



The Accounting Podcast Series

S02E02: Patrick Ferguson

Transcript

A: Albie Brooks

P: Andrew Williams

Albie: Welcome to [TAPS](#). I'm Albie Brooks, and working with me is Abbey Treloar. Our guest today is Patrick Ferguson. Patrick completed a Bachelor of Arts in English and History and then backed that up with a BCom, both here at Melbourne. Then in 2016, Patrick headed to Harvard University to undertake his PhD in Business Administration. Close to finishing, Patrick joins us today to chat about the really interesting work he has been undertaking, particularly in the area of sport analytics and its application to business-related problems and issues. Welcome to TAPS, Patrick.

Patrick: G'day, Albie.

A: Well, there's so much to talk to you about. Let's start with how you came to be interested in the field of sport analytics. Where did that originate from?

P: I guess from a fairly young age. So like most middle-class Australian kids sort of growing up in the '90s, there was a lot of sport on television and a lot of sport in which Australia was fairly successful, particularly with regards to cricket and rugby, the two sports I probably followed most closely. And then from there it was a case of getting given books every Christmas and birthday that had a lot of sports statistics in them, and that became a way for me to sort of make sense of the teams that were playing at the time but also try to develop an understanding of the history of those different games and sort of contextualise how successful the Australian teams of my childhood were relative to the historical teams.

I have memories of sort of flipping through old copies of Wisden and thinking about how someone like Steve Waugh or Allan Border may have compared to a player like Don Bradman or Greg Chappell. And even at that young age when I didn't have a lot of sort of formal mathematics behind me, I was already going through the kind of process of thinking, "Well, how would you make those comparisons? They faced different rules, different playing conditions, the strength of the opposition varied in quality. How do you adjust for those sort of differences over time to make an apple for apple comparison rather than an apple for orange comparison?" So that was something that had already sort of piqued my interest in sport from a young age. And then when I moved into my tertiary studies, taking stats and econometrics courses as an undergrad at Melbourne Uni, I wanted a way to make that material more accessible, and sport became an avenue through which to do that.

And so rather than kind of learning some of the dry material, I devised ways in which to sort of force comparisons with sport or towards sport. And a big tool that helped me sort of make that transition was I was given a copy of the book Moneyball, and that was a big moment where I kind of realised, "Hang on, this is clearly about baseball and about sport and the business of sport, but a lot of the underlying philosophical conceptual things that play in the book were about how do you use numbers and data and statistical techniques to test theories about the world and understand how humans behave?" And that was a really interesting intellectual moment for me, this idea that sports - sorry, stats and econometrics can help us understand things about sport, but it also tells us something about how humans make decisions as individuals or within organisations.

So it could be anything from why are certain types of individuals undervalued by the market? Why doesn't the market fully appreciate maybe the performance that they provide? Why are certain attributes or traits overvalued? How can we identify those instances? How can you potentially exploit them or how can you potentially correct them if they're things that are causing losses in efficiency and equitability across the market? So that was kind of another big moment in my sort of growing interest in sports analytics. And then finally, towards the end of my undergraduate studies at Melbourne I went through that fairly standard process where you apply to a bunch of internships and grad year programs at consulting and banking firms. And it dawned on me that that wasn't really the space that I was most interested in, and I'd been fortunate enough to come across someone that had just founded a small sports analytics start-up in Melbourne.

And I basically said, "Look, I don't have anything else to do with the summer and potentially the following year. Will you bring me on board and let me play around with the data and try to help deliver some value to the start-up?" So I then proceeded to spend the sort of next two years making, yeah, very little money but getting to spend a lot of my time looking at player performance statistics and contracting data of professional rugby and AFL. And that was a really great experience because it allowed me to get my hands dirty with some of the data out in the field. It allowed me to get to know people at these different sports organisations, and also it allowed me to sort of gain what I think was an important piece of information, and that was that professional sports has become very rich with data, but fairly insight-poor.

And the sense I got was that there are lots of great measures out there in pro sport, but also in lots of industries, and only a very small part of it is actually capturing those measures and storing them somewhere. The real value comes from how you actually make decisions based off those measures. And I saw pretty quickly that that was true of a lot of sports organisations, that they'd gone crazy over measuring lots and lots of things without thinking about what decision were they ultimately trying to make with that data, and were they necessarily exploiting that data to its full potential? So that got me thinking, "Look, great, I've spent some time now in industry looking at these sports questions. How can I then maybe think more deeply and develop an even more developed set of skills to address those questions?" And that was when I thought about and ultimately did apply to a PhD program in the US.

