

Enhanced indexing and selectivity theory

Andrei Bolshakov¹ | Ludwig B. Chincarini² | Christopher Lewis¹

¹Wedge Capital Management, Charlotte, North Carolina, USA

²United States Commodity Funds, School of Management, University of San Francisco, San Francisco, California, USA

Correspondence

Ludwig B. Chincarini, Professor of Finance, University of San Francisco, School of Management, USA and Advisor to United States Commodity Fund, 101 Howard St., Suite 500, San Francisco, CA 94105, USA.

Email: chincarini@hotmail.com and lbchincarini@usfca.edu

Abstract

Investment managers that believe to have skill must choose some fraction of stocks in a benchmark to hold. Recent theory predicts that the optimal percentage of holdings for a manager with skill is between 50% and 80% of the benchmark. This theory requires a number of assumptions. Using simulations, we relax some of the assumptions to examine if the theory still holds. We find that, for the most part, the theory holds when the assumptions are relaxed. We also extend the analysis from a one-period horizon to a multiperiod horizon using the actual returns of stocks in the S&P 500.

KEYWORDS

active management, enhanced indexing, information ratio, portfolio management, simulations

JEL CLASSIFICATION

G0, G13

1 | INTRODUCTION

Past research on the topic of active versus passive investment management has broadly focused on the returns of active managers: whether there is evidence of outperformance in excess of the manager's fees and expenses (Berk & van Binsbergen, 2015; Fama & French, 2010) and what portfolio characteristics tend to correspond with outperformance (Cremers & Pareek, 2016). For managers pursuing a near-passive or enhanced index approach to investing, the

We would like to thank Viktoria Dalko, Christopher Otrók, Ekatarina Sergina, the editor, John Doukas, the participants at the 2020 Western Economic Association International annual conference, the participants at the 2021 Western Economic Association International annual conference, the participants at the 2021 Southwestern Finance Association annual conference and an anonymous referee for helpful comments. Andrei Bolshakov and Christopher Lewis are general partners of Wedge Capital Management, 301 South College Street, Suite 3800, Charlotte, NC 28202.

information ratio (IR) of their portfolio is the relevant metric to assess performance. Recent theory (Bolshakov & Chincarini, 2020) provides a novel theoretical approach to determining the optimal number of stocks a skilful manager should hold from their benchmark to maximize their IR. The theory focuses on the ratio of the number of stocks selected by the portfolio manager to the number of benchmark constituents and labels this as the “selectivity ratio” of the manager. The authors find that depending on whether stocks are purchased into the portfolio in a single, one-time, bulk selection or by sequential, discrete additions, the manager’s IR will be optimized at either a 50% selectivity ratio or at an 80% selectivity ratio, respectively. Their theoretical approach is novel and provocative but introduces certain simplifying assumptions to derive their results about optimal selectivity.

This selectivity theory could have important applications in the real world. One such application would be to infer a manager’s skill from performance data to understand how effective they are at picking stocks and whether they are picking the appropriate number of stocks. An extension of this would be to use the technology of selectivity theory to understand in more depth what type of skill the portfolio manager has, not just whether they have the skill or not. Another important application would be to determine how a chief investment officer (CIO) could combine the information of the portfolio selections of many professional money managers with a low selectivity ratio to create a more optimal portfolio with a closer-to-optimal selectivity. All of these extensions would be a test of this new theory, which if successful might explain certain practices in the industry, as well as improve current practice. However, to bring selectivity theory to real-world applications, some of the simplifying assumptions must be relaxed so as to determine whether they are limitations to the theory or not. Furthermore, the process of relaxing the assumptions might lead to techniques that can be used to apply selectivity theory to real-world problems.

This paper uses Monte Carlo simulations to examine whether selectivity theory holds even when the simplistic assumptions are relaxed. We find that the theory is quite robust even when the assumptions are relaxed. In addition, the use of simulations allows the work of Bolshakov and Chincarini (2020) to be extended to a more realistic dynamic investment environment, where portfolio managers are able to select stocks every month, rather than for one investment period. This expanded dynamic environment suggests that the optimal range of selectivity for a portfolio manager, based on various assumptions in the dynamic case, is somewhere between 50% and 80% of the manager’s benchmark.

The paper is organized as follows: Section 2 reviews selectivity theory; Section 3 discusses the data and methodology used in the paper; Section 4 discusses the methods used to relax the theoretical assumptions and the simulation results; Section 5 describes a methodology for a dynamic portfolio process and the simulation results of that process, and Section 6 concludes.

2 | THEORY

Bolshakov and Chincarini (2020) consider the case of a skillful portfolio manager seeking to maximize their IR by choosing a portfolio selectivity ratio when forming their portfolio. The manager’s skill is described as a probability weighting advantage to choosing outperforming stocks (‘winners’) versus under-performing stocks (‘losers’). For a skilful manager, this probability advantage is encompassed in ω (i.e., omega), which is greater than one for a skilled manager and equal to one for an unskilled manager. Bolshakov and Chincarini (2020) brought this entirely new approach to the portfolio management literature that involves the rarely

mentioned Fisher and Wallenius Noncentral Hypergeometric distributions. An ω greater than one means the portfolio manager has a higher likelihood of choosing a winner stock rather than a loser stock.¹ The theory predicts that when a portfolio manager selects stocks in bulk manner, that is, they decide the percentage of stocks that they wish to buy of the benchmark in advance, the process is governed by the Fisher Noncentral Hypergeometric Distribution (FNCH) and it will be optimal for managers to select 50% of the benchmark. When a portfolio manager selects stocks in a sequential manner, then the process is governed by the Wallenius Noncentral Hypergeometric Distribution (WNCH). In this case, the theory suggests that the portfolio manager should choose approximately 80% of the investment universe if such a manager has skill.

The theory makes four simplifying assumptions to arrive at these conclusions. The first assumption is that half of the stocks will outperform the index over some period in the future and half will under-perform. The second assumption is that the performance of each stock over the investment period can be expressed as a binary event. That is, either it beats the index or it does not. This implicitly assumes that the magnitude of the performance does not matter. The third assumption is that the portfolio manager has a constant ability to pick winner stocks over loser stocks, regardless of how many stocks have been chosen. The fourth assumption is that the benchmark and portfolio are equally weighted. Finally, the theory takes place in a one-period investment horizon, just like the Capital Asset Pricing Model (CAPM) (Sharpe, 1964), however, in this paper, we will also present results for dynamic choice and a multiperiod horizon.

3 | DATA AND METHODOLOGY

To relax the assumptions of the theoretical work, we make use of simulations to examine the behaviour of the optimal portfolios and selectivity. Relaxing some of the assumptions requires the use of actual real stock market data. In our work, we make use of the S&P 500 stock return data and market capitalization data from 31 December 1988 to 31 December 2018. We used monthly total stock return data that was obtained from S&P Global. We also obtained the monthly constituents of the S&P 500 from that same database. Over the entire period, company data might be missing for a particular month due to inaccurate data, bankruptcy, de-listing, merger and acquisitions, or a host of other corporate actions. We worked closely with the Standard and Poor's team to make sure that our data was void of survivorship bias and, to the extent possible, that our total returns for any stock were accurate, even if a corporate action led to the stock no longer being traded.

The summary statistics for our S&P 500 data are shown in Table 1. In this paper, we primarily consider an equal-weighted S&P 500, but we also examine a market capitalization weighted S&P 500. In Table 1, we show statistics for both benchmarks over several time periods.

The Monte Carlo simulations vary in each simulation in which we relax an assumption of the model, however, the general process is the same. Given a set of parameters and/or real data, we simulate the behavior of an individual portfolio manager in selecting a percentage of the underlying benchmark. In this paper, our benchmark is the S&P 500, and hence the benchmark always has 500 stocks. We choose every selectivity ratio from 10% (50 stocks) to 90% (450

¹The parameter ω is not equal to the probability of choosing a winner stock, but it is related. We chose this parameter, as it is standard when using the Fischer and Wallenius Noncentral Hypergeometric distributions, which are relevant to this type of portfolio manager selection.

TABLE 1 Summary statistics for the S&P 500 data

This table reports summary statistics of the S&P 500 stocks used in the analysis of this paper. μ represents the monthly average return of the stocks over the relevant sample period, σ represents the standard deviation of the monthly returns of the stocks over the sample period, Sharpe represents the Sharpe ratio for the portfolio of stocks over the sample period, Max represents the maximum monthly return for an individual stock over the sample period, Min represents the minimum monthly return for an individual stock over the sample period, % Winners is the number of stocks that have a monthly return above the average of the benchmark and Herf. represents the Herfindahl index of diversification of stocks with 20 being the most diversified and 100 being the most undiversified.

Time period	Market capitalization-weighted						Equal weighted							
	μ	σ	Sharpe	Max	Min	% Winners	Herf.	μ	σ	Sharpe	Max	Min	% Winners	Herf.
1989–2018	0.88	4.08	0.16	11.44	−16.69	50	79.07	1.00	4.63	0.16	18.48	−20.91	49	20.00
1989–1998	1.54	3.85	0.29	11.44	−14.37	48	71.67	1.37	4.17	0.23	11.38	−14.88	49	20.00
1999–2008	−0.01	4.35	−0.06	9.93	−16.69	51	90.27	0.35	5.08	0.02	11.04	−20.91	49	20.00
2009–2018	1.11	3.91	0.28	10.89	−10.40	50	75.27	1.28	4.56	0.27	18.48	−11.74	50	20.00

stocks) in 0.2% intervals, thus replicating a total of 401 selectivity ratios. For each selectivity ratio, we allow the portfolio manager with skill level, ω , to select N_S stocks for that particular selectivity ratio. Thus, for a 10% selectivity ratio, it would be 50 stocks. For the single-horizon simulations, we do this 100,000 times resulting in 100,000 portfolios chosen with the particular parameters. For the dynamic, multiperiod portfolios, we do this 10,000 times resulting in 10,000 portfolios chosen with the particular parameters.² We then stored all of the resulting information of that portfolio, including, average return, standard deviation of returns, tracking error, and most importantly the IR. We did this separately for the bulk selection method (FNCH) and the sequential selection method (WNCH).³ Finally, to plot or analyze the IRs based on different selectivity ratios, we took the average for each selectivity ratio across different simulations.

