informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–16
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

When Success Is Rare and Competitive: Learning from Others' Success and My Failure at the Speed of Formula One

Michael A. Lapré,[a,*] Candace Cravey[b]

[a] Owen Graduate School of Management, Vanderbilt University, Nashville, Tennessee 37203; [b] School of Law, University of Virginia, Charlottesville, Virginia 22903
*Corresponding author
Contact: m.lapre@vanderbilt.edu, https://orcid.org/0000-0003-2259-8739 (MAL); keu5dc@virginia.edu (CC)

Received: March 16, 2020
Revised: May 17, 2021; August 23, 2021
Accepted: August 27, 2021
Published Online in Articles in Advance: February 14, 2022

https://doi.org/10.1287/mnsc.2022.4324

Abstract. Organizations can learn from prior successes and failures to improve organizational performance. Few learning-curve studies have investigated this phenomenon at the individual level. A notable exception found that surgeons learn from their own success and others' failure. Success in surgery is common and individually independent from other surgeries. We study learning from success and failure in a context where success is rare and competitive: Formula One (F1) racing. Only one driver will win a race, preventing the other competitors from winning. Even severe failures causing drivers to abandon the race are common. We investigate two types of abandonments: car failures and driver failures. Our data set covers F1 from the start of F1 in 1950 through 2017, yielding 21,487 driver-race observations. We find that win probability follows an inverted U-shaped function of racing experience. We also find that drivers learn from their own success, teammates' success, as well as own car failures. However, drivers do not learn from their own driver failures. A teammate's win increases the probability of winning the next race by 1.8%. An own car failure increases the probability of winning the next race by 1.9%. We use two characteristics of success, frequency and competitiveness, to define a spectrum of organizational settings. Placement of our F1 findings and the surgery findings on this spectrum reveals when managers can expect benefits from their own versus others' success and failure.

History: Accepted by Charles Corbett, operations management.

Supplemental Material: The online appendix and data are available at https://doi.org/10.1287/mnsc.2022.4324.

Keywords: learning curve • quality • failure learning • Formula One racing • industry study

1. Introduction

> I had matured into an analytical professor of information relating to the conditions of the track, my physical well-being and the car, to a point where I could clearly define issues out of a blur into crystal clear focus in milliseconds.
> —Jackie Stewart, Three-time Formula One World Champion

> I think crashes are necessary for the career of any racing driver. What matters is whether you learn your lesson from them or not.
> —Niki Lauda, Three-time Formula One World Champion

Organizations can improve performance as a function of experience. This learning-curve phenomenon has been observed in many contexts for performance measures such as cost, quality, and survival (Lapré and Nembhard 2010, Argote 2013). Interestingly, organizations, teams, and individuals show tremendous variation in the rate with which they learn. Understanding what contributes to the variation in learning rates remains an active field of inquiry. Recently, scholars have distinguished successful experience from failure experience. For example, do defective outputs and good outputs contribute equally to learning by doing (Lapré and Nembhard 2010, Dahlin et al. 2018)? Most of the studies on learning from success and failure are at the organizational level. However, to better understand organizational learning, a better understanding of individual learning is required.

A notable exception at the individual level is KC et al. (2013). The authors study patient mortality in cardiac surgery, where patient survival determines the distinction between success and failure. KC et al. (2013) find that surgeons learn from their own prior success, but own prior failures worsen surgeon performance. In contrast, surgeons learn from other surgeons' failures (in the same hospital), but not from other surgeons' success. The authors explain that surgeons attribute success to themselves but failure to external factors. The authors pose the question of how their findings will hold up in contexts where success is rare as opposed to common. Research on failure has focused on settings where success is common (Haunschild and Sullivan 2002). In cardiac surgery, fortunately, success is common, whereas in contexts such as sports, auctions, and bidding for contracts, success is both rare and competitive. Winning prevents opponents from winning. In contrast, surgeons do not attempt to increase the mortality rate for other surgeons.

In this paper, we ask what is the role of learning from success and failure from own experience and others' experience when success is rare and competitive. We study individual learning from success and failure in the context of Formula One (F1) racing. Only one driver will win a Grand Prix race, thereby preventing all other drivers from winning that race—essentially a zero-sum game. So, in F1, success is both rare and competitive. In F1, even severe failures forcing drivers to abandon the race occur with high frequency. Eichenberger and Stadelmann (2009) classify abandonments in F1 as equipment dropouts or human dropouts. Accidents can have different causes, and therefore learning can be greater from some accidents than others (Baum and Dahlin 2007). Hence, it is worthwhile to investigate learning from different types of failures (Haunschild and Sullivan 2002, Baum and Dahlin 2007). In this paper, we study learning from both types of severe failures in F1: car failures and driver failures.

2. Learning in Formula One

2.1. The Context

Formula One (F1) is a series of Grand Prix races on unique tracks throughout the world. F1 has been held every year since the inaugural season in 1950 (Smith 2016). Some races are held on occasional tracks, such as the Monaco Grand Prix, which uses regular streets, whereas most races are held on permanent race tracks, such as Monza in Italy or the Circuit of the Americas in Austin, TX. A few races are held on semipermanent tracks combining regular streets with a permanent track, such as Sochi in Russia. A typical race weekend has 20 or more drivers. First, drivers use prerace practice sessions to "set up" the car. This is a complicated process unique to each track, because tracks differ tremendously in terms of straights and corners (Lauda 1977). A car setup involves many trade-offs; for example, a change in a wing setting could reduce straight-line speed but improve cornering speed. Next, prerace qualifying session(s) determine the order of the cars on the starting grid for the race. In the actual race, the drivers with the top finishing positions score points. At the conclusion of the race, the top three finishers receive trophies at a podium ceremony. The driver with the most points at the end of the season is crowned the drivers' world champion. In F1, a team (formally "entrant") is the entity who registers their car(s) and driver(s) for a race. Typically, teams enter two cars. Teams prepare and maintain the cars during the race weekend. Cars are designed and built by constructors. Cars differ significantly across constructors. In the early decades, there were both "works teams" who constructed their own cars and "privateer teams" who bought cars from constructors. Since 1981, privateers are no longer allowed to enter a race. So, nowadays, "team" and "constructor" are synonymous. In 1958, F1 added a constructors' world championship (Smith 2016). Cars with the top finishing positions score points counting toward the constructors' world championship. So, a driver win also counts as a team win. Both the drivers' world championship and the constructors' world championship are prestigious and financially important for F1 teams (Stewart 2009). F1 has become a big business with global appeal. In 2016, U.S. firm Liberty Media bought F1 for \$4.4 billion. In 2017, F1 enjoyed 352 million unique viewers worldwide (formula1.com).

As Gino and Pisano (2011, p. 70) note, "Racing may seem a long way from the world of business, but in fact it provides a perfect laboratory for research on learning." Race results are objective measures of performance, so drivers always know how their longitudinal performance stacks up against the competition. Since only one driver can win a race, racing is a great setting to study zero-sum competition in tournaments (Bothner et al. 2007). Recently, F1 has been used for empirical research (Aversa et al. 2015). Castellucci et al. (2011) studied the impact of aging on points scored by F1 drivers. The authors found an inverted U-shaped relationship between age and productivity. Eichenberger and Stadelmann (2009) and Bell et al. (2016) built empirical models to find rankings of the best F1 drivers. Castellucci and Ertug (2010) studied the role of status in exchange relationships. The authors found that high-status F1 teams can secure more engine modifications and redesigns from their engine suppliers than low-status teams can. Studying escalation of competition, Piezunka et al. (2018) found that F1 drivers of similar status (in terms of characteristics such as age and points scored) are more likely to collide. None of these papers study learning curves in F1. One notable exception is the work of Mourão (2017), who found that own prior success increases the probability of winning a race. However, Mourão (2017) did not study learning from teammates' success, own failures, or teammates' failures. Mourão (2018) found that podium positions prolong F1 driver careers, whereas failures to finish races shorten F1 driver careers. These findings further substantiate the importance of studying learning from success and failure in F1.

