

Max Darnell
Partner
Chief Investment Officer



Treading Where Alpha Does *Not* Sum to Zero

Most investors assume that alpha can only be earned by taking it away from someone else. This leads investment managers to ignore sources of alpha that are not zero-sum, thereby missing an abundant source of alpha. We offer an explanation for why this mindset is so prevalent, and give examples of non-zero-sum game transactions that lead firms like ours to deliver alpha streams that are truly different.

What we believe has a bearing upon what we do. This is no less true for investment managers, which is why all introductory meetings between investment managers and prospective clients begin with the customary statement of investment philosophy. Specifically, whether or not a manager believes that the pursuit of alpha is a zero-sum game or not will influence the decision about what alpha he pursues and how he pursues it.

Whether or not a manager believes that the pursuit of alpha is a zero-sum game or not will influence the decision about what alpha he pursues and how he pursues it.

What do investors generally believe about the pool of extractable alpha? Most readily accept the notion that any alpha captured corresponds to alpha that

someone else has lost. In other words, most believe that alpha capture is, indeed, zero-sum. Some, on the other hand, take the view that for all *practical* purposes, it is a zero-sum game, but admit that there is legitimacy, at a *conceptual* level, in the case that can be made opposing the zero-sum assumption.² These investors would, therefore, presumably behave as if alpha were zero-sum.

It is, we believe, the distinct minority who hold as a core belief the notion that alpha does not always sum to zero, that opportunities to create net positive alpha across investors exist, and who actively

seek out such alpha. The most conventional sources of alpha do, we believe, tend to sum to zero, so looking for “constructive” alpha opportunities – those transactions that lead to net positive aggregate alpha – typically means looking for alpha in places other than where most active managers tread. As a result, we should expect the investment results of those who seek constructive alpha opportunities to look different from those who don’t.

Pareto Optimality

While those with an investment mindset tend to assume away the potential for net economic gain in their own activities, finance’s closest cousin, economics, is centrally concerned with the discovery of systems, rules, and interactions that lead to, or at least encourage, net economic gain. From the economist’s perspective, the principle value that markets hold lies in putting economic agents together to trade or transact so that they may become better off in aggregate.

At the inception of the twentieth century, an Italian economist, by the name of Vilfredo Pareto introduced the concept of “Pareto Optimality.”³ Pareto Optimality is the end goal of a process of economic

¹ We’re grateful to Joanne Hill for having motivated this paper with her own, and for taking the time away from “enjoying” this weekend’s blizzard in NY to read and offer comments on this paper. Should we have succeeded in advancing the discussion about this topic by offering either additional clarity or insights, then we’ve done so on the shoulders of her earlier paper.

² See “Alpha as a net zero-sum game: How serious a constraint”, Joanne Hill, Goldman Sachs Quantitative Insights, September 20, 2005.

³ *Manual of Political Economy* (1906)

interaction between individual participants who are differentiated by their prefer-

By taking differential preferences into account, we break away from the classic definition of alpha, and allow the question about zero-sum to become meaningful.

ences, and it is defined as that state of the world where no individual can be made better off without making at least someone worse off. Such a state of the world is what is meant by a zero-sum game. To Pareto, and to economists in general, what is interesting is the potential for Pareto Efficient exchanges that occur along the path towards optimality – exchanges that lead to a net benefit rather than just a zero-sum outcome.

The fact that individuals can hold different “preferences” from each other was important to Pareto. He objected to the Utilitarian school that loosely assumed that “utility” (a measure of individual satisfaction) was somewhat universal, that what was good for any individual would be generally good for all. Pareto assumed that individuals were motivated by their preferences rather than what was necessarily “best” for them, and that this differentiated individuals.

This break with the Utilitarian viewpoint should sound vaguely familiar in form. Many investors today too casually assume universal “utility,” assuming that what looks like a good investment for one type of investor is necessarily good for another. Even investors with a high degree of similarity (e.g., similar liabilities, and similar funding levels) may have nevertheless developed different preferences. The Capital Asset Pricing

Model, from which the original definition of alpha derives, defines an alpha that is universal, and that by implication yields the same utility to all investors. All theoretical models stand upon simplifying assumptions, e.g., all investors are perfectly rational, fully and equally informed, and share identical utility functions. Where we get into trouble is when we ask important questions without reminding ourselves about those simplifying assumptions, and insuring that the answer to the question will still be meaningful if we allow those simplifying assumptions.

