t,h_2}} & \text{if } n_{t,h_2}(\overline{g}x_{t-1}-(1+r)T) \leq a\sigma^2\overline{Z} < n_{t,h_2}\overline{g}x_{t-1} \\ \frac{n_{t,h_2}\overline{g}x_{t-1}+(1+r)Tn_{t,h_1}}{1+r} & \text{if } n_{t,h_2}(\overline{g}x_{t-1}-(1+r)T) > a\sigma^2\overline{Z}. \end{cases} \tag{13}$$

i.e. three cases for the price and the sets $\mathcal{B}_t$, $\mathcal{S}_{1,t}$ $\mathcal{S}_{2,t}$.

In the above example, ordering the two types in terms of optimism reduces the number of cases to check from 5 to 3. The remaining three cases correspond to the sets $\mathcal{B}_t = \{h_1, h_2\}$, $\mathcal{S}_{1,t} = \mathcal{S}_{2,t} = \emptyset$ (h_1 and h_2 are buyers), $\mathcal{B}_t = \{h_2\}$, $\mathcal{S}_{1,t} = \{h_1\}$, $\mathcal{S}_{2,t} = \emptyset$ (h_2 buys, h_1 has a zero position), $\mathcal{B}_t = \{h_2\}$, $\mathcal{S}_{1,t} = \emptyset$, $\mathcal{S}_{2,t} = \{h_1\}$ (h_2 buys, h_1 short sells and pays the tax).

Since there are $3^H - 2^H$ cases given H types, it would be desirable to reduce the number of cases that need to be checked in any example that we face. We now show how the principle of *ranking* belief types in terms of optimism reduces the number of cases substantially, taking us from exponential in the number of types to *quadratic*. As a result, ranking types by optimism can drastically speed up discovery of the sets $\mathcal{B}_t$, $\mathcal{S}_{1,t}$ $\mathcal{S}_{2,t}$, and hence price and demands, when we have a large number of belief types. We then show how 'pruning' can further speed up the algorithm by eliminating further cases that do not need to be checked.

3.1 Algorithm

Ranking types by optimism is suggested by Anufriev and Tuinstra (2013). In the Appendix of their paper they present an algorithm for the case of many types that uses an arbitrary guess on the market-clearing price, followed by a procedure to progressively narrow down to a single region and price. The key advantage of the analytics in Proposition 1 is that the sets of buyers, zero and short-selling types are determined in terms of *beliefs* that do *not* depend on the endogenous market-clearing price x_t. As a result, we can construct a fast algorithm that can solve models with hundreds or thousands of belief types in a matter of seconds.

We start by presenting a 'simple' algorithm that is powerful because it reduces the number of cases from $3^H - 2^H$ (when all H types are unordered) to at most $H(H+1)/2$, i.e. quadratic in the number of types. To do so, we rank and re-label types according to their optimism in each period t. Thus, consider the function $v_t : \mathcal{H} \to \tilde{\mathcal{H}}_t$, where $\tilde{\mathcal{H}}_t := \{1, \ldots, \tilde{H}_t\}$ is an adjusted set of types with the property that the most optimistic type(s) in $\mathcal{H}$ get label $\tilde{H}_t$, the next most optimistic type(s) gets label $\tilde{H}_t - 1$, and so on, down to the least optimistic type(s) in $\mathcal{H}$ with label 1. Types with equal optimism get the *same* label, so $\tilde{H}_t \leq H$, which implies that $|\tilde{\mathcal{H}}_t| \leq |\mathcal{H}|$. In the case of ties in terms of optimism, the period t population share of the 'group' is the sum of the population shares of the individual types.[8]

Before presenting the algorithm, we introduce some notation that defines the measures of belief dispersion in Parts 1–2(iii) of Proposition 1 for the new set of types $\tilde{\mathcal{H}}_t$. Belief dispersion relative to the most pessimistic type, $\sum_{h\in\mathcal{H}} n_{t,h}\left(f_{t,h} - \min_{h\in\mathcal{H}}\{f_{t,h}\}\right)$, is given by $disp_{t,1} = \sum_{h=2}^{\tilde{H}_t} n_{t,h}(f_{t,h} - f_{t,1})$ (see Proposition 1, Parts 1–2). Analogously, belief dispersion relative to the marginal buyer and the marginal non-buyer are given by

$$disp_{t,k} := \sum_{h=k+1}^{\tilde{H}_t} n_{t,h}(f_{t,h} - f_{t,k}), \quad k \in \{1, ..., \tilde{H}_t - 1\}$$

where type k is either the most pessimistic buyer, the most optimistic non-buyer, or the most pessimistic non-buyer type (see Proposition 1 Part 2(i)).

Finally, for the other concepts of belief dispersion in Proposition 1 parts 2(ii)–2(iii), we define the following dispersion measures for use in our algorithm (again $k \in \{1, ..., \tilde{H}_t - 1\}$):

$$\hat{disp}_{t,k} := \sum_{h=1}^{\tilde{H}_t} n_{t,h}(f_{t,h} - f_{t,k}), \qquad \tilde{disp}_{t,k} := \sum_{h\notin[\underline{k},\overline{k}]} n_{t,h}(f_{t,h} - f_{t,k})$$

where $\underline{k}$ ($\overline{k}$) denotes the most pessimistic (most optimistic) zero-position type.[9]

We first present a simple benchmark algorithm before adding some speed improvements.

