University of Zurich[UZH]

Department of Economics

Seemingly Equivalent Firm Decision Heuristics

Draft 2016-05-31

Luzius Meisser, luzius@meissereconomics.com

Being complex systems, agent-based models can be sensitive to subtle changes in the micro-behavior of their agents. This paper evaluates different heuristics for calculating dividend payments of firm agents in a basic setting, namely an Arrow-Debreu economy with Cobb-Douglas production. While the evaluated heuristics all are equivalent in equilibrium, their resulting aggregate dynamics vary between stability, oscillations, and chaos. Mapping the candidate heuristics onto a continuous parameter space, I show that the equilibrium approach, namely equating dividends with profits calculated as income minus expenses, is exactly on the edge of systemic instability. A further complication stems from the sequential nature of agent-based models, which can lead to situations in which agents have to take decisions based on variables whose latest values have not been realized yet, forcing them to rely on estimates or earlier observed values. Again, a simulation can be sensitive to the choice of such input variables, for example whether the dividend decision is based on the latest realized expenses or the next planned expenses, even though both are identical in equilibrium. These results exemplify the high attention to detail that is necessary to build reliable agent-based models.

1 Introduction

The building blocks of agent-based models are individual, stateful agents that follow specific rules. Instead of basing their decisions on a global state, they use their own local observations to form beliefs on which they base their actions. One of the most common agent type in agent-based economics is the firm. Firms buy input goods to produce and sell output goods, trying to maximize their profit, which they distribute to their shareholders. Furthermore, they often need to find the right prices by trial and error, which also is the case in the discussed setting. As simple as this sounds, there is already a lot that can go wrong in this simple setup. This paper focuses entirely on the seemingly trivial question of how much dividends firm agents should pay out, showing that this question is not so trivial after all and that the standard apprach can fail to move the simulation towards the efficient equilibrium.

Classic equation-based models are inherently mathematical and are solved analytically or numerically. In contrast, agent-based models are inherently algorithmic and solved by simulation. Even though each algorithm can be expressed in equations and vice versa, agent-based models tend to exhibit fundamentally different qualities than equation-based models, one such quality being their rich, often chaotic dynamics. Often, these dynamics are prematurily celebrated as deep findings about the nature of economies in general, when in fact they can stem from small implementation details or even innocuous programming errors.[1]

Generally, it is easy to create agent-based models that look meaningful. It is much harder to create agent-based models that fulfill clear metrics, such as reaching an efficient equilibrium or at least a self-

[1] For example, concluding that "because of dispersed information [...], the system fails to mechanically reach a Pareto efficient [...] general equilibrium position," (Gatti et al., 2011) seems premature to me as there are agent-based simulations with dispersed information that are perfectly capable of reaching Pareto-efficiency – such as the one discussed in here.

confirming equilibrium, which is easier to verify and a more pragmatic metric for models with bounded rationality. (Fudenberg and Levine, 1993) Whenever possible, one should verify agent-based models with classic equilibrium benchmarks, as recommended by LeBaron (2001). The model used in this paper is simple enough to do so, providing a clear metric for evaluating different dividend heuristics. Knowing that a particular heuristic yields good results in simple models, it can be employed with more confidence in complex settings that are harder to verify.

Other authors of agent-based models have chosen a wide variety of different dividend heuristics. There is no established consensus about what works best, and the subject generally lacks discussion. Often, the source code - if available - is the only reference. Some models simply distribute all profits as dividends.[2] Some models – for example the Crisis Economics model behind the publications by Cincotti et al. (2012) and Farmer et al. (2012) – apply an exogenously given stochastic dividend heuristic that is detached from the firm's fundamentals.[3] Gatti et al. (2011) let firms pay out a varying fraction $f \leq 1$ of their profits as dividends, allowing firms to grow, but not to shrink towards their equilibrium size. Meisser and Kreuser (2015) let firms distribute all cash above an exogenously given threshold, which works for the discussed configurations, but lacks flexibility. The most reasonable existing approach is probably that of the May 2016 version of the model by Seppecher (2012). There, firms distribute somewhat more or somewhat less than the net profits as dividend, depending on whether they want to grow or shrink.[4] Depending on how the firm determines its target size, this can be seen as a generalization of the heuristics analyzed in the paper at hand.