4 | EXAMINATION OF ASSUMPTIONS

4.1 | Baseline model

Before relaxing the assumptions, we use the Monte Carlo simulations to replicate the base case presented in Bolshakov and Chincarini (2020). In that paper, the authors assumed that the benchmark was composed of 50% winner stocks and 50% loser stocks. They also assumed that the return of good stocks was 10%, whereas the return of bad stocks was -10% . Figure 1 shows the average IRs from the 100,000 simulations along with the theoretical prediction. As would be expected, the simulations correspond to the theoretical predictions in the baseline case.

The figure also shows a fourth-degree polynomial function fit through the simulated IRs.⁴ The optimal selectivity ratio based on the polynomial function is marked with a vertical, black line at 79.0%. At that level of selectivity, the resulting IR is 0.855. It should be noted that as the selectivity ratio approaches 100%, the IR converges to zero as there is no added value by the portfolio manager. This is true in all of our figures, even though we do not explicitly show the 100% selectivity point.⁵

In the simulations that follow, where we relax some of the assumptions from the theoretical model, we provide a summary of the major differences in each simulation in Table 2. In most of these simulations, we use a value for ω of 1.1. The nature of the results remains the same even if we choose other values of ω . When ω is very high (larger than 2), we need to use the downside IR, rather than the regular IR for reasons discussed in Bolshakov and Chincarini (2020).⁶

²The reason we limited the dynamic simulations to just 10,000 is due to the time-intensive nature of the simulations. The qualitative results of the simulations would not change if we did more simulations.

³Appendix A describes the simulation process for the Fisher distribution, which is more complicated than the Wal-lenius to simulate.

⁴We estimate a fourth-degree polynomial using linear regression of the form, $y = ax + bx^2 + cx^3 + dx^4$. The fitted polynomial function achieves a correlation of 0.9997 with the theoretical IRs from the WNCH distribution. Also, as the benchmark return is by construction 0%, each IR is calculated as the ratio of expected returns to standard deviation of returns for simulations at a given selectivity ratio.

⁵Also, one should note that when ω is unrealistically high, say a value of 10, then the optimal selectivity converges to the percentage of winners in the benchmark.

⁶One should also refer to Bolshakov and Chincarini (2020) for the importance of using downside risk in the Fisher calculations due to an exploding IR when ω is very high.




FIGURE 1 Baseline simulations. This figure shows the baseline simulations of each selectivity ratio and its corresponding information ratio (IR). The dots represent the average IR from 100,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages and the theoretical Wallenius IRs for each selectivity ratio. The optimal selectivity ratio is at 79%. The parameters for the simulations are $N = 500$, $\omega = 1.1$, $r_g = 0.10$, $r_b = -0.10$ and $r_{BM} = 0$. WNCH, Wallenius Noncentral Hypergeometric Distribution

4.2 | Concentration of winners

One of the assumptions in Bolshakov and Chincarini (2020) was that 50% of the stocks in the index would outperform and 50% would under-perform in a given period. This assumption is simplistic and does not take into account the positive skewness of stock returns as the upside potential of a stock is unlimited, whereas the downside loss of a stock is limited. Thus, there could be fewer winner stocks than loser stocks. In this simulation, we keep all other assumptions the same, however, we vary the percentage of winners in the actual universe from 20% to 50%.⁷

Table 3 shows the benchmark return, the optimal selectivity ratio, the expected return at the optimal selectivity ratio, and the IR at the optimal selectivity ratio for different proportions of winner stocks in the benchmark universe. As in the theoretical model for the sequential selection method and the baseline case, the optimal selectivity is around 80% of the stocks in the universe. Decreasing the number of winners from 50% to a more realistic 48% does not alter the optimal ratio, and only slightly diminishes the IR achieved at that point. Even more extreme scenarios, where there are only 20% or 30% winners, only slightly shifts the optimal selectivity point towards holding more

⁷To compute the IRs for these simulations the numerator is the difference between the expected portfolio return and the constructed benchmark return. The denominator retains its prior definition as the cross-sectional standard deviation of returns within a given selectivity ratio, as the benchmark return is constant across all simulations.

TABLE 2 Summary of parameters in the different simulations

This table reports the parameters for each simulation that relaxes the assumptions of selectivity theory. N represents the number of stocks in the benchmark portfolio, Winners (m_1) represents the number of winner stocks in the benchmark, Losers (m_2) represents the number of loser stocks in the benchmark, Returns describes whether the simulation uses 10% for winners and -10% for losers or uses the actual stock returns of S&P 500 stocks normalized as a Z-score, Weight is for the mechanism to weight stocks, where EW is for equal-weighted, Holding Period is the typical holding period after a portfolio selection is made, ω represents the skill of the portfolio manager, N_{sims} is the number of simulations, and S^* is the optimal selectivity. † This number will vary depending on the cut-off point where the manager's skill ends. We chose to use 30% as the cutoff point. †† This is only true for reasonable values of ω . For a very high ω , where the portfolio manager is extremely skilled, this optimal selectivity ratio will converge to the saturation point. ††† The optimal selectivity will converge to 79% if the appropriate portfolio is chosen based on the weighting of the benchmark.

Parameters		Winners		Losers		Returns	Weight	Holding period	ω	N_{sims}	S^*
Simulation name	N	m_1	m_2								
Baseline	500	50%	50%	Varies	50%	+ / - 10%	EW	One horizon	1.1	100,000	Near 79%
Concentrated winners	500	Varies	Varies	Varies	Varies	+ / - 10%	EW	One horizon	1.1	100,000	Near 79%
Nonbinary returns	500	Varies	Varies	Varies	Varies	S&P 500 Z-Scores	EW	One horizon	1.1	100,000	Near 79%
Steadily declining ω	500	50%	50%	50%	50%	+ / - 10%	EW	One horizon	1.1 to 1	100,000	Near 50%
Jack-knife	500	50%	50%	50%	50%	+ / - 10%	EW	One horizon	1.1	100,000	Near 30%†
Ability saturation	500	50%	50%	50%	50%	+ / - 10%	EW	One horizon	1.1	100,000	Near 75%††
Uncertainty	500	50%	50%	50%	50%	+ / - 10%	EW	One horizon	$\omega \sim N(1.1, \sigma_\omega^2)$	100,000	Near 79%
Non-equal weight	500	50%	50%	50%	50%	+ / - 10%	Varies	One horizon	1.1	100,000	Near 79%†††
Dynamic	500	Varies	Varies	Varies	Varies	S&P 500 Stock Returns	EW	Multihorizon	1.1	10,000	Near 79%

TABLE 3 Portfolio characteristics as concentration of winners changes

This table reports the results from 100,000 simulations for each selectivity level for the sequential selection method. The parameters for the simulation are $N = 500$, $\omega = 1.1$, $r_g = 10\%$ and $r_b = -10\%$. The percentage of winners in the benchmark is altered in each of the simulations. Thus, 50% corresponds to a benchmark with 250 winner stocks and 250 loser stocks, whereas 20% corresponds to a situation where the benchmark has only 100 winner stocks and 400 loser stocks and so on. The benchmark return is provided for convenience showing that when there are less than 50% winner stocks, the benchmark return is naturally negative. AR is for average return and IR is for information ratio.

	Percentage of winners in benchmark				
	50%	48%	40%	30%	20%
Benchmark return (%)	0.000	−0.400	−2.000	−4.000	−6.000
Optimal selectivity ratio	79.0	79.0	79.4	79.4	79.6
AR at optimal selectivity	0.197	−0.203	−1.813	−3.839	−5.878
AR at 79% selectivity	0.197	−0.203	−1.812	−3.836	−5.876
IR at optimal selectivity	0.855	0.854	0.851	0.791	0.697
IR at 79% selectivity	0.855	0.854	0.846	0.795	0.694

stocks.⁸ The maximum achievable IR shrinks more drastically though; as winners become scarcer within the benchmark, the manager's skill in selecting winners has fewer opportunities to display itself. In the case where winners make up 30% of the benchmark, small selectivity ratio changes are so immaterial that the original 79.0% point shows a higher IR within its 100,000 simulations than the polynomial-fitted optimal location of 79.4% does.⁹

One may also notice the negative expected return values. This is because, with very few winners in the universe, the winner stock returns of 10% are overwhelmed by the loser stock returns of −10%. Thus, the benchmark return and the portfolio return are both negative given our parameters.

4.3 | Binary return distribution

The second assumption in selectivity theory is that the magnitude of the returns is not important. This is surely not a fair assumption, as a very good stock pick with a very high return can afford a manager many bad picks of lower magnitude poor returns. To relax this assumption, we use the actual distribution of stock returns, while holding all else constant to examine the impact of this assumption. Figure 2 shows a time-series, cross-sectional average of

⁸Some studies have shown that just a few winners drive the bulk of stock returns (Bessembinder, 2018). Although not shown in the paper, even when we reduce the saturation point of winners, to as little as 4%, the optimal selectivity ratio is still near 80%. Naturally, the average return declines tremendously, but relative to the benchmark, the portfolio still has a positive IR. For example, in a simulation of 4% winners in the benchmark, we obtained an optimal selectivity ratio of 80% with an $\omega = 1.1$, with an average return of −9.171 and an IR of 0.345.

⁹This is just the approximation error from fitting a polynomial function to a series of points.




FIGURE 2 Average cross-sectional Z-scores for S&P 500 stock returns. This figure shows the distribution of Z-scores of stock returns for stocks in the S&P 500 from December 1988 to December 2018. In each period, cross-sectional Z-scores are created based on the returns of each stock in the S&P 500. This is repeated for every month in the sample period and the average Z-scores compute along with the frequency of occurrence. A normal distribution is fitted to this histogram

stock return Z-scores from the S&P 500.¹⁰ Figure 2 illustrates what we already know. Stock returns are not binary and vary substantially within a stock universe.