Algorithm 1

1. Find the set $\tilde{\mathcal{H}}_t$ and the population shares $n_{t,h}$ for $h = 1, ..., \tilde{H}_t$. Compute $disp_{t,1}$. If $disp_{t,1} \leq a\sigma^2\overline{Z}$, then $x_t = x_t^*$ is the date t price, compute the demands $z_{t,h} \geq 0$ for $h = 1, \ldots, \tilde{H}_t$ and move to period $t+1$ and repeat. If $disp_{t,1} > a\sigma^2\overline{Z}$, move to Step 2.

2. Guess there is 1 non-buyer. Check if $\max\{disp_{t,2}, disp_{t,1} - (1+r)T\sum_{h=2}^{\tilde{H}_t} n_{t,h}\} \leq a\sigma^2\overline{Z}$. If so, $x_t = \tilde{x}_t^{(1)}$ is the price (Proposition 1 Part 2(i), $\mathcal{S}_{1,t}^* = \{1\}$), compute the demands $z_{t,1} = 0$ and $z_{t,h} \geq 0$ for $h = 2, \ldots, \tilde{H}_t$ and move to period $t+1$. If the above condition is not met, check if $\hat{disp}_{t,2} + (1+r)Tn_{t,1} \leq a\sigma^2\overline{Z} < \hat{disp}_{t,1} - (1+r)T(1-n_{t,1})$. If so, $x_t = \hat{x}_t^{(1)}$ is the price (Proposition 1 Part 2(ii), $\mathcal{S}_{2,t}^* = \{1\}$), compute the demands $z_{t,1} < 0$ and $z_{t,h} \geq 0$ for $h = 2, \ldots, \tilde{H}_t$ and move to period $t+1$. Else, move to Step 3.

3. Guess $k = 2$ non-buyers. If $\max\{disp_{t,k+1}, disp_{t,k} - (1+r)T\sum_{h=k+1}^{\tilde{H}_t} n_{t,h}\} \leq a\sigma^2\overline{Z} < disp_{t,k}$, $x_t = \tilde{x}_t^{(k)}$ is the price (Proposition 1 Part 2(i), $\mathcal{S}_{1,t}^* = \{1, \ldots, k\}$), compute the demands $z_{t,1}, \ldots, z_{t,k} = 0$ and $z_{t,h} \geq 0$ for $h = k+1, \ldots, \tilde{H}_t$ and move to period $t+1$.

[8] From a computational perspective, grouping together types with equal optimism is not necessary and can impair computation speed. Nevertheless, we choose to present our approach under the 'grouping' assumption because it is conceptually simpler to have a strict ranking of beliefs, i.e. $f_{t,1} < f_{t,2} < \ldots < f_{t,\tilde{H}_t}$.

[9] Note: $\overline{k} \geq \underline{k}$. If the set of zero-position types $\mathcal{S}_{1,t}$ is a singleton, then $\underline{k} = \overline{k}$. Otherwise, $\overline{k} > \underline{k}$. The variables $\underline{k}$ and $\overline{k}$ are period-specific; however, to ease exposition we refrain from including a t subscript here.

Else if $\hat{disp}_{t,k+1}+(1+r)T\sum_{h=1}^{k} n_{t,h} \le a\sigma^2\overline{Z} < \hat{disp}_{t,k}-(1+r)T(1-\sum_{h=1}^{k} n_{t,h})$, $x_t = \hat{x}_t^{(2)}$ (Proposition 1 Part 2(ii), $\mathcal{S}_{2,t}^* = \{1,\dots,k\}$), compute the demands $z_{t,1},\dots,z_{t,k} < 0$, $z_{t,h} \ge 0$ for $h = k+1,\dots,\tilde{H}_t$ and move to period $t+1$. Else, move to Step 4.

4. Consider all *relevant partitions* of $\{1,\dots,k\}$ into $\mathcal{S}_{1,t}$ and $\mathcal{S}_{2,t}$ and let $\underline{k} := \min \mathcal{S}_{1,t}$, $\overline{k} := \max \mathcal{S}_{1,t}$ $(= k)$. Check if $\max\{\tilde{disp}_{t,\overline{k}+1}, \tilde{disp}_{t,\underline{k}} - \tilde{T}_t\} \le a\sigma^2\overline{Z} - (1+r)T\sum_{h=1}^{\underline{k}-1} n_{t,h} < \min\{\tilde{disp}_{t,\overline{k}}, \tilde{disp}_{t,\underline{k}-1} - \tilde{T}_t\}$ for each partition $\underline{k},\overline{k}$ (Proposition 1 Part 2(iii)).[10] If this condition is met, then $\mathcal{S}_{2,t}^* = \{1,\dots,\underline{k}-1\}$, $\mathcal{S}_{1,t}^* = \{\underline{k},\dots,\overline{k}\}$ and $\mathcal{B}_t^* = \{\overline{k}+1,\dots,\tilde{H}_t\}$, price is $x_t = \overline{x}_t^{(\underline{k},\overline{k})}$ (Proposition 1 Part 2(iii) for $\mathcal{S}_{1,t}^*, \mathcal{S}_{2,t}^*, \mathcal{B}_t^*$ above), compute demands $z_{t,1},\dots,z_{t,\underline{k}-1} < 0$, $z_{t,\underline{k}},\dots,z_{t,\overline{k}} = 0$, $z_{t,h} \ge 0$ for $h = \overline{k}+1,\dots,\tilde{H}_t$, move to period $t+1$.