A: Yeah. Well, excellent. You're quite right about the availability of data in sport in analytics and performance metrics which is mirrored, of course, inside organisations as well, where we have so much in the way of analytics and metrics that, in effect, we can be swamped with the data rather than, you know, using it to lead to better decision-making, which is clearly what you've tried to do with sports.

P: Yeah, yeah. Absolutely. And I think that's the key connection to make, that really none of the phenomena - at least what I think about in my research in sport - is unique to that setting. I use sport as a sort of lab to answer questions about broader organisations and firms and the challenges they face. And I think that's kind of key there, is that these are pretty fundamental features of how people make decisions and how organisations function.

A: Yeah.

P: And you can port that across to healthcare or finance or whatever that may be. Retail. There's a bunch of operational accounting strategy questions that are really, when you abstract away, largely the same across all these different settings.

A: Yeah. Now, as you alluded to, you were able to sort of build on some of this early work and make it the focus of your PhD studies. Could you tell us just a little bit about the program you've undertaken and how you set about acquiring the skillset? Clearly you had a really good skillset to start with, but how you built on that skillset to undertake the kind of work that you've been doing?

P: Yeah, so I was fortunate enough to be encouraged to apply to PhD programs at business schools in the US, and then fortunate enough to get into the business program at Harvard Business School. So I'm currently in my fourth year of the PhD there, and my primary concentration is in accounting in management, and I have a fairly strong focus as well on what I'd call kind of more general managerial economics. Now, the great thing about the program that I'm in and I'm beginning to wrap up now is there's a strong focus on training individuals to produce world-class research and then go off and kind of continue do that research in academic roles around the world. Now, as we've touched on earlier, my research heavily relies on data from professional sport, but that's always had an intention behind it to generalise to more economic questions.

So one of the strengths of that program at Harvard is that there's a really great sort of two-year, two-and-a-half-year component of the PhD which is just coursework. And I wouldn't go so far as to call it hazing, but they definitely stretch you in terms of your intellectual and academic capabilities. So there's a big onus on, "Look, here's a firehose. Basically try to drink as much as you can from it in this short timeframe to develop your technical skills, whether that be microeconomic theory, your econometric and statistical modelling skills, and then also you're exposed to reading a lot of academic research in finance and accounting and management more generally." So that provided a really intense sort of burst of exposure to really great courses, building up technical skills and then building a lot of research broadly in a range of areas, which was really great for kind of knowing where to look to access this sort of stuff, and then having guidance from faculty there that were teaching.

And then concurrently to the coursework, there's also a strong focus on producing your own research straight away, and this idea that you should tackle important, challenging questions that are potentially risky and uncertain, but that's what you're there to do. That's the whole idea of being an academic at one of these well-funded institutions is you should go after questions that maybe in industry don't always get answered because there's incentives, whether that be more short-term focused, you can't dive into these problems. And that was always the sort of philosophy that was sitting over the work and research we were encouraged to do at Harvard. And then finally, these PhD programs have kind of what I would call a "hidden curriculum", and that's the fact that there's not formal coursework, but there is a bunch of stuff, a bunch of skills that you need to acquire to do rigorous, empirical research.

And that really is only developed by throwing yourself into real-world research projects and having to work with the data. So that just involved pairing up with other PhD students and faculty and identifying interesting research questions, and then going out there and working out, well, how do I get data to try to answer that question? I have learn how to basically write programs to scrape that data, I need to be able to write code to analyse that data. I need to come up with ways to share those results and collaborate with faculty. And that was a very rewarding portion of the project - sorry, the PhD. And that's something that you continue to acquire as you tackle these sorts of research questions, is keeping the sort of best practices in place by collaborating on

these projects. And that's been a hugely rewarding component of the studies as well.

A: Yeah. It certainly seems to have put you in a position of being - preparing you perfectly for some of the work that you have undertaken. You'd mentioned having sort of nice, interesting research questions and exploring them out in the field, so I'd like to turn our focus just to one of your current papers. And it's titled, "Consuming Contests: Outcome Uncertainty, Information Disclosure, and Spectator Demand". Really interesting field of inquiry. Could you just elaborate on how that work came about?