Among the tested heuristics, I find most stable results for setting dividends $d = (1 - \lambda)R$ with profit share $1 - \lambda$ and revenue R. It is independent of costs and robust for different choices of measuring R. Upgrading the previously presented agent-based simulation (Meisser and Kreuser, 2015) with this new heuristic allows to significantly extend the range of parameters for which it yields stable results, namely for to returns to scale of close to 1 instead of only up to 0.6.

Section 2 specifies the problem in more detail, describing the two fundamental decisions firm agents face, namely how much money to allocate on inputs and how much on dividends. The subsequent section 3 describes the rest of the model in more detail and can be safely skipped by readers not interested in reproducing the presented results. Results section 4 illustrates how the presented variants of the dividend decision can lead to dramatically different dynamics.

2 Problem

Firms in agent-based simulations buy inputs, transform them into outputs, and try to sell these outputs at a profit, which they pay out to their owners. Sometimes, they also use leverage, accumulate capital, and do research. But the core premise remains the same, revolving around the two decisions of how much money to spend on inputs, and how much money to return to their owners. For simplicity, I ignore leverage and capital, and also assume that there is a way of determining market prices in place, allowing to focus entirely on the aformentioned two decisions.

Furthermore, I assume Cobb-Douglas production with decreasing returns to scale. Decreasing returns to scale ensure the existence of a competitive equilibrium with multiple firms. And Cobb-Douglas functions have the nice property of constant revenue shares for each factor as well as for profits, which

[2] For example the Computational Economy by Wolffgang (2015)

[3] See method `FirmStubg.setNextDividend()` on github.com/crisis-economics.

[4] See `BasicFirm.payDividend()` in github.com/pseppecher/jamel.

greatly simplifies the discussion. The discussed problem should not fundamentally differ for other production functions.

2.1 Spending Decision

Regarding the decision of how much to spend on the acquisition of inputs, I resort to the trivial heuristic of simply spending a constant ratio $s = 0.2$ of the available cash. Under decreasing returns to scale, this assumption does not restrict the economy's capability of reaching the efficient equilibrium as long as prices are allowed to adjust and each firm is allowed to save or dissave to reach the cash levels that imply optimal spending. At the macro-level, choosing this spending heuristic can be seen as a form of price normalization, with equilibrium prices being directly proportional to money supply m and to the spending fraction. Interpreting the spending ratio as the velocity of money $v = s$, this matches the classic monetary equation: $pt = mv$ with prices p and real trade volume t.

Note that the choice of the spending ratio s only affects the nominal equilibrium price level, and not real prices. Assuming that no other form of price normalization is applied and that nominal price levels are considered irrelevant as usual, any value $s \in (0, 1]$ can be chosen without loss of generality. Considering the dynamics, it makes sense to choose a value that allows firms to keep a cash buffer to guard against random fluctuations. I use $s = 0.2$. A nice side-effect of this particular spending heuristic is that it cannot bankrupt firms as they never spend more than they have.

2.2 Dividend Decision

The firm's decision of how much dividends d to pay out is probably the most important of all. It controls the growth of a firm. As long as the equilibrium rule of simply paying out all profits π is in place, a firm cannot increase its cash balance or otherwise reinvest its earnings, and therefore cannot grow, regardless of all other decisions.[5] Thus, any reasonable firm implementation must somehow regulate its dividends beyond just equating them with profits. It is essential to choose a dividend heuristic that is not only correct in equilibrium, but also exhibits benevolent dynamics out of equilibrium. In the following, I will construct a one-dimensional parameter space of ways to calculate profits given revenue R and cost C under Cobb-Douglas production. They all are correct in equilibrium, but can deviate off the equilibrium. This allows to systematically analyze and test them.