5. If the condition in Step 4 is not met, increase k by 1 and repeat Steps 3–4 until a solution is found. Once a solution is found, move to period $t+1$ and repeat.

Algorithm 1 uses the adjusted set of types $\tilde{\mathcal{H}}_t$ in conjunction with Proposition 1. The case where all types are 'buyers' is first proposed as a solution; if this guess is rejected, the algorithm proceeds by guessing sequentially, starting from $k = 1$ non-buyers. The algorithm is 'simple' in the sense that it starts 'at the bottom' – by assuming there is one non-buyer (the most pessimistic type only) – and then increases the guess k in steps of 1 until a solution is found. Although this solution algorithm can easily be improved upon, it is *quadratic* in the number of types H in the worst-case scenario, which is a big improvement on the $3^H - 2^H$ cases when types are left unordered. Thus, even this 'simple' algorithm can deliver very large speed gains in numerical simulations with moderate or large numbers of types H.[11]

The worst-case is the maximum cases traversed by Algorithm 1 (naive) before finding a solution. Figure 1 shows there are $\tilde{H}_t(\tilde{H}_t+1)/2$ cases, where $\tilde{H}_t \le H$. There are $\tilde{H}_t$ initial 'branches', each representing a different number of non-buyers k (from 0 up to $\tilde{H}_t - 1$). The number of subsequent branches depends on the permutations of non-buyers into short-sellers and zero positions. For zero non-buyers (Algorithm 1 Step 1), we have an empty set and no subsequent branches (1 case); for 1 non-buyer (Algorithm 1 Step 2) there may be either 1 zero position or 1 short-seller (2 cases); for 2 non-buyers (Algorithm 1 Step 3–4), there may be 2 zero positions, 2 short-sellers, or 1 zero position and 1 short-seller (3 cases).[12] In general, for k non-buyers there are $k+1$ subsequent branches ($k+1$ cases), where $k \in \{1,\dots,\tilde{H}_t-1\}$, such that summing over all cases gives $\tilde{H}_t(\tilde{H}_t+1)/2$ cases in total, as shown in Figure 1.

It is instructive to consider some numerical examples. For $H = 5$, the number of cases is reduced from $3^5 - 2^5 = 665$ to at most $5(5+1)/2 = 15$, while for $H = 20$ we go from around 3.5 billion cases(!) to at most 210. These examples make clear the computational advantages

[10] Here, $\tilde{T}_t = (1+r)T\sum_{h\notin[\underline{k},\overline{k}]} n_{t,h}$. By 'relevant partitions' we mean all those partitions of $\{1,\dots,k\}$ for which $\min \mathcal{S}_{1,t} > \max \mathcal{S}_{2,t}$, i.e. $\mathcal{S}_{2,t} = \{1\}$, $\mathcal{S}_{1,t} = \{2,\dots,k\}$; $\mathcal{S}_{2,t} = \{1,2\}$, $\mathcal{S}_{1,t} = \{3,\dots,k\}$,..., up to $\mathcal{S}_{2,t} = \{1,\dots,k-1\}$, $\mathcal{S}_{1,t} = \{k\}$. Here, we use the fact that short-sellers must be more pessimistic in terms of their price beliefs than zero-position types; see, for example, Equation (5).

[11] Algorithm 2 (see below) is a faster algorithm that exploits more the analytical results in Proposition 1.

[12] Note that the less optimistic type, 1, cannot have a zero position if type 2 has a short position (see (5)).

