the journal of PORTFOLIO management

INVITED EDITORIAL COMMENT

Pulling the Goalie: *Investment Lessons from Watching Hockey*

CLIFFORD ASNESS AND AARON BROWN

Portfolio management is tough. Markets defy prediction. Managers must hew rigorously to timeless principles, while taking immediate notice of changes in market behavior. Risk and return never quite seem to be in balance. When this all gets to be too much, it's time to head out to Madison Square Garden, where the New York Rangers play hockey, and where taking a hit carries more pain than just the psychological kind described by prospect theory.

Like portfolio managers, hockey coaches face many risk-versus-return problems. The most dramatic of these dilemmas in hockey is when to pull the goalie.[1] The coach making this decision must weigh dozens of factors that are specific to the team and the game. At the same time, the coach should be aware of some unchanging math. We do not have the requisite hockey experience to give useful advice about when a goalie should be pulled in any given game. However, in cases where many other considerations may matter,[2] we do have a model that we trust for an estimate of when goalies should be pulled, on average.

The model uses five inputs; two of these inputs reflect the game status:

1. The current goal differential (who's winning and by how much).
2. The time remaining in the game.

And the remaining three inputs entail the probabilities of scoring goals in these situations:

3. You score with your goalie in place.
4. You score with your goalie pulled in exchange for an extra attacker for you.
5. The other team scores, when your goalie was pulled in exchange for an extra attacker for you.

Proceeding in this way creates a simple model with attractive intuition. Using data from the 2015–2016 NHL season, we estimated values of 0.65% per 10-second interval for input 3, 1.97% for input 4, and 4.30% for input 5.[3]

[1] Here's a lightning explanation for readers misguided enough to have ignored hockey. Hockey teams have six players each. Normally one player, designated as goaltender, guards the net to prevent the other team from scoring, and five "skaters" move around the ice trying to put the puck in the other team's goal, and to interfere with the opposing team's skaters' efforts to score. When a team is trailing late in the game, it usually pulls its goalie in order to have six skaters. This increases the team's chance of scoring, but it also increases the chance that it will be scored upon. Despite this imbalance, pulling the goalie is often a very good strategy and this is precisely the point that we are studying.

[2] For instance, our general results could require modification based on the relative skills of the teams playing, or based on specific game conditions (e.g., you'd be far less likely to pull the goalie when there's a face-off in your own end).

[3] Our results are fairly robust to reasonable perturbations of these inputs—but obviously very different estimates will change things. For those so inclined, a more wonky version where we derive the optimal strategy, and present some additional analytics, is available on SSRN.

This immediately shows that pulling the goalie is a negative expected value move in terms of goals (as 1.97% is a lot lower than 4.30%). The team that pulls its goalie multiplies its opponent's chance of scoring by almost seven, while multiplying its own chance of scoring by just a little over three. That certainly seems like a bad idea! But the number of expected goals is actually not the appropriate criterion. What matters here is the expected number of standings points.[4] A team down a goal with short time remaining in the game will gain a lot by scoring and will lose little if the other team scores, because losing by two goals is no worse than losing by one. (Admittedly our model doesn't consider pride). Perhaps the first lesson we can apply outside of hockey is that sometimes what seems to be the right criterion is, in fact, not the correct choice at all, and selecting the right criterion can make an important difference. We will discuss this more when we pivot to investing.

To solve for the optimal strategy, we use backward induction. A tie is worth 1.5 standing points,[5] and a loss is worth zero. If we ignore the tiny probability of multiple goals within a 10-second period, a team trailing by 1 goal with 10 seconds to go has a 0.65% chance of tying the game with their goalie in place; the chance rises to 1.97% if the team pulls its goalie. That means that its expected standing points are $1.5 \times 0.65\% = 0.0098$ if it doesn't pull, and $1.5 \times 1.97\% = 0.0296$ if it does, so it should definitely pull that goalie in this situation. (Note the chance the other team will score doesn't change a thing here, as you still get zero standing points).

