Index Funds, Asset Prices, and the Welfare of Investors*

Martin Schmalz† William Zame‡

December 29, 2022

Abstract

We present a general equilibrium model in which heterogeneous investors choose among bonds, stocks, and an Index Fund holding the market portfolio. We show that, under standard assumptions, an equilibrium exists. We then derive predictions for equilibrium asset prices, investor behavior, and investor welfare. The presence of the index fund (or a decrease in the fee charged by the index fund) tends to increase stock market participation and thus increase asset prices and decrease expected returns from investing in the stock market. As a result, few - if any - investors benefit from the availability of cheap market indexing.

JEL Classification: D14, D53, G11, G12, G23

Keywords: portfolio choice, asset pricing, ownership, indexing, inequality, household finance

*We thank seminar audiences at Yale University's Cowles Foundation, Carnegie Mellon University, and University College London. Schmalz acknowledges support from Deutsche Forschungsgemeinschaft under Germany's Excellence Strategy – EXC 2126/1–390838866. Zame acknowledges the hospitality and support of the Geneva Finance Research Institute and the Einaudi Institute for Economics and Finance, and the hospitality of University College, London. Wentao Li and Wenchuan Zhao provided excellent research assistance.

†University of Oxford Saïd Business School, CEPR, ECGI, CESIfo, and C-SEB, martin.schmalz@sbs.ox.ac.uk

‡Department of Economics, UCLA, zame@econ.ucla.edu

1 Introduction

Vanguard, and other index funds, are celebrated for allowing small investors to diversify their equity portfolios and enjoy market returns. Indeed, this was the expressly stated objective of Vanguard when it was first introduced in 1975. This view of index funds has dominated academic and policy discourse for the past four decades. However, this *partial equilibrium* view tacitly treats index funds as small, so that they have a negligible effect on asset prices. But index funds are no longer small: At the present time, Blackrock, Vanguard, and State Street (the three largest funds) own 23% of the S&P 500 companies Amel-Zadeh et al. (2022), and it is estimated that the totality of index funds own almost half of the equity of all publicly traded firms Chinco and Sammon (2022). Indeed, index funds have grown so large that legal scholars and policy makers have begun to propose wide-ranging regulations that would limit the size of funds and the influence of passive asset managers. A proper analysis of effect of such regulations would seem to require abandoning the partial equilibrium view that might have been appropriate when index funds were small in favor of a general equilibrium view which acknowledges that index funds are large. To take a step towards such an analysis, we analyze what the effects of indexing are on asset prices, and as a result for the welfare of investors across the distribution of wealth and risk aversion.

In this paper, we offer a *general equilibrium* model of index funds in which heterogeneous investors choose among investments in individual firms, an index Fund that holds the market portfolio, and a risk-free bond. Holding shares in an individual firm exposes investors to both idiosyncratic firm-specific risk and market-wide aggregate risk; holding shares in the Fund shields investors from idiosyncratic risk but not aggregate risk – but requires paying a cost-covering fee charged by the Fund. In our stylized model, all firms are *ex ante* identical, so sell for the same price in equilibrium. At the equilibrium price, investor's choices are determined by their wealth and their attitude toward risk; different investors make different choices and experience different effects from the presence of the Fund and by the fee it charges.

The presence of the index Fund leads investors to shift wealth from bonds into stocks through the Fund in order to benefit from an increase in expected returns, and from individual stocks into the Fund in order to benefit from a decrease in risk. At the individual level, these shifts are welfare-improving. However, in the aggregate, these shifts increase the demand for stock, which in turn increases the price of stock. Because firm earnings remain constant, expected returns fall.

For reasonable ranges of parameters, we find that the welfare of many – or even *all* – investors *falls* when the index Fund enters and continues to fall as the fee charged by the Fund (the cost of indexing) falls. This general equilibrium effect is reminiscent of the strategic effects in both the Prisoner's Dilemma and the Tragedy of the Commons.

Our model also predicts portfolio holdings as a function of wealth and risk aversion. For the same ranges of parameters, our model predictions are consistent with the data as reported in Bach et al. (2020); Fagereng et al. (2020); Beutel and Weber (2022). In particular, our model predicts a non-monotonic relationship between wealth and the equity allocation in individuals' portfolios, and it predicts that the very rich don't diversify much, believing that their undiversified investments will yield greater returns; this matches survey evidence (Bender et al., 2022).

The present paper is sharply distinguished from extant theories of stock market equilibrium and asset pricing. In contrast to the literature which assumes that firms are price takers (e.g. Hart (1979)), we allow firms to have arbitrary objectives. In contrast with the literature that argues that investor heterogeneity doesn't matter for asset prices (Panageas et al., 2020), we find that, in the presence of index funds, investor heterogeneity matters a great deal. In contrast to the literature on the effect of intermediaries, (e.g., He and Krishnamurthy (2013); Haddad and Muir (2021), we study the effect of an unlevered intermediary on asset prices. In contrast to the literature on cross-sectional differences (e.g., Jiang et al. (2022)), we focus is on the general level of the equity market. Indeed, our study examines the effect of textbook indexing, and does not engage with non-market variations of indexing that occur in practice. Bond and Garcia (2022) examine the effect of a decrease in the cost of indexing on price efficiency and the welfare of heterogeneously informed investors. In contrast to the seminal paper by Rotemberg (1984) and more recent theories proposed in the common-ownership literature (e.g., López and Vives (2019); Antón et al. (2022)) our model features endogeneous asset prices.[1] Piccolo and Schneemeier (2020) endogeneize ownership by diversified investors in a model in which ownership affects firm behavior. By contrast, the present paper endogenizes ownership whereas ownership does not influence firm behavior.[2] Azar and Vives (2021); Eeckhout and Barcelona (2020); Philippon et al. (2021); Azar

[1]Indexing causes some degree of common ownership of industry rivals in theory. However, in practice, much common ownership is driven by active portfolio choice (Amel-Zadeh et al., 2022). Furthermore, no study has measured common ownership at the fund level. As a result, it is not known in how far indexing drives common ownership.

[2]Our companion paper Schmalz and Zame (2022) extends the present model to allow analysis of firm behavior within industries.

et al. (2021) are complements to the present paper, as they debate equilibrium aspects of common ownership when the ownership of firms affects product and labor markets, but not asset prices.

2 Model

We build a model with a focus on aggregate qualitative predictions. To that end, we oversimplify in a number of dimensions. In particular, we consider a setting in which investment decisions are made at date 0 and consumption takes place at date 1. (In the simulations that follow, we think of the interval between date 0 and date 1 as 20 years.) We assume a large number of identical firms, operating in small industries; thus firms make positive profits, which are subject to both firm-specific idiosyncratic shocks and a market-wide random shock. We assume that the number of firms is sufficiently large that the firm-specific idiosyncratic shocks wash out in the aggregate. Because the index Fund is completely diversified across the whole market, investment in the Fund is immune to the firm-specific idiosyncratic shocks but still experiences the market-wide random shock. (In the simulations, we take the number of firms as 5,000, which was roughly the number of publicly traded firms in 1980.)

2.1 Industries and Firms

There are N_0 identical industries. Within each industry, there are $m \geq 1$ identical firms, so the total number of Firms in the market is $N = mN_0$. We think of m as small – so we allow for monopoly, duopoly, oligopoly – and N_0 – and hence N – as large. Within each industry, firms compete. Rather than specifying the objective of each firm, the mode of competition, and the (equilibrium) behavior of firms, we take the shortcut of specifying the *profit* function Π of each firm.

The profit of each firm is subject to two kinds of random shocks: an idiosyncratic, firm-specific shock ϵ and a market-wide aggregate shock Δ; the total profit of a firm is the sum of a deterministic profit π and the two shocks; $\Pi = \pi + \epsilon + \Delta$. We might view ϵ as arising from as a shock to the firm's cost of production and Δ as arising from a shock to the demand structure of the entire economy.[3]

[3] In reality, each of π, ϵ, and Δ might depend on the firm's ownership, but for our present purpose we ignore this possibility. To see how such a dependence might arise, consider the very simple case in which $m = 1$ (so that each firm is a monopolist within its industry), demand within each industry is $Q = 1 - P$, the firm's expected cost

2.2 Index Fund

There is a single index Fund.[4] The Fund does not maximize profits; rather it charges a fixed fee $k \geq 0$ which simply covers its operating expenses.[5] The Fund invests all of its Assets Under Management AUM to buy an equal share λ of all firms. Because firms are identical, at equilibrium they all trade at the same price p so the Fund's expenditure is $[(1+k)p](\lambda N)$. This must equal the Fund's Assets Under Management:

$$\text{AUM} = [(1+k)p](\lambda N)$$

(In reality, index funds charge an annual fee that is a fraction of the current value of the investor's portfolio. Our method of accounting for the Fund's fee should be thought of as a convenient two-date proxy for the reality.)

Because the Fund is perfectly diversified, the firm-specific idiosyncratic shocks wash out but the market-wide shock does not. Hence *the Fund's realized revenue* is

$$R_F = \lambda N[\pi + \Delta]$$

is c, and that there is a mean-zero cost shock ϵ_0 and a mean-zero demand shock Δ_0. Assume the firm maximizes expected profit. The obvious simple calculations show that the firm produces the quantity $q = (1-c)/2$, that the expected price is $p = (1+c)/2$ and that the expected profit is

$$\pi = \frac{(1-c)(1+c)}{4} - \frac{c(1-c)}{2} = \frac{(1-c)^2}{4}$$

However, because there are shocks to both demand and cost, the *realized profit* is

$$\Pi = \frac{(1-c)^2}{4} + (\epsilon_0 + \Delta_0)\left(\frac{1-c}{2}\right)$$

Thus, realized profit depends on the idiosyncratic shock, the market shock *and* the endogenously determined equilibrium quantity. If the expected cost c depends on the fraction of the firm owned by the index fund, then the equilibrium quantity $(1-c)/2$ depends on the fraction owned by the index fund, and hence the idiosyncratic and aggregate shocks to profit also depend on the fraction owned by the index fund.

The firm's cost might depend on the fraction of the firm owned by the Fund for a number of reasons. Perhaps the simplest reason is that the Fund *votes* the shares it holds, and hence exercises some oversight on the managers of the firm. This oversight might lower (expected) cost by deterring managers from looting the firm or by incentivizing them to bargain harder with suppliers or workers.

[4]Our examples are calibrated to the period of time in which Vanguard was the only operating index fund. At present, there are more than 20 such funds, of which Blackrock, Vanguard and State Street are the largest.

[5]This is an accurate description of Vanguard and of most of the other index funds. We allow for the extreme $k = 0$ because it is a useful benchmark in simulations.

and so the Fund's *revenue per dollar of investment* is

$$\begin{aligned} r_F &= \frac{R_F}{\text{AUM}} \\ &= \frac{\lambda N[\pi + \Delta]}{(1+k)p\lambda N} \\ &= \frac{[\pi + \Delta]}{(1+k)p} \end{aligned}$$

The Fund makes 0 profit so the *investor's yield per dollar of investment in Vanguard* is

$$Y_F = r_F \tag{1}$$

2.3 Stocks

An investor who owns stock in an individual firm obtains a share of the profits of that firm. The realized profit of a random firm is $\Pi = \pi + \epsilon + \Delta$ so the *(random) yield per dollar of investment in stocks* is

$$Y_S = (1/p)\Pi \tag{2}$$

Note that different investors own stocks in different firms so experience different idiosyncratic shocks. The assumption that each investor can only hold one stock is a proxy for the idea that the costs of creating and maintaining a diversified portfolio comprising many stocks is prohibitively expensive. In reality, most investors held a small number of individual stocks, not entirely washing out idiosyncratic risk.