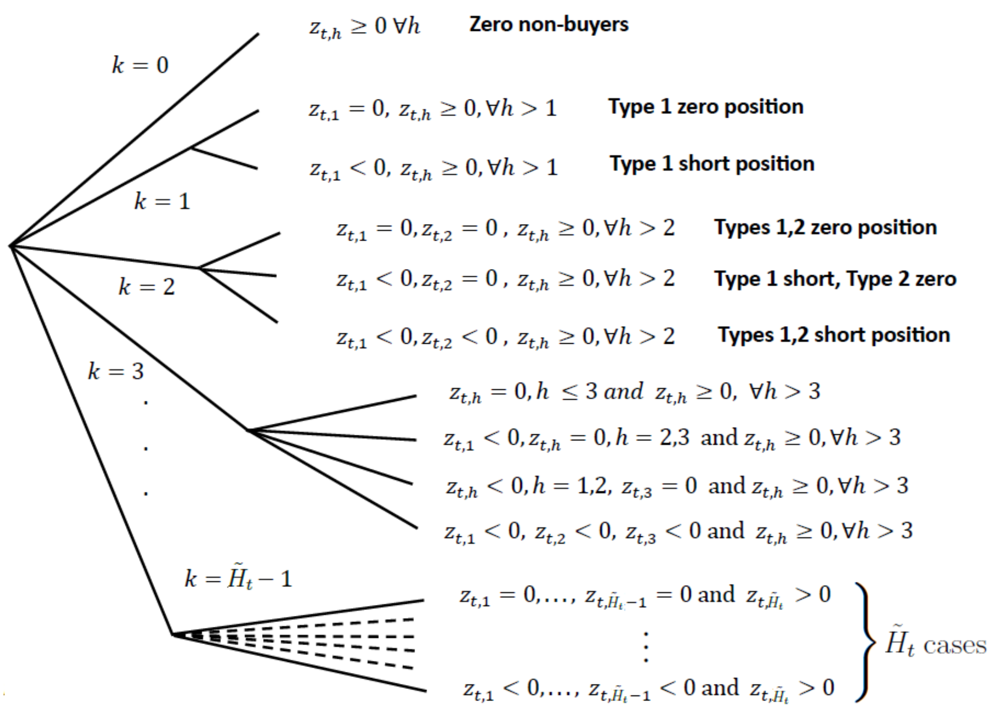

Figure 1: Tree diagram of the $\tilde{H}_t(\tilde{H}_t + 1)/2$ cases with ordered types, where $\tilde{H}_t \leq H$. Each main branch corresponds to a number of non-buyers (k); subsequent branches are sub-cases.

of ordering beliefs in terms of optimism as in Algorithm 1. However, for large numbers of belief types such as $H = 500$ or $H = 1,000$, there are still many cases and a further speed up is desirable. We now show how to do so by making further use of Proposition 1.

Algorithm 2 (fast)

1. Find the set $\tilde{\mathcal{H}}_t$ and the population shares $n_{t,h}$ for $h = 1, ..., \tilde{H}_t$. Compute $disp_{t,1}$. If $disp_{t,1} \leq a\sigma^2\overline{Z}$, then $x_t = x_t^*$ is the date t price, compute the demands $z_{t,h} \geq 0$ for $h = 1, \ldots, \tilde{H}_t$ and move to period $t+1$ and repeat. If $disp_{t,1} > a\sigma^2\overline{Z}$, move to Step 2.

2. Find the largest h such that $z_{t,h}^* = (a\sigma^2)^{-1}(f_{t,h} + a\sigma^2\overline{Z} - (1+r)x_t^*) < 0$, say k_0, where x_t^* is the price if a short-sales tax were absent (see Proposition 1 Part 1). If desired, k_0 may be updated in an iterative manner by updating the price and then updating k_0.[13]

3. Run Steps 3–5 of Algorithm 1, starting from $k = k_0$ (see Step 2). Continue until a solution is found, then move to period $t+1$ and repeat.

Algorithm 2 builds on Algorithm 1 by using 'pruning' – in place of sequential search starting from $k = 1$ non-buyers – in order to cut computation time. In particular, Step 2

[13]Our algorithm allows the user to 'turn on' such updating, and we find a non-trivial improvement in computation times for very large numbers of types, such as $H = 1,000$ or more.

avoids checking the first $k_0 - 1$ branches (i.e. $(k_0 - 1)k_0/2$ cases) when they cannot be market-clearing outcomes. The key point is that the market-clearing price is *higher* when there are some short sellers or zero-position types (see Proposition 1), so if some type is a non-buyer at the no-tax price x_t^*, they *must* be a non-buyer at the market-clearing price $x_t > x_t^*$, giving us a lower bound k_0 on the number of non-buying types. Algorithm 2 proceeds from this value to find the 'right' initial branch (i.e. the equilibrium no. of non-buyers), followed by the 'right' subsequent branch (the split between short and zero positions). Clearly, if Algorithm 2 eliminates many initial branches in Figure 1, computation time will be much reduced.

3.2 Discussion

We have set out analytical expressions that determine the sets of buyers, short-sellers and zero positions at market clearing as a function of beliefs and the short-selling tax. We used these results to construct a fast solution algorithm that is based on ordering types by optimism to eliminate irrelevant cases (Algorithm 1) plus additional *pruning* to avoid irrelevant equilibrium search strategies among the set of ordered types (Algorithm 2).

The approach used in Algorithms 1 and 2 has some similarity to the 'branch and bound' method used to solve integer or mixed-integer programming problems (Land and Doig, 1960). The branch and bound approach solves optimization problems by breaking them down into smaller sub-problems and then eliminating sub-problems (i.e. branches) that cannot contain the optimal solution. By comparison, we are able to eliminate certain permutations of the set of investor types into buyers, zero-position types or short-sellers because, given the known price beliefs of different types, some permutations cannot be consistent with equilibrium. The key to this result is that for any given price, demands cannot decrease as price beliefs increase (see (5)), so more optimistic types must have weakly larger equilibrium positions than less optimistic types, such that many 'branches' can be eliminated *a priori*.

Algorithm 1 uses this approach to reduce the number of branches (i.e. cases) at any given date t from $3^H - 2^H$ to at most $H(H+1)/2$ (for H investor types); this amounts to a big reduction in the number of cases for moderate or large values of H.[14] Algorithm 2 allows the remaining $H(H+1)/2$ cases to be cut further by using the no-tax price solution x_t^* (see (8)) to eliminate branches that contain too few non-buyers to be an equilibrium. In particular, Algorithm 2 uses two specific features of the problem at hand: (1) any types h whose untaxed demand is negative at price x_t^* must be non-buyers at the equilibrium price $x_t > x_t^*$ by Proposition 1 (i.e. have $z_{t,h} \leq 0$); (2) an updated price can be computed by counting the number of non-buyers in (1), and using this to generate an initial guess for the number of non-buying types in equilibrium (i.e. the k_0 in Step 3 of Algorithm 2).