Given a trivial Cobb-Douglas production function $x(h) = Ah^{\lambda}$ without capital, profit maximization results in a labor share λ and a profit share $1 - \lambda$. In case of multiple input goods, the labor share can still be denoted as $\lambda = \sum_c \lambda_c$ with λ_c being the elasticity of input c. Both profits π and costs C can be expressed as a fraction of revenue, with $\pi = (1 - \lambda)R$ and $C = \lambda R$, implying $\pi = \frac{1-\lambda}{\lambda}C$. Together with the standard profit function $\pi = R - C$, this results in three ways to calculate profits π:

$$\pi = R - C = (1 - \lambda)R = \frac{1 - \lambda}{\lambda}C$$

Linearly combining them, any choice of finite coefficients a_R, a_C in equation 1 must lead to the same optimal result in equilibrium.

$$\pi = (1 - a_C - a_R)(R - C) + a_R(1 - \lambda)R + a_C\frac{1 - \lambda}{\lambda}C \qquad (1)$$

[5] In the absence of investments, the size of a firm is simply its cash balance.

This equation is somewhat redundant and can be transformed into the linear combination

$$\pi = b_R R + b_C C \tag{2}$$

whose coefficients b_R and b_C have to fulfill:

$$\lambda(b_C + 1) = 1 - b_R \tag{3}$$

While in equilibrium, any choice of b_R yields the same result, the agent-based model usually is slightly off the equilibrium, leading to different dynamics depending on the choice of b_R. For some values, it will not converge at all. In order to converge towards the efficient equilibrium, firms must not pay out all the profits when they are below their optimal size, and must pay out more when they are above their optimal size.

2.3 Observing Variables in a Sequential World

To complicate matters further, it is not obvious how a firm should measure revenue R and costs C. Agent-based models are inherently algorithmic and therefore sequential by nature, not allowing the circular dependencies often found equation-based models. Here, the latest values R_t and C_t have not been realized yet when the firm has to decide about dividends d_t, forcing it to rely on older values or other estimates.

Dividend payments take place at beginning of each day in the discussed model and not at the end. Both variants are equivalent, but distributing them in the morning ensures that dividends d_t are both distributed and spent on day t, reducing potential confusion and making clear that it is indeed the firm decisions that drive the dynamics and not those of the consumers. Simulation days are structured as follows:

1. Consumers are endowed with 24 man-hours each.
2. Firms distribute dividends as calculated by their dividend heuristics.
3. Firms post asks to the market, offering yesterday's production in accordance with their individual price beliefs; for example *"we sell 79 pizzas for 7.30$ each"*.
4. Firms calculate their budget using their spending heuristic and post bids in the form of limit-orders to the market, for example *"we buy up to 50 man-hours for 13$ each"*.
5. In random order, consumers enter the market and optimize their utility given the offers they find, selling man-hours and buying output goods.
6. The market closes and each firm updates its price beliefs based on whether the relevant orders were filled or not.
7. Firms use all acquired man-hours to produce the outputs to be sold tomorrow. In equilibrium, all money resides with the firms again at this point in time, although not necessarily equally distributed.

At the point in time at which the dividends d_t are determined, no trade has taken place yet. Neither R_t nor C_t is not known yet. Instead, the dividend decision could for example be based on the known R_{t-1} and C_{t-1} from yesterday. Another variant is using expected revenue $E[R_t]$ given the firm's price beliefs and goods in stock. As costs, the planned spendings $C_{plan,t}$ according to the spending heuristic could be plugged in. Another interesting variant is to calculate and use optimal spending $C_{t,opt}$ given current price beliefs. There is a multitude of additional thinkable options, resulting in a zoo of of slightly different dividend heuristics. In here, I focus on the aformentioned ones. Again, they are all identical in equilibrium.

3 Test Environment

This section describes the architecture of the simulation in which the firms are tested. It can be safely skipped by readers not interested in the exact specification. The simulation is built according to the principle of least surprise and should not contain anything special. It is based on the simulation presented earlier in Meisser and Kreuser (2015), also making use of exponential search and sensor prices for better accuracy. However, it differs in the way prices are normalized and in its spending and dividend heuristics. The simulation was written in Java and can be found on heuristics.meissereconomics.com.

