

Paul-Erik Veel 416-865-2842 pveel@litigate.com

The First Three Months of the SCC Leave Project: A Successful Start

Three months ago, we launched our Supreme Court of Canada Leave Project. Part of that project is a machine learning algorithm that provides predictions of the likelihood of different cases getting leave to the Supreme Court. Since launch, we've made fourteen weekly predictions for leave applications to the Supreme Court of Canada. Our model has provided predictions for the likelihood that leave would be granted from 123 decisions of Courts of Appeal across the country.

Since we've launched, we've heard two questions. First, how does the model perform? Second, it's a neat idea, but is it actually useful? This blog post answers those two questions.

HOW DOES THE MODEL PERFORM?

Very, very well.

As we've stressed since we launched, our model gives probabilities, not certainties. It's not that a 90% predicted probability case is a slam dunk, and a 10% predicted probability case is hopeless. Rather, a case with a 10% predicted probability of leave will get leave one time out of 10, while a case with a 90% predicted probability of leave won't get leave one time out of ten. What that means is that our model can't be judged by its performance in predicting any single case, since there will always be outliers. Rather, it has to be evaluated based on its performance over a range of cases. Now that we've made three months' worth of predictions, we have that experience to give some preliminary results. And the results are good.

The easiest way to evaluate the performance of our model is by looking at how different categories of predictions turned out. If you've been following our model, you know that we group leave application cases into four categories, depending on their predicted probability:

 Cases to Watch – These are cases where our model predicts greater than a 25% chance that leave will be granted. Over the last three months, our model identified 10 Cases to Watch.



- Possible Contenders These are cases where our model predicts between a 5% and 25% chance that leave will be granted. Over the last three months, our model identified 42 Possible Contenders.
- Unlikely Contenders These are cases where our model predicts between a 1% and 5% chance that the case will get leave. Over the last three months, our model identified 36 Unlikely Contenders.
- Long-Shots These are cases where our model predicts a less than 1% chance that the case will get leave. Over the last three months, our model identified 35 Long-Shots.

So how did our predictions do relative to the outcomes? Essentially spot on:

- Among the 10 Cases to Watch (>25% predicted probability of leave), 5 got leave (50%).
- Among the 42 Possible Contenders (5-25% predicted probability of leave), 7 got leave (17%).
- Among the 36 Unlikely Contenders (1-5% predicted probability of leave), 1 got leave (3%).
- Among the 35 Long-Shots (<1% predicted probability of leave), 0 got leave (0%).



For each of the four categories of predictions, the proportion of cases that got leave was roughly in the middle of the predicted



probability range. This is exactly what you'd expect in a model performing accurately: on average, over its first three months of its performance, the predictions of the model were spot on.

IT'S A NEAT IDEA, BUT IS IT ACTUALLY USEFUL?

Yes. 100%. Our model is useful for any client considering whether it's worth the time, effort, and cost of seeking leave to the Supreme Court of Canada.

Any lawyer advising their client on the likelihood of a successful leave application will tell their clients that most leave applications are unsuccessful. Our data shows that over the last few years, only 7% of leave applications are successful. However, for most clients considering a Supreme Court leave application, they also have to think, to use a poker concept, about the "pot odds" of a leave application. What that means, put simply, is that it can make sense to take a bet on a low probability outcome if the cost is relatively low and the potential benefit is relatively high. Take a typical civil or commercial case where a client is considering a leave application: by that point, the client has already gone through a contested trial or application as well as an appeal, representing years of time and effort. If they're applying for leave, they've been unsuccessful below. That means they have either not received money they think they should have (if they're plaintiffs) or been ordered to pay money that they think they shouldn't have (if they're defendants). They have also probably been ordered to pay costs. That means that they may have the potential to gain quite a bit from a successful leave application. That's why clients in so many cases (several hundred per year, in recent years) are willing to take a shot at a leave application, despite the low chances of getting leave.

While most clients understand that leave applications are an uphill battle, what they want to know is whether they have a decent shot at getting leave, or whether their case has very little chance. Rational clients may be willing to take a chance on a case with a decent shot, but not be willing to spend the time and money on a complete long-shot. That's where the model's predictions come in handy: the model provides an easy and objective way of distinguishing between those.

As an example, imagine we take the threshold for a case having very little chance as anything below a 5% chance. Among the 71 cases with a predicted probability below 5%, only one got leave (1.4%). That seems pretty hopeless. By contrast, among the 52 cases with a predicted probability above 5%, 12 got leave (23%). That's by no means a certainty, but the odds are likely good enough to justify the cost of bringing a leave application. By quickly classifying a case as



having either a decent chance or a very low probability, the model allows clients to make an easy and informed decision as to whether to seek leave or not, with limited work needed to reach that conclusion.

One comment we've heard is that the model isn't useful, because a sophisticated lawyer can already give their client a sense of the likelihood of getting leave. No doubt, there's value in having a skilled and experienced lawyer to advise you, and the model can't and doesn't intend to replace that legal advice. But it can complement conventional legal advice in various ways.

We see four important ways that the model can complement legal advice.

- If you already have a great lawyer advising you, think of the model as providing a second opinion. If the model and the lawyer agree that a case has a decent shot, that's a good sign that it really does have a decent shot. By contrast, if the model and the lawyer disagree, that might be a reason to dig a little deeper and understand why.
- The model can help correct for biases that lawyers might bring to the situation. There is a host of psychological research showing that human decisionmaking suffers from all kinds of biases, particularly when it comes to estimating probabilities. Lawyers of course try to correct for those when giving legal advice, but that doesn't mean that we don't all sometimes fall prey to reasoning biases. The model can help identify those spots where lawyers might have either fallen a bit too in love with their argument, or who might be overly pessimistic. By relying purely on objective datasets, the model provides a check, free from human biases. Moreover, even if a lawyer doesn't fall prey to those biases, it can be hard for the client to know that; the model provides an objective, corroborating data point that can help reassure clients that they're getting appropriate advice from their lawyers.
- The model can play a screening function that, in at least some cases, can help clients make decisions at lower cost. For example, 35 of the cases predicted by our model (28%) were predicted as Long-Shots, meaning that they were predicted to have less than a 1% chance of leave. Among the Long-Shots we predicted, not a single one received leave. It might be that if the model predicts your case as a Long-Shot, you might not even



pay for an opinion or further advice as to whether it's worth seeking leave, since the probability is so remote.

• The model can give clients a quantitative prediction, which some clients find valuable. While many clients are content to know whether a case has a decent chance or has very little chance, some sophisticated clients want quantitative predictions of their chances. As lawyers, we're often asked to give these types of assessments, though they're often fuzzy and impressionistic. By contrast, a quantitative model can give a concrete probability to clients who want that.

Will the model replace lawyers? No. And it was never meant to. But it does provide a valuable complement: a good lawyer and the model, working together, will provide better advice at lower cost than conventional legal advice alone.