In F1, learning is imperative. Each year, teams introduce new cars. During the year, teams and drivers develop their cars. For each race, they need to learn the best car setup for the driver-car-track combination. Performance in a race depends critically on this driver-car-track specific setup of the car. For example, "[no] two Grand Prix tracks call for the same optimum transmission gearing" (Lauda 1977, p. 66). In the early decades of F1, two-time world champion Graham Hill "used to keep a small black book, in which he recorded every race and lap time, every mechanical detail of the cars he drove – spring ratings, valve settings on the dampers, every roll stiffness. Nothing was omitted" (Stewart 2009, p. 143). Nowadays, F1 cars use 140 sensors yielding 15 gigabytes of data each race weekend. Regardless of the era in F1, every race weekend, each team is figuring out the optimal car setup to maximize performance of the driver-car-track combination (Stewart 2009).

2.2. Hypotheses

Our F1 performance measure is the probability of winning a Grand Prix race. Winning races is the ultimate goal for F1 drivers (Prost 1990, Aversa et al. 2015). Three-time world champion Jackie Stewart describes the importance of winning:

> But for most sporting champions even second place is regarded as just another form of losing. If anybody offered me a million pounds today to tell them how many times I finished second in a Grand Prix, I wouldn't know the answer because it doesn't matter. On the other hand, I remember each of my wins. (Stewart 2009, pp. 52–53)

Winning a race is a measure of performance relative to all the other drivers in a race. The U-shaped learning-curve literature has studied relative measures of performance at the organizational level (Ingram and Baum 1997, Baum and Ingram 1998, Ingram and Simons 2002, Lapré and Tsikriktsis 2006). This literature expects absolute measures of performance such as efficiency to continue to improve as a function of experience. Yet, building on the notion of competency traps (Levitt and March 1988), this literature argues that relative measures of performance—such as organizational survival, profitability, and customer dissatisfaction—will follow a U-shaped function of experience. Initially, relative performance improves as a function of operating experience as organizations learn to perfect their routines. However, in the long run, routines can become obsolete as the organizational environment changes. Baum and Ingram (1998), for example, documented U-shaped learning curves for organizational failure in the hotel industry. Similarly, Lapré and Tsikriktsis (2006) observed U-shaped learning curves for customer dissatisfaction in the airline industry.

As an organization consists of individuals, we expect a U-shaped relationship between experience and relative measures of performance for individuals as well. Haltiwanger et al. (1999), for example, documented lower levels of individual sales per employee in firms with higher proportions of older workers (above the age of 55). So, at higher ages, individual performance can decline. In F1, Castellucci et al. (2011) found that points scored by F1 drivers followed an inverted U-shaped function of age. Even though age had a positive effect on points scored, age squared had a negative effect on points scored. Since learning curves capture performance as a function of experience—as opposed to age (Argote 2013)—we expect win probability to follow an inverted U-shaped function of racing experience.

Several reasons contribute to the eventual decline in win probability in F1. Graham Hill won world championship titles in 1962 and 1968. He won 14 races between 1962 and 1969. In the U.S. Grand Prix of 1969, he had a bad accident and suffered serious leg injuries (Smith 2016). Although he would recover and continue to race until retiring from F1 in 1975, he never won another race again. Many successful, experienced drivers end up with less competitive teams at the end of their careers for various reasons. Emerson Fittipaldi, for example, won world championship titles with Lotus in 1972 and McLaren in 1974, but, by 1975, "his commitment and input weren't always at his usual level" (Ménard et al. 2006, p. 33). In 1976, Fittipaldi chose to race for a less-competitive team that he had founded himself with his brother. Whereas Fittipaldi had won 14 races from 1970 through 1975, he would not win a single race from 1976 until the end of his career in 1980. Damon Hill peaked in 1996, the year he won the world championship driving for Williams. In his 67th race, he notched his 21st career win in the last race of 1996. At the end of 1996, Williams informed Hill that he was not being retained (Ménard et al. 2006). In his final three years, he drove for less competitive teams Arrows and Jordan, winning only one race. He retired from F1 after his 115th race at the end of 1999. Because of factors including declining driver skills, injuries, and moving to less competitive teams at the end of a career—voluntary or involuntary—we hypothesize the following.

Hypothesis 1. *For F1 drivers, win probability follows an inverted U-shaped function of racing experience.*

KC et al. (2013) found that surgeons do not learn from other surgeons' success in the same hospital. We argue the opposite in F1, where success is rare. When success is rare—and consequently failure is common—success by others becomes beneficial to learn from (Baum and Dahlin 2007). Any learning-by-doing effect builds on repetition. Setting up a car requires experimentation with many variables (gear ratios, camber and toe for tires, suspension settings, etc.). Variability in car settings can create noise, which makes it hard to learn from experiments (Bohn 1995). The simplest way to overcome this problem is to increase the sample size of useful observations (Bohn and Lapré 2011). Success by others can augment own success to accumulate useful experience more quickly. So, when success is rare, success from others increases the number of successes that an individual can learn from.

Learning from others, however, is not straightforward. Lapré et al. (2000) found that transfer of a successful practice within the same organization requires both a proven practice (in F1 the setup of a winning car) and a causal understanding of the principles behind the practice. Cars are very different from team to team. However, a teammate's car is identical, meaning it has the same chassis (with its many components), the same engine, and the same tires. Moreover, all components originate from the same suppliers. Hence, the engineering principles used in the setup (Lauda 1977) of a teammate's winning car have high validity because the teammate's success was achieved with the same car.

Will teammates actually share winning car setups? In F1, teams have a major incentive to share such knowledge because points scored by both drivers count toward the constructors' world championship. Yet, each driver wants to win races themselves. So, are drivers on a team cooperative or competitive? We provide both driver and team anecdotes of cooperation on F1 teams. Jackie Stewart described knowledge sharing with his teammate François Cevert:

> I always made a conscious effort to help [François Cevert] develop, and from the start we discussed everything. The relationship worked on both sides. . . . We drove essentially the same car in 1971, 1972 and 1973, and, in testing, practice, qualifying and before the race, I was happy to share everything with him, from my overall strategy to the gear ratios I would use, to the gears I was planning to take at each and every corner, and all my braking distances. There were no secrets. (Stewart 2009, pp. 269–270)

The Austrian Grand Prix of 1986 illustrates the benefits of working together. Keke Rosberg and Alain Prost were the two drivers for McLaren. In a divide-and-conquer strategy, during practice Rosberg tried out the traditional settings, whereas Prost experimented with some new settings. During the warm-up before the race, Prost was unhappy with the new setup of his car. The team changed the setup for Prost's car in a classic manner, approximating the settings used by Rosberg (Prost 1990). The car behaved perfectly, and Prost won the race.