2.4 Bonds

An unlimited supply of riskless bonds is available with *yield per dollar of investment*

$$Y_B = (1 + \rho) \tag{3}$$

where $\rho \geq 0$.

2.5 Investors

There are a continuum of investors. It is convenient (and involves no loss of generality) to index the space of investors by $T = [0, \infty)$. Investor t is characterized by its wealth $w^t \in (0, \infty)$, its choice set $X^t \subset \mathbb{R}^3_+$ and its utility function U^t for *random* consumption (i.e., for lotteries over consumption bundles). For simplicity, we assume that investors maximize expected utility with respect to some Bernoulli utility function for consumption; i.e. $U^t = E[u^t(c)]$, where u^t is continuous and strictly increasing. A choice $x \in X$ is a triple

$$x = (x_S, x_F, x_B)$$

where x_S is the number of firm shares held directly, x_F is the number of firm shares held through the Fund, and x_B is the number of bonds held. (In our formulation, there is a single share in each firm, so investors hold fractional shares and the price p is the price of the entire firm.) In view of our calculations of yields for Stock, the Fund, and Bonds, the utility of a consumer who holds $x = (x_S, x_F, x_B)$ is

$$\begin{aligned} U^t(x) &= U^t\Big(px_SY_S + px_FY_F + (1+\rho)x_B\Big) \\ &= E\left[u^t\Big(px_SY_S + px_FY_F + (1+\rho)x_B\Big)\right] \end{aligned}$$

(Keep in mind that the yields of Stock and the Fund are *random*; but that the yield of bonds is not.)

We assume X^t is a closed cone (i.e., if $x \in X$ and $\alpha \in \mathbb{R}_+$ then $\alpha x \in X$), that u^t is continuous and strictly increasing, and that the map $t \mapsto (w^t, X^t, u^t)$ is measurable.[6] Note that choice sets X^t need not be convex; in particular we allow for the possibility that investors can invest in stock or in the Fund or in Bonds – but cannot hold a mixed portfolio.

The distribution of investor characteristics is given by a positive, non-atomic measure ϕ on $[0, \infty)$ of total mass M.[7]

[6] All that is really needed is that, given other parameters, the map from investors to optimal choices is measurable.
[7] Note that we deviate from the usual convention and do not normalize the mass of investors to one.

3 Equilibrium

An *equilibrium* consists of

- A *price* $p > 0$ for the firms (all firms have the same price)
- A (measurable) investor choice function $x : T \to \mathbb{R}^3_+$

$$x^t = \left(x^t_S, x^t_F, x^t_B\right)$$

  where

  - x^t_S = number of shares of stocks purchased individually
  - x^t_F = number of shares of stock purchased through the Fund
  - x^t_B = number of bonds purchased

such that

- Each investor maximizes utility

$$U^t\left(px^t_S Y_S + px^t_F Y_F + x^t_B Y_B\right) = E\left[u^t\left(px_S Y_S + px_F Y_F + (1+\rho)x_B\right)\right]$$

  subject to the feasibility constraint

$$x^t \in X^t$$

  and the budget constraint

$$px^t_S + px^t_F + x^t_B = w^t$$

- The market for stocks clears

$$\int x^t_S d\phi(t) + \int x^t_F d\phi(t) = N$$

  (The first integral is the total of stocks purchased directly by investors, the second integral is the total of stocks purchased through the Fund. In the model, all shares are held by small investors.)

Note that, at equilibrium, the Fund's share of stocks is

$$\lambda = \frac{1}{N} \int x_S^t(p) d\phi(t)$$

Theorem *An equilibrium exists.*

We defer the proof to Appendix A.

4 Numerical Simulations

It seems difficult – perhaps impossible – to solve the Model in closed form, even for simple specifications of the distribution of wealth and utility functions. Instead, we offer numerical simulations. Our choice of baseline parameters is suggested by data from the US stock and bond markets circa 1980, at a time when Vanguard was just established but was the only significant index fund. (We remind the reader that Vanguard was explicitly designed as a non-profit-making institution; it charged a fee that represents its cost of operation, rather than a fee designed to maximize profits.) Aside from this choice of parameters, we have made little effort to calibrate the model to real data. Our objective in these simulations is to provide intuition about the effect of index funds on asset prices, portfolio holdings, and investor welfare. We believe that some of the results are surprising – even counter-intuitive.

Our procedure is as follows. First, we choose a candidate price for equities. Second, we calculate what each investor would choose given that price. Third, we check whether the market clears (i.e., whether the demand for stock equals the supply of stock). We iterate this procedure until we find a price at which the market clears - an *equilibrium price.* Having found an equilibrium price, we graph the choices and utilities and choices of 1000 investors, at the equilibrium price.[8]

4.1 Timing

We view date 0 (the date at which investors make investments) and date 1 (the date at which investors realize the returns on investments) as approximately 20 years apart.

[8]In principle, our model might have multiple equilibria, but we do not find multiple equilibria in any of our simulations,

4.2 Bonds

Throughout, we assume that $\rho = 0.5$. This represents a real rate of return of roughly $0.02 = 2\%$ per year, compounded over the 20-year period.

4.3 Firms

Throughout, we assume the total number of (publicly traded) firms is $N = 5000$ and that the expected profit of each firm is $\pi = \$500$ Million. (Keep in mind that this is profit over a 20-year period.) We assume that idiosyncratic risk $\epsilon = \pm 0.5\pi$, each occurring with probability 0.5 and that aggregate market risk is $\Delta = \pm 0.5\pi$, each occurring with probability 0.5. Thus, the distribution of realized profits for each firm is

$$\Pi = \begin{cases} \$1,000 \text{ Million} & \text{with probability} \quad 0.25 \\ \$500 \text{ Million} & \text{with probability} \quad 0.50 \\ \$0 & \text{with probability} \quad 0.25 \end{cases}$$

These parameters imply that, *ex-ante*, the probability that a given firm will go bankrupt during the 20-year period is 0.25. Lest this seem unreasonably large, note that the average bankruptcy rate of publicly traded firms is actually in the vicinity of 2% per year. In our model, a firm will go bankrupt only if it experiences a negative idiosyncratic shock *and* the market experiences a negative aggregate shock; if the market shock is positive, no firm will go bankrupt. However, from the *ex ante* perspective of investors, what matters is that the return on a stock investment will be 0 with probability 0.25.

4.4 Investors

Throughout, we assume investors' utility functions display constant absolute risk aversion (CARA) and that investor t's coefficient of risk aversion is t. However, we make two normalizations:

- In order to facilitate utility comparisons across investors with different risk aversions, we normalize by dividing investor t's Bernoulli utility function by $1 - e^{-t}$.

- We measure wealth and consumption in units of \$10,000.

Thus, if investor t's realized consumption is c (measured in units of \$10,000), its utility is

$$u^t(c) = \frac{1 - e^{-tc}}{1 - e^{-t}}$$

Consumers maximize expected utility so $U^t = E[u^t(c)]$.

Throughout, we assume the total mass of investors (the number of investors) is $M = 100$ Million and the total invested wealth is $W = \$2$ Trillion.[9]

4.5 Baseline Scenario

In our Baseline Scenario, we assume that

- the wealth of investors is exponentially distributed;
- the risk aversion of investors is uniformly distributed on the interval $[0, 5]$;
- investor choice sets are $X^t = \mathbb{R}^3_+$.

Thus investor t's wealth is Ce^{-t}, where the constant C is chosen so that total wealth of all investors is W; and investor t's Bernoulli utility function is, as above, $u^t(c) = (1 - e^{-tc})/(1 - e^{-t})$; and investors can (and some do) hold mixed portfolios that contains individual stock *and* the Fund *and* bonds.

Note that the richest investors are the least risk-averse, and that the richest 20% of investors hold roughly 62% of all invested wealth. Conversely, the poorest investors are the most risk averse, and the poorest 20% of investors hold roughly 2% of all invested wealth.

We also simulate a number of alternative scenarios, allowing for different distributions of wealth and risk aversion, and also for the possibility that investors are not allowed to hold mixed portfolios (i.e., constraining investors to hold *only* stock or *only* the Fund or *only* bonds). Allowing for different distributions of wealth and risk aversion provides a robustness check on our findings in the Baseline Scenario; not allowing investors to hold mixed portfolios illustrates the intuitions more sharply. We relegate all these simulations to Appendix B.

[9] In 1980, total capitalization of all publicly traded US firms was approximately \$1 Trillion and the total capitalization of the US bond market is estimated to have been in the range of \$0.5-1.5 Trillion, so $W = \$2$ Trillion seems reasonable.

4.6 The Fund

As in the Model, we assume there is a single Fund, so it is completely specified by k, the fee charged by the Fund. In the simulations, we consider various values of k; for the Baseline Scenario described above, we simulate outcomes for values of k in the interval $[0, 1]$ in increments of 0.01; this allows us to demonstrate that price is strictly decreasing in k. (Keep in mind that k represents the *total fee* charged over a 20-year period and that funds typically charge a fee based on the *current value* of assets under management, not on the initial investment.) We also simulate outcomes for $k = \infty$, which represents the setting in which no Fund is available to investors.

5 Findings

Here we report the findings of our simulations of the Baseline Scenario, for various values of the fee k. (As noted above, Appendix B reports the findings of our simulations in various other scenarios.) We begin with the effect of indexing on the equilibrium price of equities.

5.1 Equilibrium Prices

In Figure 1 we show the equilibrium price of equity as a function of the fee k charged by the Fund. The Figure plots prices for $0 \leq k \leq 1.0$, in increments of 0.01. Note that the equilibrium price is strictly monotonically decreasing in k. We also note that the equilibrium price when $k = \infty$ (i.e., no fund is available) is $p = 0.079$, which is even lower than when $k = 1.0$.

5.2 Portfolio Choices

In Figures 2 - 7 we show the portfolio choices of investors as a function of their risk aversion, for values of $k = \infty, 0.4, 0.3, 0.2, 0.1, 0.0$. As k decreases (so indexing becomes cheaper), more invested wealth flows into the Fund. When $k = 0$ indexing is free so of course no one chooses to hold individual stocks. Surprisingly, most risk-averse investors invest heavily in individual stocks and in the Fund, which are risky, but not in bonds, which are risk free. This is a consequence of the fact that the most risk-averse investors are also the poorest, and have little to invest.