In the next section we illustrate our algorithm using a numerical example. We report measures of computation speed and accuracy and compare against the 'standard case' of *no* short-selling tax where the price follows the no-tax solution x_t^*; see (8).

[14]Recall our previous numerical example: for $H = 20$, $3^H - 2^H \approx 3.5$ billion, while $H(H+1)/2 = 210$. Even for $H = 10$ the difference is very large: 58,025 cases (no ordering) versus 55 when using Algorithm 1.

4 Numerical example

We consider a version of the Brock and Hommes (1998) model with a large number of belief types and a short-selling tax. Accordingly, the demands of types $h \in \mathcal{H}$ are given by (5) and dividends $d_t = \overline{d} + \epsilon_t$ are IID with $\tilde{E}_{t,h}[d_{t+1}] = \overline{d}\ \forall t, h$. The fundamental price is $\overline{p} = (\overline{d} - a\sigma^2 \overline{Z})/r$, where $\overline{Z} > 0$ is the outside supply of the risky asset. The deviation from the fundamental price is $x_t := p_t - \overline{p}$. Investors have linear predictors of the form:

$$\tilde{E}_{t,h}\left[x_{t+1}\right] = b_h + g_h x_{t-1}. \tag{14}$$

Here, $b_h \in \mathbb{R}$ is the 'bias' of type h and $g_h \geq 0$ is their trend-following parameter. Type h is a pure fundamentalist investor if $b_h = g_h = 0$, while larger values of g_h or $|b_h|$ imply, respectively, stronger trend-following and stronger forecast bias.

Fitness $U_{t,h}$ is a linear function of past profits net of predictor costs $C_h \geq 0$. Profits at date t are given by the realized excess return $R_{t,h} := x_t - (1+r)(x_{t-1} - \mathbb{1}_{\{z_{t-1,h}<0\}} T) + a\sigma^2 \overline{Z} + \epsilon_t$ scaled by demand $z_{t-1,h}$, where ϵ_t is the IID dividend shock, and we abstract (for simplicity) from memory of past performance. For all $t > 1$ fitness and population shares are given by

$$U_{t,h} = R_{t,h} z_{t-1,h} - C_h, \qquad n_{t+1,h} = \frac{\exp(\beta U_{t,h})}{\sum_{h \in \mathcal{H}} \exp(\beta U_{t,h})}, \qquad \text{where } \beta \in [0, \infty). \tag{15}$$

Profits $U_{t,h}$ determine the population shares $n_{t+1,h}$ via a discrete-choice logistic model with intensity of choice β. The intensity of choice determines how fast agents switch to more profitable predictors. For $\beta = 0$ no switching occurs; increasing the value of β implies more switching to relatively profitable predictors. Following Brock and Hommes (1998), this profit-based *evolutionary competition* mechanism has been widely studied.

We use the same parameters as in Section 3.1 of Anufriev and Tuinstra (2013): $\overline{Z} = 0.1$, $a\sigma^2 = 1$, $r = 0.1$ and $\overline{d} = 10$, giving a fundamental price $\overline{p} = \frac{\overline{d} - a\sigma^2 \overline{Z}}{r} = 99$. In their model there are two types: a fundamentalist type with $\tilde{E}_{t,f}\left[x_{t+1}\right] = 0$ and cost $C = 1$, and a chartist type with $\tilde{E}_{t,c}\left[x_{t+1}\right] = \overline{g} x_{t-1}$, where $\overline{g} = 1.2$, and cost 0. We consider many types with predictors described by (14), population shares $n_{t,h}$ given by (15), and predictor costs C_h depending on the 'closeness' of beliefs to a pure fundamentalist (see above).

We start by looking at some individual numerical simulations, along with computation time and accuracy, before presenting a numerical bifurcation diagram. Figure 2 plots the price deviation x_t under *four* different scenarios labelled 1 to 4, both for a two-type model (as in Anufriev and Tuinstra (2013)) and for $H = 100$ heterogeneous belief types. In the latter case, we add heterogeneity in fundamentalist and chartist types by giving 50 types a trend-following parameter g_h linearly spaced on the interval $[1, 1.4]$, and the remaining 50 types a zero trend parameter $g_h = 0$ but bias b_h linearly spaced on the interval $[-0.2, 0.2]$ and cost $C_h = 1 - |b_h|$ which is decreasing in the bias of their predictor; these relatively small heterogeneities in beliefs have non-trivial implications for price dynamics. The four scenarios differ only in terms of the intensity of choice β and the short-selling tax is $T = 0.1$.

