The Continuous Capital Corporation

Luzius Meisser*, luzius.meisser@uzh.ch

2021-01-31

Abstract

Switzerland recently did what other European countries are still working on: creating explicit legal foundations for the issuance of blockchain-based cryptosecurities. These enable companies to directly interact with their investors and shareholders at scale, bypassing the established financial system and potentially creating their own, decentralized stock markets. I investigate the pricing rules companies should adhere to when using this newly available technology to establish markets for their own shares. In particular, I provide a deterministic pricing function that allows companies to make a market for their own shares and that provably supports the economically efficient outcome under a wide range of circumstances.

Disclaimer

This is a preprint. So far, it has neither been peer-reviewed nor formally been accepted in an academic publication. The first review - for the 2021 conference of the European Finance Association - is pending.

1 Introduction

Blockchain technology is often seen as a means to (potentially) make the existing financial markets more efficient, thereby mainly enabling a quantitative, but not a qualitative evolution of the markets. However, the nascent field of Decentralized Finance shows that Blockchain technology is also likely to enable new types of markets and new forms of financing.

My favorite example of such a new type of market is the decentralized exchange Uniswap. Uniswap is based on completely deterministic pricing rules, with anyone being able to take part in the market making by passively contributing to so-called liquidity pools. The daily trading volume on Uniswap is currently approaching one billion dollars and Uniswap itself is valued by the market with roughly four billion. [Uniswap, 2021] [Coinmarketcap, 2021] I want to investigate whether this kind of market could not only be used to create liquidity for speculative cryptocurrencies, but also serve the same purpose for the equity of real-world companies.

A determining feature of Decentralized Finance is the use of so-called smart contracts. These are small pieces of software that are deployed to a blockchain-system and that can interact with

*I would like to thank the participants of the Financial Economics Doctoral Seminar at the University of Zurich for their valuable inputs. Also, I would like to thank the the participant of the SFI Research Days 2020, most notably session chair Philipp Krueger and discussant Sidita Hasa. Also, I would like to thank Thorsten Hens for his guidance in finding the right model to tackle the problem at hand, as well as Alexander Wagner for some early inputs. Special thanks go to my wife, Verena, for her relentless support and enabling me to work hard on interesting problems.

the financial instruments present within that system. In October 2020, the Swiss parliament approved the legal foundations necessary to embody traditional securities as so-called cryptographic tokens within a blockchain, and other countries are following suit. [Bundesrat, 2020] This opens the possibility to write smart contracts that do automated market making for securities. In contrast to traditional market making, these smart contracts act with full public transparency and according to deterministic, predetermined rules. This has interesting consequences for how the market functions, in particular when the issuer itself is involved.

In theory, financial markets are supposed to play a crucial role in the capital allocation of the real economy. While this might indeed be the case for credit markets and private equity, there is only a weak connection between public stock markets and the capital allocation in the real economy. The vast majority of capital flows on the stock market are between investors, and only very rarely between the issuer and the investors. Dow and Gorton [1997] show that as long as there is no direct link between stock markets and the economy, an efficient stock market can guide management decisions to a certain degree through its informational value, but it is not sufficient for economic efficiency.

When a company uses a smart contract to kickstart a market for its own shares, providing the necessary liquidity itself, it becomes what I call a *continuous capital corporation*. Financially, it engages in the continuous issuance and buyback of its own shares at a minimal spread. This gives the markets much more direct control over the capital allocation in the real economy. Buying shares of a firm injects capital, selling the same shares extracts capital from the firm. The main challenge when creating a continuous capital corporation is the design the firm's pricing function so that a free and efficient market consisting of rational investors automatically drives the capital allocation between firms towards the efficient level, without explicit financing operations initiated by the firm's management.

The pricing mechanism proposed herein fulfills three important properties. First, it is shown to lead to the efficient capital allocation in a market with rational investors, at least in the stationary case. Second, it provides a price-point at which, under a wide range of circumstances, both the sale of new shares as well as the repurchase of old shares is Pareto-improving when transacting with an informed investor. These circumstances include interest rate changes, technology shocks, and – under the assumption of Cobb-Douglas production – also input price changes. Third, the pricing mechanism is simple enough to support automated market-making with a blockchain-based smart contract, obsoleting the financial intermediaries that are traditionally needed to run financial markets.

Like Tinn [2017], I do not focus on efficiency gains enabled by blockchain technology, but on the new possibilities it enables. Also similar to Tinn and unlike many authors in the corporate finance literature [Biais et al., 2013], I disregard agency issues and assume symmetric information.

This paper is structured as follows: after introducing the reader to the technical and legal background in section 2, I state the problem to be tackled in section 3. Section 4 presents an economic model of equity financing through the continuous issuance and repurchase of shares and derives a pricing function for the continuous capital corporation. The model is then extended in section 5 and the market interactions of the continuous capital corporation compared to traditional market making in 6. Finally, I conclude with 7.

2 Background

This section provides some background about the technical and legal developments that enable the continuous capital corporation.

2.1 Technical Developments

Since Nakamoto [2008] invented Bitcoin, the underlying blockchain technology has been improved by various others such that it can today not only serve as a payment system, but as a general-purpose, unstoppable "Internet-computer". The most popular general-purpose blockchain is Ethereum as formally specified in Wood et al. [2014]. It comes with its own programming language that allows anyone to formulate so-called smart contracts and then deploy them to the system for a fee, paid to an anonymous collective of system operators. Smart contracts are small computer programs anyone can interact with in the ways their author defined. Since the underlying system is completely decentralized, there is no authority that could ever stop a smart contract unless the author explicitly added a function to do so. Thanks to these properties, systems like Ethereum provide reliable foundations for digital ownership registries that keep track of financial assets. The process of registering securities in a smart contract is called *tokenization* and the individual securities then often referred to as *security tokens*. Unlike other digital assets, such tokens cannot be copied and are thus a suitable replacement for physical paper certificates. No centralized securities depository or other intermediary is necessary to issue security tokens, enabling disintermediation of financial markets. The resulting ecosystem can be seen as an "Internet of Finance" and is often referred to as *Decentralized Finance* or short *DeFi*.

