

Is Exclusion Effective?

September 2019

David Blitz and Laurens Swinkels

David Blitz is Head of Quantitative Research at Robeco, the Netherlands

d.c.blitz@robeco.com

Laurens Swinkels is Senior Researcher at Robeco, the Netherlands and Assistant Professor of Finance at Erasmus University, the Netherlands

l.swinkels@robeco.com

lswinkels@ese.eur.nl

- Excluding firms with the worst sustainability profiles boils down to a transfer of ownership. It is not obvious how this is supposed to lead to changes for the better in society.
- We challenge the effectiveness of four frequently used arguments in favor of exclusion policies: increasing a firm's cost of capital, pushing a firm out of business, improving investment performance, and signaling stakeholders to change behavior.
- Altogether, the effectiveness of exclusions is questionable. Investors may well achieve more by exerting influence as an active owner, through voting and engaging with firms.

Accepted for publication in

Journal of Portfolio Management

<https://doi.org/10.3905/jpm.2020.46.3.042>

Important disclaimer: the views expressed in this article are solely those of the author and are not necessarily shared by Robeco or its subsidiaries

Is Exclusion Effective?

Abstract

Many believe that investors can contribute to a more sustainable world by divesting from firms with the worst sustainability profiles. However, exclusion comes down to a transfer of ownership from sustainability-minded investors to other investors, and it is not obvious how this is supposed to lead to changes for the better in society. This article critically examines the arguments for exclusion and concludes that the effectiveness of exclusion policies is questionable. Investors may well achieve more by exerting influence as an active shareholder, through voting and engaging with firms.

Investors increasingly wish to incorporate sustainability considerations into their investment policy. One way to go about sustainability integration is to obtain the ESG (Environmental, Social, Governance) score of each stock and require that the portfolio has a significantly higher average ESG score than the benchmark. Instead of looking at a broad ESG measure, which typically consists of dozens or even hundreds of underlying variables, one can also target some very specific metrics, such as the carbon footprint of the portfolio. Yet another approach is to make an allocation to a dedicated impact portfolio, which invests exclusively in certain sustainability themes, such as renewable energy technology.

An unavoidable question with sustainability integration is what to do with the worst offenders, i.e. the firms that are particularly problematic from a sustainability perspective. Classic examples of this are firms in the tobacco, weapons, gambling and alcohol industries. In the literature these are referred to as 'sin stocks', because their business models are based on the exploitation of human vice and weakness. Next to these classic sin sectors, sustainability-minded investors also struggle what to do with firms involved with, to name a few, thermal coal, palm oil, cannabis, child labor, or animal testing. Investors increasingly choose to black-list the firms with the worst sustainability profiles, i.e. to explicitly exclude such firms from their eligible investment universe. Although the other forms of sustainability integration do not explicitly exclude any particular firm, they implicitly force certain firms out of the portfolio. For instance, if a substantial improvement of the portfolio ESG score or carbon footprint is imposed, then this effectively means that some firms with very bad scores on these metrics need to be sold off.

Many believe that investors can contribute to a more sustainable world by divesting from firms with the worst sustainability profiles. However, exclusion comes down to a transfer of ownership from concerned investors to other (perhaps less, or not at all concerned) investors, and it is actually not obvious how this is supposed to lead to changes for the better in society. This is a crucial realization: any improvement in the sustainability profile of the portfolio achieved by an individual investor, i.e. at the micro-level, has zero direct implications at the macro-level, because the excluded securities will end up being owned by someone else. In other words, if one investor divests from a sin stock, another investor will be overweight this sin stock; if one investor improves the ESG score of the portfolio, another investor will have a deteriorated ESG score; and if one investor lowers the carbon footprint of the portfolio, another investor will have a higher carbon footprint, *by definition*.