At this point, we use this distribution of stock returns in our simulations to investigate the optimal selectivity ratio. Thus, in every simulation, a portfolio manager selects a certain number of stocks. If the manager picks n_w winner stocks in his basket of n stocks, then we randomly assign a Z-score to each winner stock based on one of the winner Z-scores in Figure 2.¹¹ We do the same with his loser picks. We then compute the IRs and plot them against selectivity ratios.

Figure 3 shows the results of the simulations for the sequential selection method. Once again, the maximum IR is achieved very close to the theoretical selectivity ratio of 79.0%. However, the maximum IR value declines by 27% compared to the base case, from 0.855 to

¹⁰To construct this figure, we compute the Z-scores of stock returns in each month from 31 December 1988 to 31 December 2018. See Chincarini and Kim (2006) for the detailed methodology on computing stock Z-scores. We then rank stocks by highest Z-score in each period to lowest Z-score so that there are 500 ranked Z-scores. We then average over time, each ranked Z-score with its corresponding Z-score in all other months. Of course, these will not necessarily be the same stock. We then plot the histogram of average Z-scores versus frequency and also fit a normal distribution curve to the data. Although we used Z-scores for simplicity, we could have also used stock returns. We could have also just used one specific random month of data, rather than average over time. Neither of these alternative methods would change the nature of the results.

¹¹The Z-scores are assigned according to the truncated distribution of positive Z-scores. Thus, if the portfolio manager chose 20 winners, then we assigned Z-scores based upon the actual distribution of the 250 winner Z-scores, that is, higher Z-scores would occur less frequently.




FIGURE 3 Selectivity ratio and information ratio (IR) for empirical S&P 500 returns. This figure shows each selectivity ratio and its corresponding IR when returns are not binary for good and bad stocks but are based on the cross-sectional Z-scores of actual S&P 500 individual stock returns. The dots represent the average IR from 100,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages and the theoretical Wallenius IRs for each selectivity ratio. The parameters for the simulations are $N = 500$, $\omega = 1.1$, $r_i = \bar{z}_i$ and $r_{BM} = 0$

0.624, as random luck now plays a larger part in the manager's posted results, as their skilful selection of a particular winner could be either a big winner or a mild winner.¹²

4.4 | Constant manager skill

In the basic model of selectivity, Bolshakov and Chincarini (2020) assume that the manager's skill is constant regardless of how many stocks he picks. This assumption might be criticized as one might believe that the manager's skill worsens over time as more stocks are picked.¹³ One might also imagine a world where the manager believes he has the skill but is uncertain as to how much skill he has. That is, the skill might only be for a small handful of stocks and then his ideas run out. Another manager might worry even whether he has skill on any given pick.

¹²This 27% reduction in IR can be directly explained by the increased volatility in the return distribution. In the binary case of returns (+/− 10%), the coefficient of variation (i.e., the standard deviation divided by the absolute mean [CV]) is 1.0. Figure 3 shows a return distribution with $CV = 1.38$ which proportionately reduces the denominator of the IR resulting in a 27.5% ($1 - 1/1.38$) discount which is very close to the reduction realized in the simulations.

¹³Worth noting is that even with a fixed skill parameter ω , the WNCH distribution exhibits a gradual decline in the probability of selecting winners. As more stocks are added into the portfolio, the number of winners left in the benchmark shrinks on average and, therefore, the probability that the manager picks a winner for their next stock tends to shrink as well, all the way towards a 0% probability if all the winners have already been chosen.

That is, maybe the manager believes he has skill on average, but there is an uncertainty with his skill in any given period or on any given stock.

In this section, we attempt to relax the assumption of constant manager skill regardless of how many stocks are picked and examine how the optimal selectivity changes. If the portfolio manager believes he has skill, then some of these results will be relevant. If the manager has no skill, then the manager should not be in the stock-picking game and should be managing an index fund.

4.4.1 | Steadily declining manager skill

Traditional equity factor analysis, where the researcher uses historical data to periodically build theoretical portfolios based on some objective characteristic, can be thought of as an extension of these declining skill cases. For example, the portfolio returns from sorting stocks into deciles based on their cashflow-to-price ratio might produce a gradual decrease in returns for every successive decile (i.e., cashflow-to-price getting smaller with each decile).

To simulate the effects of a decline in skill, the initial parameter of $\omega = 1.1$ was allowed to gradually decline as each stock was picked from the benchmark of stocks. Thus, before the first stock pick, $\omega = 1.1$ and linearly declines to $\omega = 1$ by the 500th stock pick. The declining ω leads to a straight decline in the expected return of the portfolio at each selectivity ratio, whereas the standard deviation of the portfolio declines in a convex fashion versus selectivity. The resulting IR is optimal near the 50% level (see Figure 4), reminiscent of the bulk selection method rather than the sequential method. In Figure 4, we also show the selectivity ratio and IR when skill does not decline (it's the line on the left most of the figure that extends beyond the top of the graph—theoretical WNCH IRs). Interestingly, in the linearly declining skill case, the sequential selection method converges to the bulk selection method (i.e., theoretical FNCH) at a 50% selectivity ratio.

4.4.2 | Jack-knife skill

Another way to imagine skill not being constant is a portfolio manager that has good stock picking ability ($\omega = 1.1$) only for a portion of the stock universe and then no ability after that. For illustration purposes, we arbitrarily chose a manager with a selection ability of up to 30% of the universe and then no ability ($\omega = 1$) for any stock after that. One reason that we use this so-called 'jack-knife' scenario is to illustrate a more dramatic departure from the steadily declining skill situation. In that situation, although skill is declining, the manager always has skill. In the case of the 'jack-knife', at a certain selectivity point, skill completely disappears. Selectivity theory would not apply in this case, as the manager would not select beyond the point at which his or her skill vanishes. This can be seen in Figure 5, where there is a clear break from the theoretical WNCH function and the optimal selectivity is maximized at the arbitrarily chosen 30% point.

For a manager to know exactly how and when their skill declines is unrealistic, but the two prior examples are helpful as they suggest a more general conclusion for sequential skill decay functions. It seems to be the case that the standard deviation of portfolio returns does not shift by much compared to the base case. Therefore, as different functions are incorporated to reflect a more rapid skill decline, the optimal selectivity ratio will shift towards holding fewer stocks




FIGURE 4 Selectivity with steadily declining manager skill. This figure shows the information ratios (IR) and selectivity levels when the manager's skill is steadily declining from $\omega = 1.1$ on the first stock pick to $\omega = 1$ on the last stock pick. The dots represent the average IR from 100,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages and the theoretical Wallenius IRs for each selectivity ratio. The optimal selectivity ratio is at 50%. The parameters for the simulations are $N = 500$, $r_g = 0.10$, $r_b = -0.10$ and $r_{BM} = 0$. FNCH stands for the Fisher Noncentral Hypergeometric Distribution and WNCH stands for the Wallenius Noncentral Hypergeometric Distribution

depending on how punitive that function is on the manager's expected returns. The maximum attainable IR will decline significantly though, so instead of shrinking their portfolio a skilful manager may be better served trying to hire like-minded analysts to their team, or devote time to educating new recruits in their successful methods.

4.4.3 | Ability saturation

Another way of adjusting the constant skill assumption is to assume that the manager believes to have skill only for a subset of the benchmark population. After that, the manager has no skill. In the real world, such a scenario might reflect a manager's familiarity with certain sectors, or a quantitative model's inability to classify certain stocks such as initial public offerings or mergers that lack relevant financial data for the new company. Table 4 shows how the optimal selectivity changes with the percentage of the benchmark that the manager has skill for, from 100% to 20% (the portfolio manager only has stock picking ability for 20% of the stock universe). This is accomplished in the simulation by allowing the manager to pick stocks sequentially until the manager reaches the saturation level. That is, we do not allow him to pick any more winners once the saturation level is reached.

Having fewer identifiable winners available to the manager lowers the IR that the manager can generate. The course-correcting nature of the WNCH distribution is once again effecting




FIGURE 5 Selectivity with jack-knife skill. This figure shows each selectivity ratio and its corresponding information ratio (IR) when the portfolio manager's skill completely vanishes at a certain selectivity ratio. That is, the manager has skill until a certain selectivity and then it reverts to no skill. The dots represent the average IR from 100,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages and the theoretical Wallenius IRs for each selectivity ratio. The parameters for the simulations are $N = 500$, $\omega = 1.1$ for $\phi = n/N < 0.30$, $\omega = 1$ for $\phi = n/N \geq 0.30$, $r_g = 0.10$, $r_b = -0.10$ and $r_{BM} = 0$. WNCH stands for the Wallenius Noncentral Hypergeometric Distribution

TABLE 4 Portfolio characteristics with winner saturation

This table reports the results from 100,000 simulations for each selectivity level for the sequential selection method. The parameters for the simulation are $N = 500$, $\omega = 1.1$, $r_g = 10\%$ and $r_b = -10\%$. The percentage of winners that the portfolio manager can identify is altered in each of the simulations. Thus, 50% corresponds to a situation where the portfolio manager can only identify 50% of the winner stocks, whereas 20% corresponds to a situation where the portfolio manager can only identify 20% of the winner stocks. Thus, if there are 250 winner stocks, the portfolio manager can only identify 50 of those and no more. AR is for average return and IR is for information ratio.