Bypassing the middle-men and red tape, Decentralized Finance lowers the barriers to enter financial markets. Whereas a traditional IPO can cost legal and other fees in the seven-digit range, it is possible for a firm to create a small market for its own shares for a few hundred Euro using blockchain technology. Besides enabling the trading of shares, a blockchain could also enforce the drag-along clause of the shareholder agreement or automate other corporate actions that previously required considerable paperwork. Thanks to these efficiency gains, it becomes feasible for small companies to have hundreds of shareholders without losing control over shareholder relations.

Of particular interest in our context are smart contracts that perform fully automated market making for a pair of tokens. The most well-known examples are Bancor, Kyber and Uniswap. With Uniswap, anyone is free to contribute to the capital reserves of the market maker and will in turn proportionally participate in its gains and losses. A detailed discussion of the properties of Uniswap has been published by Angeris et al. [2019]. An important source of inspiration for the continuous capital corporation was Rosenfeld [2017], where I first encountered the idea of continuously creating and redeeming tokens for the purpose of market-making.

2.2 Legal Developments

On the legal side, I want to point to three notable local developments that helped making the continuous capital corporation possible.

First, new accounting rules have come into force in 2013 that reclassified a firm's transactions with its own shares. Previously, buying its own shares low and selling them high would have caused the firm to make a profit. Recognizing that this is economically equivalent to a buyback and reissuance and following international accounting standards, it is now possible to book transactions with own shares in a profit-neutral way. [Schnell Luchsinger and Montavon, 2018] [Court, 2019]

Second, Switzerland is about to introduce a so-called *capital band*, allowing the general assembly of a corporation to authorize the board of directors to issue new or destroy old shares much more easily, creating up to 50% additional shares or taking up to 50% of the old shares out of circulation. This is still not as flexible as the rules of the United States and many other European countries, but a significant step towards enabling a higher frequency of capital adjustments. [Forstmoser and Küchler, 2020]

Third, the Swiss parliament has passed a law to provide explicit legal foundations for tokenized securities. [Bundesrat, 2020] It is already possible to contractually bind uncertificated shares to blockchain-based tokens under the current law, as von der Crone et al. [2019] point out. But having a law that essentially says "the token *is* the share" provides a much higher level of legal certainty. This has been set into force on February 1st 2021. While for example Germany is also warming up to the topic, Switzerland seems to be ahead of most other European countries. [Bundesregierung, 2020]

Each country seems to have its own legal obstacles that prevent novel types of financial markets from emerging. For example, the United States has a rule that hits companies with the full weight of public reporting requirements once they reach 2000 shareholders, whereas Germany's tax laws make employee participation plans with real shares very unattractive. A review of local obstacles for the evolution of a Decentralized Finance for securities might be worth its own article.

3 Problem Statement

Traditionally, both the issuance of new shares as well as the repurchase of old shares are one-time decisions driven by the management of a firm. But when a firm does its own market making using a smart contract, thereby engaging in the deterministic and continuous issuance and repurchase of its own shares at a narrow spread, this is turned around and the capital allocation decision put into the invisible hand of the free market. The question that this paper aims to answer is: at which price should a continuous capital corporation offer and repurchase its shares in order to support the efficient capital allocation?

The according pricing function shall fulfill the following three criteria:

1. Efficiency: In a competitive market with rational investors, the efficient outcome should be reached. It must not be possible for market participants to exploit the firm's market making. The value created for the shareholders when investors buy newly issued shares should outweigh the dilution effect. And the capital outflow when repurchasing shares should be compensated for by an according increase in the value per outstanding share. Under these conditions, all transactions will be Pareto-improving as they will increase the wealth of the existing shareholders and not take place at all if they decreased the wealth of the rationally acting counterparty.

2. Attractiveness: The efficient outcome should be reachable in incremental steps with each marginal purchase or sale of a share making sense in itself. In more colloquial terms, I want to rule out situations in which buying two shares would make the buyer better off than before, but buying only one share does not. This ensures that each self-confirming equilibrium is also an efficient equilibrium and allows agents following what computer scientists call a *greedy algorithm* to find the efficient outcome. [Cho and Sargent, 2008]

3. Simplicity: It should be feasible to fully automate the market making, for example in a blockchain-based smart contract that mechanically applies the pricing function. This implies that the price must only depend on easily observable variables and that the function should be of low computational complexity.

4 Basic Model

To build an economic model of the continuous capital corporation, we first need a model of equity financing. Most classic models avoid the explicit modeling of share issuance and resort to a work-around instead. For example, Diamond [1967] proposes to model equity financing by separating it into a first step of transactions between shareholders and a second step of all shareholders adding capital to the company in proportion to their holdings, which essentially constitutes a negative dividend. Hens et al. [2019] is a notable exception. They show how equity issuance can be explicitly modelled in a general equilibrium setting. In their model, households treat shares like bonds, namely seeing them as a security with a fixed coupon, disregarding dilution and other secondary effects their buying of newly printed shares might have.

For the purposes of this paper, I will depart from the classic approach by letting the investors know about their impact of buying newly issued shares or returning existing shares to the company. These effects are two-fold: first, there is a dilution effect. The more shares there are, the smaller the dividend per share. Second, there is a productivity impact, the more shares the firm sells, the higher its capital and therefore also its profits. The investors are assumed to be aware of both.

Based on this premise and the assumptions outlined below, this section derives the participation constraint of the firm acting in the interest of its existing shareholders as well as that for new investors. I show that the efficient capital allocation is reached in equilibrium when following these constraints. The participation constraint of the firm is later used to provide the price for the market making by the firm. By definition, it provides the point at which the shareholders are indifferent about a marginal issuance or repurchase of shares.

4.1 Assumptions

I assume a production function $f(K)$ that depends on capital alone and that satisfies the Inada conditions ($f(0) = 0$, $f'(K) \geq 0$, $f''(K) \leq 0$). Time is split into discrete steps, with the fruits of production from capital K_t becoming available in the subsequent period $t + 1$. Later on, the production function will be extended to include other inputs and technology. Without other inputs and without debt, the firm's profits only depend on capital: $\pi(K) = f(K)$. With equity financing, the capital costs do not appear in the accounting and do not reduce profits. The "costs" of raising capital are hidden in the dilution that all existing shareholders suffer from when new shares are issued. To account for that, firms should not maximize profits $\pi(K)$, but the wealth of their current shareholders as later expressed in equation 2.