Teams want drivers to cooperate because of the constructors' world championship. Drivers can clearly benefit from cooperation, as the Rosberg-Prost example shows. However, each driver also wants to win. So, for teams, keeping a spirit of knowledge sharing while still allowing drivers to race for individual success is key. Across four decades, Ross Brawn worked for several F1 teams in roles ranging from engineer to technical director to team principal. In these roles, he contributed to 10 drivers' world championships (seven with Michael Schumacher) and 10 constructors' world championships for Williams, Benetton, Ferrari, and Brawn (his own team). When asked about how to handle the psychological battle between two teammates, Brawn replied:

> You can't avoid it completely, but the blatant and ugly stuff I think I generally managed to avoid. It can be destructive to the team, because it can seep into the mechanics, into the engineers. I always wanted a competitive spirit between all the crew, but it was a balancing act of then pulling them back together and saying we are all in this together. This is one team. So if you do something to benefit your driver at the expense of the other driver, that's unacceptable. (Brawn and Parr 2017, p. 36)

Based on the sample-size argument, the engineering principles behind a teammate's winning car setup, and the motivation to share knowledge within the team, we hypothesize the following.

Hypothesis 2. *In F1, drivers learn from teammates' prior successes.*

KC et al. (2013) argued that surgeons do not learn from their own failures due to attribution theory—individuals attribute failures to external factors. Fortunately, in healthcare, success is common and failure is rare. In F1, on the other hand, failure is common. Moreover, F1 allows us to distinguish between two different types of reasons causing drivers to abandon a race: car failures, such as blown engines and broken suspensions, and driver failures, such as accidents and collisions. Consequently, F1 provides a context to address the call for research in which we can differentiate failures (Baum and Dahlin 2007). Dahlin et al. (2018) provide an excellent review of the literature on failure learning. The authors identify three mechanisms to learn from failure: opportunity, motivation, and ability. Failure learning is difficult, as it requires all three mechanisms (Dahlin et al. 2018).

The opportunity to learn from failure comes from the information about previous failures. Failures that are larger in magnitude, more frequent, and salient have greater information content and thus provide learning opportunities (Chuang and Baum 2003, Baum and Dahlin 2007, Madsen and Desai 2010, Dahlin et al. 2018). Car failures are significant, as they are costly both in terms of money and sporting losses. Car failures are frequent. From 1950 to 2017, only 4.5% of all driver-race observations represent wins, whereas 30% of all driver-race observations represent car failures. Car failures are salient, as they prohibit a driver from competing. As the adage in F1 goes—to finish first, first you've got to finish. Clearly, in F1, car failures are large, frequent, and salient and provide learning opportunities. Complex problems provide a greater opportunity to learn from the information about prior failures (Stan and Vermeulen 2013). An F1 car is very complex. Car failures have many different causes: engine, gearbox, transmission, suspension, tires, brakes, and so on. Heterogeneous causes benefit learning about complex systems, resulting in more in-depth analysis of underlying causes as opposed to blaming an operator (Haunschild and Sullivan 2002). Thus, own car failures allow for in-depth root-cause analysis. Car failures can often be traced back to engineering, manufacturing, machining, assembly, or suppliers (Stewart 2009). For example, Stewart had to retire from the Monaco Grand Prix in 1969 because the drive shaft failed. "An investigation subsequently revealed the problem was directly caused by a bad batch of universal couplings that had slipped through the outside supplier's inspection" (Stewart 2009, p. 202).

Motivation to learn from failure is the desire to invest in reducing failure frequency (Dahlin et al. 2018). A failure is the result of performance falling short of an aspirational level. A shortfall in performance triggers a search for solutions (Baum and Dahlin 2007). When success is rare and competitive, many failures can intensify the search for solutions to become competitive because competitors aspire to win. Car failures "are frustrating, but they happen. You have to accept it, resolve it, regroup and move on" (Stewart 2009, p. 263). When own success experience is rare, the only substantial own experience left to analyze is failure experience. Attributing failures to external factors in order to avoid blame can hinder learning from own failure (KC et al. 2013, Desai 2015). However, as we have mentioned, own car failures can be investigated with objective root-cause analysis, preventing the driver from being blamed. Attributing failure to the car as opposed to the driver enhances the willingness to learn from car failures.

Ability to learn from failure refers to identifying failure, understanding failure, and implementing solutions to prevent future failures (Dahlin et al. 2018). In F1, identifying failure is trivial. F1 teams continuously monitor their cars and will immediately notice if a driver abandons a race due to a car failure. After-event reviews enhance failure understanding (Ellis et al. 2006). In F1, there are both comprehensive reviews an hour or so after the race as well as race debriefs on the day after the race (Brawn and Parr 2017). Next, we illustrate how F1 drivers are able to learn from car failures in the closing stages of a race. Three-time world champion Niki Lauda explains:

> You should only drive as fast as is necessary to win. If I am lying first and [my pit crew] show me + 10 [seconds lead over the driver in second place] and I still have ten laps to go, then I will drop back about one second per lap . . . and that is what the others also do if they are in front, particularly the really cool and intelligent drivers like Fittipaldi. What looks like the "drama" of the finish is often nothing more than the leader's deliberating slackening off, to spare his own car. (Lauda 1977, p. 197)

Toward the end of a race, the driver lying first should focus on bringing the car home for the win. The leader should not unnecessarily push a complex F1 car with all its components to the limit. Any component is susceptible to failure. In the closing stages of the race, the leader can effectively take the time to reflect on past car failures, listen for any unusual sounds indicative of any potential car problems, and "nurse the car home" (Lauda 1977). Since own car failures provide opportunity, motivation, and ability to learn from prior failures, we hypothesize the following.

Hypothesis 3. *In F1, drivers learn from their own prior* car *failures.*

Next, we compare learning from abandonments due to driver failures versus car failures. Driver failures are less frequent than car failures. Of all driver-race observations, 12% are driver failures, whereas 30% are car failures. So, lower frequency of driver failures reduces leaning opportunities.

If failures are less concentrated and more broadly dispersed, then there will be a search for more thorough knowledge regarding causal and contributing factors (Desai 2015). The distribution of car failures is less concentrated, as there are many possible reasons for car failures (engine, gearbox, transmission, suspension, brakes, tires, etc.). On the other hand, the distribution of driver failures causing abandonment is more concentrated. Single-driver accidents and multidriver collisions are typically caused by driver error or driving style. Hence, the difference in distribution of failures reduces learning from driver failures.

Heterogeneous causes in car failures can shift focus away from blaming the driver toward in-depth root-cause analysis (Haunschild and Sullivan 2002). Conversely, driver failures consist of accidents and collisions. In multidriver collisions, a driver can place blame with another driver. In single-driver accidents, drivers may at least partially attribute accidents to external factors. "In motor racing . . . , you often hear people complaining about their bad luck" (Stewart 2009, p. 202). Consequently, as drivers attribute driver failures to external factors, drivers reduce their willingness to learn from their own driver failures. Dahlin et al. (2018, p. 270) note that "using external attributions to avoid altering ones' method of working is an individual-level defense mechanism demonstrating both lack of motivation and inability to learn."

Piezunka et al. (2018) found that F1 drivers of similar status are more likely to collide. The authors quote Damon Hill: "If I am pushed, I will push back, that is the way I am. I am very British. We don't like to be pushed around. When the chips are down we might have to step into grey areas" (Piezunka et al. 2018, p. E3362). The escalation of competition into conflict among F1 drivers suggests that F1 drivers have less motivation to learn from collisions. Given the lower levels of opportunity and motivation to learn from driver failures compared with car failures, we hypothesize the following.