Although exclusions have zero direct implications, they might contribute to a more sustainable world via other, indirect channels. In this article we critically examine the various arguments for exclusion. We first explore the notion that exclusion can increase the cost of capital of a firm, leading to reduced consumption of its products, or forcing it to adopt more sustainable business practices. We next discuss the more ambitious objective to cut a firm off from capital markets entirely, or to even push it out of business. A third motivation for exclusion may be to improve investment performance. The fourth and final argument which we discuss is that exclusion may be an effective signaling mechanism. Altogether, we conclude that the effectiveness of exclusion is questionable, and that investors may well achieve more by exerting influence as an active shareholder, through voting and engaging with firms.

INCREASING A FIRM'S COST OF CAPITAL

We begin by discussing whether exclusion can increase the cost of capital of a firm. A higher cost of capital can be argued to result in reduced consumption of a firm's products, because it needs to charge higher prices, or spend less on marketing and advertising. Also future consumption may be reduced, because with a higher cost of capital fewer projects will have a positive net present value. Alternatively, it may force a firm to adopt more sustainable business practices. Theoretical models predict that if many investors exclude a certain firm, the cost of capital of this firm will indeed go up; see, for example, Heinkel, Kraus, and Zechner [2001] and Hong and Kacperczyk [2009]. But do such theoretical arguments hold up in the real world?

The cost of capital is a weighted average of the cost of debt capital and the cost of equity capital. Determining the cost of debt capital is relatively easy, as it is directly observable in the corporate bond market. The empirical evidence appears to be mixed here, as Bauer and Hann [2010] find a higher cost of debt financing and lower credit ratings for firms with high environmental risk, while Fabozzi et al [2019] find that the most shunned firms experience lower rather than higher expected financing rates in the corporate bond market. In any case, the impact of the cost of debt capital tends to be relatively small compared to the impact of the cost of equity capital, which is the main determinant of the total cost of capital.

A popular approach for determining the cost of equity capital is to simply use the CAPM. Graham and Harvey [2002] find that the vast majority of CFOs use the CAPM to determine the cost of equity capital for their firm, Bancel and Mittoo [2014] find that also most analysts use the CAPM for determining the cost of equity capital, and Berk and van Binsbergen [2017] infer from the capital allocation decisions of mutual fund investors that also investors rely primarily on the CAPM to determine the discount rate. In the CAPM, the only firm-specific determinant of expected return is the market beta of a firm, which reflects its degree of covariation with the market. The covariance between the returns on a firm and the market is unlikely to be affected by the one-off exclusion decisions of a certain group of investors. This implies that if the cost of equity capital is based on the CAPM, an exclusion approach is unlikely to be effective at increasing it.

The CAPM is an elegant theoretical model, but its empirical track record is rather poor. Therefore, another way to determine the cost of capital is by looking at long-term realized returns of firms with certain characteristics. Fabozzi, Ma, and Oliphant [2008], Hong and Kacperczyk [2009], and Statman and Glushkov [2009] examine the historical performance of sin stocks, i.e. the typical exclusion candidates, and find strong evidence for positive abnormal returns, which is commonly interpreted as an exclusion premium. The idea here is that if many investors prefer to stay away from sin stocks, an extra financial incentive is needed to compensate the ones who are willing to take on this reputational risk. If long-term investor returns can be equated to the cost of capital, then these results suggest that sin stocks have, indeed, a higher cost of capital. However, Blitz and Fabozzi [2017] find that the outperformance of sin stocks can be fully explained by their factor exposures, in particular their exposure to the quality and low-risk factors. In other words, the historical performance of sin stocks is not

significantly different from the performance of non-sin stocks which have similar quality and low-risk characteristics. Based on these results they conclude that there is no evidence for a distinct sin premium. This means that there is no empirical support for a higher required cost of capital for sin stocks compared to non-sin stocks with similar characteristics.