	Percentage of winners manager can identify				
	100%	80%	60%	40%	20%
Optimal selectivity ratio	79.0	78.2	77.0	76.2	75.0
AR at optimal selectivity	0.197	0.161	0.124	0.084	0.042
AR at 79% selectivity	0.197	0.157	0.116	0.079	0.040
IR at optimal selectivity	0.855	0.681	0.508	0.334	0.161
IR at 79% selectivity	0.855	0.683	0.506	0.342	0.174

the results, so the skilful manager is incentivized to hold more than 75% of the benchmark in all cases, to increase the probability that their skill shines through in the resulting portfolio. One can see from the table that in the original theoretical model, the optimal selectivity is 79%, but when the manager can only pick 20% of the winners in the universe, the optimal selectivity drops to 75%, but still far more than many managers hold.¹⁴

4.4.4 | Uncertainty in manager skill

Yet another way that manager skill might not be constant is that a manager may believe that on average he has skill, that is, $\omega > 1$, however, he is unsure about the exact value of his ω . In each of our simulations, we chose the manager's ω from a distribution. In particular, $\omega \sim N(1.1, \sigma_\omega^2)$. We allowed σ_ω^2 to vary from 0 to 0.20. The variability in a manager's skill could describe many real-life situations. For example, it could describe a manager that makes accurate predictions but the rest of the market fails to value stocks commensurate to his new information in a sufficient time period. It could also represent a period in which a manager's quantitative factor model or macroeconomic developments render the manager's skill ineffective within a particular time period. An example of such a changing period was 2018 and 2019, when the performance of momentum and value strategies was lackluster.¹⁵

The simulation was done two ways. In the first method, every time the manager selected a stock, a random draw was made to determine the skill of the manager on the next draw. This was repeated throughout the selection. In the second method, we drew the ω parameter for a particular manager and kept it the same for the entire stock-picking exercise. The results of these simulations are shown in Table 5. The first method might represent a portfolio manager that has skill on average, but for any given stock pick, or period of selection, has variability in his skill level. Sometimes the manager has bad days and sometimes good days, but on average is good. The second method could explain a manager that has bouts of good and bad streaks during his investment horizons. Over a career, this is surely possible.

The table shows that portfolio expected returns for a given selectivity ratio shrink at roughly the same rate for both methods. Overall, the methods have slightly different results. With method 1, the variability of ω has less of an effect on the IRs than it does for method 2. With method 2, an increase in the uncertainty of ω causes a large decrease in the expected IR of the manager. However, the important point is that for both methods, the optimal selectivity is still around 80%. In the end, the actual value of the uncertainty does not matter. We could increase the uncertainty further, but it would simply increase the dispersion amongst different simulations and lead to a lower average IR, but the mean results of optimal selectivity would not change.

4.4.5 | A simple and practical application of declining skill

Most practitioners are more concerned with practical, straightforward empirical relationships. To provide something along those lines, we use the Fama–French factor cash-flow-to-price as a

¹⁴If ω is unrealistically high, then this is no longer true. With an extremely high ω , the manager's maximum point might come much earlier with a selectivity close to the saturation level. This is because the manager would be close to a perfect stock picker, which is unlikely in practice.

¹⁵It also could be a consequence of crowding in these factors (Chincarini, 2012).

TABLE 5 Portfolio characteristics when manager skill is stochastic

This table reports the results from 100,000 simulations for each selectivity level for the sequential selection method when the skill of the manager is uncertain. The parameters for the simulation are $N = 500$, $\omega \sim N(1.1, \sigma^2)$, where σ varies in the table from 0.00 to 0.20, $r_g = 10\%$ and $r_b = -10\%$. The benchmark return is 0% for information ratio calculations. The standard deviation of skill, σ , is used in two ways. In Panel (a), a random draw of skill occurs before a manager selects his group of stocks for each selectivity level. In Panel (b), the manager's skill is selected before each simulation of all selectivity levels. AR is for average return and IR is for information ratio.

	Standard deviation of skill (σ)				
	0.00	0.05	0.10	0.15	0.20
(a) Skill varies for each stock selection					
Optimal selectivity ratio	79.0	78.8	79.0	79.0	79.2
AR at optimal selectivity	0.197	0.196	0.188	0.177	0.161
AR at 79% selectivity	0.197	0.196	0.188	0.177	0.163
IR at optimal selectivity	0.855	0.841	0.816	0.766	0.701
IR at 79% selectivity	0.855	0.851	0.816	0.766	0.705
(b) Skill varies for each portfolio simulation					
Optimal selectivity ratio	79.0	79.0	79.0	79.2	81.0
AR at optimal selectivity	0.197	0.195	0.187	0.173	0.150
AR at 79% selectivity	0.197	0.195	0.187	0.178	0.160
IR at optimal selectivity	0.855	0.788	0.627	0.474	0.354
IR at 79% selectivity	0.855	0.788	0.627	0.483	0.354

quantitative factor to separate portfolios.¹⁶ The factor portfolios are divided into deciles by their value of the cash-flow-to-price ratio every year and the monthly equal-weighted returns are computed. It has been shown that cash-flow-to-price can be considered a 'value' factor since companies with higher values tend to have higher average returns on average.

Traditional equity factor analysis, where the portfolio manager uses historical data to periodically build portfolios based on some objective characteristic, can be thought of as an extension of these declining skill cases. For example, the portfolio returns from sorting stocks into deciles based on their cashflow/price ratio produces a presumed decrease in returns for every successive decile. Although that pattern means the first decile has offered the best returns historically, a quantitative manager might still consider owning additional deciles to reduce tracking error against their defined universe. Moving from 10% to 20% selectivity would do just that, but it also introduces a deterioration of skill commensurate with the step down in returns. Further expansion of the selectivity ratio would lead to additional—but uneven—declines in both the numerator and denominator of the IR.

In Table 6, we produce the annualized returns of each decile from July 1951 through May 2019. One can see the implicit selectivity ratios in that, each decile moving from one to 10 is implicitly moving towards more and more stocks according to that particular selection criterion

¹⁶Source: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

TABLE 6 Cashflow-to-price factor and selectivity

This table reports the historical returns from July 1951 to May 2019 of the different deciles of the Fama–French cashflow-to-price variable. The data is from https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. The portfolios are constructed as the ranked deciles of the cashflow-to-price variable. Portfolios are formed on CF/P at the end of each June using NYSE breakpoints. The cashflow used in June of year t is total earnings before extraordinary items, plus equity's share of depreciation, plus deferred taxes (if available) for the last fiscal year-end in $t - 1$. P (actually ME) is price times shares outstanding at the end of December of $t - 1$. All NYSE, AMEX, and NASDAQ stocks for which we have ME for December of $t - 1$ and June of t , and cashflow for calendar year $t - 1$. Panel (a) represents the historical performance of the Fama–French decile portfolios. Average return (AR) is the annualized average return of the decile over the entire period. Marginal decline in AR represents the annual percentage return decline by moving from decile i to decile $i + 1$. Thus, while decile 1 has a 19.25% annualized return, decile 2 has an 18.03% return, which is 1.22% less than decile 1. Panel (b) represents the cumulative decile performance. That is if an investor buys the stocks in all preceding deciles as well as the current decile. For example, if the investor purchases only decile 1, he achieves a return of 19.25% per year, if he buys both decile 1 and 2, he achieves a return of 18.64% per year, and so on and so forth. Decile 10 represents an investor that buys all 10 deciles and holds them in equal weighting. Thus, decile 1 would be the equivalent of a 10% selectivity ratio and decile 10 would be the equivalent of a 100% selectivity ratio using this factor. AR is the annualized return of each combination, excess return is the return of the group of deciles minus the returns of the benchmark, which is the entire universe of stocks (i.e., holding all 10 deciles), Tracking Error is the annualized standard deviation of returns of the particular portfolio or group of deciles versus the benchmark, and the information ratio is the information ratio of the group of deciles or selectivity ratio versus the benchmark. For example, holding deciles 1 and 2 (selectivity ratio of 20%) leads to an AR of 18.64% per year compared to the benchmark (holding all deciles) of 15.17% per year with an excess return of 3.47% and a tracking error of 4.41%, and an information ratio of 0.82.

(a) Individual decile performance										
	1	2	3	4	5	6	7	8	9	10
Average return	19.25	18.03	16.57	16.56	15.28	14.87	14.17	13.51	12.43	11.02
Marginal decline in AR	-	1.22	1.46	0.01	1.28	0.41	0.71	0.66	1.08	1.40
(b) Cumulative decile performance										
	1	2	3	4	5	6	7	8	9	10
Average return	19.25	18.64	17.95	17.60	17.14	16.76	16.39	16.03	15.63	15.17
Excess return	4.08	3.47	2.78	2.43	1.97	1.59	1.22	0.86	0.46	0.00
Tracking error	5.86	4.41	3.39	2.84	2.39	2.01	1.67	1.34	0.85	0.00
Information ratio	0.70	0.79	0.82	0.86	0.82	0.79	0.73	0.64	0.54	-

or that factor, that is, cash-flow-to-price. The ‘best’ stocks would be considered Decile 1, whereas the worst would be considered Decile 10. If you think of all the deciles as the benchmark portfolio, then, each additional decile is closer to the benchmark.

In Table 6 we also show the IR historically of different selectivity ratios. In this particular case, the 40% selectivity obtained the highest historical IR. One way to view these results is to consider the real world of factor selection as a combination of linearly declining skill and jack-knife skill as we presented earlier. In this case, the results might naturally lie at 40% selectivity.

In a more general sense, the cashflow yield factor also represents one factor over one sequence of history.¹⁷ In practice, quantitative portfolio managers do not buy all deciles. They might buy a portion of the first decile and wait a period, then buy or sell depending on the movements of stocks according to that factor. Thus, their behavior is more dynamic than a one-period investment choice.¹⁸

4.5 | Equal-weighted benchmark

The fourth assumption in Bolshakov and Chincarini (2020) is that the benchmark is an equal-weighted benchmark. This is not a strong assumption, in the sense, that many benchmarks are equally weighted. In this section, we consider how the results would change if we use a market-capitalization-weighted benchmark like the S&P 500. In practice, there is considerable size variation to companies within a benchmark, particularly for large-cap benchmarks where the largest stock can be hundreds of times the size of the smallest stock. Many popular benchmarks such as the S&P 500 weigh their holdings according to their market capitalizations. That methodology can lead to a handful of stocks having an outsized influence on the benchmark's return. Figure 6 shows the returns of S&P 500 stocks for the month of January 2019. The stocks are listed in order of smallest return to largest return and each company's market capitalization is depicted as a circle proportional to its market capitalization. The underperformance of a handful of very large stocks such as Microsoft, Apple, and Alphabet pulled the market-cap-weighted average return down below the equal-weighted average return. The return for the actual S&P 500 index that month was 8.01%, whereas an equal-weighted version of the same index returned 9.85%.

