Simple Market Protocols for Efficient Risk Sharing

Marco LiCalzi and Paolo Pellizzari
Dept. of Applied Mathematics
University of Venice

SIAM Meeting
11 July 2005 - New Orleans
http://www.dma.unive.it/∼paolop

Overview

We study the performance of four market protocols:

– batch auction, continuous double auction, specialist, hybrid.

We evaluate them by running (computerized) experiments.

Overview

We study the performance of four market protocols:

– batch auction, continuous double auction, specialist, hybrid.

We evaluate them by running (computerized) experiments.

All protocols achieve an efficient allocation in finite time.

Overview

We study the performance of four market protocols:

– batch auction, continuous double auction, specialist, hybrid.

We evaluate them by running (computerized) experiments.

All protocols achieve an efficient allocation in finite time.

We refine the comparison using other performance criteria.

And the winner is...

The model

We identify three components for our (simulated) exchange markets:

— the environment (including agents' preferences and endowments);

— the market protocols;

— the behavioral assumptions.

The environment

There are two assets: stock and cash.

The rate of interest is zero. Cash is the numeraire.

The stock pays no dividends and has a final (random) value Y.

There are n traders, who believe that Y is normally distributed with mean $\mu \geq 0$ and precision $\tau = 1/\sigma^2 > 0$.

No new information is ever released. Therefore, traders' beliefs about Y are homogeneous and never change until uncertainty resolves.

Each trader has an initial endowment of cash $c_i \geq 0$ and shares $s_i \geq 0$.

Each trader has "cara" preferences over his final wealth, with a coefficient of risk tolerance $k_i > 0$.

Trader i's excess demand function for stock (net of his endowment s_i) is linear and decreasing:

$$q_i(p) = k_i \tau(\mu - p) - s_i. \quad (1)$$

The total amount of cash and stock is $C = \sum_i c_i > 0$ and $S = \sum_i s_i > 0$.

Let $K = \sum_i k_i$. The unique efficient risk-sharing allocation requires that trader i holds $s_i^* = (S/K)k_i$ shares of the stock; that is, the efficient allocation is proportional to the coefficient of risk tolerance.

The competitive equilibrium

The competitive equilibrium achieves the efficient allocation.

The zero aggregate excess demand condition implies

$$p^* = \mu - \frac{S}{\tau K}. \tag{2}$$

At price p^*, trader i's net demand

$$q_i(p^*) = \left(\frac{S}{K}\right) k_i - s_i$$

is exactly filled and his final allocation $q_i(p^*) + s_i$ equals $s_i^* = (S/K)k_i$.

How can we reasonably implement the (equivalent of a) competitive equilibrium? (Just imagine Walrasian tâtonnement.)

The market protocols

We compare the performances of four *simple* market protocols:

— the batch auction (B);

— the continuous double auction (C);

— a nondiscretionary specialist dealership (D);

– a hybrid protocol (H).

The first protocol is simultaneous, the other three are sequential.

Common features

A protocol is organized in trading sessions.

Agents can exchange at most one unit in each trading session.

The sequence in which agents place their orders is randomly chosen for each trading session. Each agent selects randomly one side of the market where he attempts to place a trade.

Prices are quoted on a finite grid.

The books are completely cleared at the end of each trading session.

Batch auction

In each session, after traders submit their orders, the exchange price p^* is determined by equating the aggregate excess demand to zero.

Shares and payments are exchanged between traders who submitted bids not lower than p^* and asks not higher than p^*.

Traders who placed orders exactly at price p^* may be rationed.

This protocol is also known as the k-double auction, with $k = 1/2$.

Continuous double auction

In each session, traders place their orders on the selling and buying books.

Orders are immediately executed if they are marketable; otherwise, they are recorded on the books with the usual price-time priority.

Orders are canceled only when a matching order arrives or the trading day is over.

Automated dealership

A specialist dealer posts bids and asks valid for a unit transaction.