Hypothesis 4. *In F1, drivers learn less from their own prior* driver *failures than from their own prior* car *failures.*

3. Data and Method

We collected data on all F1 Grand Prix races from the start of F1 in 1950 through 2017. Our main data sources were formula1.com, Ménard et al. (2006), and Wikipedia: WikiProject Formula One. All three sources document for each race, the starting position for each driver on the grid and the race result, which is the finishing position or a "did not finish" (DNF) if the driver did not finish the race. Eichenberger and Stadelmann (2009) classify DNFs as human dropouts or technical dropouts. Similarly, we use the DNF reasons and race reports documented on statsf1.com, race-database.com, grandprix.com, Wikipedia: WikiProject Formula One, and Ménard et al. (2006) to classify each DNF as Driver DNF or Car DNF. (We describe the classification procedure in the online appendix.) A Driver DNF is typically caused by an accident or a collision: 95.7% of Driver DNF observations are single-driver accidents or multidriver collisions. The few remaining Driver DNF observations concern disqualifications (e.g., for ignoring a flag) or drivers being physically unfit to continue a race (e.g., due to exhaustion). A Car DNF is the result of a failure due to a broken engine, gearbox, transmission, suspension, tire, brake, and so on. Unlike Driver DNF, there are many possible causes for a Car DNF. After a car failure, cars are fixed before the next race, either by repairing/replacing the broken part or by rebuilding the car.

From 1950 through 1960, the Indy 500 race counted toward F1, even though the Indy 500 was the only race not held under FIA rules (the FIA is the governing body of F1). Regular F1 drivers ignored the Indy 500, and Indy 500 drivers ignored the regular F1 races. The Indy 500 results were essentially irrelevant for the F1 championship (Smith 2016). Excluding the eleven Indy 500 races, our data set includes 965 races and 656 drivers for a total of 21,616 driver-race observations.

In the early years of F1, drivers would sometimes "share a drive." At some point during a race, a driver would come into the pits and turn the car over to another driver. This drive-share practice typically happened during long, hot races when drivers would get tired. From 1950 through 1957, drivers who shared a drive and finished in the points, shared the points. In 1958, drivers could no longer score points with drive shares. Teams quickly stopped sharing drives. Drive shares account for 0.7% of our observations. For 129 of these observations, drivers drove multiple cars. If a driver crashed, then the driver could take over a teammate's car and get another chance in the same race. Limiting our data set to drivers who only drove one car in a race, we end up with 965 races and 655 drivers for a total of 21,487 driver-race observations. In robustness tests excluding all drive shares, none of our findings changed.

Out of all 21,487 observations, 29.8% are Car DNF observations and 12.1% are Driver DNF observations. In contrast, only 4.5% of the observations represent wins. Clearly, in F1, success is rare and severe failure is common.

3.1. Variables

Our main dependent variable is $Win_{dr} = 1$ if driver d won race r, and 0 otherwise. We also use an alternative dependent variable $Podium_{dr} = 1$ if driver d finished first, second, or third in race r, and 0 otherwise.

Our main independent variables measure different types of experience. Let $Cumulative\ Races_{dr}$ be the number of races started by driver d prior to race r. We include $(Cumulative\ Races_{dr})^2$ to test for an inverted U-shaped learning-curve effect. To capture own success experience, we define $Cumulative\ Wins_{dr}$ as the number of races won by driver d prior to race r. Similarly, we calculate $Cumulative\ Car\ DNF_{dr}$ and $Cumulative\ Driver\ DNF_{dr}$ to capture own failure experiences. We measure teammates' success experience with $Cumulative\ Teammate\ Wins_{dr}$ which sums the wins by drivers who at the time of the win were on the same team as driver d prior to race r. Likewise, we calculate $Cumulative\ Teammate\ Car\ DNF_{dr}$ and $Cumulative\ Teammate\ Driver\ DNF_{dr}$ to measure teammates' failure experiences. At the end of a season, several drivers might change teams. As an example of driver churn, the appendix shows the teams and teammates for Alain Prost's career. The example illustrates which teammates contributed to Prost's cumulative teammate experiences. We define cumulative podium variables analogous to the cumulative win variables.

We use several control variables. Let $Grid\ Position_{dr}$ be the position for driver d on the grid at the start of race r (1 for the first driver, 2 for the second, etc.). The control variable *Grid Position* is determined by prerace qualifying session(s). The faster the driver-car combination on the track, the better the *Grid Position*. There is variation in driver ability and car performance both within a race and across races. Note that *Grid Position* is a succinct control variable for how fast the driver-car-track combination is (Piezunka et al. 2018). The combination of these three elements (driver, car, and track) varies from observation to observation, and this variation can be very different from merely three fixed effects held constant for the entire data set. First, consider a driver moving from a great team to a mediocre team. In 1996, Damon Hill drove for Williams—the best team in 1996. Hill won the world championship, and his teammate Jacques Villeneuve was the runner-up. In 16 races, Hill's average grid position was 1.44. In 12 classified finishes (covering 90% of the race distance), his average finishing position was 1.75 (including eight wins). In 1997, in stark contrast, Hill drove for Arrows, a team which had never won a race. His average grid position was 11.8. He had six DNFs. In 10 classified finishes, his average finishing position was 9.3 (no wins). Grid position controls for the car-driver differences between Hill driving for Williams in 1996 versus Hill driving for Arrows in 1997. Moreover, teams' competitive performance varies significantly as well. Since 2005, Williams has only won a single race (in 2012). Second, consider variations in car-track combinations. Some cars perform comparatively better on certain tracks. In 1980, Renault was the only team with turbo engines. Three races were held at altitude (Brazil, South Africa, and Austria), where Renault had an advantage (Ménard et al. 2006). The average starting position for the two Renault drivers, Jabouille and Arnoux, at the three altitude races was 2.2 compared with 8.5 at all other tracks. The two Renault drivers combined for three wins in 1980, all obtained at the three altitude tracks. At the other tracks, Renault had more DNFs than classified finishes. Grid position controls for the differential car-track performance of the 1980 Renault at altitude versus sea level. In subsequent years, other teams switched to turbo engines. By 1985, every team used turbo engines and Renault's altitude advantage had vanished.

Using dates of birth and race dates, we determine Age_{dr} as the age of driver d at the time of race r (Castellucci et al. 2011). We define $Home\ Edge_{dr} = 1$ if race r was held in the home country of driver d, and 0 otherwise (Castellucci et al. 2011). Finally, we use race controls: X_{ir} is the value for the ith control variable for race r. The race controls are wet weather conditions, occasional track, permanent track, and number of drivers starting (Castellucci et al. 2011). Both Smith (2016) and Wikipedia document the data for all of these race controls.

For robustness tests, we introduce several variables. We create individual driver dummy variables for each of the drivers who won three or more world championship titles. These triple world champions are considered the greatest in the sport—Juan Manuel Fangio, Jack Brabham, Jackie Stewart, Niki Lauda, Nelson Piquet, Ayrton Senna, Alain Prost, Michael Schumacher, Sebastian Vettel and Lewis Hamilton. To measure failures by competitors, we calculate *Cumulative Competitor Driver* DNF_{dr} as the number of driver failures by other teams' drivers in races started by driver d prior to race r. In the online appendix, we explain how we code the Driver DNF observations as single-driver accidents and multidriver collisions. Subsequently, we create *Cumulative Single-driver DNF* and *Cumulative Multidriver DNF* variables for the focal driver, teammates, and competitors. In the online appendix, we define additional variables for robustness tests. Table 1 shows the summary statistics.