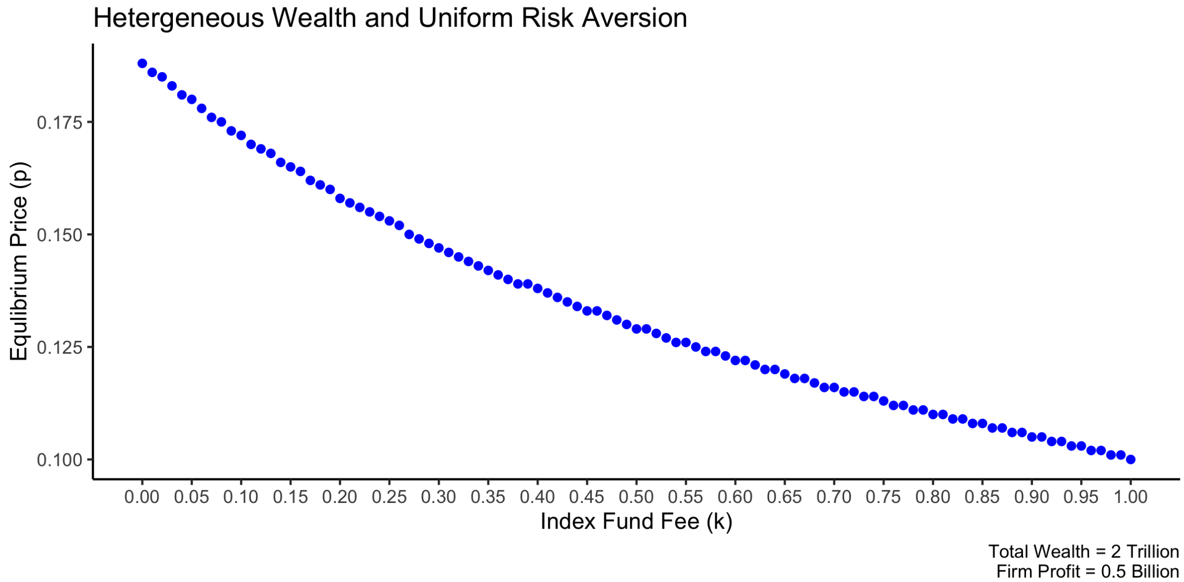

Figure 1. Equilibrium Price of Stock as a Function of Fee

5.3 Welfare Comparisons

In Figure 8, we provide welfare comparisons (in percentage terms). The benchmark for comparison is welfare when $k=\infty$; i.e., when no fund is available. Remarkably, the availability of the fund *reduces welfare for all investors*, and *the reduction in welfare increases as the fund becomes cheaper*. Because asset prices rise when the Fund is available and continue to rise as investing in the Fund becomes cheaper, it is no surpise that the richest investors – who invest most heavily in individual stocks – suffer the largest losses. (The curves representing utility losses for investors whose risk aversion is below about 2.5 are not distinguishable in the Figure because the losses for these investors are not very sensitive to the cost of the Fund.)

5.4 Welfare of the Marginal Investor

As Figure 8 demonstrates, the welfare of investors *falls* when the Fund becomes available, and continues to fall as the Fund becomes cheaper. This is the general equilibrium effect which have already discussed: the availability of the Fund drives up asset prices and the negative effect of this increase in asset prices is greater than the positive effect of the diversification that the Fund provides.

To emphasize the difference between the conclusions that follow from the (correct) general