Agents check sequentially the dealer's quotes.

If an agent accepts the dealer's quote, the exchange takes place at the quoted price.

After a transaction is completed, the two dealer's quotes for bid and ask move by one in the obvious direction.

The size of the bid-ask spread stays fixed over time, so the price is never unique.

Hybrid market

Separate books hold quotes from the specialist and from the public.

The dealer revises her quotes as in the automated dealership.

Agents check sequentially the books.

Their orders are executed at the best price available (possibly different from the specialist's) if they are marketable; otherwise, they are recorded on the traders' book with the usual price-time priority.

Behavioral assumptions

We concentrate on the characteristics of the protocols, by making general-purpose assumptions on traders' behavior.

There are three established assumptions in the literature:

1) traders are restricted to trade one unit at a time;

2) budget constraints hold;

3) each trader has a constant valuation for each unit traded.

We maintain the first two assumptions, but relax the third one.

The demand function $q_i(p) = k_i\tau(\mu - p) - s_i$ of a trader is decreasing.

When the current endowment is s_i, we invert the demand function and derive a valuation for the next unit to trade as

$$p(\pm 1) = \mu - \frac{s_i \pm 1}{k_i\tau}, \quad (3)$$

depending on which side of the transaction he decides to attempt.

We assume that each trader chooses either side with equal probability.

Truthtelling (TT): a trader offers a price equal to his valuation.

Alternative behavioral assumptions

We test the robustness of our conclusions under two different sets of behavioral assumptions, representing different tradeoffs between immediacy and efficacy.

Zero intelligence (GS): a trader bids a random price never worse than his valuation; f.i., a buyer bids a price uniformly drawn from $[0, p(+1)]$.

Mental accounting (MA): People show higher propensity to risk in the presence of a prior gain. We accordingly modify traders' offers when their past transactions have been profitable.

At time t, the valuation of a buyer is his price $p(+1)$ increased by his current trading surplus G. The valuation of a seller is $p(-1) - G$.

The exemplar case

	Parameters		Initialization
Global	n	$=$	$1,000$
	μ	$=$	$1,000$
	σ^2	$=$	120
	t	$=$	2,500
	Δ	$=$	1
Trader	k_i	$=$	divisors of σ^2 in $[10, 40]$
	c_i	$=$	$50,000$
	s_i	$=$	permutation of $2k_i$

The equilibrium price is $p^* = \mu - 2\sigma^2 = 760$ for all simulations.

The simulation *exactly implements* TT.

Exact implementation

Assume TT. Valuations for one unit obey $p(\pm 1) = \mu - \frac{s_i \pm 1}{k_i \tau}$.

Since the tick is $\Delta = 1$, ticked prices must be integers.

When traders have integer valuations, the market protocol need not round bids and asks to a ticked price.

The most important consequence of exact implementation is that *exact* convergence to the equilibrium price is possible.

The appeal of this is mostly theoretical, but experimental testing should approximate ideal conditions.

The distribution of the coefficients of risk tolerance is independent of the initial endowments.