P: Yeah. So this is a question I've sort of been working on for probably the last 12 months, maybe even 18 months, in various iterations. And this largely came about by just following any major sports league or sporting contest around the world, so moving between Australia and the US, you start to see these sort of commonalities in terms of how these leagues' contests are run. And one of the big things that jumps out is this focus on competitive balance, as an attribute quality of sporting contests that captures a lot of popular interest. So you'll often see in the US and you'll see in Australia but less in Europe, a lot of administrators talk about, "Well, we need to have a fair competition. We need to have a competition where strong teams can win but also weak teams have the opportunity to, over time, improve themselves and become competitive." And this objective of competitive balance has been pursued by the leagues through the implementation of a range of policies.

So it could be things like salary caps, drafts, revenue-sharing. And these policies have been shown to indeed make contests more well-balanced and competitive, but they've also faced some resistance from player associations, from owners, on the grounds that they violate a range of kind of rules and laws. They're seen as anti-trust violations. They're seen as infringements upon the fair and free competition of the market. They're not always as embraced and as welcome by all the stakeholders in a league as they are by the administrators. So my first question here was, well, yeah, we've seen some opposition. The leagues nearly always push back by saying, "Well, yeah, maybe players do suffer a little under these regulations, maybe the owners don't extract the full rents they could, but this is all done in the benefit or in the interests of the consumer, of the spectator. We should have an even contest because that's exactly what a spectator wants. When they decide to show up at a game, it's driven by the fact that they want some uncertainty around the outcome of that upcoming game."

Now, to me this seemed anecdotally obvious, intuitively fairly obvious that, like, if I'd spoken to someone at the footy this would largely be what they'd say. But if you go and look at empirical research on this topic, the results are all over the place and you don't actually see this obvious anecdotal answer borne out in the research. What you see is across different leagues, over time, across different countries, the results are mixed. Some leagues' competitive balance and uncertainty of outcome drive attendance; other leagues there's no relationship with attendance. Even in some instances it seems like it's negatively related to attendance, as if spectators want games that are very certain, whether that be very certain that their own team wins, or even one study showed very certain that their own team loses, which made very little sense to me.

A: Of course, I understand.

P: Yeah. So I guess the way I came in towards this project was this question of your average punter on the street has this idea that, yeah, spectators care about uncertain contests. That makes sense. The literature has said, no, this is all over the place. This has not held up. So where myself and my co-author approached this question was understanding, well, maybe why is the evidence not as clear-cut as we would expect? So we dug around and looked at prior studies, and the hunch that we had is that they're informative and useful because they are data points along the way, but they're all subject to a pretty major empirical limitation. And that's that they're not relying on what I would argue to be a robust, well-identified empirical strategy to establish a causal effect.

So what I mean by this is they don't just isolate the relationship between the uncertainty of an upcoming game

and demand for that game. They partly capture that, but they also capture lots of other things that move with game outcome uncertainty and attendance. So this could be things like promotions and advertising differences across games. This could be pricing differences across games; this could be the timing of the game during the season; this could be weather. All these sorts of things that both simultaneously affect the outcome uncertainty of a game, but also attendance. Now, some of these prior studies have tried to address these confounding factors, but some of them are just unobservable. You can't actually put a control in for all these things. There's lots moving around and, as researchers, we can never fully know the full set of things that determine attendance.

So where my co-author and I have decided to direct the project was towards thinking, "Well, where can we find natural experiments in sports leagues or sports events that cause random variation in the uncertainty of an upcoming game?" In effect, we want shocks that are uncorrelated with anything else other than outcome uncertainty, and we want to use those shocks to then identify the causal effect of this thing on attendance, which in theory sounds really neat and appealing, but it's very messy because there's a range of ways in which these sorts of natural experiments basically can be violated as identifying instruments. But, having spent some time sort of thinking carefully about the institutional details of different sports leagues, it occurred to me that the AFL provided a really nice setting to think about this.

So basically in the AFL prior to games being played, teams must disclose the line-up of their team several days prior to a game. So what I thought about this was that it's a nice opportunity to exploit potential changes in teams' line-up that are a shock to the general market, a shock to the general public. And these most likely arise due to injuries that are random or exogenous or occur with very little control or ability to be influenced by the teams themselves. So basically what my co-author and I did was look at a huge number of AFL games that have been played over the last 10 years or so, and for each game scrape all the line-up changes that were made in advance of that game, then scrape all the betting data about those games to trace the effect on the financial markets of these disclosures of injuries, and then go and look at how attendance changed when these random injury shocks made upcoming games closer or less close.

So using this design, we were largely able to then hold constant all the other things that also move with game outcome uncertainty and also move with attendance, things that the prior studies in the literature had not been able to hold constant. So that was sort of what we decided to set out to investigate and how we provided a novel way of thinking about that problem. But there's a reason why people really hadn't necessarily done this before, or at least several reasons. One is that there's been a more recent turn towards - I suppose in economics and management and finance and accounting more generally, there's been a more recent turn towards thinking about these sorts of causal inference problems more robustly, but also there's just the data collection challenge of doing this sort of work. It requires you to collect a lot of stuff, which is not trivial to do.