The issuance price of share number θ is denoted $p(\theta)$. It provides the price at which the continuous capital corporation issues and repurchases marginal units of a share. Its integral $P(\theta) = K(\theta)$ returns the total amount of capital raised through the net issuance of θ shares. In case of a constant price $p(\theta) = p$, capital raised is trivially $K(\theta) = p\theta$. The pricing function is assumed to be positive: $p(\theta) \geq 0$, and increasing: $p'(\theta) \geq 0$. The latter reflects the basic wisdom that higher demand should lead to higher prices.

Note that the pricing function $p(\theta)$ is path-independent. The price is fully determined by the number of outstanding shares, regardless of the path taken to reach this number. Path-independence guarantees that there is no sequence of trades that would allow an attacker to exploit the market maker as long as the sequence starts and ends at the same price.

The internal valuation of the firm is defined as:

$$V_i(K) = V_i(P(\theta)) = \theta p(\theta) \tag{1}$$

It denotes the valuation at which the firm is currently willing to issue new or repurchase existing shares. It is fully determined by the pricing function and the outstanding number of shares.

As long as the market is not in equilibrium, the internal valuation can differ from the external market valuation, which is denoted $V_e(K, r)$ and typically depends on the interest rate and other external variables or anticipated events. To anchor the optimization problem, I assume that there always is an alternative investment opportunity that pays interest rate r and that all market participants except the firm can borrow at that same rate. The way for the firm to increase or decrease its capital under this model is through the issuance or repurchase of its own shares. It cannot borrow and is not allowed to retain earnings.

A notable consequence of having a pricing function instead of a constant price is that it allows the firm to issue different shares at different prices within the same period. This generalization allows the firm to preserve path-independence and helps to fulfill the attractiveness requirement. At the same time, this also means that we are not in a traditional market setting in which the law of one price holds. Such a deviation from the Walrasian setting can be justified by considering that the n-th issued share represents something else (namely $\frac{1}{n}$ of the company at the time of issuance) than the next issued share (namely $\frac{1}{n+1}$ of the company at the time of issuance). This is something traditional market models are not designed to deal with anyway. In a Walrasian setting, the quality of a good does not depend on the number of units sold.

I further deviate from the Walrasian setting by allowing market participants to see how their buying and selling of shares affects a firm's profits. This goes against the usual premise that market participants optimize given prices alone, without any knowledge about other market participants.

4.2 The Firm's Optimization Problem

At each point in time t, the firm maximizes the profits attributable to its existing shareholders, who hold θ_{t-1} shares. Given pricing function $p(\theta)$, it chooses the right amount of outstanding shares θ_t at each point in time t.

$$\max_{\theta_t} \frac{\theta_{t-1}}{\theta_t} \pi(K(\theta_t)) = \max_{\theta_t} \frac{\theta_{t-1}}{\theta_t} f(P(\theta_t)) \tag{2}$$

The more shares are issued, the more capital is available and the higher are the profits of the firm. At the same time, issuing more shares also dilutes existing shareholders, decreasing their profit share. The optimal θ, where these two effects are in balance, is reached when:

$$\frac{d}{d\theta_t} \frac{\theta_{t-1}}{\theta_t} f(P(\theta_t)) = 0$$

This results in:

$$\theta\, p(\theta) = \frac{f(P(\theta))}{f'(P(\theta))} \tag{3}$$

This is a maximum if the second derivative at this point is negative. This is the case as long as:

$$\frac{p'(\theta)}{p(\theta)^2} < -\frac{f''(P(\theta))}{f'(P(\theta))} \tag{4}$$

The right side of this inequality is positive thanks to the Inada conditions. The inequality holds for constant prices (implying $p'(\theta) = 0$) and more generally for prices $p(\theta)$ that do not grow too quickly, whereas the meaning of "quickly" depends on the shape of the production function.

Noting that $\theta p(\theta)$ corresponds to the internal valuation of the firm as defined in equation 1, one can conclude that choosing the optimal θ implies:

$$V_i(\theta) = \theta\, p(\theta) = \frac{f(P(\theta))}{f'(P(\theta))} \tag{5}$$

4.3 The Investor's Optimization Problem

The aim of this section is to derive the investor's valuation function. It provides the current valuation of the firm from the point of view of a marginal investor and is based on potential capital gains, dividends, and the return r of the outside option. Given that capital gains and dividends are fully determined by the number of outstanding shares θ_t at each point in time t, the external valuation function looks as follows:

$$V_e(\theta_t, \theta_{t+1}, r) = \frac{p(\theta_{t+1}) - p(\theta_t)}{r}\theta_t + \frac{f(P(\theta_t))}{r} \tag{6}$$

This valuation function is shown to hold for any increasing utility function and any discount rate. In equilibrium, it coincides with the valuation function of a strategic investor. In contrast to the strategic investor, the marginal investor does not consider the impact of buying or selling a marginal share on the value of other shares the investor already holds or might hold in the future.

Investors maximize life-time utility with discount rate β and an utility function $U(c)$ with $U'(c) > 0$:

$$U = \sum_t \beta^t U(c_t)$$

At each point in time, the investor holds ϵ_t shares, whereas θ_t denotes the shares held by everyone else. As an alternative to investing his wealth W_t in shares, the investor can also deposit it with interest rate r. Labeling capital $K_t = P(\theta_t + \epsilon_t)$ and investment $I_t(x) = P(\theta_t + x) - P(\theta_t)$, the amount available for consumption at time $t + 1$ is determined by the decisions taken in the previous period and captured by the following term:

$$c_{t+1} = I_{t+1}(\epsilon_t) + \frac{\epsilon_t}{\theta_t + \epsilon_t} f(K_t) + (1 + r)(W_t - I_t(\epsilon_t)) - W_{t+1}$$

It consists of the current value of the shares bought in the previous period $I_{t+1}(\epsilon_t)$, the dividend income $\frac{\epsilon_t}{\theta_t+\epsilon_t} f(K_t)$, the part of the previous periods savings that was not used to buy shares and the interest on these savings $(1 + r)(W_t - I_t(\epsilon_t))$, minus the wealth W_{t+1} put aside for the next period. Taking the derivative of life-time utility with respect to one particular ϵ_t, all but one summands drop out and one obtains the following first order condition:

$$\frac{dU}{d\epsilon_t} = \beta^{t+1} U'(c_{t+1}) \frac{d}{d\epsilon_t} c_{t+1} = 0$$