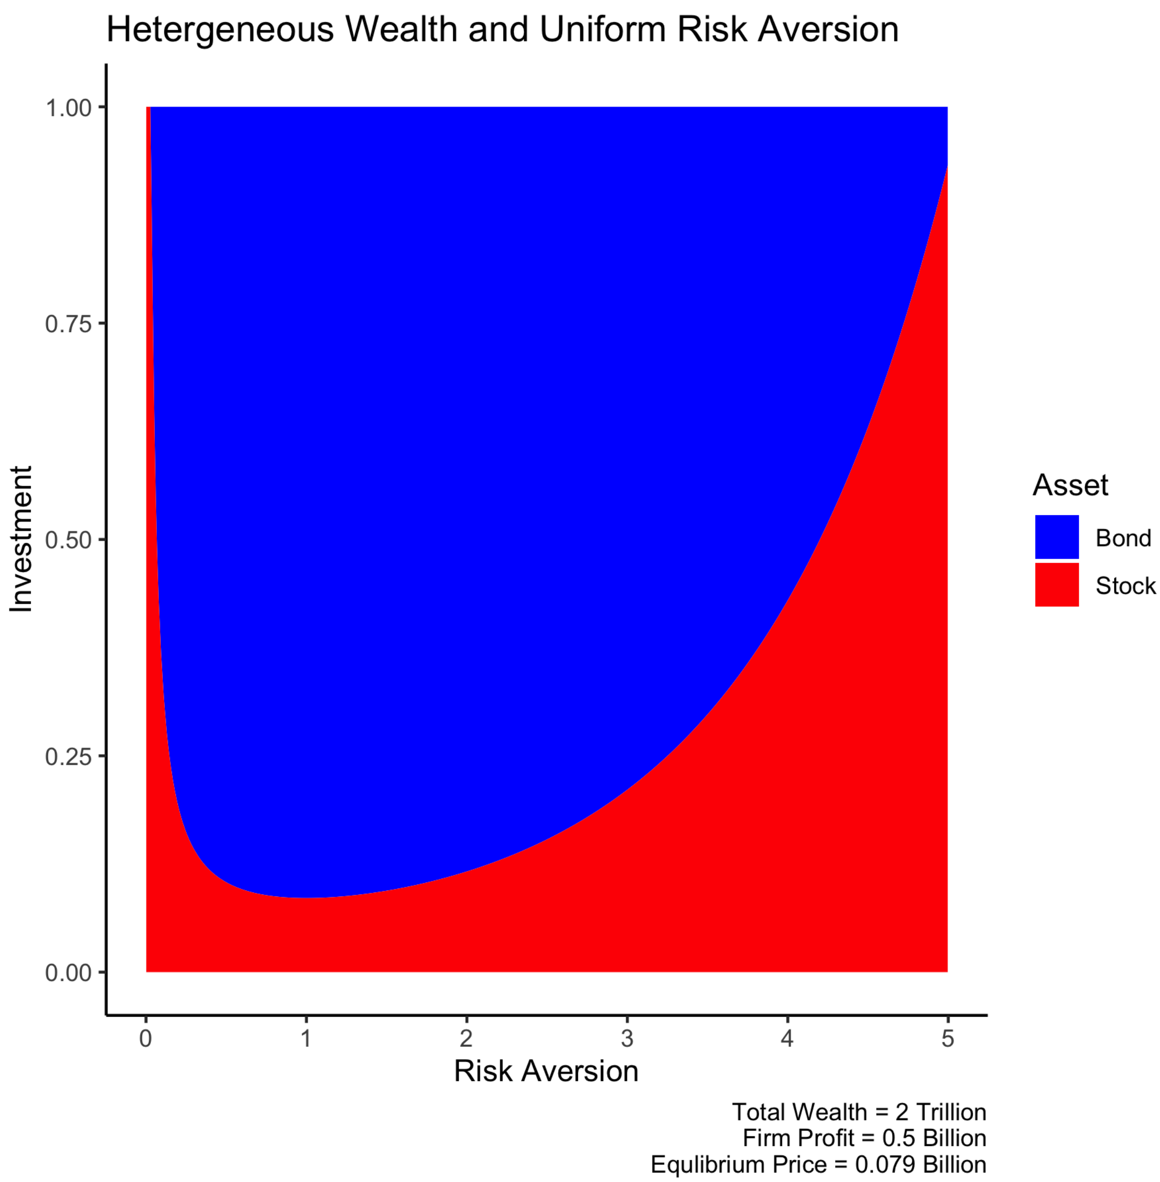

Figure 2. Portfolio Choices: $k = \infty$

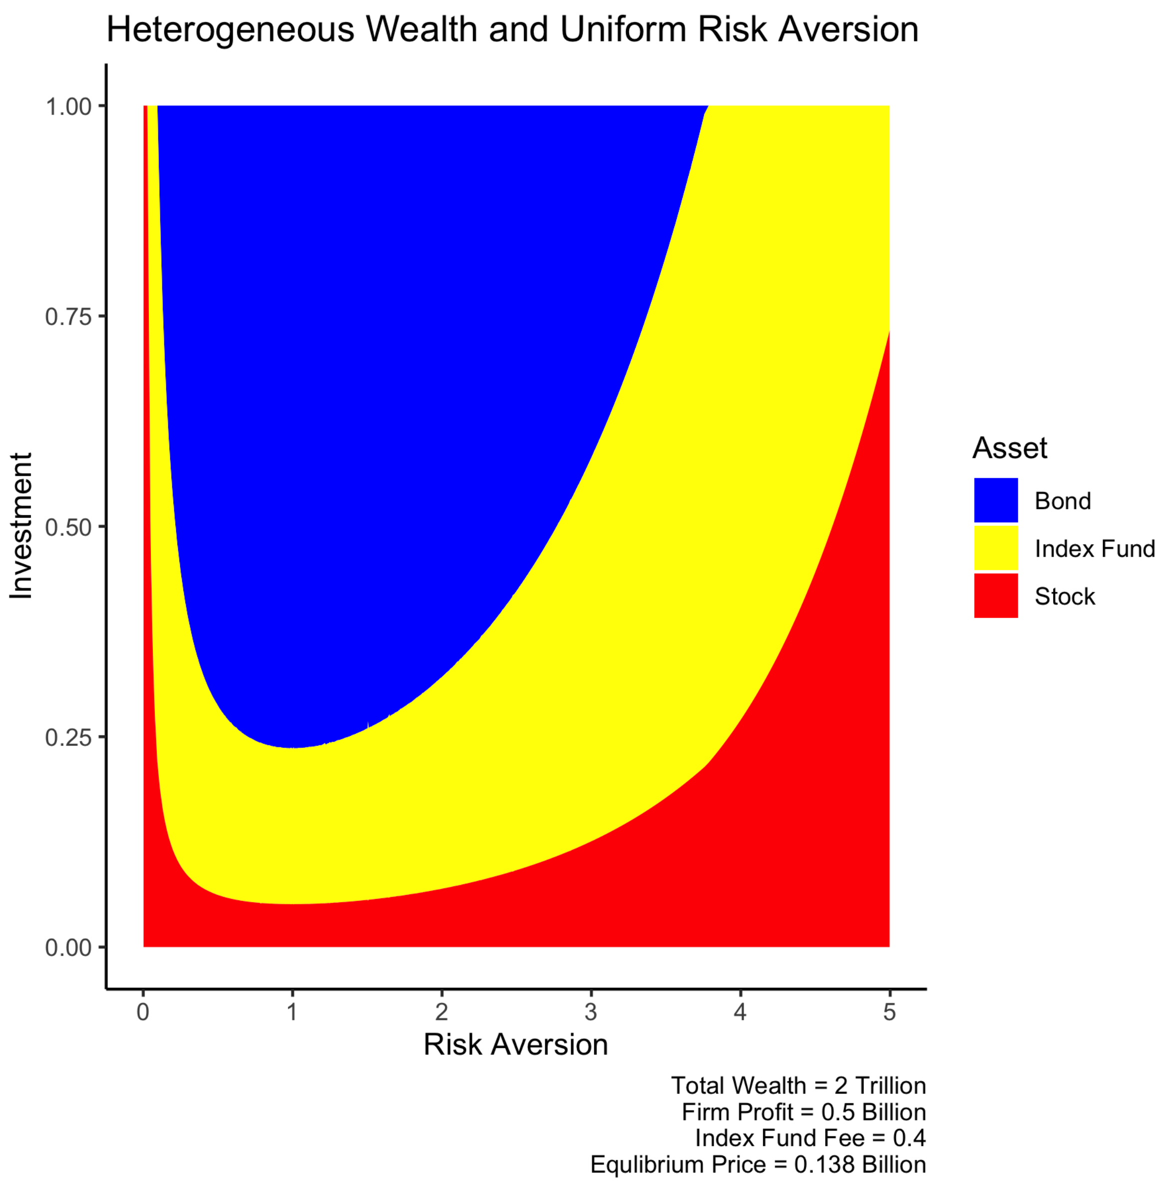

Figure 3. Portfolio Choices: $k = 0.4$

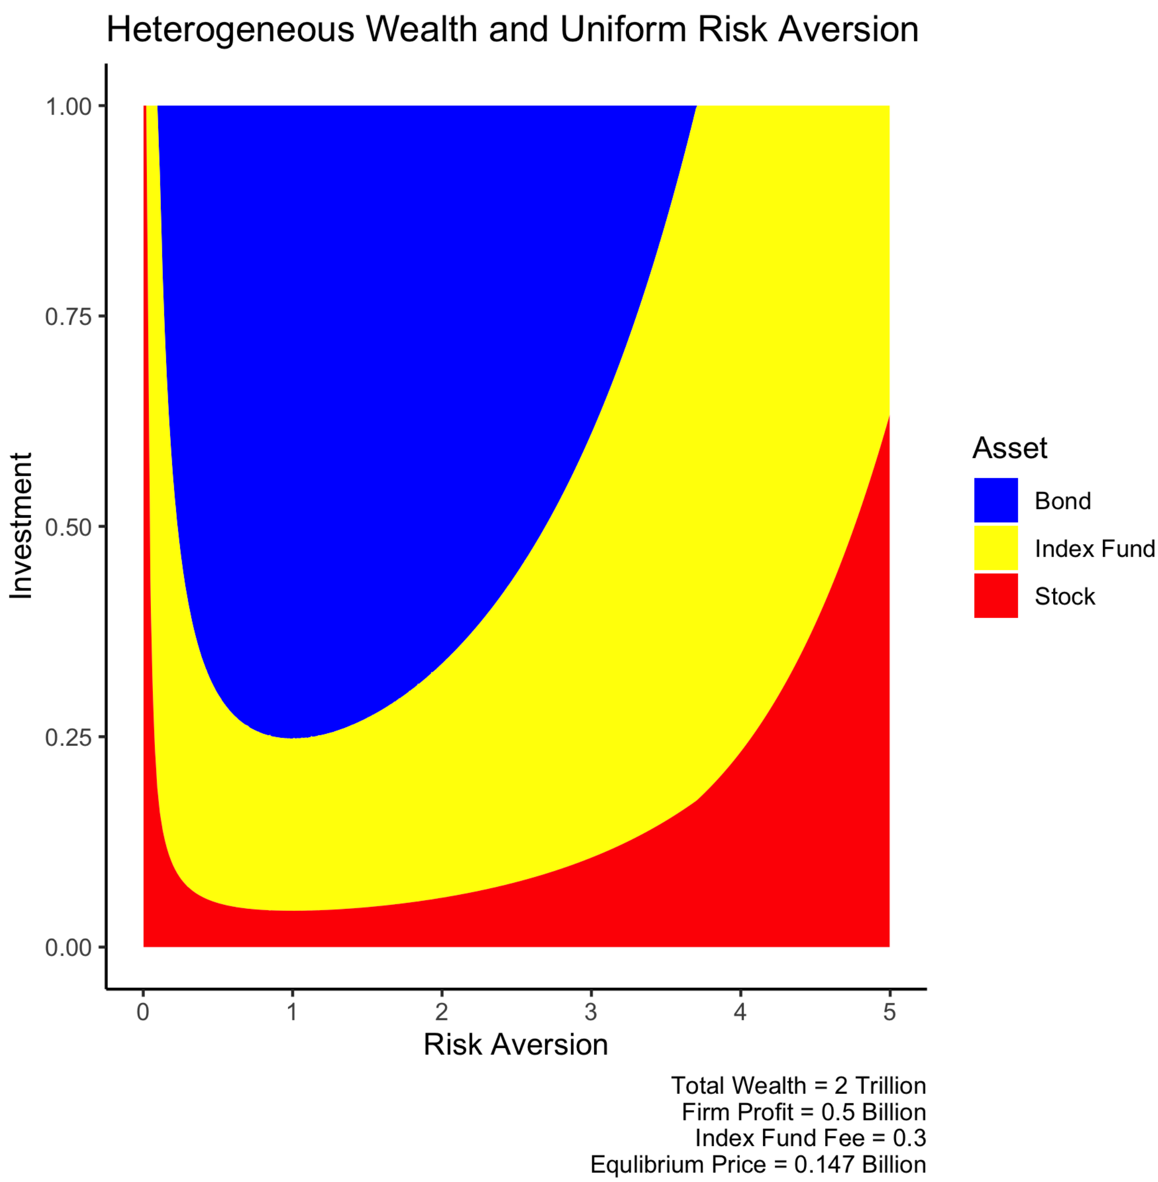

Figure 4. Portfolio Choices: $k = 0.3$

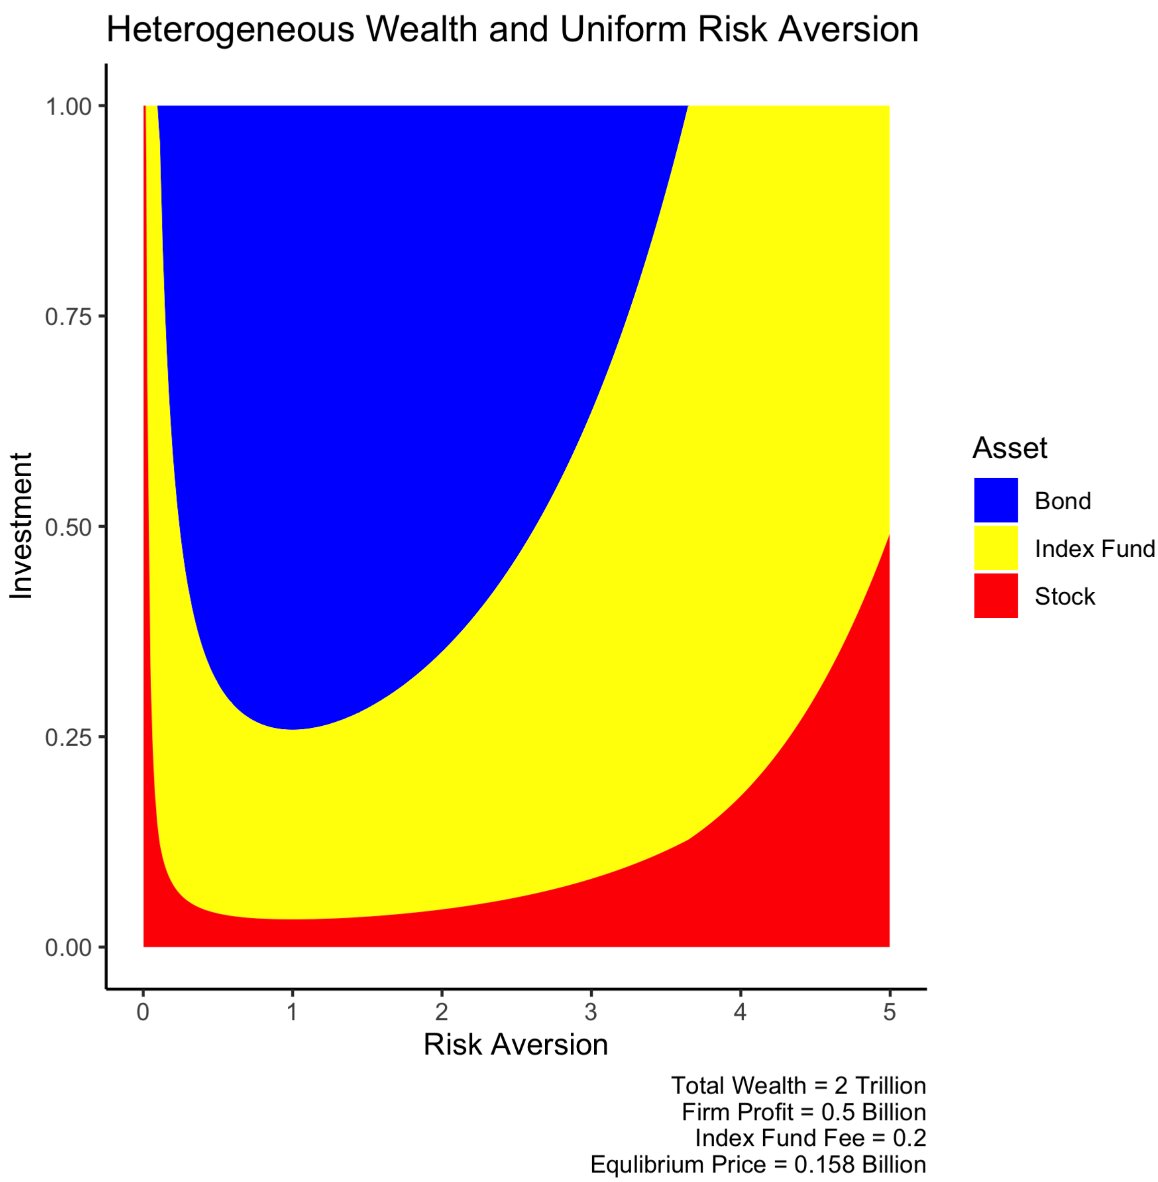

Figure 5. Portfolio Choices: $k = 0.2$

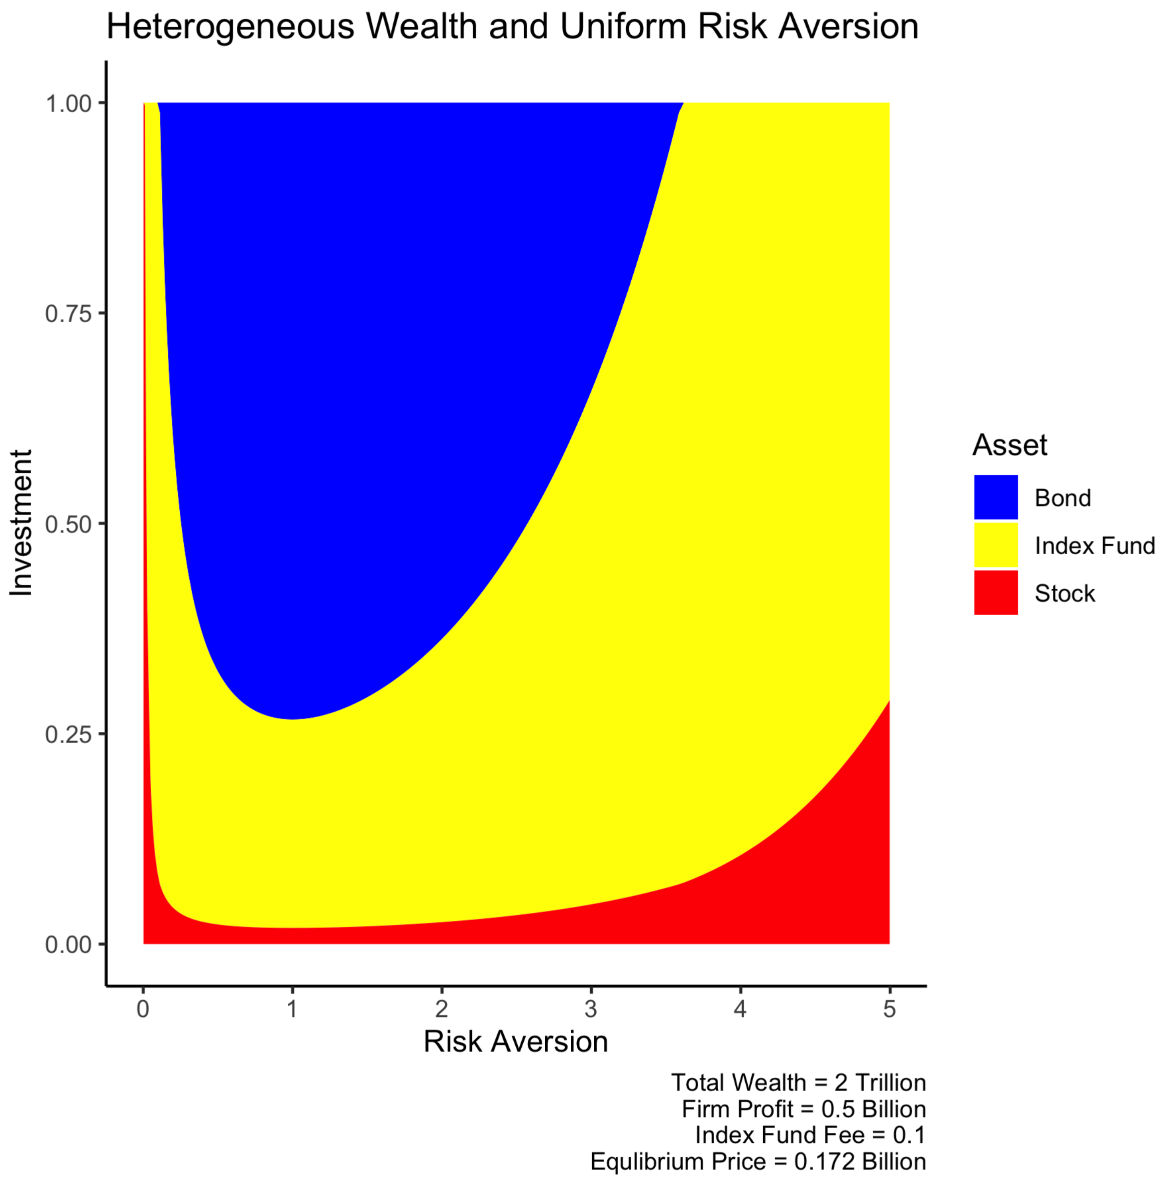