With 20 seconds to go, a team that keeps its goalie in place has a 0.65% chance of getting to the 10-second mark with a tied game, a 0.65% chance of getting to that mark down two goals (as that's the chance the other team scores), and a 98.70% (100% − 0.65% − 0.65%) chance of getting there down the same 1 goal (i.e., nothing happens). We assume that when they reach the 10-second mark, they will again act optimally. That gives the team expected standing points of $1.5 \times 0.65\% + 0 \times 0.65\% + 0.0296 \times 98.70\% = 0.0103$. If it pulls the goalie, the formula changes to $1.5 \times 1.97\% + 0 \times 4.30\% + 0.0296 \times 93.73\% = 0.0573$. So, as 0.0573 is greater than 0.0103, it should indeed pull.

Working backwards in this way, we find that a team trailing by one goal should pull with 6 minutes and 10 seconds to go (that's the time when, finally, the expected standings points of pulling is not greater than that obtained by leaving the goalie in the net).[6] Traditionally, NHL coaches waited until about one minute to pull their goalie when trailing by one goal. In recent years, some coaches have started pulling earlier, but the average when trailing by one goal is still well under two minutes, much later than what we find to be optimal.

Incidentally, with around six minutes to go, the decision curve is quite flat. It doesn't matter much for a team down by one whether it pulls from a period between, say, eight minutes and four minutes.[7] The real value of early pulling comes from the range of four minutes or less to go, and for pulling even earlier when the team is down by more than one goal. So, small differences in assumptions or parameters could make the case for somewhat later pulling times than our model recommends. However, no reasonable estimates can justify the current NHL practice.

A team that practices optimal goalie pulling gains an average of 0.04 more points per game.[8] That would

[4] Very soon we're going to question whether that's actually what coaches are optimizing for. Perhaps we should say, "One would think coaches are optimizing for expected points in the standings, but they may have other concerns and biases."

[5] NHL teams get two points for winning and one point for losing in overtime or shootout. So, if the game ends in a tie, one team will get one standing point and the other team will get two. We assume who gets one vs. two is a coin toss to arrive at one and a half.

[6] Given that this is based on a coarse model and empirically estimated inputs, nobody should take the specificity of "6 minutes and 10 seconds" too seriously. Rather, one should simply focus on the "pull a lot earlier than they do in real life" idea.

[7] One bit of confusion we encountered when circulating earlier versions of our article is that some people thought we meant a team down by 1 should pull with 6:10 to go and never reinsert its goalie. That would be foolish. The optimal strategy is to pull at 6:10 and reinsert the goalie if you subsequently tie the game.

[8] In the 2015–16 season with 1,230 regular season games, teams following our recommendations would have pulled 315 times down 1 goal, 221 times down 2 goals, 174 times down 3 goals, and 44 times down 4 goals. This only counts our first recommendation in each game, because we don't know what might have happened afterwards, so this understates the value of optimal pulling. On the other hand, it overstates it because it compares to never pulling, while NHL teams do gain some pulling advantage by pulling (too) late in the game. The total season advantage from optimal pulling is 89.17 standing points, or about 3 standing points on average for each of the 30 teams.

have been worth an average of 3 standing points in an 82-game season, over a team that never pulls the goalie. In the 2015–16 season, every team except for the Washington Capitals was closer than three standing points to either the team ahead of them or behind them in the standings. Boston could have made the playoffs with three extra standing points, and Detroit and Philadelphia would have failed to make the playoffs with three fewer standing points. So, this is a material difference in expected performance that comes without extra cost or work, but simply from acting in an optimal way.

Now let's get crazy. When down 2 goals, it pays to pull the goalie with 13 minutes to go, less than halfway through the third period. If you score to make it 1 goal down, you replace your goalie until 6 minutes and 10 seconds, as playing without a goalie any earlier is too aggressive when you are only down by 1. So, if you're still down by 1 goal at that mark (6:10), you should pull again. When down 3 goals, you should pull at 23:40, that is with 3:40 to go in the *second* period; down 4 goals or more and you can pull at any time.