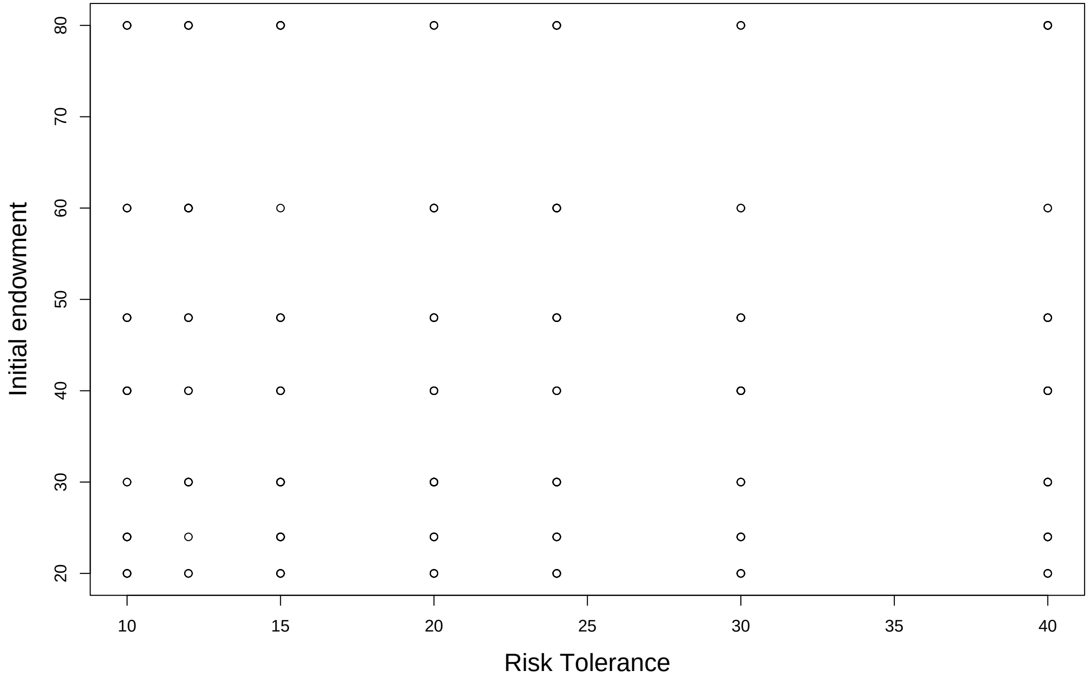

The performance criteria

1) Convergence to the efficient allocation: "do we get there?".

2) Speed of convergence: "how fast we get there?".

3) Distance from the efficient allocation: "what happens before we get there?"

4) Traded volume: "how many wasteful trades"?

5) Average certainty equivalent: "is any wealth lost"?

6) Volatility (and kurtosis) of prices: "how wild is convergence?"

The efficient allocation is proportional to the coefficient of risk tolerance. Assuming sufficient liquidity, all simulations converge.