Figure 6. Portfolio Choices: $k = 0.1$

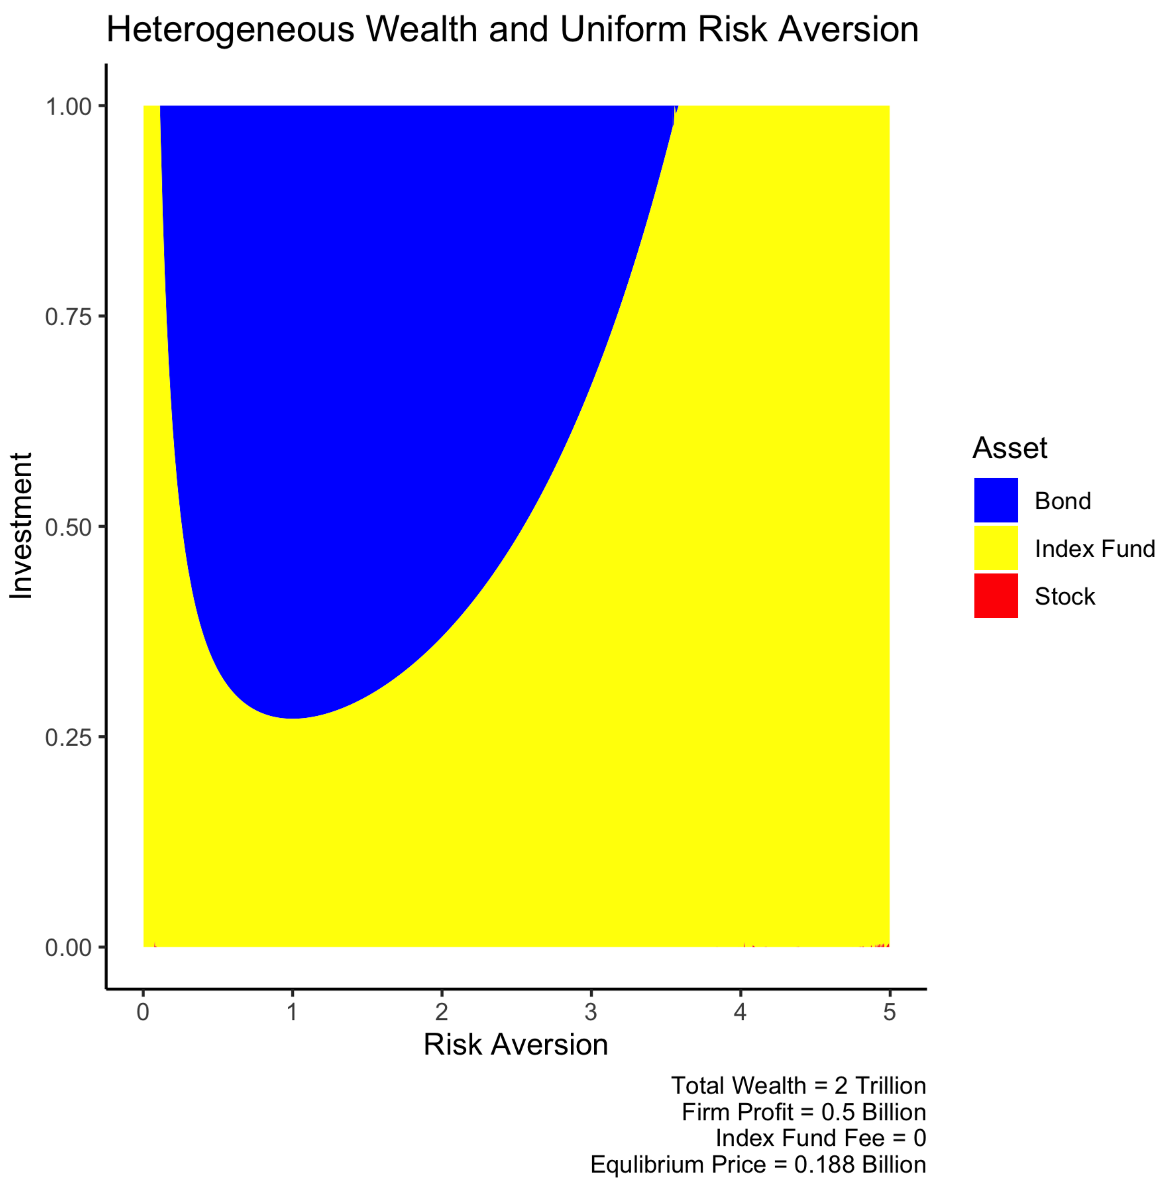

Figure 7. Portfolio Choices: $k = 0.0$

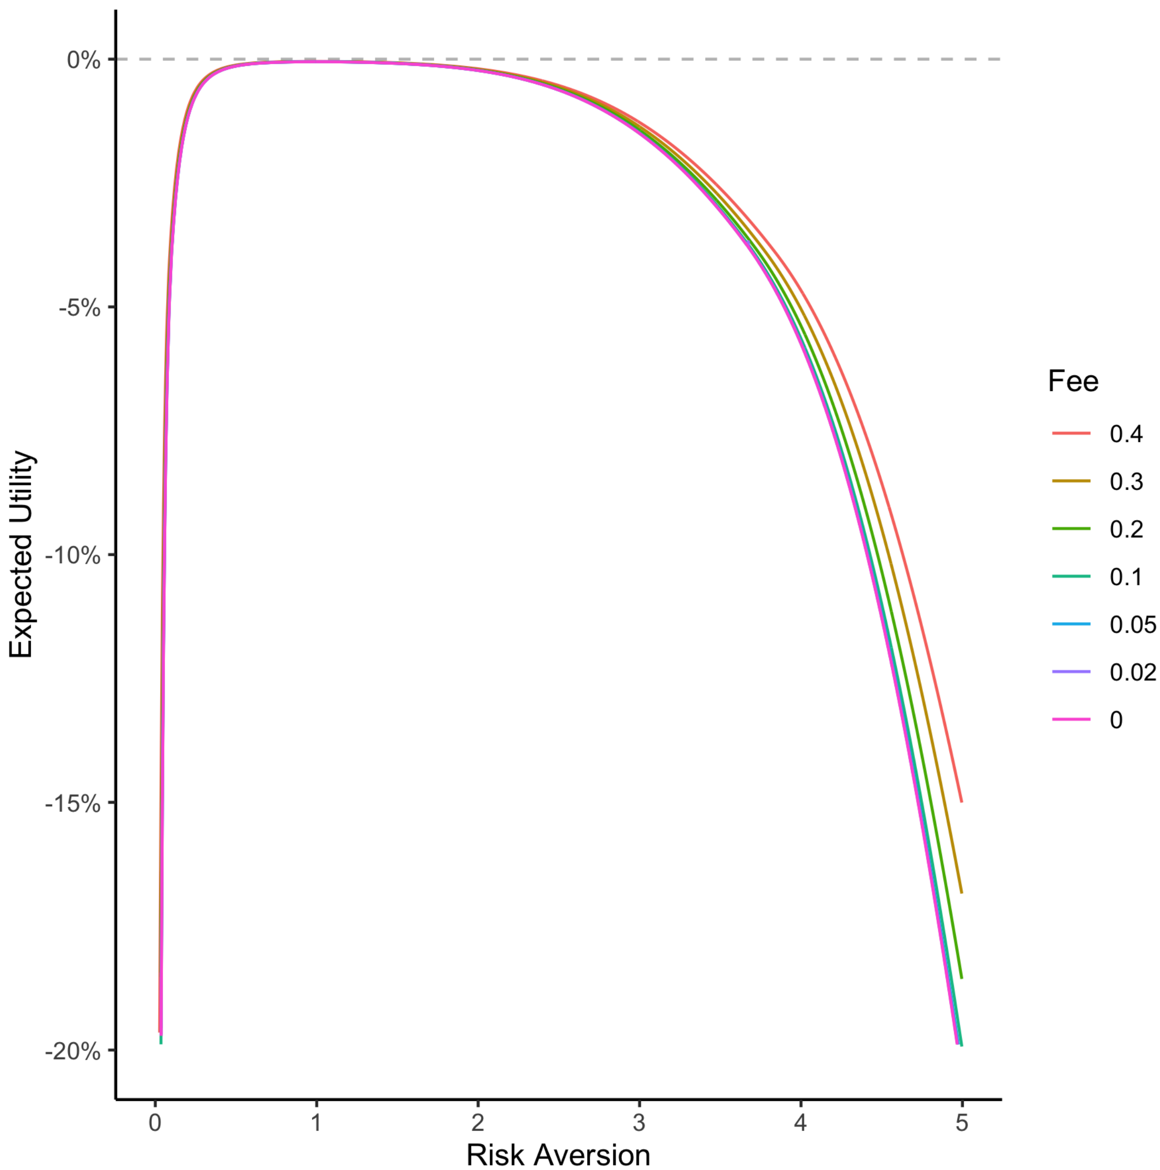

Figure 8. Welfare Losses; Benchmark = No Fund

equilibrium reasoning and the (incorrect) partial equilibrium reasoning, we can conduct a counterfactual analysis by examining the welfare consequences for the marginal investor.