Even if exclusion does lead to an increase in the cost of capital of a firm, the question is how much impact this has in the real world. If we take classic sin sectors such as alcohol and tobacco as an example, then for all practical purposes the cost of capital probably has a negligible impact on the amount of drinking and smoking. In many countries the price of a bottle of wine or a pack of cigarettes mainly consists of taxes. As such, the price for these products is effectively set by governments, rather than by firms themselves. Production and marketing costs of alcohol and tobacco producers only make up a small fraction of the price paid by the end-consumer. Thus, even if by excluding alcohol and tobacco firms investors would succeed in increasing the cost of capital of these firms, the impact of this on ultimate alcohol and tobacco consumption is dwarfed by the impact of government tax policy, which is the primary determinant of alcohol and tobacco prices. In practice, higher taxes on alcohol and tobacco appear to be a lot more effective at raising tax revenues than at reducing consumption. Perhaps the low price elasticity of these products stems from their addictive nature.

If the objective is to make a firm change its business practices, then exclusion may actually be interpreted as a sign of failure. If a majority of shareholders agrees that certain things should change, they can insist that the management of the firm acts upon this, or replace the old executive board with a new one that is willing to comply with these shareholder demands. Investors who resort to excluding a firm basically acknowledge that their view is the minority view and that they have given up on trying to get majority support for this. A constructive engagement approach may be more effective though at making a firm change its practices than a hostile exclusion approach. Evidence for the effectiveness of engagement is provided by Dimson, Karakaş, and X. Li [2015] and Barko, Cremers, and Renneboog [2017], who find significant improvements in the average Environmental, Social, and Governance (ESG) ratings of firms following shareholder engagement efforts.

STRANGLING A FIRM

The goal of exclusion may not be limited to increasing a firm's cost of capital. The more ambitious objective may be to fully cut a firm off from capital markets, or to even push it out of business entirely.

If the objective is to cut a firm off from capital markets, then it is the primary market that really matters, i.e. firms raising fresh money by issuing new stocks or bonds, needed to expand their business activities. The secondary market, i.e. the trading in stocks and bonds that are already out there, is irrelevant from a macro-economic perspective, as it merely transfers holdings from one investor to another. Although the primary market is what really matters, because here capital markets fulfill their crucial role of efficiently allocating fresh capital to firms in need of capital, one could argue that the only way to be credible with boycotting new issuance is by

refraining from investing in a firm altogether. If an exclusion approach is limited to the primary market one can expect to be arbitrated by other investors, who first subscribe to new stock or bond issuance of firms, and then simply pass on these stocks and bonds once they have become part of the secondary market. Full exclusion of a firm sends the most credible signal that one will not finance future growth of this firm by any means.

Cutting a firm off from capital markets may not be a very realistic objective though. For one, there is so much capital in search of attractive returns, that even if a large number of investors exclude certain firms, this does not make it much harder for these firms to attract fresh capital. Proponents of exclusion might counter that this only proves that exclusion at an even larger scale is needed. It is unclear, however, whether the degree of exclusion that is required to really cut firms off from capital markets is realistic. Moreover, many firms are not dependent on capital markets to finance their activities in the first place, because they have sufficient retained earnings. The popularity of share buyback programs clearly indicates that many firms have too much cash, rather than too little at their disposal. If firms anticipate that raising fresh capital in the future might be difficult, they can also be expected to reduce their dependence on external financing by hoarding cash.

The objective with exclusion may be even more extreme, namely to push a firm out of business entirely. Again the question is whether this is realistic. As long as there is demand for the products or services of the firm in question, supply is likely to find its way. Even if investors manage to drive a firm out of public markets, it may simply end up going private, like for instance the adult entertainment industry. Since private firms can operate more in the shadows and have less reporting obligations, this would make them less subject to public scrutiny. If investors really want to close a firm down, then they could consider buying up the whole firm, and then liquidating it. This capital destruction would of course be financially painful, but it would be a very tangible contribution to a more sustainable world. Again, however, this ignores that cutting off supply does not make demand disappear. As long as there is demand and as long it is legal, some profit-maximizing firm can be expected to seize that opportunity.