3.2. Methodology

Following KC et al. (2013) and Clark et al. (2018), we estimate learning-curve effects in a logistic regression framework. As these authors note, this log-linear form is identical to the theoretically derived learning-curve model by Lapré et al. (2000) and Lapré and Tsikriktsis (2006):

$$\begin{aligned}\ln\frac{Pr(Win_{dr}=1)}{1-Pr(Win_{dr}=1)} &= \alpha_0+\alpha_1 Grid\ Position_{dr}\\ &+\alpha_2 Age_{dr}+\alpha_3 Home\ Edge_{dr}\\ &+\sum_i \alpha_{ir}X_{ir}+\beta Cumulative\ Races_{dr}\\ &+\gamma(Cumulative\ Races_{dr})^2+e_{dr}.\end{aligned}$$

A positive value for β and a negative value for γ would support Hypothesis 1. In the full model, we replace *Cumulative Races* with own and teammates' success and failure experience:

$$\begin{aligned}\ln\frac{Pr(Win_{dr}=1)}{1-Pr(Win_{dr}=1)} &= \alpha_0+\alpha_1 Grid\ Position_{dr}+\alpha_2 Age_{dr}\\ &+\alpha_3 Home\ Edge_{dr}+\sum_i \alpha_{ir}X_{ir}\\ &+\beta_1 Cumulative\ Wins_{dr}\\ &+\beta_2 Cumulative\ Teammate\ Wins_{dr}\\ &+\beta_3 Cumulative\ Car\ DNF_{dr}\\ &+\beta_4 Cumulative\ Driver\ DNF_{dr}\\ &+\beta_5 Cumulative\ Teammate\\ &\quad Car\ DNF_{dr}\\ &+\beta_6 Cumulative\ Teammate\\ &\quad Driver\ DNF_{dr}\\ &+\gamma(Cumulative\ Races_{dr})^2+e_{dr}.\end{aligned}$$

A positive value for β_2 would support Hypothesis 2; a positive value for β_3 would support Hypothesis 3; $\beta_3 > \beta_4$ would support Hypothesis 4.

Table 1. Summary Statistics

Variable	Mean	Standard deviation	Min	Max
Win	0.0448	0.2069	0	1
Podium	0.1342	0.3406	0	1
Cumulative Races	59.85	59.86	0	321
Cumulative Wins	3.81	9.72	0	91
Cumulative Teammate Wins	3.27	6.73	0	55
Cumulative Car DNF	16.69	17.03	0	113
Cumulative Driver DNF	7.26	7.52	0	39
Cumulative Teammate Car DNF	15.69	16.29	0	92
Cumulative Teammate Driver DNF	7.30	8.27	0	45
Grid Position	11.96	6.84	1	34
Age	30.13	5.14	17.46	55.80
Home Edge	0.0828	0.2756	0	1
Wet Weather Conditions	0.1575	0.3643	0	1
Permanent Track	0.7301	0.4439	0	1
Occasional Track	0.1534	0.3604	0	1
Number of Drivers Starting	22.83	3.03	6	34

Note. $N = 21{,}487$.

In logistic regression, Hosmer et al. (2013) recommend at least 10 events per parameter to avoid overfitting. With 965 wins in our data set, Hosmer's recommendation means that we should not estimate more than 96 parameters. Thus, we cannot jointly include dummy variables for drivers, teams, and races. Fortunately, we do not have to do so. As we have explained, *Grid Position* succinctly controls for the driver-car-track performance across all observations. Furthermore, in robustness tests, we separately include dummy variables to control for several driver, team, track, and year fixed effects.

Within a single race, observations are not independent. If one driver wins the race, then the other drivers lose. There are two approaches commonly used to model correlated binary data: a random effects model (also called a cluster-specific model) and a population average model (Hosmer et al. 2013). "The cluster-specific model is most useful when the goal is to provide inferences for covariates that can change within cluster, whereas the population average model is likely to be more useful for covariates that are constant within cluster" (Hosmer et al. 2013, p. 317). Our observations are clustered by race. Most of the independent variables of interest, such as the experience variables, vary within a race, whereas some of the control variables such as the race controls are constant within a race. Therefore, we report all of our estimations obtained with the cluster-specific model. We also estimated all of our models with the population average model. Both methods yield the same findings and the same support for our hypotheses.

4. Empirical Results

Table 2 shows the logistic regression results for *Win*. The first model provides evidence for an inverted U-shaped function of *Cumulative Races*. The positive and statistically significant coefficient for *Cumulative Races* combined with the negative and statistically significant coefficient for $(\textit{Cumulative Races})^2$ supports Hypothesis 1. In the second model, the positive and statistically significant coefficient for *Cumulative Wins* is evidence that drivers learn from their own prior successes. The positive and statistically significant coefficient for *Cumulative Teammate Wins* supports Hypothesis 2. The positive and statistically significant coefficient for *Cumulative Car DNF* supports Hypothesis 3. The coefficient for *Cumulative Driver DNF* is smaller than the coefficient for *Cumulative Car DNF*. We use a Wald test to determine that the two coefficients are statistically significantly different from each other. The Wald test rejects that the two coefficients are equal ($p < 0.05$). So, Hypothesis 4 is supported. The coefficient for the teammates' driver failures is positive and statistically significant, whereas the coefficient for teammates' car failures is not significant.

Our research focus is individual driver learning. In F1, learning can also take place at the team level. Team fixed effects can control for team factors that can contribute to learning. We cannot include a team dummy for every team, because we would substantially exceed the recommended maximum of 96 parameters that we can estimate with just 965 wins (Hosmer et al. 2013). However, a small portion of teams wins a disproportionate share of races. There are 20 teams with five wins or more. In fact, these 20 teams have at least eight wins. With this cutoff, we have at least eight wins per team dummy variable. This is close to 10 wins per dummy variable—which is the preferred minimum (Hosmer et al. 2013). These 20 teams account for 97.6% of all wins. See the appendix for more information on the teams. Models (3) and (4) in Table 2 show the results when we control for these 20 team fixed effects. The results are robust—all four hypotheses continue to be supported.

Table 2. Logistic Regression Models for *Win*: Base Models

	(1)	(2)	(3)	(4)
Cumulative Races	0.0118***		0.0123***	
	(0.0018)		(0.0020)	
Cumulative Wins		0.0416***		0.0426***
		(0.0040)		(0.0043)
Cumulative Teammate Wins		0.0280***		0.0175*
		(0.0073)		(0.0078)
Cumulative Car DNF		0.0162**		0.0189**
		(0.0054)		(0.0062)
Cumulative Driver DNF		−0.0223		−0.0135
		(0.0116)		(0.0125)
Cumulative Teammate Car DNF		−0.0059		−0.0066
		(0.0049)		(0.0052)
Cumulative Teammate Driver DNF		0.0219*		0.0269**
		(0.0089)		(0.0095)
$(Cumulative\ Races)^2$	−0.00004***	−0.00004***	−0.00005***	−0.00006***
	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Team fixed effects	No	No	Yes	Yes
Grid Position	−0.5209***	−0.4798***	−0.4740***	−0.4388***
	(0.0243)	(0.0236)	(0.0251)	(0.0242)
Age	Yes	Yes	Yes	Yes
Home edge	Yes	Yes	Yes	Yes
Wet weather conditions	Yes	Yes	Yes	Yes
Permanent track	Yes	Yes	Yes	Yes
Occasional track	Yes	Yes	Yes	Yes
Number of drivers starting	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
Wald χ^2	642.30***	863.91***	736.97***	915.82***

Notes. $N = 21{,}487$. Standard errors adjusted for clustering on observations by race in parentheses.
* Significant at 0.05; ** at 0.01; and *** at 0.001.