To relax the assumption of an equal-weighted benchmark, we created simulations of selectivity, but with market capitalization-weighted portfolios and benchmarks using S&P 500 stocks. The first question becomes, what is an appropriate benchmark for the market capitalization-weighted portfolio? Given that half the stocks are winners and half the stocks are losers, how do we allocate these in our benchmark portfolio?

The expression for the return of a market capitalization-weighted benchmark is given by:

$$r_{MC}^{BM} = \sum_{i=1}^N w_i r_i = \sum_{i=1}^N \frac{MC_i}{\sum_{i=1}^N MC_i} r_i, \quad (1)$$

¹⁷Although an optimal selectivity ratio can be found in examples like the one above, its precise location will be quite sensitive to the historical data sample. Paul Samuelson used to say that even the U.S. stock market was just one sample of history. He used to suggest a thought experiment: Write down stock returns for every minute or every hour or every day for the past 200 years on pieces of paper, put them into a big bowl or hat and mix thoroughly. 'Now let's draw out a new history of 200 years', he continued. 'And let's do that so we have 20,000 different histories'. He noted that in most of these alternate histories, the probability increases that you'll be ahead, the longer you invest. 'But with the same remorseless certainty', he explained, 'there will be some long periods when you will lose, and the longer the horizon the greater the amount of loss. But of course, you don't see these alternative runs of history' Updegrave (2009). It might also be of interest to build a mathematical model that translates the empirical ability of a factor into the setting of selectivity theory as a probability of selecting a good stock.

¹⁸This idea is discussed later in the paper.



FIGURE 6 Cross-Section of individual S&P 500 stock returns for January 2019. This figure shows the returns of individual stocks in the S&P 500 for the month of January 2019. The stocks are ordered from lowest return (1) to highest return (500). The companies are also represented by a circle which is proportional to their relative market capitalization as of 31 December 2018

where MC_i is the market capitalization of company i and r_i is the return of stock i . As we wish to keep the binary nature of the returns the same as the base case, we will have 50% of the stocks with a 10% return (or any other return for that matter) and 50% of the stocks with a -10% return. Thus, we can write the return of the market capitalization benchmark as

$$\begin{aligned} r_{MC}^{BM} &= \sum_{k=1}^{n_g} \frac{MC_k}{\sum_{i=1}^N MC_i} r_g + \sum_{j=1}^{n_b} \frac{MC_j}{\sum_{i=1}^N MC_i} r_b \\ &= \frac{1}{MC} r_g \left(\sum_{k=1}^{n_g} MC_k - \sum_{j=1}^{n_b} MC_j \right), \end{aligned} \quad (2)$$

where MC is the total market capitalization of all the stocks. Thus, the return of the benchmark in this one-period setting will be the difference in the total market capitalization of the good stocks minus the total market capitalization of the bad stocks multiplied by some factor. One choice for a benchmark would be one that, similar to our equal-weight benchmark, produces an average return of zero. This would be one with 50% of the market capitalization of companies as good stocks and 50% as bad stocks. There are many ways this could be chosen, so long as we meet this criterion. We could also choose another benchmark where all of the winner stocks are the largest companies in the benchmark. This would have an average portfolio return that is positive. We could also choose a benchmark where all of the loser stocks are the largest companies in the benchmark. There is no correct benchmark. Ultimately, we must choose a market-cap benchmark where the good and bad returns are distributed amongst stocks in the benchmark.

For our analysis, we created five benchmarks. The zero return benchmark we call ‘neutral size’, the benchmark where the largest 250 stocks are winners, we call ‘strong large-cap tilt’, the benchmark where the largest 250 stocks are losers, we call ‘strong small-cap tilt’, and then we

created two portfolios that are in-between ‘strong large-cap tilt’ and ‘neutral size’ and ‘strong small-cap’ and ‘neutral size’.¹⁹ These are ‘Large-cap tilt’ and ‘Small-cap tilt’. We believe this gives us a wide variety of market-cap benchmarks from winners extremely tilted to big stocks and to winners extremely tilted towards smaller companies.

Once the benchmark was selected, the procedure for generating portfolios was the same. First, we drew from the Wallenius to determine whether we had a winner stock or a loser stock. Then we randomly assigned that winner or loser return to a company, then when we finished drawing the number of stocks for a particular selectivity ratio, we computed the market capitalization-weighted and equal-weighted return of every portfolio. We then repeated this procedure for the M simulations for every selectivity ratio. The results of the simulations are contained in Table 7.

The table shows that if the benchmark is market capitalization-weighted and the best-performing stocks are large-cap stocks, then optimal selectivity level will be far from what the theory predicts and the IR of the portfolio manager will be poor relative to the benchmark *for equal-weighted portfolios*, which is to be expected. This is reversed if the best-performing stocks are small-cap stocks. This is quite natural due to the fact that equal-weighting places less emphasis on the returns of stocks with higher market capitalizations. However, if the portfolio manager, with skill level ω , weighs his or her portfolio using a market capitalization-weighting, then the results are closer to the predictions of selectivity theory. That is, whether the best-performing stocks are large-cap (‘strong large-cap tilt’) or small-cap (‘strong small-cap tilt’), the average optimal selectivity ratio is between 70% and 83%, and all of the IRs are positive.

Similar to our discussion of the effect of return volatility on the level of IRs across all selectivity levels we can also derive the estimate of the position weights variability effect on the expected IR. Let us assume that over longer periods of time the returns of stocks are independent of their market capitalization-based weights. If this is true, we can calculate the variance of a sequence of independent random variables and estimate the effect on the IR. In Table 7, the independence between returns and position weights would result in the outcomes that reside in the middle column (‘neutral size’). The weights as of 31 December 2018 in the S&P500 index were used for the simulations in Table 7.²⁰

In summary, the least challenging assumption in the theory of selectivity is the benchmark weighting. However, we have shown that if the benchmark returns are not skewed to either large-cap or small-cap stocks, then the equal-weighted portfolio will have the highest IR.²¹

¹⁹To create a zero return benchmark, one can use optimization tools or other methods, but there will be a large number of solutions to this problem. To construct our benchmark, we simply took a random draw of market cap and assigned the stock a winner return, we then took another draw from the benchmark list without replacement and assigned a loser return, and so on and so forth, until we had a return for all N benchmark stocks. At every draw, we also computed the sum of winner market capitalizations and loser market capitalizations (i.e., $\sum MC_k$ and $\sum MC_j$ through the n th selection of stocks). If the sum of winners was greater than the sum of losers, we assigned the next winner return to a smaller company (i.e., a company in the bottom 250 of companies by market capitalization). Once again, this was randomly picked. This is only one of many ways we could have assigned stocks to create a zero return benchmark.

²⁰It is of interest to note that the coefficient of variation of market capitalization weights on that date equals 1.833. Because the binary returns of $+/-10\%$ were used for winners and losers, the coefficient of variation is 1.0. Then $\sqrt{(1.0^2 + 1.833^2)} = 2.088$ which is very close to the ratio between the Panels (a) and (b) IRs at the optimal selectivity $0.861/0.407 = 2.115$.

²¹If there is a bias to the performance in the benchmark, for example, there are more large-cap winner stocks, then it would make sense for the portfolios to be market capitalization-weighted to achieve the highest IR and vice versa. The optimal selectivity ratios for these alternatively weighted portfolios would vary from 70% to 83%.

TABLE 7 Simulation results for market capitalization-weighted benchmarks

This table reports the results from 100,000 simulations for each selectivity level for the sequential selection method using a market capitalization-weighted index. The parameters for the simulation are $N = 500$, $\omega = 1.1$, $r_g = 10\%$ and $r_b = -10\%$. The portfolios are compared against five types of market capitalization benchmark, one in which all winner stocks are the largest 250 stocks in the index ('strong large-cap tilt'), one in which all of the winner returns are the smallest 250 stocks in the index ('strong small-cap tilt'), one in which the returns are distributed such that the market capitalization-weighted returns are equal to 0 (i.e., $\sum_{k=1}^{n_g} MC_k \approx \sum_{j=1}^{n_b} MC_j$), and one in which winners are slightly tilted towards big stocks ('large-cap tilt') and one in which winner stocks are slightly tilted towards smaller companies ('small-cap tilt'). Two types of portfolios are analyzed with respect to selectivity; an equally weighted portfolio and a market capitalization-weighted portfolio. AR is for average return and IR is for information ratio.

	Type of market-cap benchmark				
	Strong large-cap tilt	Large-cap tilt	Neutral size	Small-cap tilt	Strong small-cap tilt
(a) Equally weighted portfolios					
Optimal selectivity ratio	10.0	10.0	78.8	90.0	90.0
AR at optimal selectivity	0.448	0.449	0.198	0.122	0.121
AR at 79% selectivity	0.196	0.198	0.197	0.196	0.197
IR at optimal selectivity	-5.071	-0.595	0.861	9.246	48.947
IR at 79% selectivity	-30.138	-4.559	0.851	6.260	31.848
(b) Market capitalization-weighted portfolios					
Optimal selectivity ratio	82.6	78.0	78.4	78.6	74.2
AR at optimal selectivity	7.231	1.448	0.203	-1.050	-7.032
AR at 79% selectivity	7.241	1.444	0.201	-1.055	-7.046
IR at optimal selectivity	0.586	0.397	0.407	0.406	0.611
IR at 79% selectivity	0.586	0.404	0.411	0.401	0.620
Benchmark return	7.151	1.250	0.000	-1.250	-7.151

5 | DYNAMIC SELECTIVITY

Until this point, the optimal selectivity discussion has considered a one-period horizon. A portfolio manager chooses a basket of stocks to own from a benchmark and holds that portfolio. It could be that the portfolio manager simply repeats the exercise at the next rebalancing date. However, this is not typically the way portfolio managers behave. A more likely scenario is that a portfolio manager has a method to pick stocks, whether it is fundamental or qualitative or a more automated quantitative process. For the former, as new ideas come in, the manager might purchase a new stock or group of stocks from time to time. For the latter, as new information becomes available, the portfolio manager may process his quantitative stock return model and decide to purchase new stocks. Other factors, such as tax issues and transaction costs, may also cause a manager to repeatedly adjust the holdings in his portfolio rather than buy all stocks at the same time.