Investors who are intent on pushing certain firms out of business should also remember the US government's prohibition of alcohol in the nineteen twenties. Instead of reduced alcohol consumption, the main result of that policy turned out to be a dramatic increase in crime. The same can be seen nowadays with the war on drugs, which after decades of massive efforts has done little to reduce the availability and usage of drugs, but with a huge amount of crime involved. The majority of prison inmates in the US are convicted for drug-related crimes, drug trafficking is the main source of income for various terrorist organizations, and drug cartels destabilize various developing countries. Liberals but also a conservative economist such as Milton Friedman agree that legalizing drugs would be better for society. If the conclusion is that legalizing and regulating bad stuff such as alcohol, tobacco, and drugs is, on balance, a lesser evil than banning and fighting it, then investors should acknowledge that someone should be willing, or even encouraged to invest in these kind of activities. An individual investor can opt out of course, not wanting to be associated with these activities, but this means that some other investor should be expected to take on this role.

IMPROVING INVESTMENT PERFORMANCE

Another argument for exclusion may be to improve investment performance. Empirically, however, this argument is problematic, given that the typical exclusion candidates, the so-called 'sin' stocks, have historically delivered a significant outperformance. The recent finding of Blitz and Fabozzi [2017], that this outperformance of sin stocks is fully explained by loadings on the recently established quality and low-risk factors, implies that by excluding sin stocks investors effectively go short the attractive factor characteristics that come along with these stocks. Investors might be able to restore their expected portfolio return by replacing excluded stocks with non-controversial stocks that offer similar factor characteristics. However, this is mere damage control, i.e. it still does not support the argument that exclusion is beneficial for investment performance.

The notion that exclusion improves investment performance is also inconsistent with the argument that it is aimed at increasing the cost of capital of firms. This issue is well articulated by Cliff Asness from AQR, who even goes one step further by arguing that sustainability integration should cost performance.¹ His reasoning is that the purpose of not investing in 'bad' firms is to increase their cost of capital, and that the cost of capital should be equal to the long-term expected return for investors. In other words, the more successful one is at increasing a bad firm's cost of capital by staying away from it, the higher the expected return on that stock, and so the lower the return on a portfolio which excludes this stock. The magnitude of the underperformance would then be a measure of the success of the exclusion policy.

Some argue that excluding sin stocks reduces long-term investment risk. For instance, AP4 states that *"Tobacco shares have historically generated relatively high returns. However, AP4 estimates that risks have increased and that tobacco companies long-term risk becoming so-called 'stranded assets'. Altogether, AP4 believes that the risks associated with tobacco companies is not properly priced, but these risks will negatively affect the valuation of the tobacco companies long term."*² Similar arguments have been put forward by, for instance, CalPERS and CalSTRS, who also mention the risk of lawsuits. In other words, excluding sin stocks might just be a matter of good risk management. Andersson, Bolton, and Samama [2016] also use the stranded assets argument for divesting from the fossil fuel sector.

This argument seems rather opportunistic. Tobacco and fossil fuels are clearly not the only industries for which one can be concerned about long-term profitability and survival, yet all those other industries are not excluded because of that concern. It is also unclear why, in a market which is known to be pretty efficient, this public information would not be properly priced in already. Since the stranded assets argument seems impossible to prove or disprove (falsify) with evidence, it is essentially a matter of belief. Given the remarkably strong historical performance of sin stocks, the financial risk of excluding them may well be bigger than the risk

¹ <https://www.aqr.com/Insights/Perspectives/Virtue-is-its-Own-Reward-Or-One-Mans-Ceiling-is-Another-Mans-Floor>

² <http://www.ap4.se/en/esg/ap4-does-not-own-tobacco/>

of including them. The stranded assets argument also ignores the fact that the difference between a good and a bad investment depends on the price one pays for it, i.e. anything can be a good investment as long as the price at which it can be purchased is low enough.