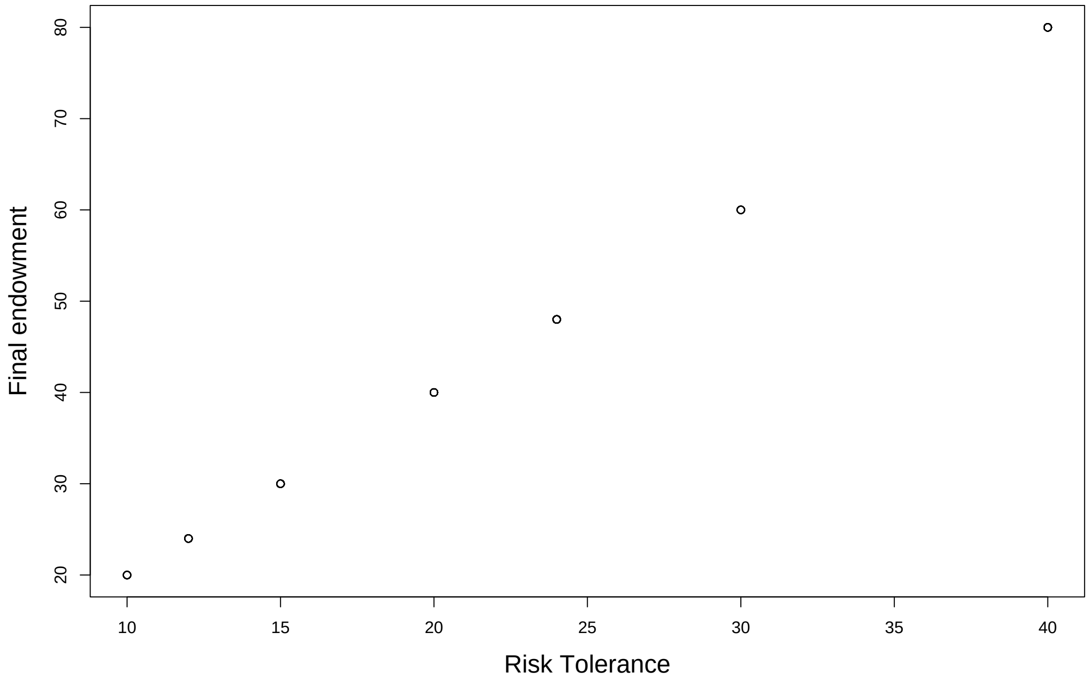

The benchmark for a comparative evaluation is Walrasian tâtonnement.

It yields the efficient allocation in one (giant) step, minimizing volume.

This table reports summary statistics over 25 different batches of simulations under TT.

Prot.	Vol	ExcV	NT	CE	Loss	Dist	SD	Kurt
B	11322.44	2.7%	148.32	88079.50	0.00%	0.00	4.57	6.288
C	28778.04	161%	149.40	88079.50	0.00%	0.00	31.72	3.405
D	11157.48	1.6%	143.56	88022.26	0.065%	0.11	4.09	0.754
H	16221.56	47%	144.88	88050.15	0.033%	0.01	2.77	-0.020

The ranking with respect to volume is $\{B, D\} > H > C$, where $>$ stands for "less volume".

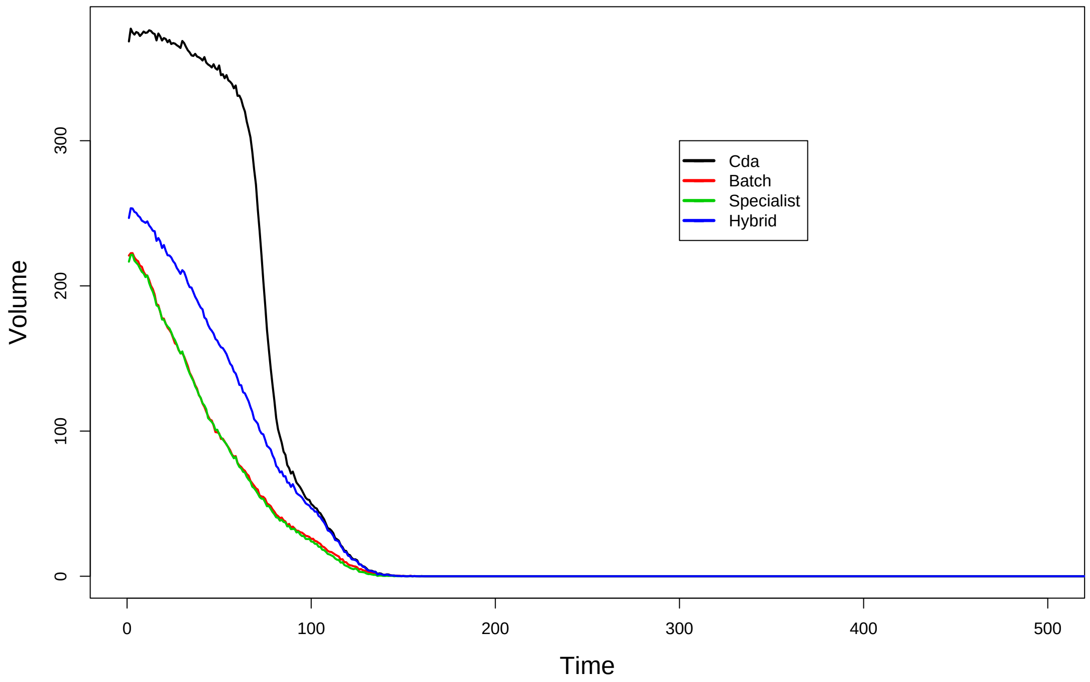

With respect to the ability of a protocol to minimize the number of wasteful trades, the ranking is $\{B, D\} > H > C$.