Since $U'(x) > 0$ and $\beta > 0$ by assumption, this implies that $\frac{d}{d\epsilon_t} c_{t+1}$ must be zero, leading to:

$$p(\theta_{t+1} + \epsilon_t) + \frac{\theta_t}{(\theta_t + \epsilon_t)^2} f(K_t) + \frac{\epsilon_t}{\theta_t + \epsilon_t} f'(K_t) p(\theta_t + \epsilon_t) = (1 + r) p(\theta_t + \epsilon_t)$$

In the case of the single, strategic investor with all other shareholders staying passive, $\theta_t = \theta_o$ a can be assumed constant. Under these conditions, the above condition simplifies to:

$$\theta_o f(K_t) + (\theta_o + \epsilon_t) p(K_t) \epsilon_t f'(K_t) = (\theta_o + \epsilon_t)^2 r p(K_t)$$

This implies that future prices do not matter for the strategic investor. Why is that? Wouldn't anticipated price changes pose an opportunity to make capital gains? No. Given a path-independent pricing function and our assumption of the strategic investor being the only active investor, any such gains would evaporate when trying to realize them. The strategic investor would be playing a zero-sum game with herself.

The first point of interest is where the strategic investor and the firm have both no incentive to buy or sell, thereby denoting a market equilibrium. This is found by plugging equation 5 into the strategic investor's first order condition, replacing $(\theta_o + \epsilon_t)p(K_t)$ with $\frac{f(K_t)}{f'(K_t)}$. After simplifying $\theta_o + \epsilon_t = \theta$, this yields the strategic investor's valuation function:

$$V_s(\theta, r) = \frac{f(P(\theta))}{r} \tag{7}$$

In order to verify that this point represents a maximum (and not a minimum) for the strategic investor, we take the second derivative of the maximization problem with regards to ϵ_t, evaluate it at the same point, and require it to be negative, obtaining:

$$\frac{\theta_t}{\theta_t + \epsilon_t}(f''(K_t)p(K_t)^2 + f'(K_t)p'(K_t)) < rp'(K_t)$$

Considering the expression in the brackets, one finds that this inequality holds when equation 4 is fulfilled, making equation 4 a sufficient condition for making the point at which the two valuation functions V_s and V_f meet a maximum for both the strategic investor and the firm.

Let us now turn to the case of a marginal investor and the question when it makes sense to buy (or sell) a marginally small amount of shares. This time, I allow the amount of shares held by the other shareholders to vary over time. The intratemporal optimization of consumption c_t tells us that buying ϵ_t shares with an investment $I(\theta_t, \epsilon_t) = P(\theta_t + \epsilon_t) - P(\theta_t)$ is worthwhile for a new shareholder as long as the capital gains from selling the ϵ_t shares at $t+1$ and the earned dividends are at least as high as the interest income that could be earned otherwise:

$$I(\theta_{t+1}, \epsilon_t) - I(\theta_t, \epsilon_t) + \frac{\epsilon_t}{\theta_t + \epsilon_t} f(K_t) \geq rI(\theta_t, \epsilon_t)$$

Since this is satisfied with equality at $\epsilon_t = 0$, we know that it is satisfied for all values of ϵ_t if its derivative is also satisfied for all values of ϵ_t. The derivative with respect to ϵ_t is:

$$p(\theta_{t+1}) - p(\theta_t + \epsilon_t) + \frac{\theta_t}{(\theta_t + \epsilon_t)^2} f(P(\theta_t + \epsilon_t)) + \frac{\epsilon_t}{\theta_t + \epsilon_t}(...) \geq rp(\theta_t + \epsilon_t)$$

Letting ϵ_t approach zero, we recognize the following condition in the limit:

$$p(\theta_{t+1}) - p(\theta_t) + \frac{f(P(\theta_t))}{\theta_t} \geq p(\theta_t) r$$

This condition says that the capital gains per share and the dividends per share must outweigh the opportunity cost of not being able to earn interest on the price of a share. It can also be expressed in terms of valuation, saying that the valuation of the company must not exceed the value of the discounted dividends and price appreciation, leading us to equation 6, which is a more general version of the strategic investor's valuation function 7.

4.4 Efficiency

The equilibrium is reached when the firm and the investors are indifferent between buying and selling additional shares. At this point, their valuations match:

$$V_e(K, r) = \frac{p(\theta_{t+1}) - p(\theta_t)}{r}\theta_t + \frac{f(K)}{r} = \frac{f(K)}{f'(K)} = V_i(K)$$

4.4.1 Stationary Case

For the stationary case with $\theta_{t+1} = \theta_t$, the market equilibrium is reached when the marginal return on capital matches the interest rate: $f'(K) = r$. This is the efficient capital allocation. Apparently, our setup leads to a Pareto-efficient outcome even if everyone only has their self-interest in mind, at least in the stationary case. Furthermore, it is not necessary to have a strategic investor to reach the equilibrium. The equilibrium can be reached incrementally with a series of marginally small investments that pay off on their own, as required in the problem statement.

Figure 1 shows how the two valuations depend on the capitalization of an example firm with production function $f(K) = K^{0.66}$. As long as the firm is selling shares below the valuation of the investor, the investor keeps buying. An additional example with Cobb-Douglas production $f(K) = K^{\alpha}$ can be found in appendix A.

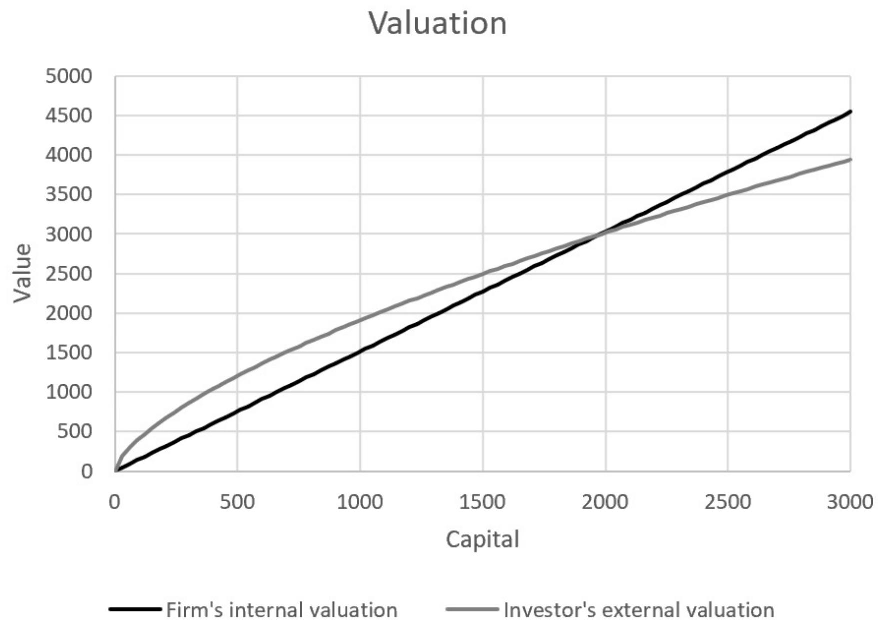

Figure 1: Illustration with $f(K) = K^{0.66}$, $r = 0.05$ and equilibrium $K^* \approx 2000$

An attentive reader might have noticed that I only derived the incentive constraints for two extreme cases of investors: a sole strategic investor and a marginally small investor. Without providing a rigorous proof, I argue that all other cases fall in between the two extremes.