In this case, the simplifying assumption that investors are homogeneous in their preferences cannot be allowed if the answer to the question is to have any importance – if it is to have any real meaning to us. The answer to the question is meant to guide us in our decision about whether to seek alpha and if so, how to do so. By taking differential preferences into account, we break away from the classic definition of alpha, and allow ourselves to ask a meaningful question about whether or not investors can trade to become better off – *according to their definition of “better off”* – without having to make someone else worse off. Should the definition of alpha fail to correspond to investor gains as investors perceive and value them, then the proclamation that alpha sums to zero becomes one of those empty observations about an investment return that is too far abstracted from the one we care about to guide our actions. It’s worth noting that Pareto’s insistence on differential preferences set the stage for the entire field of microeconomics. Important questions depended upon this change in world view.

Artificial Homogeneity

How is it that most investors so commonly accept (consciously or not) the assumption of homogeneity and from there draw the conclusion that the pursuit of

alpha is a zero-sum game? One answer resides with the proliferation of investment constraints. Constraints, by definition, limit the range of activity and outcomes, and by doing so they, in effect, homogenize our investment objectives and our investment approaches. Differentiation is critical to the existence of opportunities for constructive alpha transactions.

Constraints limit the range of activity and outcomes, and by doing so they, in effect, homogenize our investment objectives and our investment approaches.

Take, for example, the world of enhanced large cap US equities. These mandates are defined by their highly constrained approach, particularly their tight relationship to the benchmark. This occurs not only in terms of the tracking error (active risk) being very low, but also in terms of the requirement that industry, sector and style weights are either forced to mirror the benchmark, or are at least held very closely to their benchmark weights.

Investors managing such portfolios are boxed in so tightly by these constraints that homogeneity of purpose and approach are essentially forced upon those investors who enter this narrow space. It is as though they exist in an artificial laboratory environment, where the conditions are tightly controlled so as to limit the number of factors that will influence the outcome. As a result, there is no room for investors to differentiate themselves and their objectives from each other, thereby leaving no room for investors to turn their differences into their mutual advantage.

Of course, investors are *not* a homogeneous bunch. They come to the markets

with different objectives, different sensitivities to risk, different investment horizons, different liabilities, different utility functions, and different institutional or personal constraints placed upon their activities. Endowments and corporate defined benefit plans, for example, have materially different investment objectives and different appetites for risk. As a result, what may be a “prize” to one investor may not be so to another. This allows us to consider the possibility that investors may, in fact, be able to trade with each other with the result that *both* are better off!

Constructive Alpha Trades

Examples of constructive alpha transactions – even a Pareto Efficient transactions – are easy to find. The forwards market arose to allow farmers to sell their crops in advance of harvest. By paying someone with a longer investment horizon, or someone in a better position to diversify their risks, to take on the risk that crop prices would fall in the future, both the farmer, who enjoys off-loading a risk he had a very low tolerance for, was made better off in the long-run, and the “insurer” who enjoyed the premium he was paid for bearing this risk, was also made better off. To state it most generally, where and when investors are not a homogenous bunch, but instead are characterized by differences in utility functions, preferences, or differences in objectives, then the greater the likelihood that there exist Pareto efficient transactions to be made between investors.

In the example above, not only is there a net positive benefit in aggregate, but also all parties involved may end up better off. Pareto Efficiency involves a more stringent set of conditions than are necessary in that they require that no party loses. Constructive outcomes, where the sum of the benefits simply outweighs the sum of the costs, may leave some investors worse off. What we're interested

in is the more abundant sources of alpha. We need not care whether some investors are left worse off or not. What attracts us is that the sum of the gains outweighs the sum of the losses.

Where do we see other natural occurrences of constructive, possibly even Pareto Efficient trades? To start with, let's consider an example that is especially relevant today given the increased role of hedge funds in the market. Hedge funds classically work with relatively short time horizons as their investors are prepared to quickly take their money back at the first sign of material shortfall. As a result, most hedge funds exhibit a low tolerance for short-term losses and employ various forms of stop-loss mechanisms to limit such losses. These mechanisms cause them to sell assets with depressed prices to investors (such as First Quadrant) with a higher tolerance for short-term volatility and with an active interest in earning a premium, on average, over time, for taking on this risk. They, like us, may think that those assets are more likely to appreciate than not, but they can't take the risk that they are wrong and that the assets fall further before they rise. They, therefore, pay us to assume that risk.

In the example above, we observe some investors avoiding risk *after* damage has been done. In other cases we observe investors seeking to avoid risk *before* there are any signs of trouble. These latter investors classically go to the options market to buy protection, and they pay a premium (on average) to those who are willing to write a put option. There has been, for example, around a 150 basis point average spread between the price-implied volatility on the S&P 500 and its realized volatility. Admittedly, there are times where the premium grows too thin to remain attractive, and there are environments where long volatility may look more attractive than short, but this just means that a dynamic approach to earning this

premium is better than a static approach. This is a principle that applies to any premium.