These numbers may sound silly, but they make sense and they matter. A team that is down 4 goals early in the first period (we might argue in this case the goalie isn't doing any good anyway!) still expects to collect 0.09 points on average, and can increase that to 0.14 points by pulling the goalie immediately and keeping him out until the team pulls within 3 goals, and using the optimal rules thereafter.[9] The intuition is the same as pulling near the end of the game. If you do nothing when you are down four goals early, your chance of winning is very small and you have little to lose by trying to climb back into the game using this seemingly desperate, but actually optimal move (and, as with losing by two instead of by one, going down by five is not that much worse than by four). Though we would admit the assumptions and simplifications of our model are probably being pushed much harder for this "losing by a lot early" analysis, we find the intuition compelling and believe the recommendation is at least directionally accurate (pull much earlier than current practice when down by a significant margin). When you're down by a lot, you've essentially lost the game unless something radical happens. In financial terms, you're "deep out of the money." In this situation, seemingly radical actions are actually very low cost, as you're giving up very little if they backfire.

LESSONS FOR INVESTORS

Now, let's get to our real purpose (besides loving hockey and math!). There are some important risk management and investing lessons to be learned by considering this optimal hockey strategy problem. The most basic lesson is to make sure you are thinking about the right risk and return. Pulling the goalie always increases the volatility of the numbers of goals scored and is always a negative expectation in terms of the score (it's more likely that the opposing team will score than it is that you will). For those reasons, it is often used as a metaphor for a high-risk, desperation move. However, the point of hockey is not to maximize the differential between the goals your team scores during the season and the goals it gives up (if this were the case, then no one should ever pull a goalie).[10] The objective in a hockey game is to

[9] These numbers come from our model, which was not designed to simulate full games, and both numbers seem high to one of us. Unfortunately, there is not enough data on teams down by four goals early in the game to check the estimates empirically. As a sanity check, we note that in the last six NHL seasons, of the 1,169 games in which a team has fallen behind by 4 goals, the trailing team has managed 5 ties in regulation and no wins in regulation, so actual NHL strategy has resulted in an average of 0.006 standing points, not the 0.090 we suggest. But the average time of first falling behind by four goals is five minutes into the third period. It might be plausible that early in the game, with 4 times as long to go, a team has 15 times the probability of making up a 4-goal deficit. In the 3,024 NHL games in which a team has fallen behind by 3 goals, there have been 112 ties in regulation and 27 wins for the trailing team, for 0.073 standing points on average. If a team down by 4 pulls the goalie early in the game it is virtually certain that one team or the other will score, and our model implies a 0.0197/(0.0197 + 0.0430) = 31% chance that it is the pulling team. 31% of 0.073 = 0.023, which is bigger than 0.006 (if the leading team scores to get a 5-goal lead, no NHL team in the last 6 seasons has recovered to tie or win in regulation after trailing by 5 goals). A 0.017 expected standing point advantage for pulling down 4 goals with 15 minutes left in the game seems compatible with our model's claim of a 0.050 standing point advantage for pulling down 4 goals in the first period.

[10] We do have examples of sub-optimally never pulling the goalie. The Miracle on Ice of 1980 might never have occurred if the Soviet's coach Viktor Tikhonov had pulled his goalie and the fearsome Russians had 1–2 minutes of 6 on 5 (let alone the 6 minutes 10 seconds we recommend!). But he left the goalie in the nets as time expired. Some say he was philosophically against pulling the goalie (so he would have hated this article). We think perhaps he just had nearly zero practice at it, as the Russian team lost so rarely!

maximize the number of standing points (as we're still assuming, though we'll soon consider other alternatives). A team down by a single goal with short time remaining gains a lot by scoring and loses little if the other team scores—which argues for a different measure of risk and return. As we have shown, pulling the goalie actually *reduces* the risk of losing the game—and viewed this way, it's actually an insurance move.

Investors sometimes make similar mistakes when they focus on the risk of an investment stand-alone (the wrong measure), rather than the risk an investment adds to their overall portfolio (the right measure). They may also make this same mistake when they focus on the volatility of their portfolio rather than the probability of an unacceptable level of return over their risk horizon. If increased portfolio volatility comes with sufficient excess expected return, over long periods of time a portfolio can have a better chance of producing an acceptable return by taking more short-term volatility; in fact, not taking enough volatility can be very risky in terms of not hitting the required level of return.