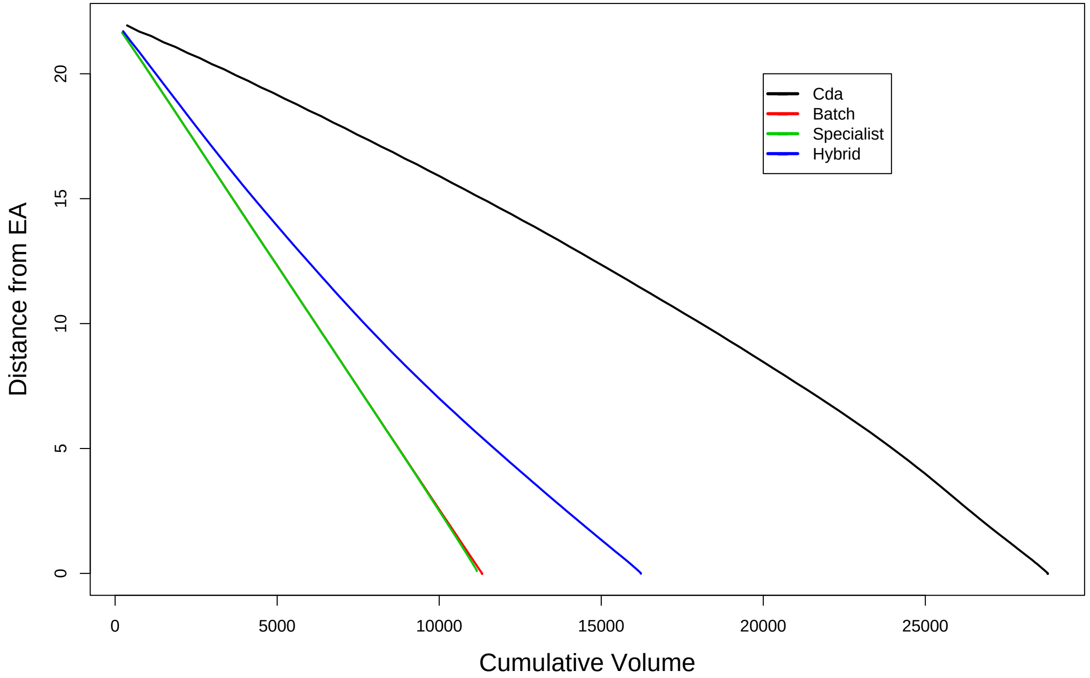

With respect to the volatility of prices, the ranking is $H > \{B, D\} > C$.