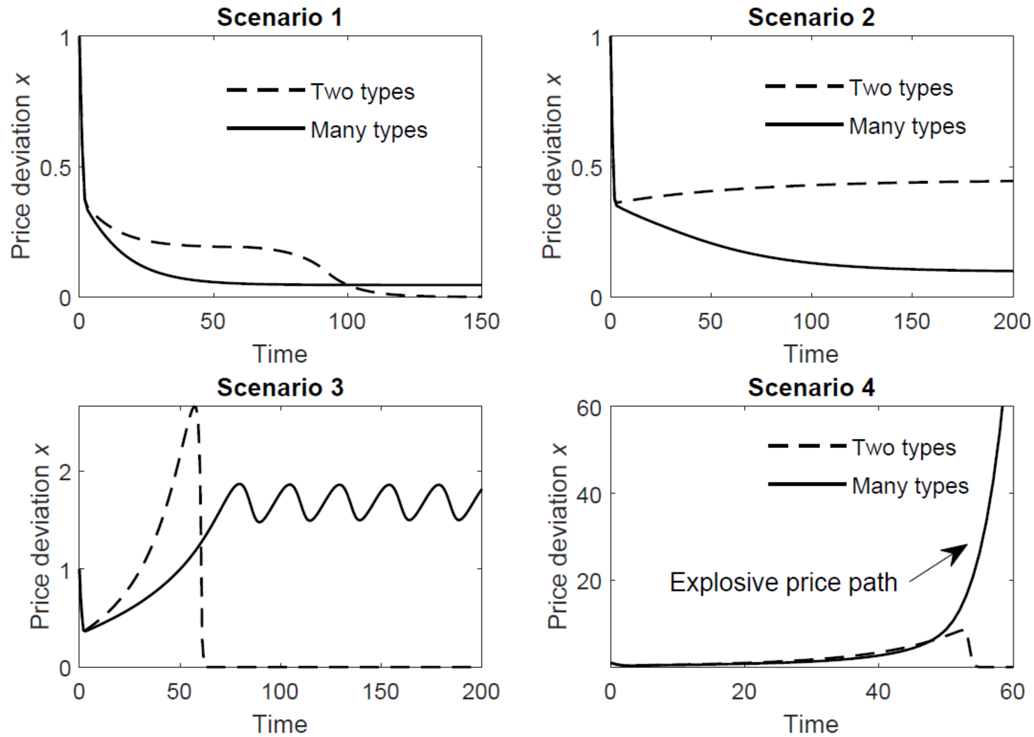

Figure 2: Four price scenarios: Two-type model vs many types ($H = 100$) when $T = 0.1$. Scenario 1: intensity of choice is $\beta = 1.6$; Scenario 2: intensity of choice is $\beta = 2$. Scenario 3 (4) increases the intensity of choice to $\beta = 2.5$ ($\beta = 2.9$). In each panel, the price deviation x_t is plotted at each t, given initial price $x_0 = 1$ and deterministic dividends $d_t = \overline{d}$ for all t.

Scenario 1 sets $\beta = 1.6$. We see that the initial overvaluation 'dies out' fairly quickly toward the fundamental price, with the main difference being the greater initial price persistence in the two-type model, due to the better relative performance of chartists with two types rather than many (Fig. 2, top left). In the many-types model, price settles at a small, positive non-fundamental price because positive bias fundamental types ($b_h > 0$) perform better. In Scenario 2, the intensity of choice increases to $\beta = 2$. Both prices now converge to a non-fundamental steady state with a sizeable overvaluation, with the latter being larger in the two-type model, reflecting the better relative performance of chartists. Scenario 3 increases the intensity of choice further to $\beta = 2.5$. In this case, price initially falls, but the trend is reversed. In the two type model, short-selling by fundamental types is discouraged by the tax and chartists outperform, so price starts increasing; however, there is sufficient short-selling by fundamental types to cause the bubble to 'burst' and price to converge near the fundamental price. With many types, the price deviation initially increases but we then see endogenous price cycles around a positive price. Finally, in Scenario 4, where the intensity of choice is $\beta = 2.9$, price converges near the fundamental price in the two-type model but diverges to $+\infty$ with many types, as belief heterogeneity means the most chartist types outperform enough to get a large population share, making price increases entrenched.

Table 1 reports computation times and accuracy for Scenario 3 with $T_{sim} = 100$ simulated periods and dividend shocks.[15] Computation times are quite fast – less than 0.5 seconds in all cases – but they increase sharply as the number of number of types H is increased, as expected for a *quadratic* algorithm. Note that both short-sellers *and* zero position types coexist in many periods (Column 4), such that our algorithm must search for the 'correct partition' between short-sellers and zero-position types (Algorithm 1, Step 4).[16] As compared to the case of *no* short-selling tax (where price is $x_t = x_t^* := \frac{\sum_{h\in\mathcal{H}} n_{t,h}\tilde{E}_{t,h}[x_{t+1}]}{1+r}$), computation times are increased somewhat but remain reasonable even for $H = 2,500$ types. Finally, the computed error in the final column is essentially zero as expected. The results for the other price scenarios in Figure 2 – see the Supplementary Appendix – tell a similar story.