The simulation is structured as a sequence of daily Arrow-Debreu spot markets with 10 competing firms and 100 consumers that are endowed with 24 man-hours per day. All firms have the same Cobb-Douglas production function

$$x_f(h_f) = max(1, A\, h_f^{\lambda}) \tag{4}$$

with h_f being the man-hours acquired by firm f and h_c the man-hours sold by consumer c, such that $\sum_f h_f = \sum_c h_c$. All consumers c have the same logarithmic utility function:

$$U(x_c, h_c) = \alpha ln(x_c + 1) + \beta ln(1 + 24 - h_c) \tag{5}$$

For the outputs, $\sum_f x_f = \sum_c x_c$ holds analogously to the inputs. The increment $+1$ in each component of the utility function serves to avoid negatively infinite utilities. Otherwise, a single consumer failing to acquire a single consumption good on a single day can spoil average utility as a metric by dragging it down infinitely. Similarily, the floor of 1 in the production function helps the firms recover after an economic meltdown, stabilizing the simulation at the boundaries. In all the tested equilibria, production is above 1, allowing to ignore the max operator when calculating the equilibrium mathematically. Furthermore, all firms can be aggregated into one in the classic view thanks to having identical parameters $A = 10$ and $\lambda = 0.7$. The same applies to consumers, who have $\alpha = 10$ and $\beta = 14$.

In the disaggregate simulation, individual agents are kept seperate. Here, each firm f and consumer c has its own state, namely its own inventory of goods and a wallet with its own money w_f or w_c. Each firm also maintains its own price beliefs, which it adjusts over night depending on trade success. Firms employ exponential search and sensor pricing for better accuracy. (Meisser and Kreuser, 2015) The presence of money as a means for exchange ensures that bilateral trade suffices to reach the efficient equilibrium. (Feldman, 1973)

Days are structured as described in section 2.2. Consumers are passive, entering the daily market one by one to work and consume optimally given the offers they find. The adaptivity of the simulation stems entirely from the firms, who adjust their individual price beliefs depending on how successful they were at buying and selling their inputs and outputs. For the test, $\lambda = 0.7$ is chosen, but equivalent results are achieved with other values.

The daily amount of inputs acquired by the firms is used as a benchmark to see how close a particular configuration gets to the unique optimum. This is a valid metric in the simple setting at hand. For more elaborate settings with multiple input goods or more agent types, more elaborate benchmarks might be necessary. The benchmark scenario runs for 5000 days, with measurements starting from day 500 to give the simulation some time to find its initial equilibrium. The relatively long measurement period of 4500 increases the chances of random deviations and is akin to averaging the results of multiple shorter simulation runs. One simulation run takes about one second to complete on my computer.