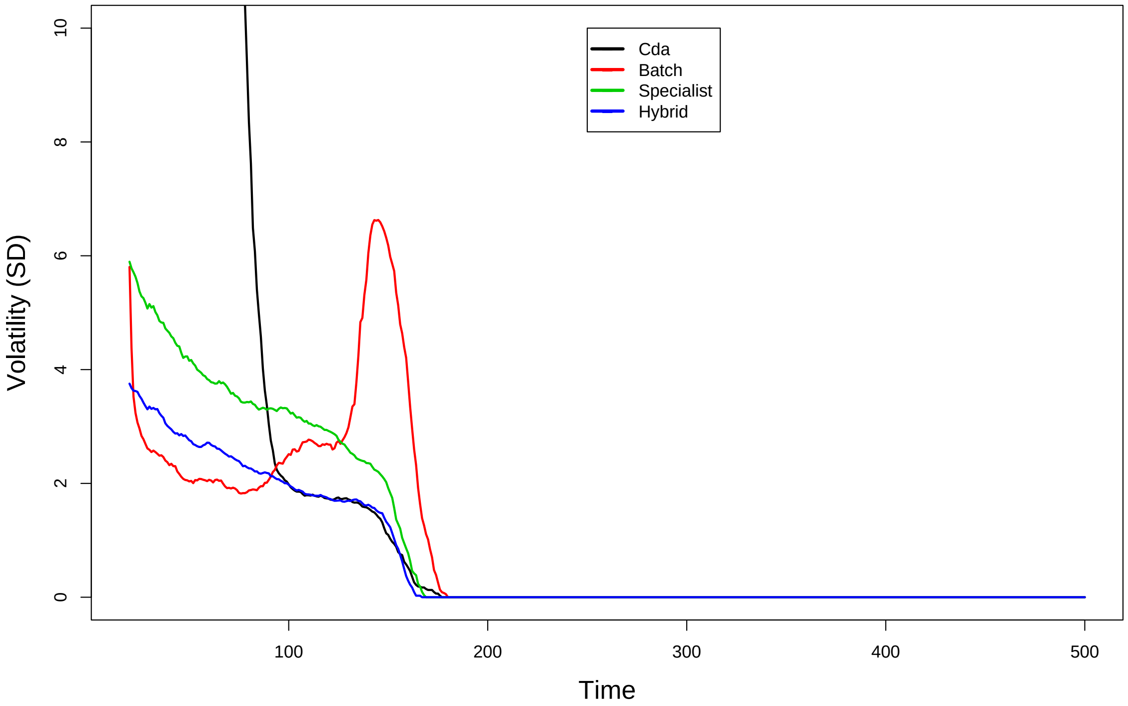

The MA assumption

When compared to TT, MA implies a less demanding search for profitable trades on the part of the agents.

Under TT, two agents agree to a trade when this improves their joint welfare w.r.t. their positions immediately *before* trading.

Under MA, their joint welfare is improved w.r.t. to their initial positions, but not necessarily w.r.t. to their positions immediately before trading.

Trading under TT is *incrementally* improving, while trading under MA is *overall* improving.

This table reports summary statistics over 25 different batches of simulations under MA.

Prot.	Vol	ExcV	NT	mCE	Loss	Dist	SD	Kurt
B	173995.40	1478%	1361.40	88079.50	0.00%	0.00	10.92	3.473
C	248730.72	2156%	785.84	88079.47	0.00%	0.018	75.12	13.150
D	106327.36	864%	925.76	87539.94	0.61%	0.016	9.64	-0.665
H	148756.60	1248%	1081.60	87757.00	0.41%	0.01	8.76	-0.771

The MA behavioral assumption is clearly less effective, but all protocols achieve allocative efficiency confirming their ability to act as partial substitutes for individual rationality.

Compared to TT, it now takes longer to reach the efficient allocation and traded volumes are uniformly much higher.

The ranking w.r.t. volume is $D > H > B > C$ (TT: $\{B, D\} > H > C$).