To this end, we consider an investor who faces the asset prices in a world in which the Fund operates (at a given cost), and we ask how much this investor could gain by investing optimally in stocks, bonds and the Fund rather than *only* in stocks and bonds. In the interests of parallelism with our analysis of welfare losses, we show in Figures 9 - 13 the *counterfactual* portfolio choices; i.e., the optimal choices at prevailing prices when the given investor unable to invest in the Fund – perhaps because she is unaware of its existence) for various values of k.

If we compare the counterfactual portfolio choices shown in Figures 9 - 13 with the actual portfolio choices shown in Figures 3 - 7, we see that the presence of the Fund causes the investor to shift some invested wealth from stocks to the Fund and also from bonds to the Fund. The former shift is motivated by the desire to avoid idiosyncratic risk; the latter shift is motivated by the willingness to take on reduced risk in return for higher expected returns.

Understanding the magnitudes of these shifts as functions of the fee charged by the fund is more subtle than it might seem. If all else were held constant, a lower fee would provoke larger shifts: because a lower fee means that it costs less to avoid idiosyncratic risk – so the Fund gains in attractiveness in comparison to stock – and it also costs less to obtain a higher expected return – so the Fund gains in attractiveness in comparison to bonds. However, this is, once again, (incorrect) partial equilibrium reasoning. As we have already seen, a lower fee leads to higher prices for and lower returns on equities, which leads both equities *and the Fund* to become less attractive. Thus, lower fees generate opposing forces and the magnitude of the net effect of these forces is not obvious.

Having said this, it is important to keep in mind that the *sign* of the net effect of these forces *is* obvious. Because the marginal investor is infinitesimal, her choices do not affect prices. Thus, when she learns about the existence of the Fund, she faces a larger choice set but the same prices, and her welfare must rise (at least weakly). Figure 14 shows the percent increase in welfare for investors who learn about the Fund and hence have the opportunity to shift invested wealth from stocks and bonds into the Fund.

Figure 14 shows the percentage gain in utility as a function of t for various values of the fee k. Several things are important to note about this Figure. The first is that the welfare gain *increases* as the fee charged by the Fund decreases. Put differently: lower fees are good for the