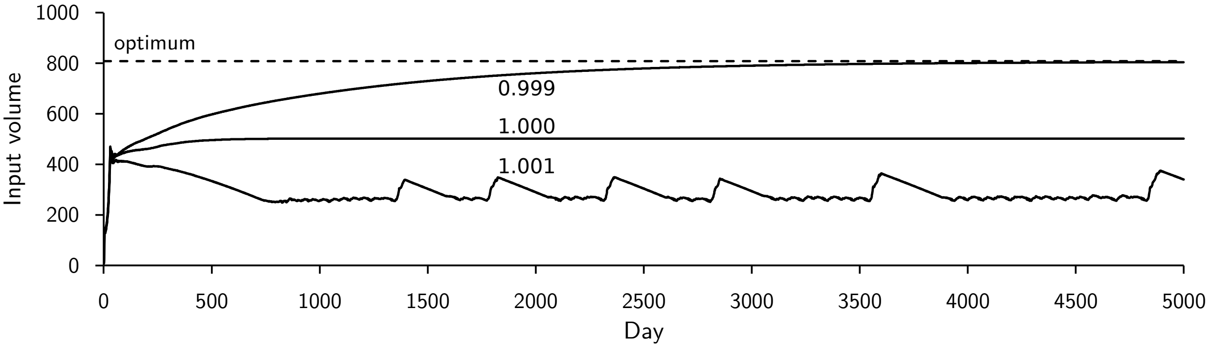

Figure 1: The standard solution $d_t = \pi_t = R_{t-1} - C_{t-1}$ (implying $b_R = 1.000$) is right at the edge of instability. Setting the parameter somewhat below ($b_R = 0.999$, implying $d_t = 0.999R_{t-1} - (1 - 0.001/\lambda)C_{t-1}$) suffices to let the simulation converge towards the efficient equilibrium. Setting the parameter somewhat above ($b_R = 1.001$) lets the simulation fall towards an inefficient state with sporadic outbreaks. For $b_R = 1.000$, the outcome is stable but the value it settles on depends on the initial conditions.

4 Results

Even though all tested heuristics are equivalent in equilibrium, they differ wildly in the resulting dynamics. All tested heuristics are based on equation 2 from section 2.2, setting dividends to

$$d = b_R R + b_C C$$

and testing the impact of changing b_R, adjusting b_C accordingly to fulfill equation 3. Orthogonal to that, three different ways of measuring R and C are tested, which are again all equal in the efficient equilibrium.

The most common way of calculating dividends is to set it equal to profits using the latest known revenues and costs: $d_t = R_{t-1} - C_{t-1}$. This is equivalent to setting $b_R = 1.0$ and adjusting $b_C = -1.0$ accordingly.[6] Doing so results in stable off-equilibrium states as shown in figure 1. The exact state the simulation settles on is not fixed and depends on the initial conditions. However, already deviating slightly to $b_R = 0.999$ allows firms to save a little cash every day as long as marginal costs are below prices, thereby approaching the efficient equilibrium. Likewise, setting $b_R = 1.001$ pushes firms away from the equilibrium into instability. This nicely shows that the standard heuristic really is at the edge of instability. In order to reach the efficient solution, one must slightly deviate from the standard heuristic by allowing firms to withhold some of their profits when they need to grow and to distribute excess cash when they are larger than optimal.

The cliff at $b_R = 1.0$ can also nicely be seen in figure 3. For lower values, the equilibrium solutions can be attained, but for higher values, the market settles on an inefficient steady state. Since the latest known R and C are used, I call this heuristic the *known* heuristic. For values of b_R that are too low, the simulation falls into a two-cycle orbit as shown in figure 2, with all 10 firms adjusting their behavior in sync even though they only interact indirectly through the market.

A second variant is the *optimal cost* heuristic, which plugs in the expected reveneus $E[R_t] = x_t p_t$ from selling the inventory x_t at price belief p_t and the level of spending $C_{t,opt}$ that would maximize profits given the current price beliefs. It fails to produce meaningful results for $b_R < 1 - \lambda$, where b_C is

[6] The firm subscripts f have again been dropped for readability. In the simulation, each firm f has its own values $d_{t,f}$, $R_{t,f}$, $C_{t,f}$, etc.

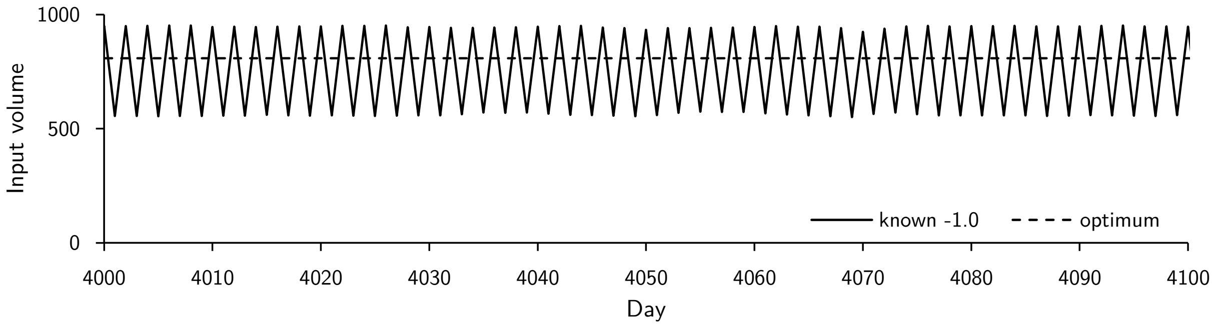

Figure 2: Setting $b_R = -1$ in the *known* heuristic implies $d_t = -R_{t-1} + \frac{2-\lambda}{\lambda} C_{t-1}$ and leads to a two-cycle.