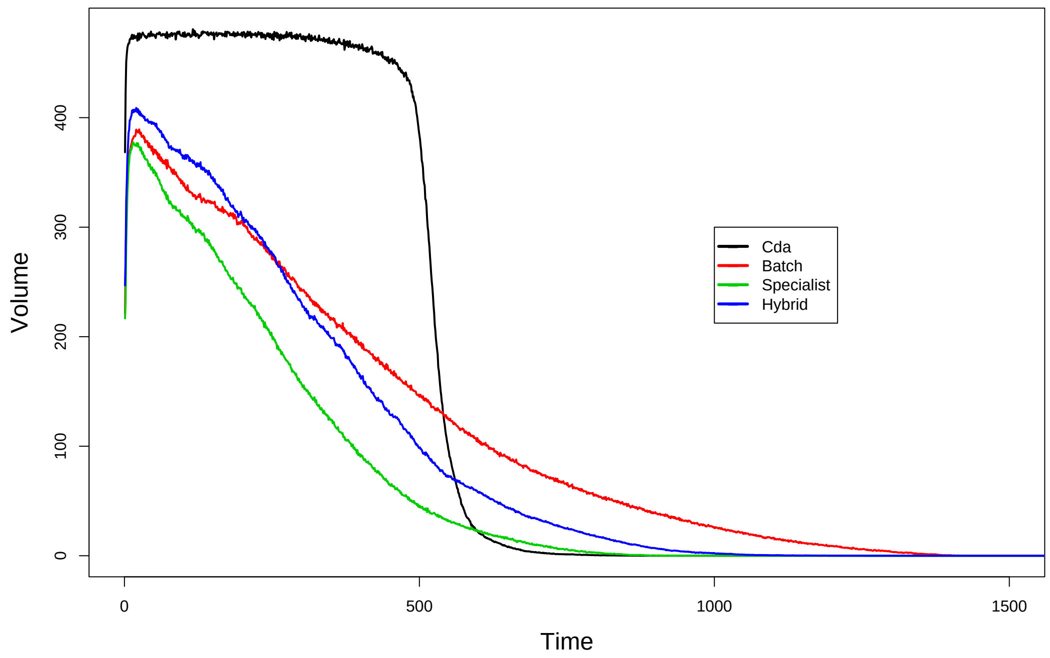

Note the sharp transition for the continuous double auction.

With respect to the ability to avoid wasteful trades, the ranking is $D > H > B > C$.

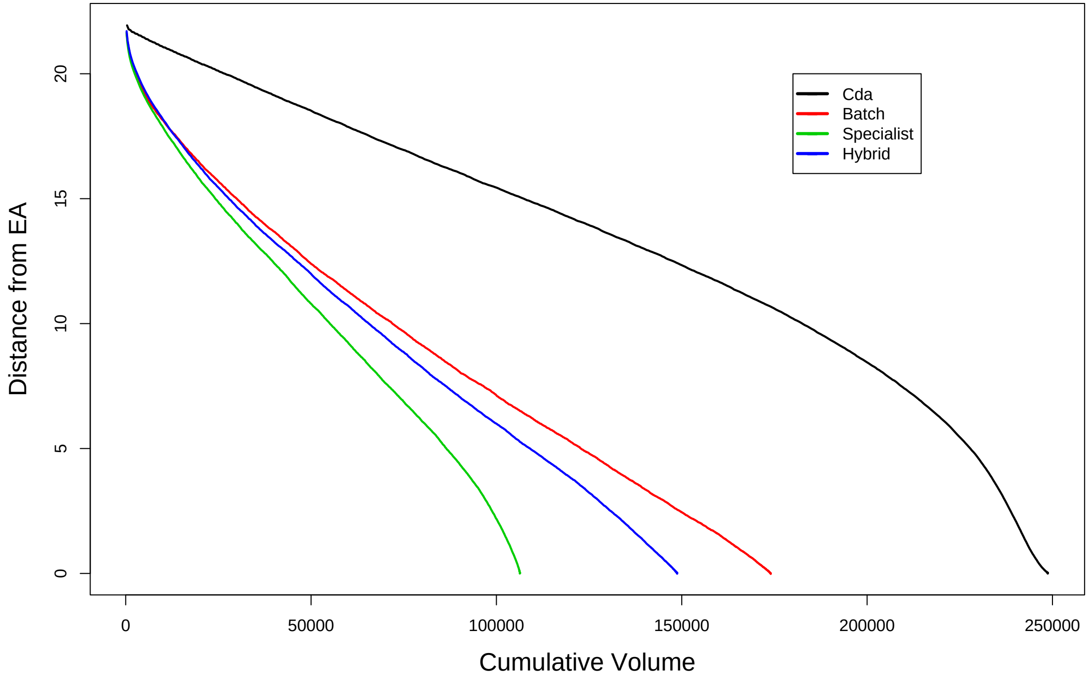

W.r.t. volatility, the ranking is $H > D > B > C$ (TT: $H > \{B, D\} > C$).