If the objective is to improve investment performance, engagement again may be a more fruitful approach. With engagement, investors can convince a firm to adopt a more sustainable business model. If this also improves the (long-term) financial health of the firm, i.e. lead to higher expected future cash-flows or reduced business risks, then this may result in a share price appreciation. Evidence in support of this hypothesis is provided by Dimson, Karakaş, and X. Li [2015] and Barko, Cremers, and Renneboog [2017], who find that firms subject to engagement tend to generate positive abnormal returns and better operating performance in the subsequent period. A lot more research is needed before it is safe to conclude that this matter is settled, but if these results prove to be robust, then it implies that investors who divest from bad firms miss out on the upside potential that may be unlocked by making these firms change for the better. From an investment performance perspective, divesting then only makes sense if one assumes that engagement is either futile, or that the adoption of more sustainable business practices will destroy shareholder value.

SIGNALING

Exclusion decisions may also be motivated by signaling objectives. The signaling may be directed at the firm itself, consumers of the firm's products, policy makers, or the clients of investment management firms.

Even if exclusion does not have any direct effects on a firm, one might argue that it does send a powerful negative signal to the firm. In this way exclusion is a way of naming and shaming, as in shouting from the rooftops that the firm in question is among the worst of the worst. This may embarrass executives and employees of the firm, and even more so when it leads to critical questions from the media or from within their own social circle. The effectiveness of this mechanism is unclear though; for all we know it might just be wishful thinking.

Exclusion may also send a signal to consumers of a firm's products. The Tobacco Free Portfolios campaign uses the argument that *"Divesting tobacco company stocks and bonds serves to challenge and change the current status quo - where most workers worldwide are unwitting owners of tobacco companies. The collective impact of many financial organizations divesting from and ceasing commercial relationships with the tobacco industry helps to 'denormalize' tobacco companies. 'Denormalization' is a key goal of both global government and public health tobacco control efforts."*³ The reasoning here is that exclusion leads to denormalization, and that denormalization in turn leads to reduced consumption. But is it plausible to assume that exclusion by professional investors sends such a powerful signal that consumers stop smoking as a result? If high taxes and gruesome pictures on packs of cigarettes do not deter people from smoking, then they are unlikely to give up the habit upon hearing that their pension fund no longer invests in tobacco firms.

³ <http://www.tobaccofreeportfolios.org/wp-content/uploads/2017/11/TFP-Toolkit-3rd-Edition-3.3.pdf>

The signal of exclusion may also be targeted at policy makers, who are in the position to change laws and regulations. Braungardt, van den Bergh, and Dunlop [2019] mention this objective in their review of the arguments in favor and against fossil fuel divestment. Again, however, we can wonder if the signal from professional investors makes much of a difference, given that policy makers already get clear signals from the media and pressure groups such as NGOs.

Next to giving a signal to firms, consumers or policy makers, asset management firms and pension funds may also want to give a signal to their own clients. From a purely economic perspective it makes sense to refrain from investing in the least sustainable firms if this is appreciated by a certain clientele, while other clients do not mind (a limited number of) exclusions. Based on survey research Riedl and Smeets [2017] find that socially responsible investors are motivated by signaling objectives and are even willing to accept lower financial returns in order to invest in accordance with their social preferences. A profit-maximizing asset manager would be foolish to ignore such client preferences. In other words, an exclusion policy may be a rational response to perceived client demand. It may also be wise from the perspective of reputation risk, because holdings in sin stocks may result in bad press. Exclusion is the course of action that makes most sense in that case, as it sends a clear, unequivocal signal to end-clients that they do not need to be concerned about potentially controversial stocks being in their portfolios.