In this part of the paper, we study the way in which a more realistic way of adding stocks to one's portfolio behaves in relation to selectivity theory.

5.1 | Portfolio formation methodology

Similar to the simulation process in other parts of this paper, portfolios are simulated by drawing stocks without replacement from a benchmark. For each individual simulation, a portfolio is created by assuming some fixed number of stocks j are added to the portfolio each month. Each individual addition will eventually exit the portfolio based on an investment horizon chosen at random. These investment horizon draws are made independent of all other aspects of the simulation. The investment horizon probability distribution can take on any form, however, for this paper, we created a holding period distribution that satisfied the following criteria:

1. The minimum holding period for any stock purchase was 3 months.
2. The maximum holding period for any stock purchase was 48 months or 4 years.
3. The average holding period for any stock purchase was 12 months.
4. The selection of stock holding period was obtained from a probability distribution that was smooth and declining for longer holding periods.

The probability distribution function that we chose for the holding period was obtained from the following function:

$$f(n) = \frac{1}{48(n-2)^x} \quad \forall n \in 3, 4, \dots, 48, \quad (3)$$

where x is a parameter to be chosen.²² One can see that this functional form leads to a much lower value of $f(n)$ when n is larger versus smaller, reflecting the idea that the manager is more likely to pick stocks with a shorter investment horizon. To make use of this as a probability distribution, we had to normalize the value such that the sum of discrete outcomes equalled one. Thus, the probability of obtaining a particular holding period, n , is given by:

$$\text{Prob}(n) = \frac{f(n)}{\sum_{i=3}^{48} f(n)}. \quad (4)$$

Finally, to establish a mean holding period of 12 months, we needed a value for the parameter x such that the mean of the holding period distribution was twelve (i.e.,

²²We chose this form so that the manager's average holding period was 12 months, the minimum was 3 months, and the maximum was 48 months and that longer holding periods were less common than shorter holding periods. There was no specific need for this particular function, other than we believed it somewhat reflected reality. In unreported results of optimal selectivity, the qualitative results still hold with other distributions. In particular, we considered a uniform distribution, where the average holding period is 24 months and every holding period is equally likely, as well as an 'inverse' distribution to the one used here. By 'inverse', we mean that the longer periods were more likely and the average holding period was closer to 48 months. In those cases, as well as other distributions that we considered, the optimal selectivity points did not change.




FIGURE 7 Holding period distribution. This figure shows the distribution of holding periods for a given stock selection by the portfolio manager. The mean holding period is 12 months, the minimum holding period is 3 months, and the maximum holding period is 48 months

$\sum_{i=3}^{48} f(n) = 12$). The value of x that achieves this desired property is $x = 1.0325$. Figure 7 shows the effective distribution function of the holding periods from which we drew each holding period for each stock selection by the manager. Besides the mean holding period of 12 months, some other characteristics are worth noting. More than 50% of the stock selections will have a holding period of less than 6 months. Only 15% of the stocks will have a holding period of longer than 2 years and only about 6% will have a holding period of more than 3 years. Generally speaking, our holding period distribution is very reflective of a portfolio manager that generally holds stocks less than 1 year. It is reasonable for some quantitative funds to have this distribution of holding periods.²³ The reason for our choice of a minimum holding period of 3 months is that it coincides with quarterly reports from companies and earnings calls, which would be a natural updating point for quantitative models, as well as for fundamental analysis.

Before every stock selection, we randomly drew from the holding period distribution. Then we determined whether the pick was a success or not, based on the Wallenius distribution. As we moved forward in the simulation, when a stock's holding period had been reached, the stock was dropped from the portfolio. In addition to the removal of a stock at the end of its holding period, we also removed stocks in the S&P 500 that were delisted, went bankrupt, were acquired, or had other corporate actions occur. When these events occurred, we attempted to remove the stock in the previous month at the best-estimated selling price of the company at

²³We also performed simulations for longer average holding periods, like 24 months, and the basic conclusions still hold. See Table 9.

that time in history. All stocks in the portfolio at any given time were equally weighted. Our benchmark was the equally weighted S&P 500.²⁴

The sequential and dynamic stock selection was done as follows. In any given month, t , the portfolio contained a subset of stocks from the S&P 500. The portfolio manager chose a new stock from stocks in the S&P 500 that were not already in the portfolio and did not have missing data for the following month. To appropriately assign winner and loser stocks to the portfolio, once the holding period was known of the particular stock that the manager chose, we separated all of the eligible stocks by their *actual return* over the holding period into winners and losers. A winner or loser was defined as a stock whose compounded return over the holding period was above or below that of the benchmark.²⁵ It was at this point that we used the Wallenius distribution and allowed the manager to pick a stock with the parameter of $\omega = 1.1$. The process was then repeated, where the portfolio manager drew a new holding period, then drew a new stock which could be a winner or a loser.

As already mentioned, our sample S&P 500 data covered the period from 31 December 1988 to 31 December 2018. However, for the dynamic simulation analysis, we used the period of data between 31 December 1992 and 31 December 2014. The first 4 years were used to build up the portfolio to a steady-state size and the last 4 years were excluded due to a lack of forward return information.²⁶ When the portfolio is in its steady-state, the portfolio has roughly n_t stocks in the portfolio, which divided by the number of stocks in the benchmark, N , equals the selectivity ratio. In a steady-state, this will be equal to the average holding period of the stocks multiplied by the number of stocks that the portfolio manager adds per month divided by the number of stocks in the benchmark (Little, 1961). That is,

$$\bar{S}_t = \frac{\bar{n}_t}{N} = \frac{j\bar{h}}{N}, \quad (5)$$

where $\bar{S}_t$ is the average selectivity ratio, $\bar{n}$ is the average holdings of the portfolio, j is the number of stocks that are added to the portfolio by the portfolio manager each month, and $\bar{h}$ is the average holding period of each stock in the portfolio, which in our case, equalled 12.

Thus, to modify our simulation for different selectivity levels, we could either change the number of stocks the manager picked per month, j , or the mean holding period of the manager, $\bar{h}$. We did this by altering the number of stocks that the manager chose per month. To compute the IR for the portfolio manager, we computed the monthly portfolio return based on the actual return of the stocks in the portfolio and the benchmark return.²⁷ The portfolio was rebalanced every month to be equally weighted amongst all the stocks.

²⁴The equally weighted S&P 500 index was actually created by one of the co-authors in 2003 and now trades as an exchange-traded fund (ETF) with the ticker symbol, RSP. For more information, see <http://ludwigbc.com/pubs/SPEWIWhitePaper010703.pdf>.

²⁵For example, if the portfolio had only one stock, there would be 499 stocks available to trade. Suppose the randomly drawn investment horizon was 6 months, then we would compute the forward-looking total returns for every one of the 499 stocks and call those that beat the benchmark, winners, and those that did not, losers. If there was any missing return data for these stocks in any month of the holding period, we assigned the return as 0%.

²⁶To get to a steady-state, we found that the number of months that must elapse is equal to the maximum holding period of any stock.

²⁷The benchmark was the equally weighted S&P 500 rebalanced monthly.




FIGURE 8 Number of Stocks in a representative simulation path for dynamic selectivity. This figure shows the number of stocks in the portfolio from the beginning of the period (January 1988) until the final period (December 2018) for one particular simulation. The parameters for the simulation were $N = 500$, $j = 25$, $\bar{h} = 12$ and the returns are the actual returns of S&P 500 stocks for each holding period, where the holding period for any given stock is drawn from the holding period distribution

5.2 | Dynamic simulation results

Before jumping to the results of the simulations, we first depict how a sample portfolio looks over our sample period. Figure 8 shows one of our simulations, where $j = 25$, $\bar{h} = 12$, the steady-state average number of stocks is 300, and hence the average selectivity ratio is 60%.²⁸

The figure also shows the 4 years that we exclude from the beginning of the analysis and at the end of the analysis. From January 1992 to December 2014, the portfolio fluctuates around the mean of 300 stocks and/or a 60% selectivity ratio.

Figure 9 plots the average IRs for each selectivity ratio along with the polynomial function fit through them. For each selectivity ratio, we conducted 10,000 simulations and each simulation resulted in a different IR. Figure 10 depicts the distribution of IRs for each selectivity ratio. The simulated maximum is technically achieved at 33 new additions per month, which, using Equation (5), corresponds to a selectivity ratio of 79.2%, just as the theoretical, one-period Wallenius distribution would recommend. Table 8 displays IR metrics for the 79.2% case as well as other selectivity ratios. We show the average, median, maximum and minimum IR across simulations along with the polynomial fit of the mean IRs across selectivity ratios.

The impact that different skill levels and holding period assumptions have on the time series simulation framework is presented in Table 9. Each column summarizes the results at

²⁸The dynamic simulation is computationally expensive. We used the R programming language ‘parallel’ package to construct our 10,000 simulations for each selectivity ratio. We reduced the number of simulations from 100,000 to 10,000 due to the computational time required to complete the simulations.




FIGURE 9 Selectivity and average information ratio (IR) for dynamic selectivity. This figure shows the simulations of each selectivity ratio and its corresponding IR for the dynamic simulations. The dots represent the average IR from 10,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages. For these parameters, the optimal selectivity ratio is at 80%. The parameters for the simulations are $N = 500$, $\omega = 1.1$, $j = 125$, $\bar{h} = 12$ and the returns and stocks are chosen from the actual returns and stocks of the S&P 500

the optimal selectivity ratio for a separate round of simulations across selectivity ratios. The left three columns increase the skill parameter while holding period assumptions remain unchanged, whereas the right three columns increase the average holding period while leaving the skill parameter unchanged. Panel (a) summarizes results for the sequential selection approach, whereas Panel (b) shows the results for a bulk-selection approach at each month in the simulation.²⁹

In all cases, the simulated distributions fall within the theoretical range of 50% to 79%. At 1-month holding periods, the optimal ratios reach those endpoints exactly, as these cases are essentially monthly replications of the one-period simulation tests conducted earlier. For both the sequential and bulk selection frameworks, as skill increases and as the average holding period decreases the level of IR achieved at the optimal selectivity ratio increases. In all sequential selection scenarios, the optimal selectivity ratio remains fixed at the theoretical endpoint, just under 80%.³⁰ For bulk selection cases, the optimal selectivity ratios are a few percentage points above the theoretical value of 50%. This shift towards the Wallenius side of

²⁹To conduct the bulk selection simulations, the process was changed such that approximately all monthly j picks were made simultaneously and for the same average holding period, $\bar{h}$, using the computational logic presented in Appendix A.