A similar observation can be made for the amount an investor is willing to pay for a marginal share outside the equilibrium. In the example from figure 1, a marginal investor is willing to buy shares at a valuation of roughly 1200 given a capital level of 500, assuming the investor does not know whether others will buy further shares. In contrast, a strategic investor would be willing to buy at a valuation of up to 2000, anticipating further purchases up to the optimal level. Everyone else would be willing to make a bid in between.

4.5 Outstanding Shares and Their Price

One might wonder how many shares $\theta(K)$ the market participants are getting for the invested capital K. Anchoring the number of issued shares at $\theta(K_0) = 1$, function $\theta(K)$ is:

$$\theta(K) = \int_{K_0}^{K} \frac{1}{p(k)} dk + \theta(K_0) = \int_{K_0}^{K} \frac{f'(k)}{f(k)} \theta(k) dk + 1 = \frac{f(K)}{f(K_0)} \tag{8}$$

In combination with the internal valuation function 5, this implies that the price per share depends on capital as follows:

$$p(K) = \frac{f(K_0)}{f'(K)}$$

4.5.1 Anticipated Price Changes

So far, we have seen that the efficient outcome is reached in the stationary case, with interest rates and the parameters of the production function being constant. In the next section, I will show that the outcome stays efficient under a wide range of unanticipated shocks. However, one major weakness of the model is that it can produce inefficient capital in the face of anticipated price changes from one period to the next.

In case of an anticipated price increase with $\theta_{t+1} > \theta_t$ and therefore $p(\theta_{t+1}) > p(\theta)$, the investors will over-invest, injecting more capital than necessary into the firm. Similarly, an anticipated price decrease leads to under-investment. Figure 2 illustrates an example with a series of anticipated technology shocks.

A simple way to restore efficiency would be to allow the firm to invest excess capital at the same interest rate r as everyone else (or to borrow at that rate in case of under-investment). This, however, would go against the article's premise that debt financing is not available and the idea of putting capital allocation into the invisible hands of the market.

Another argument one could make is that long-term anticipated price changes are rather rare and if there are any, the company might be the first to know and therefore be able to adjust its price before the market reacts. However, this would erode the deterministic nature of the pricing function.

I have a number of further ideas for how this problem could be alleviated, but for now, I leave it unresolved. After all, it is just of temporary nature. As soon as the anticipated change has materialized, the efficient capital allocation is reached again. Before spending too much time on elaborate schemes to counter this temporary problem, I'd like to turn our attention to how robust the pricing function is under various model extensions.

5 Extensions and Robustness

The outstanding property of valuation function 5 is its robustness against a wide range of shocks. The robustness is owed to its simplicity and the fact that it only depends on internal observables, namely capital K and the shape of the production function $f(K)$. The valuation function does not budge in the presence of technology shocks or interest rate changes. When extending the function

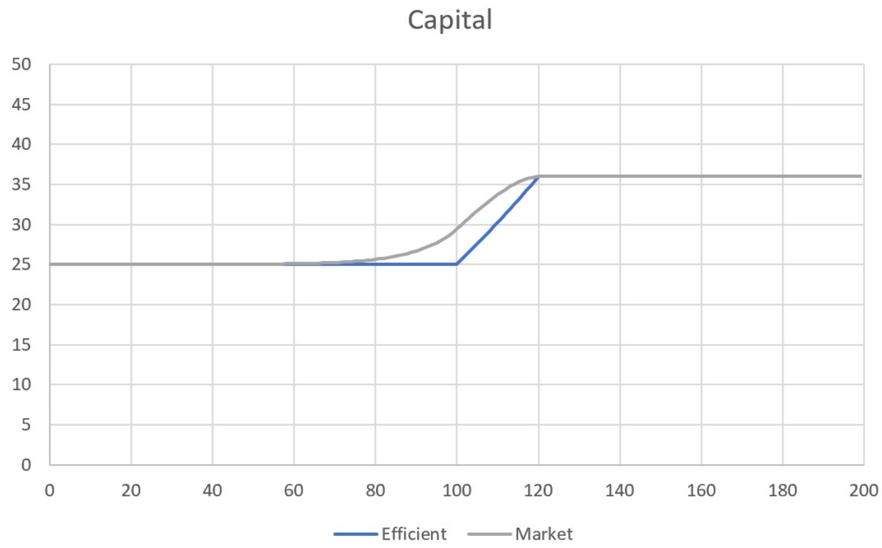

Figure 2: Temporary over-investment in an example with interest $r = 0.1$, production function $f(K) = A\sqrt{K}$ and an anticipated, gradual increase of the technology factor A from 1.0 to 1.2 over the course of 20 periods.

to include additional input goods in the style of Cobb and Douglas, it is even indifferent about input price changes! This makes our valuation function an ideal candidate for fully automated market making. A numerical example for what happens during an interest rate change can be found in appendix B. Finally, some considerations on how dividends should be handled are made in section 5.3, introducing *capital drift*.

5.1 Technology Shocks

Valuation function 5 is clearly robust against interest rate shocks, as r does not even enter the equation. But what about technology shocks?

To investigate this, I extend the production function to probabilistic technology shocks of the form proposed by Diamond [1967]:

$$y(X, K) = g(X)f(K)$$

Here, y is the stochastic output of a firm that depends on a random variable X and a capital level K. As the stochastic effect is multiplicative, it drops out of the equation:

$$V_i(X, K) = \frac{y(X, K)}{y'(X, K)} = \frac{g(X)f(K)}{g(X)f'(K)} = \frac{f(K)}{f'(K)}$$

Therefore, the market maker does not need to adjust its price when confronted with technology news. An exception are technological changes that impact the shape of $f(K)$, for example a change in α with Cobb-Douglas production $f(K) = K^\alpha$. Nonetheless, it is great to see that the valuation function V_i is robust against all kinds of multiplicative changes.