The credit markets give us another interesting example in the form of credit default swaps. Bonds are priced to compensate investors for the various risks that they carry. Most importantly, they carry interest rate risk, liquidity risk, credit risk, and default risk. What the credit default swap does is to strip out the last of these risks so that investors who may be holding the bond can earn a premium for the first three risks, but pay someone else a premium to insure against the last risk. This is one of the clearest indications that different investors have different appetites for specific types of risk.

The examples above all relate to differences in risk appetite on the part of investors. Returning to the subject of investment constraints, differences in the constraints that are either externally- or self-imposed also lead to opportunities for constructive alpha. While most stock investors will choose to cling tightly to the benchmark in terms of tracking error, sector/industry weights, style weights (e.g., value and growth), others are more willing to take risks that differentiate their returns from, and their exposures to, the benchmark.

In our management of equity strategies at First Quadrant, for example, we find that by crossing style and industry boundaries, we get paid for underweighting expensive segments of the market and overweighting cheap segments, while other investors chose to locate all of their risk in idiosyncratic (stock specific) dimensions. We are, therefore, trading with both passive and active managers who are willing to pay us a premium to buy from us stocks in those segments of the market that have become expensive so that they can avoid the risk that their return deviates materially from the benchmark.

One is tempted to argue that our gain comes at the expense of those who are more constrained, but that's paramount to claiming that their preferences for those constraints is damaging to their returns. Pareto's objection to the Utilitarians was (in part) that people were motivated by their preferences, not by a universal rule on what enhances their "well-being." We too must take seriously the fact that there isn't a universal utility function for investors, which means that we can't ignore the individual preferences expressed by, and assumed benefits derived from, such constraints.

Benchmark agnostic investors should derive the most benefit from engaging in constructive trades with those who are highly constrained. Benchmark agnostic strategies come in both long-only and long-short form, the latter – typically cast as absolute return strategies – having even more to gain from their differentiation from highly constrained long-only benchmark hugging investors. Portfolios that are 130% long, 30% short (or any variant on those magnitudes) free themselves only partially from the typical constraint set, but because they do, they also stand to benefit modestly in a similar fashion.

Finally, global macro investors provide the best example of investors who live with very little in the way of investment constraints, and are, therefore, best positioned to develop constructive alpha trades with more constrained investors. Global macro investors trade across asset classes, countries, currencies, and investment instruments (e.g., cash securities, futures, forwards, options, swaps). As they are typically free to shift their risk taking in a multistrategy context from one strategy to another if they see better opportunities ahead, they are also best positioned to take advantage of *new* opportunities that arise in our ever-changing world.

In Brief

Do zero-sum "games" exist in the competition for alpha? There's no question about it. Where do they exist? They are most likely to be found where either real homogeneity or artificially created homogeneity is found. Constraints play a material role in the creation of artificial homogeneity. The more narrowly the game has been defined, the more investors' objectives and tactics necessarily resemble each other, and therefore the less opportunity there can be for mutually beneficial trades to take place.

The more narrowly the game has been defined, the more investors' objectives and tactics necessarily resemble each other, and therefore the less opportunity there can be for mutually beneficial trades to take place.

It's interesting, isn't it, that the professional, institutional investing world spent so many years boxing in investment decisions more and more narrowly. For an increasing number of investors, alpha capture was, indeed, becoming zero-sum. Even more interesting is the fact that the constraints are being increasingly loosened today, with fewer and fewer investment constraints being applied, allowing for the potential that investment managers be more creative and more opportunistic in finding ways to generate alpha. Opportunities for constructive alpha transactions are on the rise for an increasing number of investors. So long as these managers continue to believe that alpha capture is zero-sum, however, they will fail to capitalize on these opportunities.

Our willingness to be paid to bear risks that other investors are willing to pay us to bear is one of the reasons that our investment results tend to have low correlations with the results that other managers generate.

Because our own beliefs at First Quadrant differ with the typical investor on this point, we do capitalize on these opportunities. Because we have always sought to impose fewer constraints on our investment process, we have had the freedom to do so. This has resulted in our alpha tending to behave differently from other managers' alphas, and our description of what we do tending to sound different. To those of you who know in some detail our investing approach, this may lend further clarity to the fact that for much of what we do, we don't describe what we do as "forecasting," but rather describe facets of our approach as the systematic harvesting of premiums that other investors are willing to pay us to absorb risks that they would like to insure against.

Note that we are neither suggesting that other investors stay away from zero-sum sources of alpha, nor are we implying that we avoid them ourselves we too compete in spaces where the alpha is zero-sum). The area where we think we most differentiate ourselves from the bulk of other investors is in exploiting the non-zero-sum alpha opportunities. Given our willingness to be paid to bear risks that other investors are willing to pay us to bear – in fact, our active pursuit of such premiums – it should come as no surprise that our investment results tend to have low correlations with the results that other managers generate.