³⁰For the case of the 24-month average holding periods the optimal selectivity ratio of 76.8% is actually the largest ratio below 80%, there are just fewer selectivity ratio points to choose from as the number of new adds each month shrinks to accommodate the longer holding periods.




FIGURE 10 Distribution of information ratios (IRs) at each selectivity ratio for the simulations. This figure shows the distribution of IRs at each selectivity ratio for all of the 10,000 simulations. The parameters for the simulations are $N = 500$, $\omega = 1.1$, $j = 125$, $\hat{h} = 12$ and the returns and stocks are chosen from the actual returns and stocks of the S&P 500

the spectrum reveals that there are some course-correcting properties to resampling from the same benchmark over time. Even though each monthly sampling exercise has no memory between draws, once a given winner has been selected into the portfolio, it cannot be bought again. Thus, if a manager gets excessively lucky for a few months in a row then their future binomial draw processes will be excessively tilted towards picking losers, similar to how the Wallenius distribution recalculates selection probabilities after each intramonth selection. In practice, the actual way stocks are selected into the portfolio will fall somewhere in between the

TABLE 8 Portfolio characteristics of dynamic portfolio selection

This table reports the results from 10,000 simulations for a different number of stocks acquired by the portfolio manager each month. The parameters for the simulation are $N = 500$, $\omega = 1.1$, $\bar{h} = 12$ and stock returns are obtained from the Z -scores of actual S&P 500 returns. IR is the information ratio. Average is the average of cross-sectional mean IRs from the simulations, Polynomial Fit represents the values from a polynomial fit across average IRs, median, maximum and minimum represent the respective values of the IRs across simulations.

Number of monthly stock purchases	17	21	25	29	33	37
Average selectivity ratio (%)	40.8	50.4	60.0	69.6	79.2	88.8
IRs						
Average	0.465	0.516	0.553	0.593	0.606	0.586
Polynomial fit across mean IRs	0.471	0.520	0.561	0.589	0.599	0.583
Median	0.463	0.514	0.554	0.595	0.606	0.585
Maximum	1.296	1.176	1.144	1.218	1.195	1.226
Minimum	-0.293	-0.199	-0.174	-0.021	-0.111	0.002

TABLE 9 Portfolio characteristics of dynamic selectivity with altered parameters

This table reports the results from 10,000 simulations for a different number of stocks acquired by the portfolio manager each month. The parameters for the simulation are $N = 500$ and the other parameters are listed in the table, and individual stock returns are obtained from individual stock returns of actual companies in the S&P 500. IR is the information ratio.

	Skill parameters			Average holding period (j)		
Skill (ω)	1.05	1.10	1.20	1.10	1.10	1.10
Average Holding Period (j)	12	12	12	1	12	24
Number simulations	10,000	10,000	10,000	10,000	10,000	10,000
(a) Sequential selection						
Optimal selectivity ratio (%)	79.2	79.2	79.2	79.2	79.2	76.8
IR at optimal selectivity	0.353	0.599	1.050	1.973	0.599	0.444
(b) Bulk selection						
Optimal selectivity ratio (%)	55.2	55.2	55.2	50.0	55.2	57.6
IR at optimal selectivity	0.273	0.482	0.870	1.237	0.482	0.383

bulk and sequential selection, which will result in optimal selectivity somewhere between 50% and 80%. However, as Table 9 clearly demonstrates, shifting the portfolio selection process more towards the sequential selection method increases both the IR and optimal selectivity.

In the process of confirming the general results of Selectivity Theory by relaxing the initial assumptions, we discovered a couple of interesting relationships that we think will further enhance our research on this topic. When simulating a manager's performance with real stock return data, a given selectivity ratio, and a certain skill level, we find that there seems to be a

consistent relationship between monthly outperformance, the manager's average holding period, and the manager's ability. Although we are still trying to understand the results, the implications would be very profound. It would imply that we could estimate the expected IR of a manager's portfolio if we knew the manager's omega (ω) and that we might be able to estimate an implied omega (ω) of a manager if we have the manager's portfolio holdings through time.

6 | CONCLUSION

The work in this paper attempts to further test and examine assumptions about optimal selectivity in an enhanced portfolio (Bolshakov & Chincarini, 2020). The specific purpose of this paper was to relax assumptions about the investment environment using simulations to investigate the impact of the assumptions. When we relax the assumptions in the original theoretical paper, we find the general implications of the theory still hold. However, when there are certain limits to skill, the optimal selectivity differs from the theory as would be expected. In other words, we find that even with the relaxed assumptions, it makes sense for an enhanced index manager to hold somewhere between 50% and 80% of the benchmark provided the manager has skill. If the manager has no skill, then the best course of action, is obviously, to fully index. If the manager has varying degrees of declining skill, that is, a limited ability to pick stocks, then the optimal point will be achieved at a lower selectivity ratio.

We also examined the impact of dynamic portfolio management on the optimal selectivity ratio. That is, rather than a portfolio manager choosing a group of stocks in a one horizon setting, we examine a portfolio manager that selects new stocks every rebalancing period and selects stocks with different holding periods. We find that as long as the portfolio manager has skill the optimal selectivity is somewhere between 50% and 80%. This naturally points to new and innovative directions for future research.

First, it might be possible to use these techniques to identify skilful managers from a data set of professional manager returns and selectivity ratios. We are already exploring a method to infer a portfolio manager's ω from their historical performance. Early evidence indicates an empirical relationship between a portfolio manager's performance and the implied ω . We have made progress on understanding it and we believe a practical connection might be made, but we are still in the process of thoroughly understanding the mapping.

Second, it might be enlightening to examine actual historical return data in combination with some objective skill-like criterion for selecting stocks to determine whether optimal selectivity ratios are confirmed with very practical implementation methods.

Third, some of these ideas may prove useful to investment consultants to examine the consistency between a manager's expected ω -based IR and his actual realized IR. For example, if ω implies a higher IR than is actually realized, one potential explanation would be that the manager experienced a draw of bad luck on the return distribution side (i.e., had unusually big losers or, not enough big winners). Another use of these ideas would be for a CIO to invest in several managers with high estimated omega, rather than one individual manager.³¹ This last point is quite important.

³¹All of these ideas are fruitful areas that we are currently exploring. There is a host of evidence that institutional or professional managers have skill in selecting stocks (see Amihud and Goyenko (2009), Pan et al. (2019), and Dong and Doukas (2020)).

A curious reader may ask how one can reconcile this theory with the empirical fact that many active managers hold a much smaller percentage of stocks than 50% to 80% of the benchmark. One reason could be that they have ability saturation or jack-knife skill, which we have discussed in this paper. With ability saturation, the manager with a realistic level of ω should still choose a very high number of stocks, however, a manager with jack-knife skill, will hold a much smaller percentage of the benchmark depending on what percent of the universe he believes he has skill for. Thus, if a typical manager only has skill up to 10% of the benchmark, then the optimal selectivity ratio would naturally be lower. However, as mentioned above, one can imagine a CIO that is not limited as an individual manager is. That is, a CIO could conceivably choose the holdings of many talented portfolio managers to keep $\omega > 1$ for a larger percentage of the universe. Hence, this CIO could combine the portfolio manager selections so as to create a portfolio near the optimality of 50% to 80%. If the theory is right, this combined portfolio would be a better performing portfolio than any of the individual portfolios.

In addition to the promising application for investing, we also believe that these ideas provide an alternative to potentially harming effects of total index-based management or crowding in passive asset management. As of the writing of this paper, 50% of total U.S. equity assets are controlled by index investing organizations (DiBenedetto & Lauricella, 2019). The frustration with poorly performing portfolio managers has led to this dramatic shift to passive or index-based investing, but perhaps there is room for a more optimal investing style that selects a large portion of the benchmark, yet still cares deeply about performance and high IRs.

Clearly, this interesting new theory needs to be examined and scrutinized more thoroughly. However, it is comforting that as a first step, these simulations confirm some of the original findings of the theory when the assumptions are relaxed. Another very interesting aspect of selectivity theory is the conclusion that even an active manager with superior skill might want to invest in the majority of stocks in the underlying index.

REFERENCES

- Amihud, Y., & Goyenko, R. (2009). Mutual fund's R2 as predictor of performance. *Review of Financial Studies*, 26, 667–694.
- Berk, J., & van Binsbergen, J. (2015). Measuring skill in the mutual fund industry. *Journal of Financial Economics*, 118, 1–20.
- Bessembinder, H. (2018). Do stocks outperform treasury bills? *Journal of Financial Economics*, 129, 440–457.
- Bolshakov, A., & Chincarini, L. B. (2020). Manager skill and portfolio size with respect to a benchmark. *European Financial Management* 26, 176–197.
- Chincarini, L. B. (2012). *The Crisis of Crowding. Quant Copycats, Ugly Models, and the New Crash Normal*, Wiley/Bloomberg.
- Chincarini, L. B., & Kim, D. (2006). *Quantitative equity portfolio management. An active approach to portfolio construction and management*. McGraw-Hill.
- Cremers, M., & Pareek, A. (2016). Patient capital outperformance: The investment skill of high active share managers who trade infrequently. *Journal of Financial Economics*, 122, 288–306.
- DiBenedetto, G. & Lauricella, T. (2019). A look at the road to asset parity between passive and active U.S. funds. A breakdown of passive and active funds' asset parity in 5 charts. Morningstar Blog. <https://www.morningstar.com/insights/2019/06/12/asset-parity>
- Dong, F., & Doukas, J. (2020). When is fund management skill more valuable? *European Financial Management*, 26, 455–502.
- Fama, E., & French, K. (2010). Luck versus skill in the cross-section of mutual fund returns. *The Journal of Finance*, 65, 1915–1947.
- Little, J. (1961). A proof for the queuing formula: $L = \lambda W$. *Operations Research*, 9, 383–387.