So, like I said earlier, we need to scrape data from the sports betting market, we need to get line-up change data, we need to get attendance data, we need to get player statistics from the field. So all that sort of stuff is a pretty onerous thing to gather, and it creates a fairly high barrier to entry in this space. So we decided fairly early on if we want to do a really good job of answering this question we have to incur some costs, and those costs were going out to collect this data that's not straightforward to get. And that was kind of the strategic decision we made with this research was to go after this question and to invest in the data collection up-front.

A: Yeah, so a huge data collection project, particularly as you've just indicated. You went back - was it 10 years, you said? You went back through 10 years of data?

P: Yeah, we did. Yeah, around 10 years. The final paper I think used a slightly smaller sample, but yeah, we

effectively had - so if you think about it, each observation in our study is a game. So for each game that's played in each round for, what, 10 seasons, we were getting all the line-up announcements, we were getting the performance statistics of the players on the field, we were getting things like the weather, the attendance figures. We were getting a range of things that we were scraping from numerous sources to ultimately throw into a big pooled data set that would allow us to really tightly kind of pin down the causal relationships we were interested in.

A: Okay, so let's go to those. So what is it? What did you find, having collected all this data?

P: So basically, we kind of took a two-step approach. The first thing we did was - before we used our kind of new novel empirical strategy to look at the casual outcome between game uncertainty and attendance, we said, "Why don't we just firstly do what the literature has historically done in this space and use their techniques to estimate the relationship between these variables?" So we went out there and we looked at the traditional model, we ran that. And that said a one standard deviation increase in the uncertainty of the outcome of an upcoming game in the AFL led to a four to five per cent increase in attendance for that game. So that's an extra 2000 people to a game on a, whatever, Saturday at the MCG. That's what the literature's traditional techniques had said, so a fairly small increase.

Now, we went and used our more explicitly-identified strategy to estimate this same construct, and what we found was that a one standard deviation increase in game outcome uncertainty was associated with a 12 to 15 per cent increase in attendance at a game. So that's closer to sort of 4000-, 5000-person increase in attendance. So almost a threefold larger effect size than had been previously documented. So the takeaway from this, really, is yes, people have a very strong appetite for uncertain contests. our marginal spectator does want to show up to see a game where the outcome is highly uncertain. And the prior literature had really heavily underestimated the strength of those preferences, which to a regular person on the street is not terribly surprising, but the neat part here was that we were able to say why prior work had underestimated it and then go and arrive at this more robust estimate.

Now, to the league this is also a powerful finding to have because it provides quite a lot of support for their aggressive pursuit of competitive balance over the last few seasons. So the AFL obviously faces some flack when they want to implement revenue sharing or they want to change how the draft works to create a more level playing field. What our data helps to say and argue is that, yeah, the consumer on average will benefit from the implementation of these policies as it's creating a product that's more desirable for them. And that was the primary takeaway from the paper. So I think, yeah, it has important implications for the league. It also allows us to address a fairly well-established academic literature, and it also allows us to, I suppose, take advantage of a way of thinking about this problem that is kind of unique to people that spend a bit of time thinking about finance and accounting questions.

So we thought - the way we kind of came at our strategy to identify these causal effects was through a disclosure event, which was the announcement of team line-ups, and then using a capital markets measure - betting odds - to then trace out game outcome uncertainty. So a neat thing here is we were able to draw on a set of skills thinking about accounting and finance problems to answer a consumer behaviour question in sport. And I think that's a kind of neat synthesis of inter-disciplinary techniques and topics that is partly the reason why I find this sort of academic research so rewarding and interesting.

A: Well, I don't blame you. It's fascinating. I think that combination of, as you say, the sport analytics, business, accounting, consumer behaviour-type disciplines that are all involved in this has really resulted in some ground-breaking work here, and I think you should be applauded for that, along with your co-authors and supervisors on

encouraging you to undertake work that really is a bit on the edge for many people, and to be able to combine, as you say, highly robust statistical methodologies with really interesting questions and at the same time be able to connect nicely business and accounting-related issues is, well, really neat. You should be commended for that.

P: Thank you, Albie.