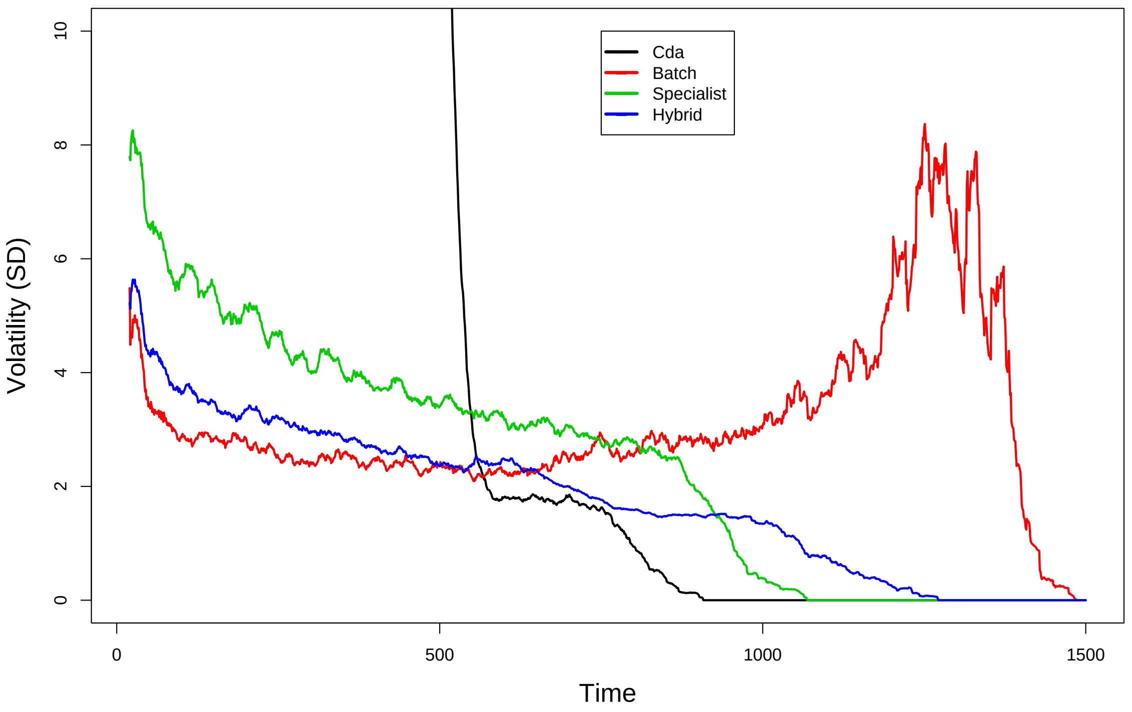

Conclusions

We extend Gode and Sunder's argument for the allocative efficiency of C to two additional behavioral assumptions and to other protocols.

Conclusions

We extend Gode and Sunder's argument for the allocative efficiency of C to two additional behavioral assumptions and to other protocols.

C is an protocol inferior to B w.r.t. volume, speed of convergence, and volatility of prices.

Conclusions

We extend Gode and Sunder's argument for the allocative efficiency of C to two additional behavioral assumptions and to other protocols.

C is an protocol inferior to B w.r.t. volume, speed of convergence, and volatility of prices.

D has a lower time to convergence and never performs worse than B, but it drains a tiny amount of wealth from traders.

Conclusions

We extend Gode and Sunder's argument for the allocative efficiency of C to two additional behavioral assumptions and to other protocols.

C is an protocol inferior to B w.r.t. volume, speed of convergence, and volatility of prices.

D has a lower time to convergence and never performs worse than B, but it drains a tiny amount of wealth from traders.

H has lower volatility and wealth drain, but is otherwise inferior to D.

Conclusions

We extend Gode and Sunder's argument for the allocative efficiency of C to two additional behavioral assumptions and to other protocols.

C is an protocol inferior to B w.r.t. volume, speed of convergence, and volatility of prices.

D has a lower time to convergence and never performs worse than B, but it drains a tiny amount of wealth from traders.

H has lower volatility and wealth drain, but is otherwise inferior to D.

The all-round ranking has D in first place and C in fourth place.