5.2 Input Price Changes

So far, we have assumed a simplistic production function $f(K)$ that only takes capital as input, resulting in profits $\pi(K) = f(K)$. Now, I extend the model with an additional input good x that costs price p_x and impose Cobb-Douglas production. Besides having nice mathematical properties, Cobb-Douglas production functions have an economically plausible shape. [Jones, 2005] In that case, profits are:

$$\pi(K, x) = f(K, x) - p_x x = K^\alpha x^\beta - p_x x$$

The price of the output good has been normalized to 1 without loss of generality. It is assumed that $\alpha > 0$, $\beta > 0$ and $\alpha + \beta < 1$.

Given capital K and price p_x, the firm chooses x^* to maximize profits, which happens at $\frac{d}{dx} f(K, x) = p_x$, leading to:

$$x^* = \beta^{\frac{1}{1-\beta}} K^{\frac{\alpha}{1-\beta}} p_x^{\frac{1}{\beta-1}}$$

Plugging the optimal x^* back into the profit function reveals:

$$\pi(K, x^*) = K^{\frac{\alpha}{1-\beta}} p_x^{\frac{\beta}{\beta-1}} \beta^{\frac{1}{1-\beta}} (\frac{1}{\beta} - 1) = K^\gamma p_x^\delta b$$

with $0 < \gamma < 1$, $\delta < 0$ and $b > 0$. Both p and b drop out when applying the internal valuation function:

$$V_i(K) = \frac{\pi(K, x^*)}{\pi'(K, x^*)} = \frac{K^\gamma}{\gamma K^{\gamma-1}} = \frac{K}{\gamma}$$

So at least for Cobb-Douglas production, our valuation function is indifferent to input or output price shocks. Adding further input factors would not change that either. Generally, when having a production function of the form $f(K, X) = K^{\alpha_0} \prod x_i^{\alpha_i}$ with many input factors, the valuation function is:

$$V_i(K) = \frac{1}{\alpha_0} K - \frac{\sum \alpha_i}{\alpha_0} K + K \tag{9}$$

For the special case of constant returns to scale ($\sum \alpha_i = 1$), this simplifies to $V_i(K) = K$. For the special case where K is the only input factor, it simplifies to $V_i(K) = \frac{1}{\alpha} K$.

Interestingly, equation 9 could be used to argue that there is a direct connection between the returns to scale of the production function of a company and its price-to-book ratio. It implies that a firm with returns to scale α should in the absence of debt have a price-to-book ratio of $\frac{1}{\alpha}$. It might be worthwhile to explore this relation more deeply in a separate publication.

5.3 Capital Drift

The basic model implies that all profits magically move from the firm to the shareholders as they happen. In practice, profits accrue within the company before they are paid out in a periodic dividend to the shareholders. Economically equivalent, but more flexible, are share buyback programs. In case of a continuous capital corporation, a buyback program can be implemented by letting the firm continuously adjust the price at which it sells and repurchases its own shares, thereby setting the incentives for capital inflows or outflows in accordance with the accrued profit or loss. This again moves control from the firm to the market.

When a firm with capital K_t makes profits π_t without the hitherto assumed immediate payout to the shareholders, its capital increases to K_{t+1} in the beginning of the next period. This in turn increases the valuation according to the internal valuation function $V_i(K)$, but not the number of shares, leading to inconsistency $V_i(K_t) = p(\theta_t)\theta_t = p(\theta_{t+1})\theta_{t+1} \neq V_i(K_{t+1})$ as $\theta_{t+1} = \theta_t$. This can be corrected by letting the pricing function depend on time t, extending it to $p(\theta, t)$ such that:

$$\frac{p(\theta, t+1)}{p(\theta, t)} = \frac{V_i(K_{t+1})}{V_i(K_t)}$$

In case of Cobb-Douglas production, this simply implies that the share price should be increased in proportion with the accrued capital:

$$\frac{p(\theta, t+1)}{p(\theta, t)} = \frac{K_t + \pi_t}{K_t}$$

For example, a continuous capital corporation with Cobb-Douglas production and a return on equity of $\pi/K = 10\%$ should let the price $p(\theta, t)$ drift upwards by also 10% over the course of a profit period. Doing so will continuously attract sellers, pulling profits out of the company as they happen and making the firm stay at the efficient level of capital. While the valuation $V_i(K)$ also stays the same, the shareholders gain π in cash, the number of outstanding shares declines, and the price per share increases.

Essentially, capital drift is a share buyback program. But unlike traditional buyback programs, the exact extent and timing of capital outflows are not driven by the managers and the traders they hire, buy by the open market and its participants. The continuous capital corporation provides the incentives for the sellers, but the sales themselves are coming from the market.

Companies could also use capital drift as a tool to adjust the capital level of a company and to bring it back onto the optimal capitalization path. Capital drift could even be negative, creating an incentive to inject more capital into the firm. We will make use of this tool in the next section, when applying our insights to real-world scenarios.

6 Relation to Traditional Market Making

This section discusses the most important differences between the market making a continuous capital corporation engages in and that of traditional market makers, looking at the inventory risk and the pricing risk.

6.1 Inventory Risk

Garman [1976] formally describes the problem of the market maker in a very similar setting as ours: the market maker is assumed to be a price-setter, defining the relation between demand and the prices. However, what complicates Garman's model significantly is the requirement that the market maker has a limited inventory and must never run out of shares to sell. The same applies to the model of Amihud et al. [1986]. Among other things, they show that the market maker's behavior is to a significant extent driven by the risk of running out of stock and that the replenishment costs play a significant role.

In contrast to a traditional market marker, the continuous capital corporation does not fear running out of shares, as it can always print new ones. Furthermore, unlike owning another company's shares, owning one's own shares does not come with a market risk or a liquidity risk.

A share owned by the firm itself is economically equivalent to a share that does not exist at all. It cannot vote and it does not lead to an outflow of dividends. According to the latest accounting standard, it is not even a position on the active side of the balance sheet, but a negative position on the passive side, leading to an overall situation that is equivalent to these shares not existing at all.