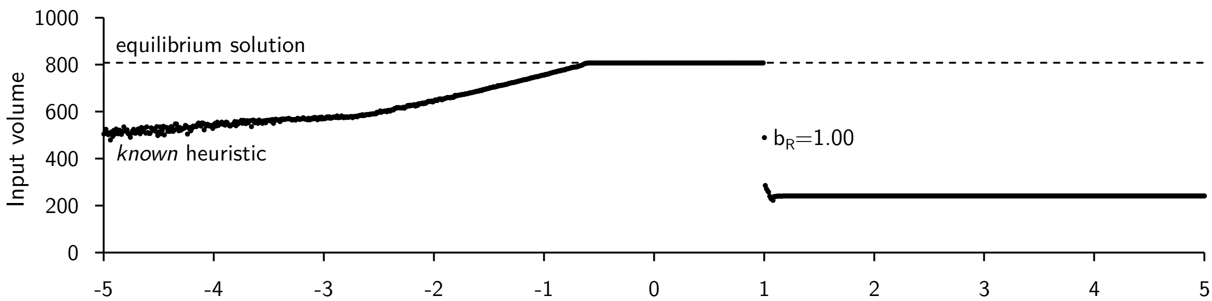

Figure 3: *Known* heuristic: each dot represents one of 10000 simulation runs with revenue weight $b_R \in [-5, 5]$ using the latest known R_{t-1} and C_{t-1}. The setting $b_R = 1.0$ implies the standard rule $d_t = \pi_t = R_{t-1} - C_{t-1}$, which results in a stable, off-equilibrium outcome. Here, it settled on 501, but that varies depending on the initial conditions.

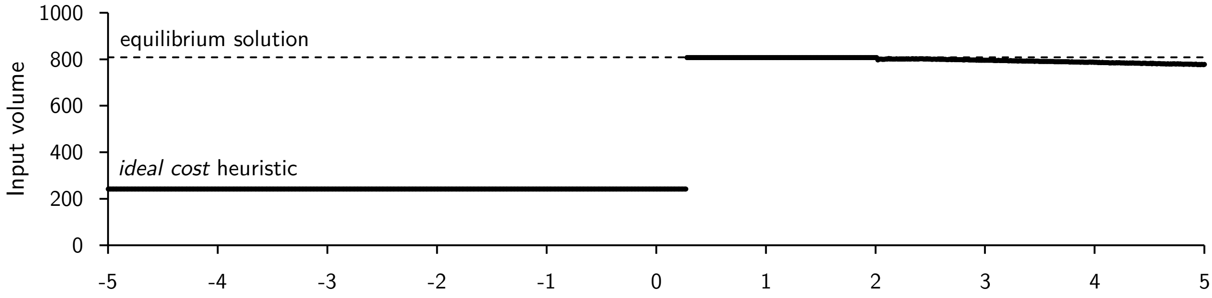

Figure 4: The *ideal cost* heuristic bases the dividend decision on expected revenue $E[R_t]$ and ideal costs $C_{t,opt}$ given price beliefs. It is attracted to the efficient solution for $b_R \geq 1 - \lambda$, but detoriates again for very large parameter values. Each dot represents one of 10000 simulation runs with $b_R \in [-5, 5]$.