This issue of selecting the proper risk measure also relates to valuing a stock option. If, for example, a call option is well out-of-the-money with little time left to expiry, an investor would prefer that the volatility of the underlying asset increase because the new higher volatility increases the value of the option and thus reduces the risk of losing (as it reduces the likelihood of the option expiring worthless). This is true even if it comes at the cost of a lower or negative expected return on the underlying asset. Pulling the goalie is effectively like option holders being able to increase volatility on their own (at the cost of lowering the expected return if measured in terms of goals, but not in terms of the truly relevant measure of points). There's no management lesson in this one, because, unless we've missed something, investors can't just turn up an asset's volatility on their own. So, it's simply a neat analogy. But seen in this context, you can think of hockey coaches as out-of-the-money option holders who indeed *can* turn up the volatility on their own in their own games!

Another lesson comes from asking *why* NHL coaches don't pull goalies earlier. This question relates to a well-documented phenomenon in sports. Examples include: basketball coaches who were very slow to have players attempt enough three-point shots; football coaches who don't go for it on fourth down often enough, nor do they attempt enough two-point conversions; and baseball managers who were very slow to appreciate the value of walks, the cost of outs, and the utility of the excruciatingly annoying radical infield shift. All of these coaches face tremendous pressure to win, and they have the benefit of many repetitions of game situations from which to obtain precise statistics that can guide the right choices. Yet, they consistently fail, often for decades, even after every numerate fan has figured things out. (Many of these sports phenomena have changed over time, some quite recently, but hockey seems to be more frozen in the past.)[11]

Two reasons have been advanced for this common failure to act optimally across these major league sports. First, coaches are not actually rewarded for winning (or our metric of "standings points"). Rather, they are rewarded for being perceived as good coaches. Obviously, the two are closely related, but they are not exactly the same thing. If a basketball coach gets his team to execute crisp offensive plays with few turnovers that lead to textbook two-point baskets on 50% of possessions, he or she is deemed an excellent coach. If his team still loses 100–102, well, the players just weren't quite good enough and they didn't get the breaks. If the same coach encourages the team to "run and gun"—shooting a ton of threes, with lots of turnovers and misses, but scoring on 35% of possessions, the coach clearly has lost control of his team. If they win 105–102, it may be perceived as mere luck, because everyone knows that was a crazy strategy. Essentially, winning ugly is undervalued compared to losing elegantly; and losing ugly can be career suicide. Once again, the way you measure risk matters in making the optimal decision.

This is a problem in portfolio management as well. A chief investment officer (CIO) who runs a tight ship, allocating funds to low-fee index funds and moderate-fee active managers who beat their benchmarks, is perceived as an excellent manager (more so these days than in earlier times when perhaps the opposite, a CIO investing only in high fee ex post successful stock pickers, was conventional wisdom). Although those are good things, they are nonetheless sometimes not good enough, in which case a CIO should look into alternative choices

[11] Pun intended, sorry!

and new types of risk, accepting some leverage risk for the benefit of additional portfolio diversification, for example, or taking some liquidity risk in exchange for higher expected return.[12] But if these things don't work out, the CIO can lose his reputation for competence; and if they do work out, well, everyone knows it was luck because those things were "risky." As John Maynard Keynes pointed out, "Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally." This applies to coaching sports and portfolio management alike.

We can also apply this lesson to security selection. Cheap stocks (using simple ratios like price-to-book, earnings, or sales) tend to outperform expensive stocks over the long term. But they also tend to be issued by "worse" companies—companies with less exciting prospects and more problems than others. Portfolio managers who own the expensive subset of stocks can be perceived as prudent, while those who own the cheap stocks seem rash. Even if the data say otherwise.

The second reason coaches shy away from actions with short-term risk is that sins of commission are far more obvious than sins of omission. The hockey coach who pulls his goalie down 0–2 with 10 minutes to go and loses 0–5 will face harsh criticism from every quarter. A coach who quietly loses 1–2, pulling his goalie only in the final minute, can hold his head up and say his guys played hard, but the puck didn't roll their way tonight; it was a close game, and they'll work even harder to get the breaks tomorrow (giving 110%, of course).