When a firm does its own market making, the inventory risk vanishes, leading to a significant simplification of the process. What is left, however, is the risk of paying too much for a share or selling a share for too little. This is discussed in the next section.

6.2 Pricing Risk

From an accounting point of view, transactions with own shares are profit-neutral. They are economically equivalent to a capital increase or decrease. However, just with the capital increase or decrease, it is still possible to hurt the shareholders by either unduly diluting them (when selling too cheap) or by unduly reducing the value of the company (when buying back for too much). For the traditional market maker, this means that she has to live in constant fear to trade against a better informed trader who knows more about the true value of the firm than the market maker does. If there are too many informed traders, market making can even become unprofitable regardless of the spread, making the market illiquid.

While the pricing risk cannot be eliminated with the continuous capital corporation, it can be significantly reduced and the risk of trading against better informed traders even turned into an opportunity. When a traditional market maker sells for price p even though the fundamental value is $p_f > p$, she has made a loss of $p_f - p$. However, in case of a continuous capital corporation, the proceeds of the sale are used to capitalize the company, thereby creating additional value for everyone and making the transaction Pareto-improving when selling for a price at least as high as valuation function 5 dictates. Of course, there is still the opportunity cost of not having sold the share at the highest possible price, but the fact that value is being created through the addition of capital turns the zero-sum game of the stock market into a mutually beneficial scenario where everyone can win.

The same holds when repurchasing shares at a price that does not exceed that given by valuation function 5. Here, the added value stems from the informed seller being able to reinvest the capital at a higher return elsewhere, thereby also making the transaction Pareto-improving.

7 Conclusion

7.1 Learnings

What have we learned from all of this? The conclusions are three-fold.

First, we have found a valuation function that fulfills the criteria from the problem statement in section 3, supporting the efficient outcome under a wide range of circumstances. It provides guidance for companies that want to make use of the new legal and technical possibilities to create a small market for their own shares. A company that does so is considered a continuous capital corporation. Its capital is continuously adjusted as the market price of its shares moves up and down.

Second, having a rigid mechanism that links capital allocation to market prices shifts the capital allocation decision into the invisible hand of the free market, potentially obsoleting dividend payments and other management-driven financing activities. Whether this is a good idea depends on how much trust you have in the management of a firm to allocate the optimal amount of capital versus the trust you have in the market to allocate the capital where it is most productive.

As a small-scale investor, I like the thought of actually investing in a company when buying its shares. Today, buying a Tesla share on the stock market just sends money to its previous holder. There is no direct connection to the company itself. However, if Tesla was a continuous capital corporation, buying its shares would actually inject cash into the company, maybe allowing it to speed up the roll-out of its next great car. The continuous capital corporation would allow impact investing to actually have an impact.

Third, there remains a lot of unresolved work to make the continuous capital company function in practice. In order to apply even the simple valuation function 5, one would need to know the shape of the production function of a firm and its current capital level. But in practice, even the capital of a company might be non-trivial to determine as there are various forms of capital that might show up or not show up on the balance sheet depending on the applied accounting methods. As clean and simple as equation 5 looks, it is not straight-forward to apply.

Furthermore, even if all problems of practical applicability were solved, there would still be the fundamental issue that the suggested pricing function represents the participation constraint of the current shareholder such that by definition, all created surplus goes to the new investors. In reality, firms will try to sell at a higher price than what the valuation function suggests and repurchase at a lower price in order to capture their fair share of the created surplus when the capital allocation of the economy is improved.

7.2 Outlook

Whereas it is questionable whether we will see a pure continuous capital corporation anytime soon in the wild, I expect an increasing number of companies to make use of the new legal and technical possibilities and to experiment with approaches between the one extreme of having a market that is detached from the firm's capital and the other extreme of having a market that fully determines a firm's capital.

A firm could achieve this relatively easily by choosing a pricing function that is steeper than what the presented model suggests. That way, the market would have some impact on the capital allocation, but the firm would still depend on some traditional management-driven financing to fully reach the efficient level. This also alleviates the surplus allocation problem as a larger fraction of the created surplus would end up with the initial shareholders. In general, having an incremental path towards an innovative new form of financial markets is invaluable to its adoption.

I hope that many European firms will recognizes these new opportunities and start creating markets for their own shares, thereby unlocking eliminating the illiquidity discount, [Damodaran, 2005] enabling founders or other significant shareholders to diversify their assets, and enabling seed stage investors to exit a firm that has grown out of the seed stage, using the proceeds for other seed stage investments. Further, if this could enable more firms to tap into a broader financial market it would broaden the investment universe for the common investor again, countering the decline in the number of listed companies, an unfortunate trend described by Doidge et al. [2018] and first postulated by Jensen [1989].

As the proof of the pudding is in the eating, I am currently testing the presented ideas with my own company. Its shares can be bought and sold on its own website by anyone. The price is deterministically adjusted in accordance with a linear valuation function. [Aktionariat, 2021] How that turns out remains to be seen.

References

Uniswap. Uniswap protocol analytics. 2021. URL info.uniswap.org/home.

Coinmarketcap. Uniswap. 2021. URL coinmarketcap.com/currencies/uniswap.

Bundesrat. Anpassung des bundesrechts an entwicklungen der technik verteilter elektronischer register. 2020. URL https://www.parlament.ch/fr/ratsbetrieb/suche-curia-vista/geschaeft?AffairId=20190074.

James Dow and Gary Gorton. Stock market efficiency and economic efficiency: is there a connection? *The Journal of Finance*, 52(3):1087–1129, 1997.

Katrin Tinn. Blockchain and the future of optimal financing contracts. *Available at SSRN 3061532*, 2017.

Bruno Biais, Thomas Mariotti, and Jean-Charles Rochet. Dynamic financial contracting. *Advances in economics and econometrics*, 1:125–71, 2013.

Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008. URL https://bitcoin.org/bitcoin.pdf.

Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.

Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An analysis of uniswap markets. *arXiv preprint arXiv:1911.03380*, 2019.

Meni Rosenfeld. Formulas for bancor system. 2017. URL https://drive.google.com/file/d/0B3HPNP-GDn7aRkVaV3dkVl9NS2M/view.

Monique Schnell Luchsinger and Pascal Montavon. Der erwerb eigener anteile durch die ag und die gmbh – 2. teil: Steuerliche aspekte. *Der Treuhandexperte*, (5):284–297, 2018. URL https://carlicahn.com/our_letter_to_tim_cook.