- Pan, X., Wang, K., & Zykaj, B. B. (2019). Does institutional ownership predict mutual fund performance? An examination of undiscovered holdings within 13(f) reports. *European Financial Management*, 25, 1249–1285.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19, 425–442.
- Updegrave, W. (2009). A nobel laureate's priceless stock tip. *Money*.

How to cite this article: Bolshakov, A., Chincarini, L. B., & Lewis, C. (2021). Enhanced indexing and selectivity theory. *European Financial Management*, 1–35.
<https://doi.org/10.1111/eufm.12330>

APPENDIX A: FISHER SIMULATION METHODOLOGY

In terms of simulations, the Fisher distribution, that is, where managers pick in bulk, is much paradoxically more complicated than the Wallenius distribution, that is, where managers pick sequentially. In this appendix, we describe the steps that we implemented to simulate the sequential selection method, as well as present the results for the relaxation of assumptions of Bolshakov and Chincarini (2020) that were presented for the Wallenius distribution in the main body of this paper.

A.1 | Description of simulation steps

Step 1: Construct-appropriate benchmark

Similar to the Wallenius Monte Carlo simulations, a theoretical benchmark is constructed from which the manager will randomly draw stocks. The benchmark can be as simple as the original base case, where 250 winning stocks return +10%, 250 losing stocks return −10%, and all constituents are equal-weighted. Or alternatively, the benchmark can be tailored to different winner/loser populations, more complex return distributions, or market capitalization-weighted schematics.

Step 2: Calculate probabilities for all selectivity ratios

One of the primary distinctions between the Fisher and Wallenius distributions is that in the former the number of stocks selected (and therefore the selectivity ratio) is not known in advance; the total number of stocks held is itself a random outcome from independent binomial processes. If one attempts to plug in the base case WNCH probabilities for $\omega = 1.1$ to these binomial processes (i.e., 0.523 for winners and 0.476 for losers), then the resulting FNCH simulated portfolios will be tightly clustered around a 50% selectivity ratio. It would take a massive number of simulation trials to generate simulated observations of portfolio sizes far away from 50% selectivity.

However, just as a quantitative manager can tailor their model to be more or less strict in its rules, so can the binomial probabilities assigned to the winner and loser populations be shifted to focus the FNCH simulations around a desired selectivity ratio. Because the binomial distributions are independent, the expected number of stocks selected from each subpopulation

will be additive (see Equation (8) below). In the algebra that follows, X_1 and X_2 are the numbers of winners and losers randomly drawn, m_1 and m_2 represent the number of winners and losers in the benchmark, and p_1 and p_2 signify the binomial probability of drawing a winner or drawing a loser, respectively, from those subpopulations. Thus,

$$E[X_1] + E[X_2] = m_1 p_1 + m_2 p_2. \quad (6)$$

If one is targeting a given selectivity ratio S , then the binomial expected value in Equation (6) can be rewritten as:

$$S = \frac{m_1 p_1 + m_2 p_2}{m_1 + m_2}. \quad (7)$$

For an assigned skill level ω there is a direct relationship between p_1 and p_2 . As each of the two subpopulations follows a binomial process of either (a) select a winner versus select nothing, from the subpopulation of winners; or (b) select a loser versus select nothing, from the subpopulation of losers, the skill parameter can be expressed as the following ratio:

$$\omega = \frac{\frac{p_1}{1-p_1}}{\frac{p_2}{1-p_2}}. \quad (8)$$

The ratio in Equation (8) can be reorganized in terms of p_2 :

$$p_2 = \frac{p_1}{\omega(1-p_1) + p_1}. \quad (9)$$

At this point, p_2 can be substituted into Equation (7). We are left with the following expression.

$$S = \frac{m_1 p_1 + \frac{m_2 p_1}{\omega(1-p_1) + p_1}}{(m_1 + m_2)}. \quad (10)$$

As we also choose the selectivity level, S , we can solve for p_1 using standard computer optimization solvers.³²

With p_1 in hand, one then uses Equation (9) to find p_2 , which will correspond with the chosen ω and the chosen S .

At this point, a complete Fisher distribution can be simulated by looping through all possible S values to create a series of p_1 and p_2 probability pairs to use, one pairing for each distinct selectivity ratio.

³²There is an analytical solution for p_1 , but it's very cumbersome and involves many terms.

Step 3: Conduct independent binomial draws in each subpopulation

With the probability pairs computed, random binomial draws are drawn independently on the winner and loser subpopulations to determine how many winners and losers are selected into the portfolio. The expected selectivity ratio will remain true to the targeted levels found in Step 2, but each sampled portfolio size may deviate from that size. Nevertheless, all simulations are grouped according to their targeted selectivity ratio rather than their realized selectivity ratio.

Within each subpopulation, it is assumed that each stock is equally likely to be selected. For example, if the binomial draws indicate 25 winners and 20 losers are to be chosen, then 25 winners are randomly drawn from the winner subpopulation assuming a uniform random distribution, and the losers are randomly drawn from their subpopulation in the same way. These selections are made without replacement, which allows for the real-world complexities to be introduced, as each stock within the benchmark can be tagged with its unique market capitalization or return characteristics as desired.

Step 4: Calculate portfolio returns

Lastly, portfolio returns for a given simulation are calculated based on the stocks identified from the benchmark and their pre-assigned returns. Holdings are assumed to be equally weighted unless noted otherwise. Steps 1–4 are then repeated within each targeted selectivity ratio for all simulations. Following completion of the simulations, cross-sectional performance metrics are calculated within each targeted selectivity ratio, and then a polynomial function is fit across those cross-sectional average results.

A.2 | Base case Monte Carlo simulations

The same assumptions used in the Wallenius (WNCH) base case presented in the main paper are applied to the above Fisher simulation structure, with 100,000 simulations conducted at every possible targeted selectivity ratio between 10% and 90%. The polynomial-fitted line from these simulations is then compared to the theoretical FNCH expected returns and information ratios (IRs) as discussed in Bolshakov and Chincarini (2020). The correlation between the fitted line and actual simulations is 0.9985 when comparing IRs across all selectivity ratios. Figure 11 plots the resulting IRs from these simulations, alongside the theoretical IRs.

A.3 | Relax assumption 1: More losers than winners

Figure 12 shows the selectivity ratio and IR in the case of 40% winners alongside the polynomial function fit through simulations of the same population. With fewer opportunities for their skill to shine, the manager's optimized IR declines; however, their optimal behavior does not. Similar to the findings for WNCH simulations, the optimal selectivity ratio does not deviate materially from the base case (which for Fisher is 50%).

A.4 | Relax assumption 2: Empirical distribution of returns

Substituting the empirical distribution of returns for the binary $\pm 10\%$ returns does not alter the optimal selectivity ratio either. Using the S&P 500 Z-scores described in the WNCH simulations




FIGURE 11 Baseline simulations: Fisher Noncentral Hypergeometric Distribution (FNCH) selection method. This figure shows the baseline simulations of each selectivity ratio and its corresponding information ratio (IR) with the FNCH distribution method described in Appendix A. The dots represent the average IR from 100,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages and the theoretical Fischer IRs for each selectivity ratio. The optimal selectivity ratio is at 50%. The parameters for the simulations are $N = 500$, $\omega = 1.1$, $r_g = 0.10$, $r_b = -0.10$ and $r_{BM} = 0$




FIGURE 12 The Fisher Noncentral Hypergeometric Distribution (FNCH) selection method: when the manager only has the ability to find winners in 40% of the universe. This figure shows the results from 100,000 simulations for each selectivity level for the bulk selection method. The parameters for the simulation are $N = 500$, $\omega = 1.1$, $r_g = 10\%$ and $r_b = -10\%$

produces a similar parabola as seen in the base case Fisher distribution, although the IR for each selectivity ratio shifts down because of the greater variability within the benchmark's returns (see Figure 13).

A.5 | Relax assumption 3: Decaying manager skill

Most of the skill decay scenarios explored for the Wallenius distribution cannot be tested under Fisher conditions as the draws must occur simultaneously to create the Fisher distribution. The only case that can realistically be simulated is to vary the skill across portfolio simulations, mimicking a manager facing different regimes that periodically favour or punish their long-term skill. Although this addition of extra variability would lower the manager's optimized IR, the manager would continue making their construction decisions the same way as in the base case, using average omega ($\bar{\omega}$) instead of ω as their assumed skill parameter. As long as $\bar{\omega} > 1$, that construction decision will remain optimized at a selectivity ratio of 50%.




FIGURE 13 Selectivity ratio and information ratio (IR) for empirical S&P 500 returns for the Fisher Noncentral Hypergeometric Distribution (FNCH) selection method. This figure shows each selectivity ratio and its corresponding IR when returns are not binary for good and bad stocks but are based on the cross-sectional Z-scores of actual S&P 500 individual stock returns. The dots represent the average IR from 100,000 simulated portfolios at that particular selectivity ratio. Also shown is a fitted line across the simulation averages and the theoretical Fischer IRs for each selectivity ratio. The parameters for the simulations are $N = 500$, $\omega = 1.1$, $r_i = \bar{z}_i$ and $r_{BM} = 0$

A.6 | Relax assumption 4: Market Capitalization Weighted returns

Although we do not show the results, building a market-cap-weighted benchmark also leads to similar conclusions as those found with the WNCH distribution. The standard deviation of returns for the manager's portfolio expands, which lowers their potential IR. If larger capitalization stocks tend to outperform, the manager will have an ex-post desire to have weighted their portfolio according to market caps; conversely, if smaller capitalization stocks outperform, the manager will wish they had equal-weighted their holdings.

On an ex-ante basis though, across the different regimes that can materialize, the manager maximizes their IR by weighting their holdings proportional to their market caps and targeting a selectivity ratio slightly below 50%.