For the average F1 driver, the estimates imply that the decline starts at the 122nd race, which is after the mean cumulative number of races plus one standard deviation. The inverted U-shape for experience is observed in the sample, as there are 51 drivers with more than 122 races. The average *Pr*(*Win*) starting from an average of 12th place is 0.0072. Learning from a teammate's win increases *Pr*(*Win*) by 0.00013, or 1.8%, on average. Learning from an own car failure increases *Pr*(*Win*) by 0.00014 (1.9%) on average. Note that *Grid Position* effectively controls for the performance of the driver-car-track combination. The negative and statistically significant coefficient for *Grid Position* means that the higher position number (i.e., farther away from the front) implies a lower probability of winning the race. Figure 1 illustrates that learning effects depend heavily on *Grid Position*. Starting from first place, an additional teammate's win (own car failure) can increase *Pr*(*Win*) by as much as 0.0044 (0.0047).

We conducted several additional robustness tests. Including individual dummies for all 655 drivers would substantially exceed the recommended maximum of 96 independent variables. Moreover, for 559 winless drivers representing 45% of the observations, driver dummies would perfectly predict the dependent variable. Software packages deal with perfect prediction in one of two ways: (i) drop perfectly predicted observations, or (ii) retain all variables and produce numerically unstable estimates (Hosmer et al. 2013). Either approach biases the estimates. Instead, to investigate the potential impact of driver effects, we introduce several driver-related control variables. First, we include driver dummy variables for each of the drivers who won three or more world championship titles. These exceptional drivers account for 42% of all race wins. The first model in Table 3 shows that our results are robust when we include triple world champion fixed effects. In the second model in Table 3, we include both triple world champion fixed effects as well as the aforementioned team fixed effects. Again, our results are robust. All four hypotheses continue to be supported. Second, Online Appendix Tables 9–10 show the robustness of our results when we include additional driver-related variables (driver priority on the team, quality differences between teammates, number of cars entered by the driver's team, grid penalties, and protecting the lead in the drivers' championship standings toward the end of the season).

Sporting regulations in F1 have evolved over time. For many seasons, the FIA has modified point scoring systems, qualifying procedures, and technical specifications such as engine size. To control for changes in sporting regulations, we introduce year dummies. Our results (reported in Online Appendix Table 12)

Figure 1. (Color online) Single Learning-Curve Effects

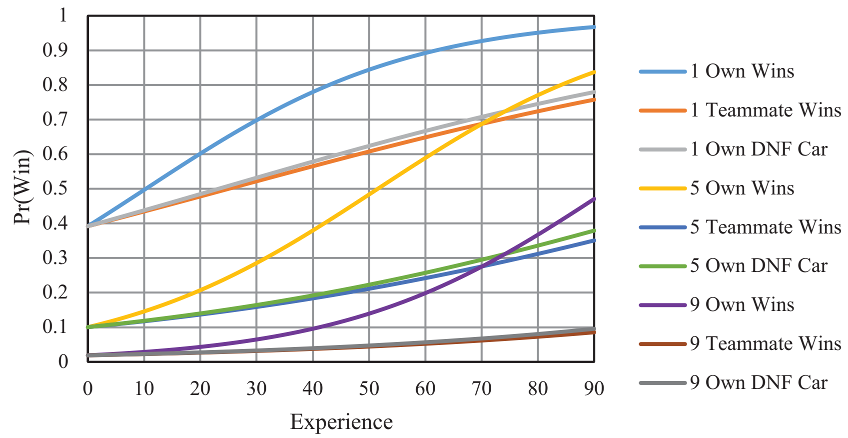

Notes. Rewriting the learning-curve model with just *Grid Position*, a single experience variable, and the estimated parameters, we get $\Pr(\text{Win}) = e^{\hat{\alpha}\ \text{Grid Position} + \hat{\beta}\ \text{Experience}} / (e^{\hat{\alpha}\ \text{Grid Position} + \hat{\beta}\ \text{Experience}} + 1)$. The *1 Own Wins* curve shows this estimated *Pr*(*Win*) as a function of *Own Wins* for a novice driver starting from first, holding all other experience variables at zero. The top (middle, bottom) three curves show such single learning-curve effects starting from first (fifth, ninth) for the *Own Wins*, *Teammate Wins*, and *Own Car DNF* experience variables. Actual increases in *Pr*(*Win*) will be less because $(\textit{Cumulative Races})^2$ has a negative coefficient.

are robust when we include year dummies. Similarly, our results are robust when we include track dummies (Online Appendix Table 13).

Models (1) and (2) in Table 3 show that drivers do not learn from their own driver failures or from teammates' driver failures. When failures are attributed to external factors, ideas from outside the organization might be necessary to change behavior (Baum and Dahlin 2007, Dahlin et al. 2018). In model (3) in Table 3, we introduce competitors' driver failures. The coefficient for *Cumulative Competitor Driver DNF* is not significant. However, we obtain additional insight in model (4), where we replace all *Driver DNF* variables with *Single-driver DNF* and *Multidriver DNF* measures. Consistent with model (3), drivers do not learn from either their own driver failures or their teammates' driver failures. Drivers do not learn from competitors' single-driver accidents either. Conversely, the coefficient for *Cumulative Competitor Multidriver DNF* is positive and significant, indicating that drivers do learn from competitors' multidriver collisions. Our hypotheses continue to be supported, even when we control for competitor failure experience and type of driver failure (single-driver accident vs. multidriver collision).

To examine whether endogeneity is a concern, we use three approaches. First, we examine whether sample selection bias is a concern as it relates to high-performing teams recruiting the highest-performing drivers to drive for their teams. The clearest example is Michael Schumacher driving for Ferrari and winning five consecutive world championship titles. There are only six pairings of the highest-performing drivers driving for the same high-performing team resulting in three or more world championships. See Table A.3. These six collaborations represent 663 out of 21,487 observations, which is just 3%. Yet, these six collaborations account for 25% of all the wins in the data set. So, could these six collaborations between the highest-performing drivers driving for the same high-performing team lead to any sample selection bias? In robustness tests, we omit these six collaborations representing 3% of our sample yet 25% of all the wins. The results reported in Online Appendix Table 14 show that our findings are robust. All of our hypotheses continue to be supported.

For the second approach to examine possible endogeneity, we follow the procedure used by Clark et al. (2018) to assess reverse causality for experience and the dependent variable. We estimate a regression model with *Cumulative Races* as the dependent variable and lagged *Win* as one of the independent variables. We also include the top team dummies and *Age*. The results reported in Online Appendix Table 18 show that lagged *Win* is not significant. Hence, racing experience does not seem to be endogenously determined by prior wins.

For the third way to examine endogeneity, we follow the approach by Muthulingam and Agrawal (2016). To break the potential mechanical relationship between the experience variables and the dependent variable, we estimate our learning-curve models with increased lags (two, three, and four races) for our experience variables. The models including both team-fixed effects and triple world-champion fixed effects are robust. The only experience variables that are positive and significant are the same as in Table 3: own wins, teammate wins, own car failures, and competitors' multidriver collisions. Moreover, the results reported in Online Appendix Tables 19 and 20 show that all of our hypotheses continue to be supported.