Unfortunately, CIOs face this issue as well. If a portfolio is hammered by overexposure to equity risk in a crash, well, the market went down, what can you do? The investment team worked hard, but the markets didn't roll their way this quarter (they failed conventionally), they'll work even harder to get the breaks next quarter (though hopefully CIOs don't use terms like "110%"). But a CIO who chooses an alternative investment that disappoints will face sharp criticism, even if the portfolio now has a better expected long-term ratio of risk to return (and is likelier to succeed unconventionally).

[12] Disclosure: the first one, leverage risk, is one of us talking his book.

We can view the question of why coaches don't pull early enough in yet another way that has an investing analogy. Investors have been shown to be reluctant to sell their losers (part of the so-called "disposition effect") presumably as selling is psychologically "locking in" a loss. Might the extreme reluctance toward pulling the goalie, when say down two with more than 10 minutes left, be the result of a similar cause? Pulling the goalie earlier may be the best action in terms of expected points, but it runs a very high probability of going down three goals, and thereby almost certainly "locking in" the loss.

Finally, our model absolutely ignores the fact that the game of hockey is a form of entertainment. We are cold-hearted quants after all. Perhaps, more speculatively, fans enjoy the last one to two minute drama very much, and even will stay tuned to a game with a late two-goal differential. But, pulling the goalie very early when down two goals leads to a very large chance of being down three thereafter. That would kill the entertainment. Might coaches feel a desire, or even tacit pressure, not to ruin the fun? In other words, our model is myopic, only caring about maximizing the points from the game. Perhaps coaches, and the league in general, are better long-term present value maximizers than our model if providing more entertainment maximizes the "franchise value," even if it is slightly costly in terms of points. In particular, if all teams do it, the overall entertainment value of hockey is increased with no ex ante advantage or disadvantage to any one team, as they're all acting alike.[13] Now our analogy in business has drifted to monopolistic collusion!

There is one bright spot. Coaches slowly do come to adopt better strategies. It can take decades, but things ultimately move in a rational direction. When this happens, it proves that the stats geeks were right all along, but, more importantly, it also shows that it is possible over time to surmount the social and behavioral factors that too often sabotage the optimal strategy.

[13] This is actually another analogy with investment. Initially, the advantage to the sole coach who follows our advice will be substantial (recall from earlier the three expected standings points on the season). But, over time, if all coaches read our article and follow our prescription, no team will get an advantage in the end. They would then have to act optimally just to keep up with the rest of the league. Again, hockey sounds a lot like the financial markets!

Suboptimal strategy in sports causes no net harm, and gives quants something to feel smug about, even if the jocks are still more popular. Suboptimal strategy in investing, however, can be a real problem for society. In addition, CIOs don't have the same advantages as coaches, unfortunate as that may be. They don't see the same situations repeated thousands of times. Rather, they must try to make sense of a single investment history in rapidly changing financial markets. They don't get to start fresh at 0–0 every night; instead, they need to think about 5-year, 10-year, and even longer histories, where a 5% loss this quarter means 5% less money forever.

One way to level the playing field is for CIOs to study risk carefully in laboratories like sports, where the statistics are much clearer but the human cognitive biases are essentially the same. That's why investors should definitely care about when goalies are pulled.

And, coaches, you're not pulling your goalies anywhere near early enough; well, at least not yet.

ACKNOWLEDGMENTS

We thank Malcolm Gladwell, Ronen Israel, John Liew, Michael Mendelson, Tobias Moskowitz, Chris Schindler, and Rodney Sullivan for thoughtful comments and advice. Special thanks to Michael Lopez, Director of Data and Analytics at the National Football League, who caught a data error in an earlier version. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR.

Clifford Asness *is a co-founder of AQR Capital Management in Greenwich, CT.*
cliff.asness@aqr.com

Aaron Brown *is a professor at the Courant Institute for the Mathematical Sciences in New York, NY.*
aaron.brown@eraider.com