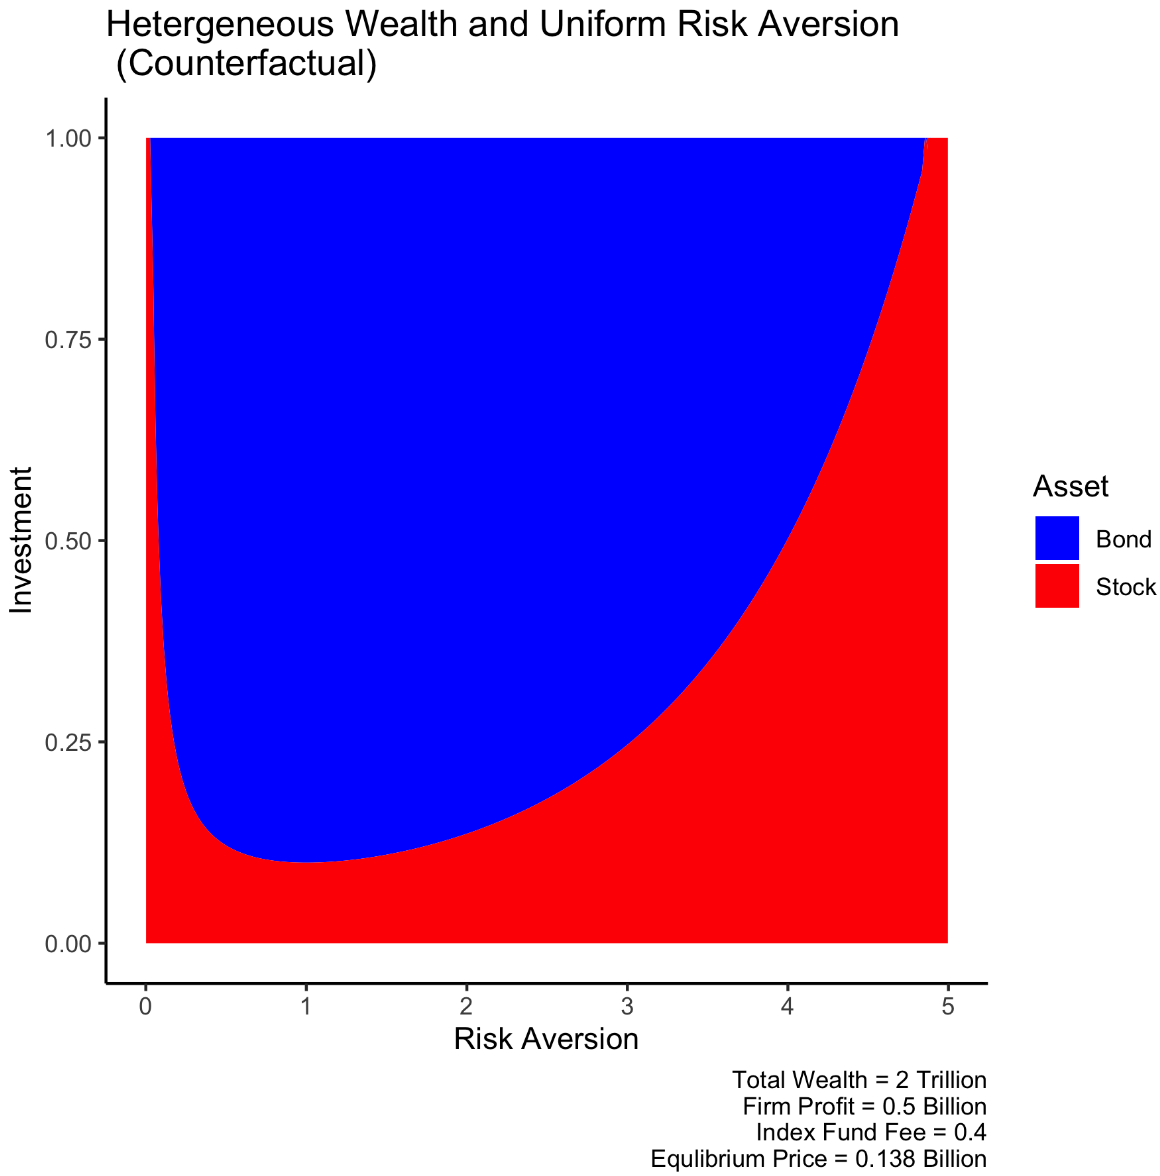

Figure 9. Portfolio Choices: $k = 0.4$

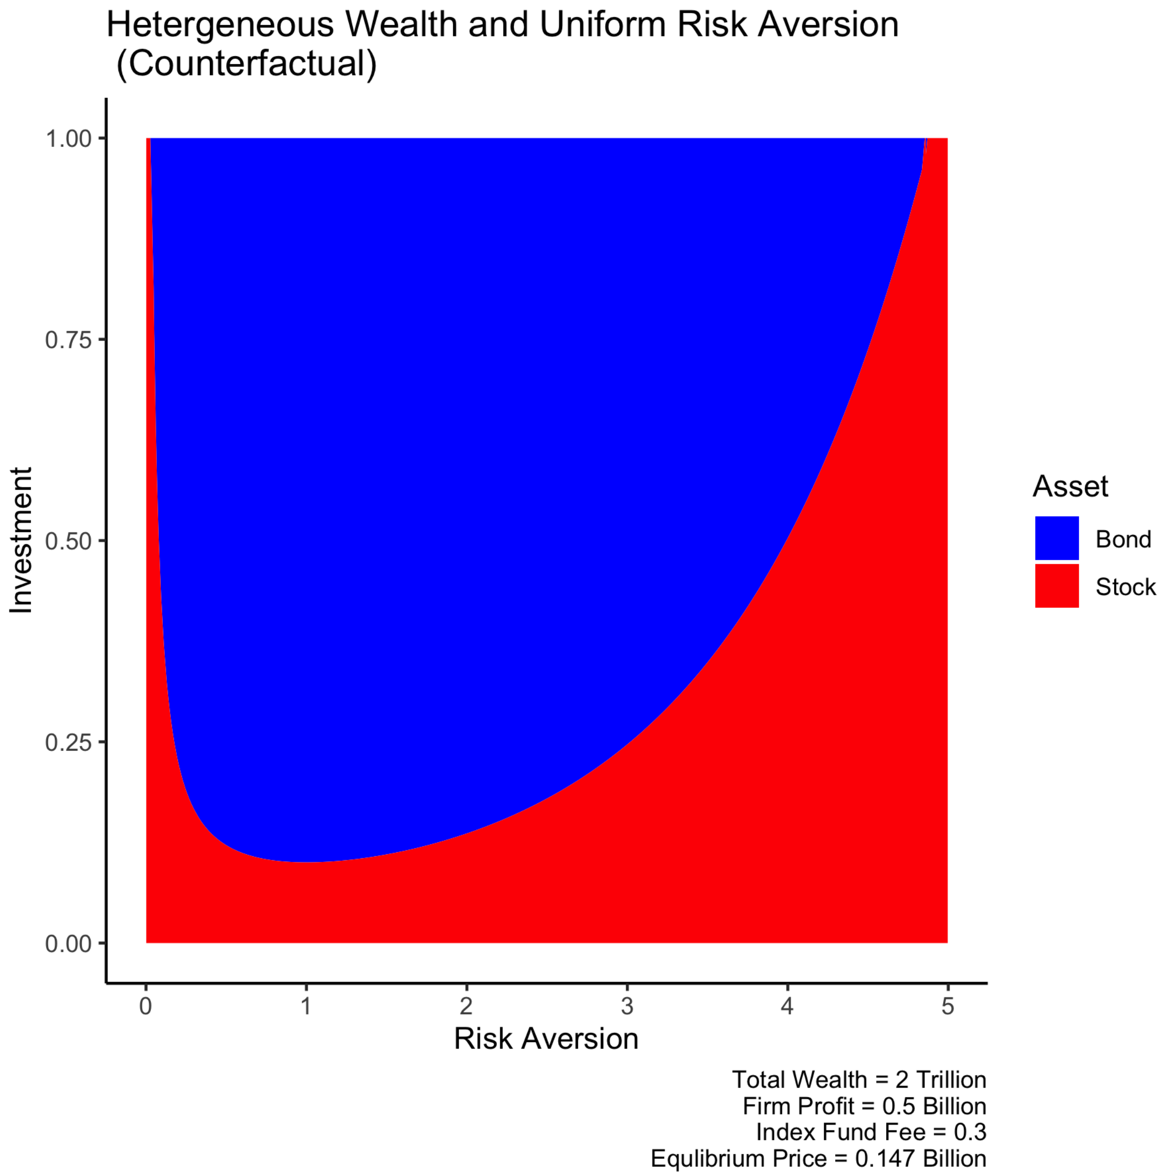

Figure 10. Portfolio Choices: $k = 0.3$

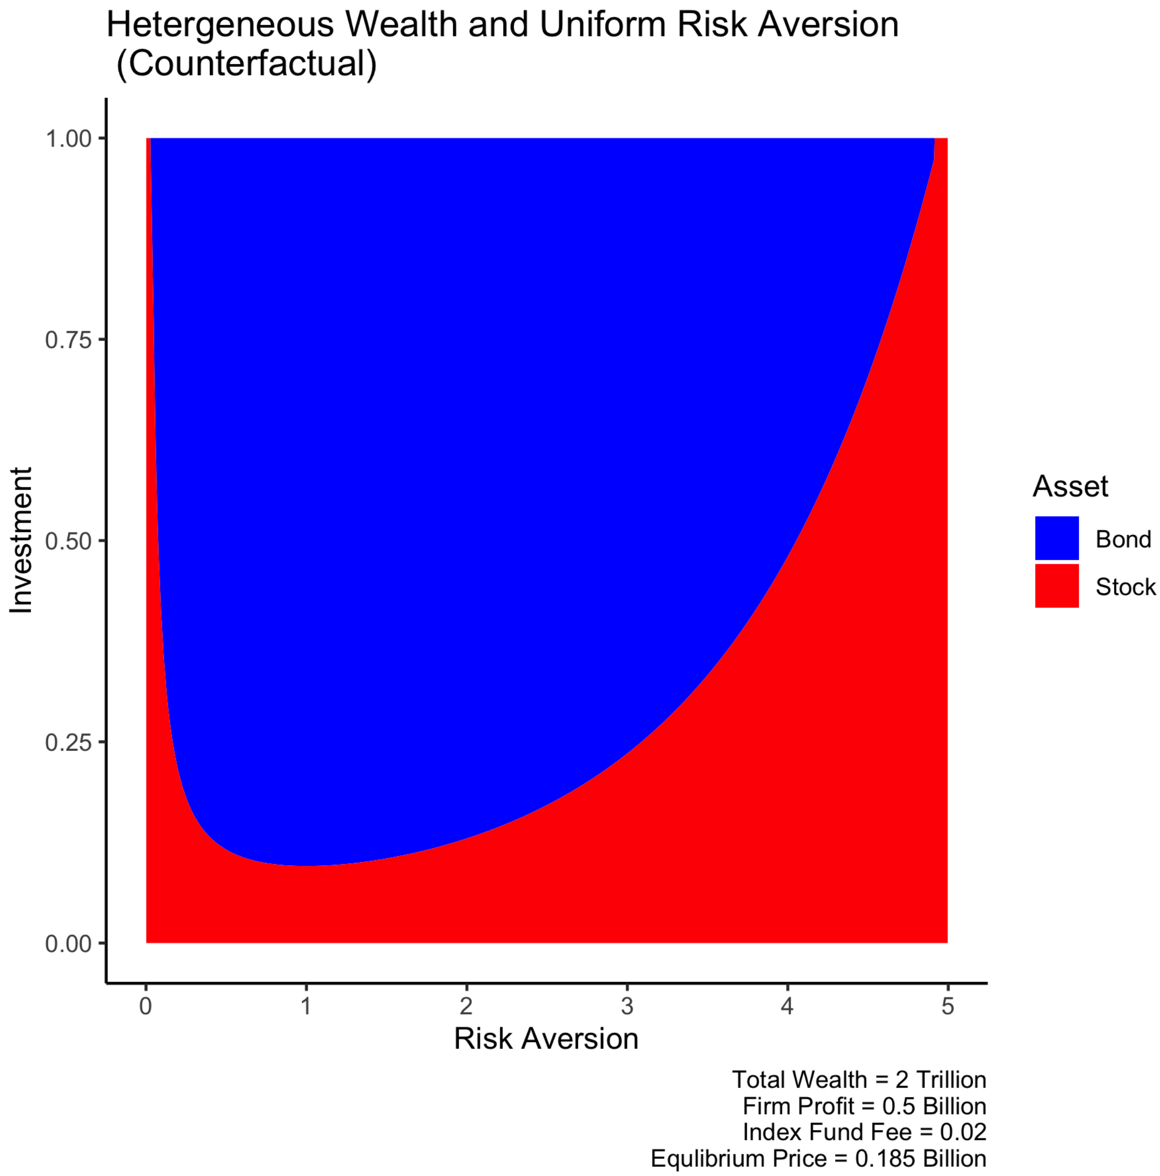

Figure 11. Portfolio Choices: $k = 0.2$

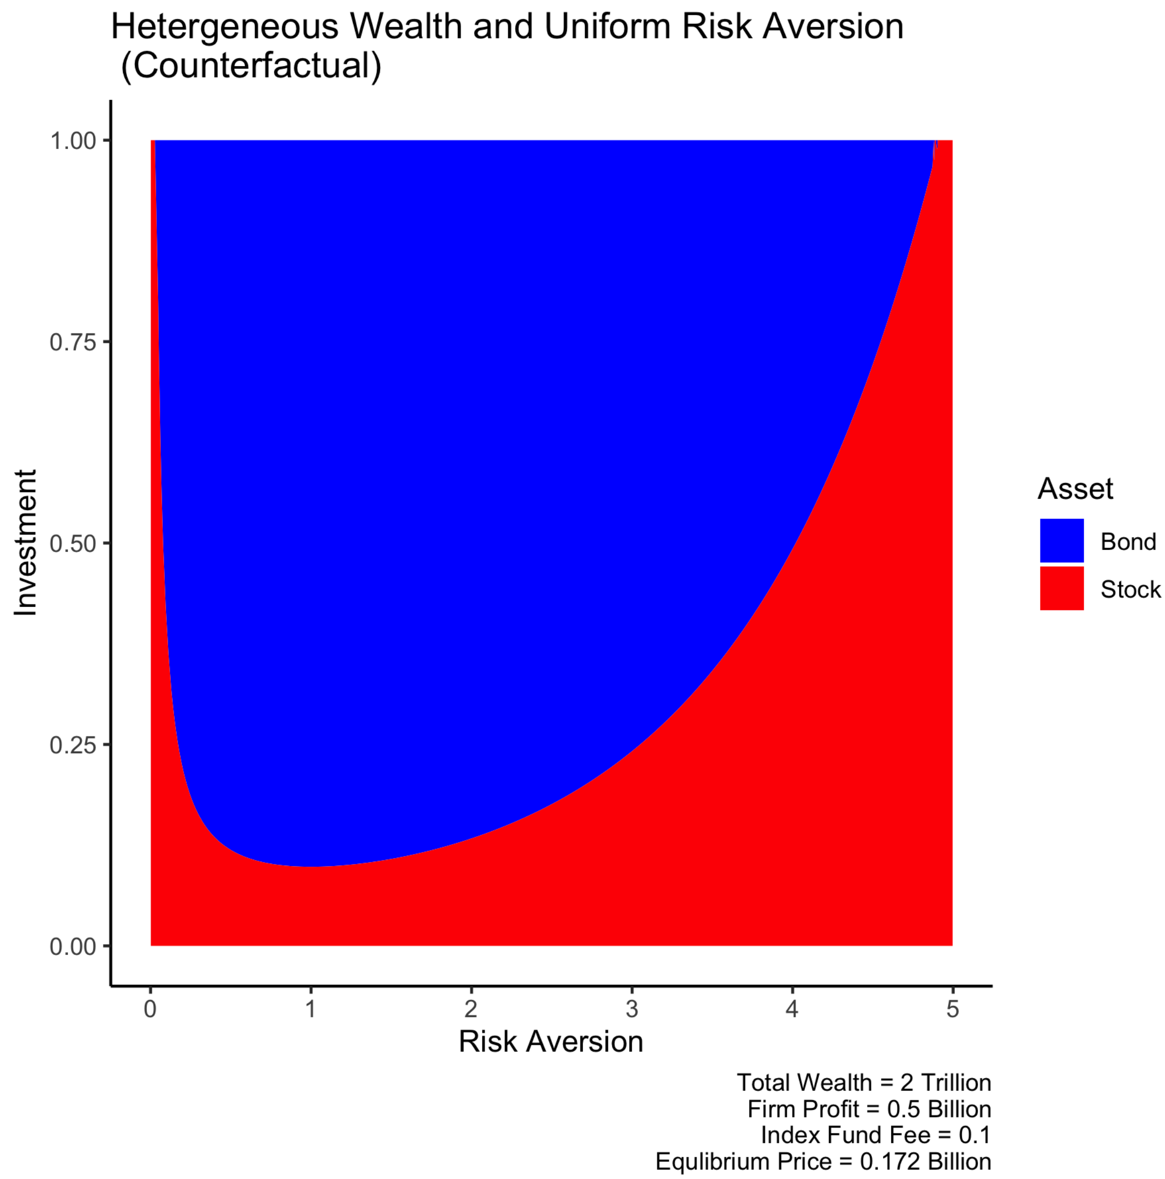

Figure 12. Portfolio Choices: $k = 0.1$

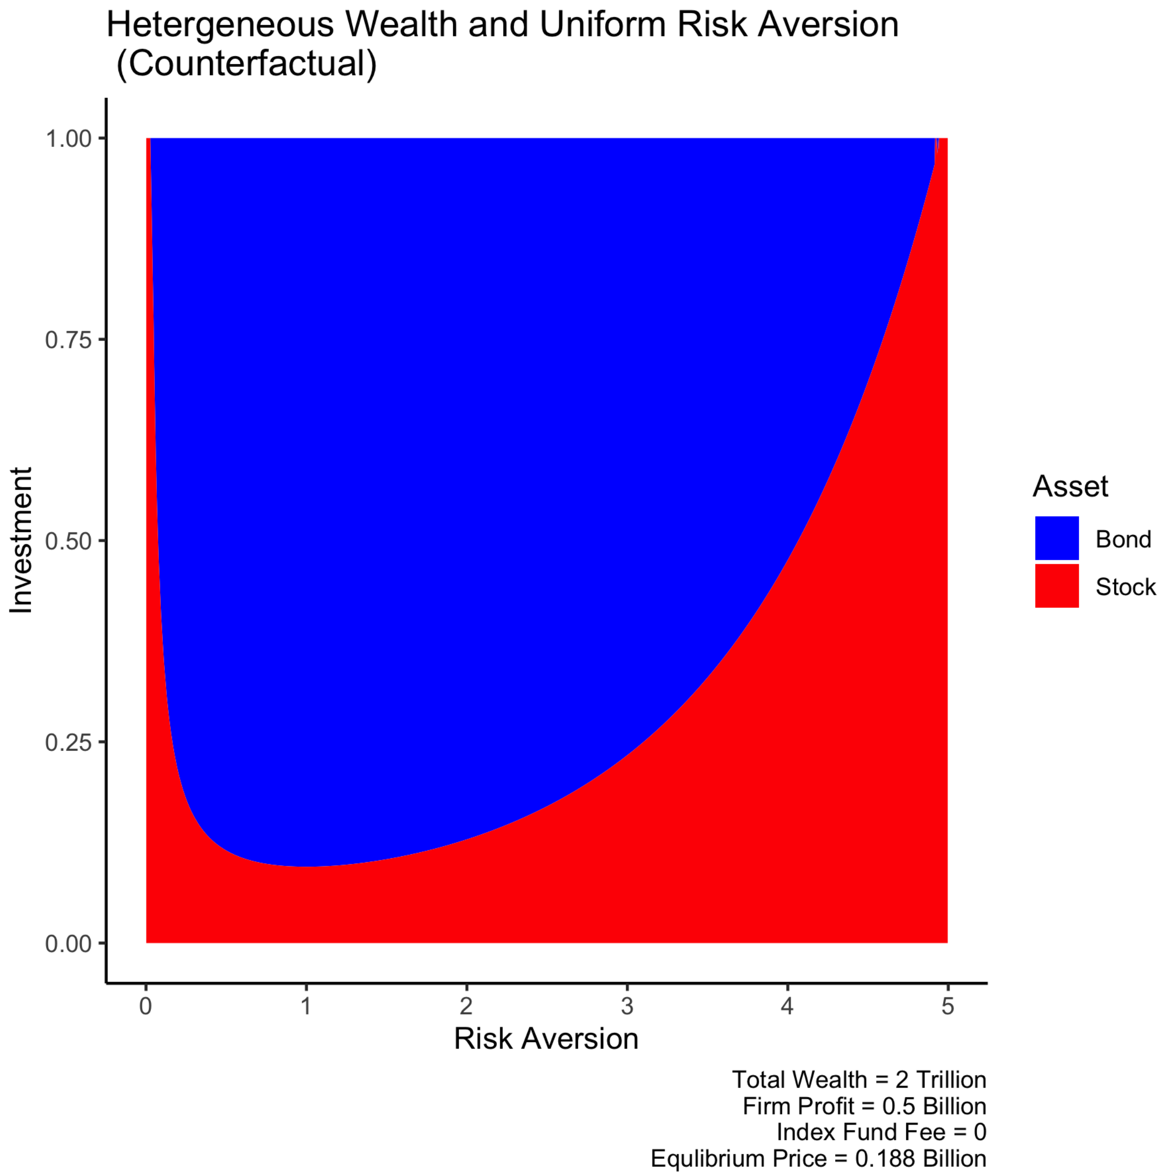

Figure 13. Portfolio Choices: $k = 0.0$

marginal investor. (Of course this confirms the obvious partial equilibrium analysis, but it is in stark contrast to the general equilibrium conclusions.) The second is that the percentage welfare gains are very different across the spectrum of investors. Roughly speaking we can say that

- Investors whose wealth is in the top 1% benefit little from the presence of the Fund because they are almost risk neutral.

- Investors whose wealth is in the top 2-10% benefit very substantially from the presence of the Fund because they shift a great deal of invested wealth into the Fund and are sufficiently risk-averse to enjoy the resulting reduction in risk.

- Investors in the upper middle class benefit little from the presence of the Fund because the gains from the reduction in risk are almost exactly balanced by the losses from the increase in equity prices.

- The remaining investors benefit substantially in percentage terms from the presence of the Fund – but, because they have little invested wealth, their benefit in absolute terms is in fact small.

6 Conclusions

This paper studies the general equilibrium consequences of the availability and cost of index funds. It shows that the availability of index funds tends to *increase* asset prices and *decrease* investor welfare, and that these effects become greater as the cost of indexing declines. While it is true that the availability of index funds allows small investors to enjoy market returns, at equilibrium, these market returns are *lower* than those that were enjoyed by investors before index funds became available. Moreover, the welfare loss that investors experience as a result of these lower market returns may outweigh the welfare gain that they experience from reduced portfolio risk; on balance, the availability of index funds may not benefit *investors*. This finding suggests that the policy debate on index funds' effects on corporate governance and consumer welfare may need to be re-framed.

Appendix A: Proof of Theorem

Proof. In principle, the price of stock could be arbitrarily large or arbitrarily small. However, it is clear that if the price of stock is too large (e.g., larger than the maximum possible profit of a firm) then all investors will prefer bonds, so the demand for stocks will be 0, and if the price of stocks is too small, then the demand for stocks will exceed the supply. Hence we can restrict our attention to stock prices in some interval $[\underline{p}, \overline{p}]$.

For each $p \in [\underline{p}, \overline{p}]$ and each investor $t \in T$, let

$$F^t(p) = \{(x_S^t, x_F^t)\} \subset R_+^2$$

be the set of investor t's optimal choices of shares of individual stocks and shares of the Fund, assuming that the price of stocks is p. (The holding of Bonds is determined by the budget constraint and the assumption that utility is strictly increasing in certain consumption.) Because investor's utility functions are continuous and choice sets are closed, $F^t(p)$ is a non-empty compact set. Moreover, for each investor t, the correspondence

$$p \mapsto F^t(p)$$

is upper-hemi-continuous. By assumption, the distribution ϕ of consumers is non-atomic. The integral of a correspondence with respect to a non-atomic measure is convex so the correspondence

$$p \mapsto F(p) = \int_T F^t(p)\, d\phi(t)$$

is upper hemi-continuous and convex-valued.

By definition, if $x = (x_S, x_F) \in F(p)$ then x_S is the number of shares of stock purchased directly and x_F is the number of shares of stock purchased through the Fund, so $x_S + x_F$ is the total number of shares purchased. Define a correspondence $G : [\underline{p}, \overline{p}] \to \mathbb{R}$ by

$$G(p) = \{x_S + x_F - N : x \in F(p)\}$$