Table 1: Computation times and accuracy in Scenario 3 (Algo. 2, $T_{sim} = 100$ periods)

No. of types	Short-selling tax	Time (s)	Freq. 1 (2)	$\max(Error_t)$
	No tax: $T = 0$	0.01	-	6.9e-17
$H = 100$	$T = 0.10$	0.02	100 (57)	7.6e-16
	$T = 1/8$	0.02	100 (41)	5.4e-16
	No tax: $T = 0$	0.02	-	8.3e-17
$H = 1,000$	$T = 0.10$	0.15	100 (60)	1.1e-15
	$T = 1/8$	0.11	100 (41)	1.0e-15
	No tax: $T = 0$	0.03	-	1.4e-16
$H = 2,500$	$T = 0.10$	0.48	100 (60)	1.1e-15
	$T = 1/8$	0.38	100 (41)	1.3e-15

Notes: $\max(Error_t) := \max\{Error_1, ..., Error_{T_{sim}}\}$, where we define the date t simulation error as $Error_t = |\sum_{h\in\mathcal{H}} n_{t,h}z_{t,h} - \overline{Z}|$. Demands $z_{t,h}$ depend on the computed market-clearing price. Freq. 1 = number of periods with $\mathcal{S}_{1,t}^* \cup \mathcal{S}_{2,t}^* \neq \emptyset$ (at least one short or zero position at date t), and Freq. 2 = number of periods with $\mathcal{S}_{1,t}^*, \mathcal{S}_{2,t}^* \neq \emptyset$ (both short and zero positions at date t).

The results in Table 1 raise the question: to what extent can the 'fast' computation times be attributed to ordering types by optimism (Algorithm 1) versus *additionally* finding a good guess for the number of non-buyers (Algorithm 2)? The short answer is that for moderate or large numbers of types H, Algorithm 1 is most important for reducing computation times, because sorting types by optimism reduces the number of cases from a very large number $3^H - 2^H$ to quadratic in H, which is a dramatic difference.[17] Indeed, searching all $3^H - 2^H$

[15] Dividend shocks ϵ_t were drawn from a truncated-normal distribution with mean zero, standard deviation $\sigma_d = 0.01$ and support $[-\overline{d}, \overline{d}]$. Simulations were run in Matlab 2023a (Windows version) on a Viglen Genie Desktop PC with an Intel(R) Core(TM) i7-7700 CPU 3.60 GHz processor and 16GB of RAM.

[16] If non-buyers are either *all* short-sellers or *all* with zero positions, rather than a mix of the two, then solution times are much faster (not shown); however, such cases do not fully test our solution algorithm.

[17] For example, for $H = 100$ the ratio $\frac{H(H+1)/2}{3^H - 2^H}$ is smaller than 1×10^{-44} (note: $3^H - 2^H = 5.15 \times 10^{47}$).

permutations of the set of investor types is not computationally tractable for large values of H, so Algorithm 1 is an essential component of Algorithm 2.

Table 2: Comparison of computation times: Algorithm 1 vs Algorithm 2

No. of types	Algo. 1 time (s)	Algo. 2 time (s)	Relative time gain
$H = 100$	0.09	0.02	77.8%
$H = 1,000$	8.25	0.09	98.9%
$H = 2,500$	146.4	0.28	99.8%

Notes: Scenario 3 with $T = 0.10$ and $T_{sim} = 100$ periods.

What the results in Table 2 show, however, is that Algorithm 2 still obtains sizeable extra reductions in computation time by finding a good initial guess for the number of non-buyers, such that simulations take a small fraction of the Algorithm 1 times.[18] We conclude that both the mass removal of branches in Algorithm 1 and the refinement in Algorithm 2 are crucial for a fast solution algorithm which makes it computationally-tractable to study simulations of many periods or bifurcation-type analyses to changes in parameter values.

Figure 3 plots a numerical bifurcation diagram as the intensity of choice β is increased; the attractors are based on *negative* initial prices $x_0 < 0$ and a particular short-selling tax, T. The price attractor for the two-type model is shown in grey, and the attractor from the many-type model ($H = 20$ types) is in black. The short-selling tax is set at $T = 0.1$.

In the *two-type* model, the fundamental steady state $x = 0$ is the only price attractor for sufficiently low intensity of choice β (grey attractor). Once a critical value of β is exceeded, the price converges to a non-fundamental steady state with negative price (i.e. undervaluation); hence the lower 'fork' seen for β between (approx.) 2.4 and 3.2 in Figure 3. Increasing β further causes the non-fundamental steady states to lose stability through a secondary bifurcation, leading to endogenous price fluctuations. The results for the two-type model are consistent with the results and bifurcation diagrams in Anufriev and Tuinstra (2013), and we found similar attractors for alternative values of the short-selling tax T.[19]

The attractor in the *many-types* model is quite different (black). For small positive values of the intensity of choice we see convergence on a non-fundamental steady state price with a small overvaluation; intuitively, this is because the fundamentalists with positive bias are best performers, such that their optimistic predictor gains a foothold in the market. As the intensity of choice β increases, price converges to a larger steady-state price since the foothold is stronger when there is more switching in response to performance. Increasing β further leads to endogenous price fluctuations which, however, exist only for a small range of β values (approx. 2.4 to 2.7). Once the intensity of choice is sufficiently high, there is a negative steady-state price (undervaluation) whenever price converges. By comparison to

[18] Slightly different hardware used relative to Table 1, so Algorithm 2 computation times are slightly lower.