Table 3. Logistic Regression Models for *Win*: Driver Effects and Competitor Failures

	(1)	(2)	(3)	(4)
Cumulative Wins	0.0301***	0.0320***	0.0296***	0.0286***
	(0.0053)	(0.0060)	(0.0062)	(0.0062)
Cumulative Teammate Wins	0.0305***	0.0200*	0.0174*	0.0202*
	(0.0075)	(0.0079)	(0.0082)	(0.0084)
Cumulative Car DNF	0.0201***	0.0223***	0.0195**	0.0208**
	(0.0060)	(0.0068)	(0.0072)	(0.0071)
Cumulative Driver DNF	−0.0240	−0.0122	−0.0234	
	(0.0125)	(0.0138)	(0.0169)	
Cumulative Single-driver DNF				−0.0433
				(0.0271)
Cumulative Multidriver DNF				−0.0267
				(0.0266)
Cumulative Teammate Car DNF	−0.0074	−0.0114	−0.0117	−0.0107
	(0.0056)	(0.0063)	(0.0063)	(0.0063)
Cumulative Teammate Driver DNF	0.0180	0.0180	0.0098	
	(0.0103)	(0.0111)	(0.0130)	
Cumulative Teammate Single-driver DNF				0.0294
				(0.0243)
Cumulative Teammate Multidriver DNF				0.0010
				(0.0204)
Cumulative Competitor Driver DNF			0.0013	
			(0.0009)	
Cumulative Competitor Single-driver DNF				−0.0012
				(0.0020)
Cumulative Competitor Multidriver DNF				0.0041*
				(0.0019)
$(Cumulative\ Races)^2$	−0.00004***	−0.00005***	−0.00005***	−0.00005***
	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Triple world champion fixed effects	Yes	Yes	Yes	Yes
Team fixed effects	No	Yes	Yes	Yes
Grid Position	−0.4707***	−0.4316***	−0.4320***	−0.4303***
	(0.0239)	(0.0244)	(0.0244)	(0.0244)
Age	Yes	Yes	Yes	Yes
Home edge	Yes	Yes	Yes	Yes
Wet weather conditions	Yes	Yes	Yes	Yes
Permanent track	Yes	Yes	Yes	Yes
Occasional track	Yes	Yes	Yes	Yes
Number of drivers starting	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
Wald χ^2	913.52***	972.18***	989.44***	1097.76***

Notes. $N = 21{,}487$. Standard errors adjusted for clustering on observations by race in parentheses.
* Significant at 0.05; ** at 0.01; and *** at 0.001.

Starting a race, the ultimate goal for a driver is to win the race. Finishing on the podium is a measure of success that is less rare (13.4% vs. 4.5%). Table 4 shows the logistic regression results for *Podium*. Compared with the win results in Table 2, the podium results in models (1) and (2) provide support for Hypotheses 1, 3, and 4. Once we control for triple world champion fixed effects in addition to the 20 team fixed effects in model (3), all four hypotheses are supported. Comparing model (3) in Table 4 with model (2) in Table 3, the only different insight is that the negative coefficient for *Cumulative Driver DNF* is statistically significant. So, in the podium analysis, drivers persist in making the same failures just like in the surgery study (KC et al. 2013). In model (4), we include competitor failures, and we replace all *Driver DNF* variables with *Single-driver DNF* and *Multidriver DNF* measures. All four hypotheses continue to be supported. Furthermore, consistent with Table 3, drivers learn from competitors' multidriver collisions. Lastly, model (4) shows that drivers' persistence in making the same failures stems from repeating single-driver accidents. In Section 5, we discuss the significance of the success and failure coefficients across Tables 3 and 4.

5. Discussion and Conclusion

5.1. Discussion of Results

The U-shaped learning-curve literature has argued that relative performance measures such as organizational survival, profitability, and customer dissatisfaction follow a U-shaped function of experience (Ingram and Baum 1997, Baum and Ingram 1998, Ingram and

Table 4. Logistic Regression Models for *Podium*

	(1)	(2)	(3)	(4)
Cumulative Races	0.0087***			
	(0.0012)			
Cumulative Podiums		0.0156***	0.0111***	0.0095**
		(0.0022)	(0.0029)	(0.0030)
Cumulative Teammate Podiums		0.0055	0.0071*	0.0067*
		(0.0030)	(0.0032)	(0.0033)
Cumulative Car DNF		0.0099**	0.0131***	0.0102*
		(0.0037)	(0.0040)	(0.0042)
Cumulative Driver DNF		−0.0176*	−0.0210**	
		(0.0075)	(0.079)	
Cumulative Single-driver DNF				−0.0442**
				(0.0146)
Cumulative Multidriver DNF				−0.0273
				(0.0153)
Cumulative Teammate Car DNF		0.0039	0.0005	0.0012
		(0.0035)	(0.0039)	(0.0039)
Cumulative Teammate Driver DNF		0.0110	0.0126	
		(0.0062)	(0.0068)	
Cumulative Teammate Single-driver DNF				0.0066
				(0.0145)
Cumulative Teammate Multidriver DNF				0.0124
				(0.0120)
Cumulative Competitor Single-driver DNF				0.0004
				(0.0011)
Cumulative Competitor Multidriver DNF				0.0024*
				(0.0012)
$(Cumulative\ Races)^2$	−0.00004***	−0.00004***	−0.00003***	−0.00004***
	(0.00000)	(0.00000)	(0.00000)	(0.00001)
Triple world champion fixed effects	No	No	Yes	Yes
Team fixed effects	Yes	Yes	Yes	Yes
Grid Position	−0.2833***	−0.2734***	−0.2712***	−0.2709***
	(0.0091)	(0.0090)	(0.0091)	(0.0091)
Age	Yes	Yes	Yes	Yes
Home edge	Yes	Yes	Yes	Yes
Wet weather conditions	Yes	Yes	Yes	Yes
Permanent track	Yes	Yes	Yes	Yes
Occasional track	Yes	Yes	Yes	Yes
Number of drivers starting	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
Wald χ^2	2,042.86***	2,305.49***	2,412.99***	2,457.58***

Notes. $N = 21{,}487$. Standard errors adjusted for clustering on observations by race in parentheses.
* Significant at 0.05; ** at 0.01; and *** at 0.001.

Simons 2002, Lapré and Tsikriktsis 2006). Our study extends this body of literature by studying win probability. Whereas survival, probability, and dissatisfaction are all indicators of competitiveness over time, winning a race is an immediate outcome of a competitive contest. So, winning versus losing a race immediately captures how well all competitors fared relative to each other. In F1, we find that win probability follows an inverted U-shaped function of racing experience.

Mourão (2017) found that a win in the previous F1 race, as well as prior podiums on the same track, increased win probability. We find that F1 drivers learn from all prior wins (on all tracks) to increase win probability. For example, while leading a race in the closing stages of a race, drivers can draw from all prior win experiences such as the race tactics described by Lauda (1977). Furthermore, Mourão (2017) found that the percent of team podiums per start did not affect win probability. Percent of team podiums incorporates all drivers who have driven for the team, including all of the drivers in the past who did not overlap with the focal driver. Focusing only on overlapping teammates can help explain why we do find learning from teammates' wins, whereas the percent of team podiums did not increase win probability for Mourão (2017).

Haunschild and Sullivan (2002) note that one issue with research on failure is that failures are generally rare events. Consequently, studying frequent failures should be promising because such "research could profit from investigating whether all failures affect learning, possibly comparing across different types of failure events" (Haunschild and Sullivan 2002, p. 382). Similarly, according to Baum and Dahlin (2007), accidents vary greatly in cause and consequence. The authors suggest that it seems likely that more will be learned from some accidents than from others. Formula One racing provides an opportunity to address these calls for research on learning from frequent failures of different types. Studying different types of failures, we find different learning effects for car failures, single-driver accidents, and multidriver collisions. First, F1 drivers learn from own car failures. Heterogeneous causes make car failures amenable to objective root-cause analysis without concerns about blaming the driver. As own car failures are frequent, there is less of a need to learn from others' car failures.