Note that $y \in G$is the market excess demand for stock. It is easily checked that G is upper-semi-continuous and has compact, convex values.

For $y \in (-\infty, +\infty)$ set

$$h(y) = \arg\max_q \{qy : q \in [\underline{p}, \overline{p}]\}$$

Note that

$$y < 0 \Rightarrow h(a) = \underline{p}$$
$$y > 0 \Rightarrow h(a) = \overline{p}$$
$$y = 0 \Rightarrow h(a) = [\underline{p}, \overline{p}]$$

and define a correspondence $H : \mathbb{R} \to [\underline{p}, \overline{p}]$ by

$$H(y, \eta) = h(y)$$

It is evident that H is an upper-hemi-continuous correspondence with compact, convex values. Finally, consider the composite correspondence $H \circ G : [\underline{p}, \overline{p}] \to \overline{p}]$. This is an upper-hemi-continuous correspondence with compact, convex values, and so (by Kakutani's fixed point theorem), $H \circ G$ has a fixed point. That is, there is some p^* such that $p^* \in H \circ G(p^*)$. In view of the definition of h and our choices of $\underline{p}, \overline{p}$, market excess demand for stock at price p^* must be 0. Hence we have found an equilibrium. $\square$

Appendix B: Additional Simulations

To be added

References

Amel-Zadeh, Amir, Fiona Kasperk, and Martin Schmalz, "Mavericks, Universal, and Common Owners-The Largest Shareholders of US Public Firms," *SSRN*, 2022.

Antón, Miguel, Florian Ederer, Mireia Giné, and Martin C Schmalz, "Common ownership, competition, and top management incentives," *Journal of Political Economy*, 2022.

Azar, José and Xavier Vives, "General equilibrium oligopoly and ownership structure," *Econometrica*, 2021, *89* (3), 999–1048.

_ , _ et al., "Reply to: Comments on "General Equilibrium Oligopoly and Ownership Structure"," *Econometrica*, 2021, *89* (3), 1061–1063.

Bach, Laurent, Laurent E Calvet, and Paolo Sodini, "Rich pickings? Risk, return, and skill in household wealth," *American Economic Review*, 2020, *110* (9), 2703–47.

Bender, Svetlana, James J Choi, Danielle Dyson, and Adriana Z Robertson, "Millionaires speak: What drives their personal investment decisions?," *Journal of Financial Economics*, 2022, *146* (1), 305–330.

Beutel, Johannes and Michael Weber, "Beliefs and Portfolios: Causal Evidence," *Chicago Booth Research Paper*, 2022, (22-08).

Bond, Philip and Diego Garcia, "The equilibrium consequences of indexing," *The Review of Financial Studies*, 2022, *35* (7), 3175–3230.

Chinco, Alex and Marco Sammon, "The Passive-Ownership Share Is Double What You Think It Is," *Available at SSRN*, 2022.

Eeckhout, Jan and UPF Barcelona, "Comments on Econometrica's Walras-Bowley lecture: Azar, J. and X. Vives "General Equilibrium Oligopoly and Ownership Structure"," 2020.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and persistence in returns to wealth," *Econometrica*, 2020, *88* (1), 115–170.

Haddad, Valentin and Tyler Muir, "Do intermediaries matter for aggregate asset prices?," *The Journal of Finance*, 2021, *76* (6), 2719–2761.

Hart, Oliver D, "On shareholder unanimity in large stock market economies," *Econometrica: Journal of the Econometric Society*, 1979, pp. 1057–1083.

He, Zhiguo and Arvind Krishnamurthy, "Intermediary asset pricing," *American Economic Review*, 2013, *103* (2), 732–70.

Jiang, Hao, Dimitri Vayanos, and Lu Zheng, "Passive Investing and the Rise of Mega-Firms," 2022.

López, Ángel L and Xavier Vives, "Overlapping ownership, R&D spillovers, and antitrust policy," *Journal of Political Economy*, 2019, *127* (5), 2394–2437.

Panageas, Stavros et al., "The implications of heterogeneity and inequality for asset pricing," *Foundations and Trends® in Finance*, 2020, *12* (3), 199–275.

Philippon, Thomas et al., "A COMMENT ON:"General Equilibrium Oligopoly and Ownership Structure" by José Azar and Xavier Vives," *Econometrica*, 2021, *89* (3), 1055–1059.

Piccolo, Alessio and Jan Schneemeier, "Ownership and Competition," *Available at SSRN 3733795*, 2020.

Rotemberg, Julio, "Financial transaction costs and industrial performance," *Massachusetts Institute of Technology Working Paper*, 1984.

Schmalz, Martin C and William Zame, "Index Funds, Asset Prices, and the Welfare of Investors and Consumers," *Working Paper*, 2022.