Cynical interpretations of this supposed client centrality are that in reality it is nothing more than opportunistic virtue signaling and green washing, or that investment managers like to emphasize sustainability issues because it is a lot easier to claim successes in this area than to beat the market with higher returns. Being tough on a small number of firms shows that one has good intentions, and that, for instance, one understands that a health insurance firm may feel uncomfortable with investing in tobacco firms given that smoking is the main cause of lung cancer. The question whether excluding tobacco stocks actually results in less deaths from smoking is hardly ever asked, because exclusion decisions are not evaluated in that way. Again, investing in firms and exerting shareholder influence might well accomplish more at the end of the day.

FINAL WORDS

Many believe that investors can contribute to a more sustainable world by divesting from firms with the worst sustainability profiles, but since exclusion is effectively a transfer of ownership from concerned to less-concerned investors, it is anything but obvious how this is supposed to lead to changes for the better in society. For instance, if investor A sells off tobacco stocks to investor B, not a single cigarette less will be lit by smokers the next day. The effectiveness of an exclusion policy is typically not evaluated by assessing its impact in the real world, even though this is what really matters. This article critically examined the various arguments for how exclusion might have a positive impact in the longer run. Since all arguments appear to be shaky, we conclude that it is questionable whether exclusion policies accomplish anything meaningful in the real world. In fact, rather than excluding firms, investors may well achieve more by exerting influence as an active shareholder, through voting and engaging with firms.

REFERENCES

- Andersson, M., P. Bolton, and F. Samama (2016). "Hedging Climate Risk." *Financial Analysts Journal* 72:3, 13-32.
- Bancel, F., and U.R. Mittoo (2014). "The Gap between the Theory and Practice of Corporate Valuation: Survey of European Experts." *Journal of Applied Corporate Finance* 26:4 (Fall), 106-117.
- Barko, T., M. Cremers, and L. Renneboog, (2018). "Shareholder Engagement on Environmental, Social, and Governance Performance." European Corporate Governance Institute Working Paper 509.
- Bauer, R., and D. Hann (2010). "Corporate Environmental Management and Credit Risk." SSRN working paper no. 1660470.
- Berk, J.B., and J.H. van Binsbergen (2017). "How Do Investors Compute the Discount Rate? They Use the CAPM." *Financial Analysts Journal* 73:2, 25-32.
- Blitz, D.C., and F.J. Fabozzi (2017). "Sin Stocks Revisited: Resolving the Sin Stock Anomaly." *Journal of Portfolio Management* 44:1, 105-111.
- Braungardt, S., J. van den Bergh, and T. Dunlop (2019). "Fossil Fuel Divestment and Climate Change: Reviewing Contested Arguments." *Energy Research and Social Science* 50 (April), 191-200.
- Dimson, E., O. Karakaş, and X. Li (2015). "Active Ownership.", *Review of Financial Studies* 28:12, 3225-3268.
- Fabozzi, F.J., A.S. Lamba, T. Nishikawa, R.P. Rao, and K.C. Ma (2019). "Does the Corporate Bond Market Overvalue Bonds of Sin Companies." *Finance Research Letters* 28 (March), 165-170.
- Fabozzi, F.J., K.C. Ma, and B.J. Oliphant (2008). "Sin Stock Returns." *Journal of Portfolio Management* 35:1, 82-94.
- Graham, J., and C. Harvey (2002). "How Do CFOs Make Capital Budgeting and Capital Structure Decisions?" *Journal of Applied Corporate Finance* 15:1 (Spring), 8-23.
- Heinkel, R., A. Kraus, and J. Zechner (2001). "The Effect of Green Investment on Corporate Behavior." *Journal of Financial and Quantitative Analysis* 36:4, 431-449.
- Hong, H., and M. Kacperczyk (2009). "The Price of Sin: The Effect of Social Norms on Markets." *Journal of Financial Economics* 93:1, 15-36.
- Riedl, A., and P. Smeets (2017). "Why Do Investors Hold Socially Responsible Mutual Funds?" *Journal of Finance* 72:6, 2505-2550.
- Statman, M., and D. Glushkov (2009). "The Wages of Social Responsibility." *Financial Analysts Journal* 65:4, 33-46.