A: So I wish you well with the paper in due course. Now, I know you're involved in a range of other pieces of work as well around sport and analytics, but I just wanted to talk to you a little bit broadly about sport analytics as we look forward. So we know that organisations and sporting organisations down to individual sporting clubs and even down to individual sports people, whether they're at professional or amateur level, continue to seek analytics. So there's that aspect, and there's obviously the research aspect that you're heavily involved in. So going forward, how do you think the nature of the analytics themselves are going to evolve?

P: Yeah, so I think this is super interesting. We're at a real inflection point in sports analytics, but also this analytics data science stuff more generally across industries. I think, as we touched on earlier, there's been an explosion in measurement. So it's become cheaper, easier to measure lots of things and store that data, and even kind of draw on that data. But we're a long way away from capturing the real value of that data. And I think this is something in sports that's going to have to begin to change. And I think one of the real opportunities for organisations to develop a comparative advantage here is around thinking very carefully about the difference between what I call causal inference problems - so saying if and how X affects Y, identifying those sorts of problems and distinguishing them from prediction problems, which is really just saying, "If X, we see this." And they're very different things, because it's basically exactly what I had to address in the paper I just described.

You can form pretty good predictions about the world that seem to hold up, at least looking retrospectively, that tell you nothing about the mechanisms and tell you little about the validity or value of a policy you may pursue. So I think where we're likely to see some developments is organisations really carefully thinking about using stats and numbers and data to make causal statements. And what I mean by this is in sport it could be anything from - let's say in cricket you'll see something like ball-by-ball data on a batsman facing a particular bowler. And they'll tell you the ball moved - whatever. When the ball pitched at this length with this amount of swing, the batsman was this likely to hit the ball to that spot on the field. Then you'll go and see a statement as if to say, "Well, if you then put the ball there again, this is the likely outcome."

When you, in drawing of inferences, neglect the fact that those deliveries historically were not randomly varying where they were pitched. So you can't make a causal statement about that. So what we're going to have to see is a move to thinking more carefully about the identifying assumptions, the assumptions you have to make for historical data analysis to tell you a casual thing. I think that's a key distinction. It's something in academia we've seen a sort of revolution towards, but I think industry is going to have to make that pivot to get the full value out of the data that they have at their hands. And I think that's kind of a key thing. So you'll see it in - if you think about regular business problems - anything from pricing to corporate governance questions around what makes for a good board member. They're all in a way causal inference questions. So if you vary this attribute, you get this outcome.

And you have to be very careful about making those sorts of statements with historical data, because lots of other things move. And I think what we're going to see is a move towards thinking about problems in this way. And what that requires is a change in mindset when working with data from thinking of this as an engineering problem. Here I can use all these hard technical skills to - whatever - fit a really convoluted model. So moving from an engineering viewpoint to a scientist in a lab's viewpoint. And that viewpoint is, "Here I am looking at the

world. I have data that tells me stuff about the world. I need to formulate and then refute hypotheses about how things in the world work." And that's - what I think about is a very experimental mindset. So if I hold constant stuff and I move this one thing, what happens to the outcome? And I think what we're going to have to see is a move towards thinking very carefully about how to take advantage of this experimental mindset more than a sort of engineering mindset which we've seen in the past.

And I think that's something that also really favours what I would consider to be people that maybe their training wasn't necessarily in the more technical areas, they're in organisations. But they have very strong institutional knowledge and very strong knowledge of their setting, and also just a very good sense about how people make decisions in the world that they inhabit. And I think that's - what you'll see is a swing back towards a value on those attributes and skills as being able to work with the data becomes easier and as many things can be automated. It's putting the onus back on people that actually carefully think about the substance of the problems, and not so much about how can you pull in the most data points possible and fit a model as efficiently as possible.

A: Excellent. Look, it really does seem to be an area going forward that provides a lot of opportunity for analysis and availability of info, whether it's in sport, business, other settings. And so, as you say, there's the collection of the data itself and then it's how we might set about trying to use it, and, as you refer to, what we infer from the information that's available to us. I think I'm going to look at my sporting contests in the future in a little bit of a kind of different way, Patrick. And I might have to go out and have another read of Moneyball as well. Look, it's been an absolute pleasure to have you joining us here at TAPS, and I appreciate that you're still in the process of completing your PhD so I understand that the work is in progress. Just fascinating to talk to you about some of the things that you've been able to investigate and some of the conclusions emerging from that work. We wish you well for the completion of your PhD and the continued work that you do in the sport analytics-business area. Hopefully some really nice publications down the track. So good luck, and thanks for joining us here at TAPS.

P: Brilliant. Thanks for having me, Albie.