[19] Figure 7 in their paper is a bifurcation diagram in the two-type model for $T = 0.1$, and their Figure 8 is a numerically-computed bifurcation diagram for $T = 0.1$, analogous to the grey attractor in Figure 3.

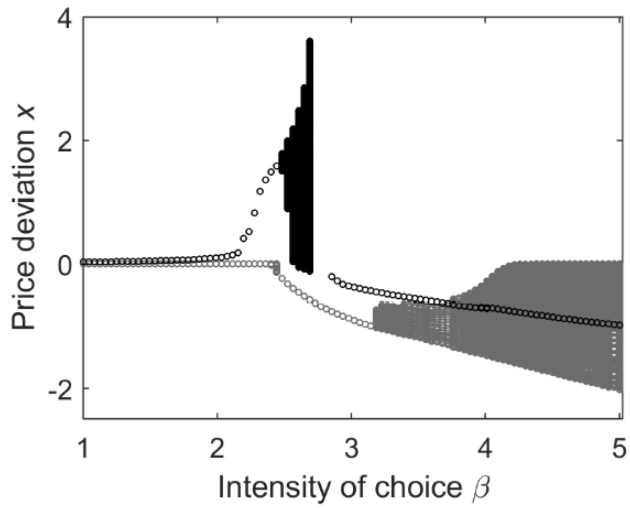

Figure 3: Numerical bifurcation diagram: two-type model (grey attractor) and many-type model ($H = 20$, black attractor) when $T = 0.1$. For each β, we ran 30 simulations from initial prices $x_0 \in (-4, 0)$ and we plot a total of 2,400 points after 3,120 transitory periods.

the two-type model, the price attractor is rather 'sparse' at higher values of the intensity of choice because a high percentage of simulations have explosive price paths, so there is an extreme form of volatility not seen in the two-type model where price paths were bounded.[20]

In short, there are non-trivial qualitative differences in the price dynamics with many versus few types, including a switch in sign of the price deviation and explosive simulations.

5 Conclusion

In this note we studied asset pricing in behavioural heterogeneous-belief models with a short selling *tax* and many belief types. We provided analytic expressions for asset prices and conditions on beliefs that determine which types take long, short or zero asset positions at the market-clearing price. These results allowed us to construct an algorithm that can solve models with hundreds or thousands of heterogeneous beliefs in a matter of seconds.

We illustrated the utility of these results using a numerical example with many different belief types in *evolutionary competition*, inspired by the two-type model in Anufriev and Tuinstra (2013). We extended the model to allow many heterogeneous beliefs around the polar chartist and fundamentalist beliefs, and we found that even small heterogeneities can have substantive implications for price dynamics and financial market volatility.

An interesting question is whether our approach could be applied to models with non-smoothness or discontinuities in demand for reasons other than short-selling regulations. Examples include market entry from the crossing of price misalignment thresholds or price

[20]For β between approx. 2.7 and 2.8, we found price was explosive in all simulations (many-types model).

beliefs that depend on such thresholds (see Tramontana et al., 2010, 2015). Extending such models to a large number of belief types might reveal new insights about effectiveness of regulatory policies or the empirical performance of this class of models.

References

Anufriev, M. and Tuinstra, J. (2013). The impact of short-selling constraints on financial market stability in a heterogeneous agents model. *Journal of Economic Dynamics and Control*, 37(8):1523–1543.

Beber, A. and Pagano, M. (2013). Short-selling bans around the world: Evidence from the 2007–09 crisis. *The Journal of Finance*, 68(1):343–381.

Brock, W. A. and Hommes, C. H. (1997). A rational route to randomness. *Econometrica: Journal of the Econometric Society*, pages 1059–1095.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22:1235–1274.

Brock, W. A., Hommes, C. H., and Wagener, F. O. (2005). Evolutionary dynamics in markets with many trader types. *Journal of Mathematical Economics*, 41(1-2):7–42.

Dercole, F. and Radi, D. (2020). Does the "uptick rule" stabilize the stock market? insights from adaptive rational equilibrium dynamics. *Chaos, Solitons & Fractals*, 130:109426.

Hatcher, M. (2024). Solving heterogeneous-belief asset pricing models with short-selling constraints and many agents. *Macroeconomic Dynamics*, pages 1–24.

in't Veld, D. (2016). Adverse effects of leverage and short-selling constraints in a financial market model with heterogeneous agents. *Journal of Economic Dynamics and Control*, 69:45–67.

Land, A. and Doig, A. (1960). An automatic method of solving discrete programming problems. *Econometrica: Journal of the Econometric Society*, pages 497–520.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of Finance*, 32(4):1151–1168.

Siciliano, G. and Ventoruzzo, M. (2020). Banning cassandra from the market? an empirical analysis of short-selling bans during the covid-19 crisis. *European Company and Financial Law Review*, 17(3-4):386–418.

Tramontana, F., Westerhoff, F., and Gardini, L. (2010). On the complicated price dynamics of a simple one-dimensional discontinuous financial market model with heterogeneous interacting traders. *Journal of Economic Behavior & Organization*, 74(3):187–205.

Tramontana, F., Westerhoff, F., and Gardini, L. (2015). A simple financial market model with chartists and fundamentalists: market entry levels and discontinuities. *Mathematics and Computers in Simulation*, 108:16–40.