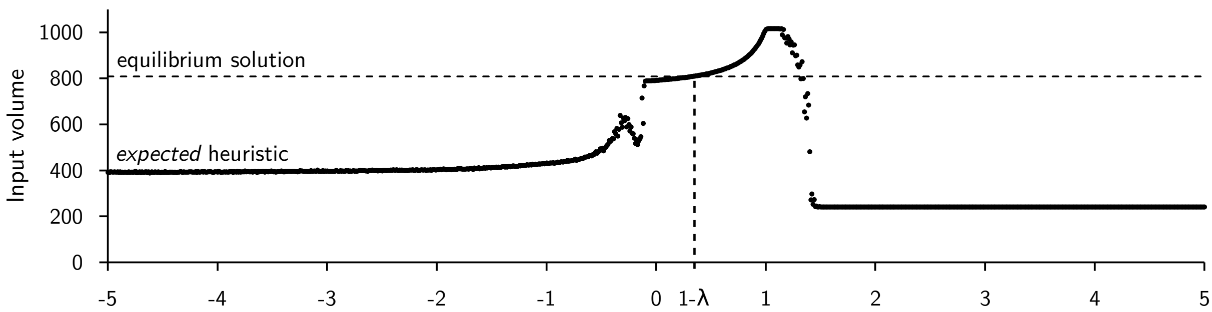

Figure 5: *Expected* heuristic: when basing the dividend decision on expected revenue $E[R_t]$ and planned costs $E[C_t]$, the efficient outcome is reached for $b_R = 1 - \lambda$, at which point $d_t = (1 - \lambda)E[R_t]$, which is the recommended dividend heuristic. Each dot represents one of 10000 simulation runs with $b_R \in [-5, 5]$.

positive. Considering that $C_{t,opt} > C_t$ holds when the firm has too little cash and $C_{t,opt} < C_t$ holds when the firm has too much cash, its behavior is not surprising. A positive b_C leads to generous dividends exactly when there is not enough cash, therefore pushing the firm away from equilibrium. The opposite is the case for negative b_C, which lets the firm restrict the dividend when its cash levels are low so it can move towards the efficient equilibrium.

The third tested variant is the *'expected'* heuristic, which is also based on $E[R_t]$, but uses planned spendings $E[C_t] = sw_t = 0.2w_t$ with w_t being the cash (wealth) of the firm. It is more adaptive than the other variants and reaches the equilibrium only for $b_R = (1 - \lambda)$, at which point it is identical to the *optimal cost* heuristic as $b_C = 0$. Thus, setting

$$d = (1 - \lambda)R$$

seems the most robust choice as it allows the simulation to find the efficient equilibrium for all three tested heuristics. It is simple, more stable than the standard profit equation, and converges reasonably fast under various conditions.

References

S. Cincotti, D. Sornette, P. Treleaven, S. Battiston, G. Caldarelli, C. Hommes, and A. Kirman. An economic and financial exploratory. *The European Physical Journal Special Topics*, 214(1):361–400, 2012.

J. D. Farmer, M. Gallegati, C. Hommes, A. Kirman, P. Ormerod, S. Cincotti, A. Sanchez, and D. Helbing. A complex systems approach to constructing better models for managing financial markets and the economy. *The European Physical Journal Special Topics*, 214(1):295–324, 2012.

A. M. Feldman. Bilateral trading processes, pairwise optimality, and pareto optimality. *The Review of Economic Studies*, pages 463–473, 1973.

D. Fudenberg and D. K. Levine. Self-confirming equilibrium. *Econometrica: Journal of the Econometric Society*, pages 523–545, 1993.

D. D. Gatti, S. Desiderio, E. Gaffeo, P. Cirillo, and M. Gallegati. *Macroeconomics from the Bottom-up*, volume 1. Springer Science & Business Media, 2011.

B. LeBaron. Empirical regularities from interacting long-and short-memory investors in an agent-based stock market. *Evolutionary Computation, IEEE Transactions on*, 5(5):442–455, 2001.

L. Meisser and F. Kreuser. An agent-based simulation of the stolper-samuelson effect. *Science Research Network Working Paper Series*, 2015. URL *http://ssrn.com/abstract=2794828*.

P. Seppecher. A java agent-based macroeconomic laboratory. 2012.

U. Wolffgang. A multi-agent non-stochastic economic simulator. In *Proc. of the 21st Int. Conf. on Computing in Economics and Finance (CEF 2015)*, 2015.