Swiss Federal Court. 2c_119/2018. 2019.

Peter Forstmoser and Marcel Küchler. Schweizerische aktienrechtsreform: Die schlussrunde ist eingeläutet! *Jusletter*, 2020. URL https://jusletter.weblaw.ch/juslissues/2020/1010/schweizerische-aktie_0417b4599b.html.

Hans Caspar von der Crone, Martin Monsch, and Luzius Meisser. Aktien-token. *GesKR*, (1): 1–17, 2019.

Deutsche Bundesregierung. Gesetz zur einführung von elektronischen wertpapieren. 2020. URL www.bmjv.de/SharedDocs/Gesetzgebungsverfahren/DE/Einfuehrung_elektr_Wertpapiere.html.

In-Koo Cho and Thomas J Sargent. Self-confirming equilibrium. *The new Palgrave dictionary of economics*, 2008.

Peter A. Diamond. The role of a stock market in a general equilibrium model with technological uncertainty. *The American Economic Review*, 57(4):759–776, 1967. ISSN 00028282. URL http://www.jstor.org/stable/1815367.

Thorsten Hens, Sabine Elmiger, et al. Economic foundations for finance. *Springer Texts in Business and Economics*, 2019.

Charles I Jones. The shape of production functions and the direction of technical change. *The Quarterly Journal of Economics*, 120(2):517–549, 2005.

Mark B Garman. Market microstructure. *Journal of financial Economics*, 3(3):257–275, 1976.

Yakov Amihud, Haim Mendelson, et al. Asset pricing and the bid-ask spread. *Journal of financial Economics*, 17(2):223–249, 1986.

Aswath Damodaran. Marketability and value: Measuring the illiquidity discount. *Available at SSRN 841484*, 2005.

Craig Doidge, Kathleen M Kahle, G Andrew Karolyi, and René M Stulz. Eclipse of the public corporation or eclipse of the public markets? *Journal of Applied Corporate Finance*, 30(1): 8–16, 2018.

Michael C Jensen. Eclipse of the public corporation. *Harvard Business Review*, (5):61, 1989. URL https://hbr.org/1989/09/eclipse-of-the-public-corporation.

Aktionariat. Investor relations. 2021. URL aktionariat.com/investors.html.

A Basic Example

This example shows what happens when a continuous capital corporation raises capital in the basic model.

The firms is assumed to have production function $f(K) = \sqrt{K}$. Interest rates are $r = 0.1$ and the firm is initially equipped with capital $K_0 = 1$ by its founders who hold $\theta_0 = 1$ shares. The optimal capital allocation is $K^* = 25$, at which point $f'(K) = r$.

When the firm starts offering shares at the valuation $V_i(K_0) = 2$ according to the internal valuation function 5, market participants immediately recognize that this is an attractive buying opportunity as the fundamental value of the firm is $V_e(K_0) = 10$, calculated using the external valuation function 7. The market participants buy shares until $V_i(K) = V_e(K)$, which happens at the efficient capital level $K_1 = 25 = K^*$.

But how many shares did the market participants get for their investment of $K_1 - K_0 = 24$ units of capital? This can be calculated with the function derived in equation 8. The market participants receive $\theta(K_1) - \theta(K_0) = 4$ shares, bringing the total number of outstanding shares to $\theta(K_1) = \sqrt{K} = 5$.

The existing shareholders started with fully owning a firm worth $V_e(K_0) = 10$ and ended with owning $\frac{1}{5}$ of a firm worth $V_e(K_1) = 50$, neither losing nor gaining anything despite the firm selling shares below the fundamental value. The new investors provided $K_1 - K_0 = 24$ in capital and got $\frac{4}{5}$ of the firm, increasing their wealth by 16. All of the surplus went to the new investors.

One might wonder what happened if the firm used the external valuation $V_e(K)$ instead of $V_i(K)$ to sell its shares. To find out, we need to derive the number of shares $\theta_e(K)$ issued when K is invested at the external valuation:

$$\theta_e(K) = \int_{K_0}^{K} \frac{1}{p_e(k)} dk + \theta(K_0) = \int_{K_0}^{K} \frac{r}{f(k)} \theta_e(k) dk + 1$$

This time, the investors only get $\theta(K_1) - \theta(K_0) \approx 1.22$ shares (solved numerically) and the initial shareholders can significantly increase their wealth. They start with fully owning a firm worth 10 and end up with owning $\frac{1}{2.22}$ of a firm worth 50, gaining about 12.5. Surprisingly, some of the total surplus of 16 still goes to the new investors. This is owed to the observation that the new investors become existing shareholders themselves as they invest and increasingly participate in the surplus assigned to the existing shareholders with each marginal trade.

B Example with an Interest Rate Shock

Given a continuous capital corporation with production function $f(K) = \sqrt{K}$, this example starts with capital $K_0 = 25$, interest rate $r_0 = 0.1$, and a total number of outstanding shares $\theta_0 = 5$ and shows what happens when the interest rate drops to $r_1 = 0.05$.

When the interest rate drops to r_1, the internal valuation does not change, but the external valuation jumps from $V_e(K_0, r_0) = f(K_0)/r_0 = 50$ to $V_e(K_0, r_1) = f(K_0)/r_1 = 100$. This attracts new investors who buy freshly printed shares until the internal and the external valuation are in balance again at the new efficient capital level $K_1 = 100$. The new investors get $\theta(K_1) - \theta(K_0) = 5$ shares for their investment, bringing the total to 10.

Before the transaction, the old shareholders fully owned a firm worth 100. After the transaction, they own 50% of a firm worth $V_e(K_1, r_1) = \frac{10}{0.05} = 200$, neither losing nor gaining anything. The new investors see their wealth increase from 75 in cash to owning shares worth 100. The transaction is Pareto-improving.

In case of a drop of the interest back to $r_2 = 0.1$, everything would symmetrically unwind and shareholders would return shares to the market maker until the initial capital level $K_2 = K_0$ is reached again. The shareholders who sell will move from owning 50% of a firm worth $V_e(K_1, r_2) = \frac{10}{0.1} = 100$ to owning 75 in cash, gaining 25 units of capital. The remaining shareholders move from owning 50% of a firm worth 100 to completely owning a firm worth 50, neither gaining nor losing anything. Again, the transaction is a Pareto-improvement.