Second, own single-driver accidents reduce the probability to finish on the podium. Driver failures with homogeneous causes are amenable to assigning blame. Just like in KC et al. (2013), attribution theory comes into play. Drivers might attribute single-driver accidents to factors outside of their control (Stewart 2009). As KC et al. (2013) suggest, individuals might not make an effort to analyze what went wrong. Consequently, drivers might not change their driving behavior and future performance suffers as a result. Staw (1981) calls this persistence "escalation of commitment" to flawed behavior. However, own single-driver accidents do not affect win probability. These results suggest that, in order to win, F1 drivers need to unlearn their disruptive escalation of commitment observed in the podium analysis. While the best drivers manage to do so, for many drivers this is hard. As Lauda (1977, p. 18) notes, "I think these crashes are necessary for the career of any racing driver. What matters is whether you learn your lesson from them or not. There are people who were driving five years ago in Formula III [a lower-level racing division] and half killing themselves and they're still driving there and they're still going off the road." Although F1 drivers do not learn from single-driver accidents to improve win probability, they have to at least learn to eliminate their disruptive escalation of commitment observed in the podium analysis.

Third, F1 drivers learn from competitors' multidriver collisions. F1 teams do not provide competitors with access to crashed cars for analysis, so we do not expect any learning from competitors' car failures. It could be hard to learn from competitors' single-driver accidents, which might depend on car setups, and drivers do not have access to competitors' car setups. It is, however, feasible to review film and study competitors' multidriver collisions. Not having been involved in these collisions, drivers can objectively analyze these collisions between other drivers. Avoiding a situation where a driver or a teammate could be blamed, drivers can objectively assess competitors and answer questions such as, Where should a driver be on the track compared with other cars? When trying to overtake, what gaps should you go for? How long can you stay in another driver's blind spot at very high speeds?

5.2. Success Characteristics and Learning Effects

We use two characteristics of success to define a spectrum of organizational settings: frequency of success and competitiveness. Table 5 shows the spectrum. Frequency of success ranges from rare to common. Competitiveness during events ranges from a zero-sum game to individually independent events. We illustrate the spectrum with our F1 win study, our F1 podium study, and KC et al.'s (2013) surgery study. The F1 win study falls on one end of the spectrum, where success is rare and essentially a zero-sum game. The surgery study falls on the other end of the spectrum, where success is common and individually independent of other events. In the long term, surgeons' performance should impact their ability to get subsequent case referrals. So, we do not imply that there is no competition between surgeons over an extended period of time. However, during a single surgery, other surgeons are not trying to prevent the focal surgeon from having a success, which is exactly

Table 5. Success Characteristics and Learning Effects

Frequency of success	Rare	…	Less rare	…	Common
Competitiveness during event	Zero-sum game	…	Less competitive	…	Individually independent
Study	F1 win (Table 3)	…	F1 podium (Table 4)	…	Cardiac surgery (KC et al. 2013)
Own success	+ *Own Wins*	≈	+ *Own Podiums*	≈	+ *Own Patient Success*
Others' success (in the same organization)	+ *Teammates' Wins*	≈	+ *Teammates' Podiums*	>	n.s. *Others' Patient Success*
Own failure	+ *Own Car DNF*	≈	+ *Own Car DNF*	>	− *Own Patient Failure*
	n.s. *Own Driver DNF*	>	− *Own Driver DNF*	≈	− *Own Patient Failure*
Others' failure (in the same organization)	n.s. *Teammates' Car DNF*	≈	n.s. *Teammates' Car DNF*	<	+ *Others' Patient Failure*
	n.s. *Teammates' Driver DNF*	≈	n.s. *Teammates' Driver DNF*	<	+ *Others' Patient Failure*

Notes. + (−) represents a positive (negative) and statistically significant coefficient estimate. n.s. means not significant. We flip the signs for estimates from KC et al. (2013), as the surgery study used failure (as opposed to success) as the dependent variable.

what F1 drivers are trying to do in the F1 win study. The F1 podium study falls in-between the F1 win study and the surgery study. Study placement on the spectrum allows us to gain insights in terms of observed learning effects.

Table 5 shows that learning from own success happens in all three studies. However, Table 5 indicates that learning from others' success (in the same organization) occurs for F1 win and F1 podium, but not in the surgery study. So, only when success is rare, do individuals turn to others in the same organization to increase the number of successes they can learn from.

Table 5 shows that learning from own car failures happens for F1 win and F1 podium. Furthermore, own driver failures do not affect F1 win, whereas such failures are disruptive for F1 podium. Similarly, own patient failures are disruptive in the surgery study. So, learning from own failures depends on the type of failure. Car failures with heterogeneous causes are amenable to root-cause analysis, which can lead to learning from own car failures when such failures are frequent and eliminating failure is required to be competitive (to finish first, first you've got to finish). In contrast, driver failures with homogeneous causes are amenable to assigning blame. Escalation of commitment to flawed behavior (Staw 1981) occurs toward the right of the spectrum in Table 5, where success is more common and less competitive.

F1 drivers do not learn from teammates' car failures. We suggest two possible explanations. First, in the development of Hypothesis 3 (drivers learning from own prior car failures), we mention that the leader can effectively take time to reflect on past car reliability problems, listen for any unusual sounds indicative of any potential reliability issues, and "nurse the car home" (Lauda 1977). It could be difficult to describe to a teammate exactly what unusual sounds a driver has learned from. Second, the frequency of own prior car failures is large, providing ample learning opportunities. Hence, drivers might not feel the need to learn from teammates' car failures.

5.3. Avenues for Future Research

It should be fruitful for future research to study learning from own and others' success and failure experience in other contexts and place such findings in Table 5. Is learning in pharmaceuticals, auctions, and bidding for construction contracts similar to F1 win? How does learning in a local restaurant market compare with F1 podium? Do learning effects in nonprofit organizations follow the pattern observed in KC et al. (2013)? Adding other contexts to Table 5 should help managers understand what types of experience might help or hurt organizational performance in different settings.

In Table 5, we identify two dimensions that impact learning from own and others' success and failure: frequency of success and competitiveness. Future research is needed to identify other dimensions, for example, direct observation of others' experience. In F1, drivers spend most of the race on a different part of the track than their teammates. Consequently, drivers have limited opportunity to observe teammates during a race. Conversely, in sports such as baseball or curling, competitors can directly observe others. How does direct observation change learning effects? What other dimensions impact learning effects across settings?

Our study demonstrates that differentiating failure is important to identify learning-from-failure effects. Win probability increases as a function of prior *car* failures, whereas podium probability decreases as a function of prior *driver* failures. One limitation of our study is that for driver failures, we cannot identify who caused multidriver collisions. Another limitation is that we do not know root-cause resolutions for car failures. Was the car failure caused by assembly, processes at a supplier, quality control? Future research should study the impact of root-cause analysis on learning effects. Do different types of root causes lead to different learning effects? How is learning from different root causes affected by supplier relationships, engineering, manufacturing, assembly, machining, and quality control?

Another promising avenue for future research would focus on differentiating success. In F1, each win is equally valuable. However, in industries for movies, gaming, and pharmaceuticals, not every successful product is equally profitable. Blockbusters are much more valuable than titles or drugs that make just small profits. How do learning effects change as a function of the magnitude of success?

Acknowledgments

The authors gratefully acknowledge Anna Lapré and Jessica Noble for research assistance during the data collection, cross-checking, and coding process. The authors thank department editor Charles Corbett, the associate editor, and three anonymous reviewers for the thoughtful and constructive feedback during the review process. The authors also thank Yasin Alan, Peter Haslag, Mümin Kurtuluş, Megan Lawrence, Tim Vogus, participants at the Carnegie School of Organizational Learning 2018 Conference, INFORMS 2018 Annual Meeting, POMS 2018 and 2021 Annual Conferences, as well as seminar participants at Georgetown University and Vanderbilt University for helpful comments.

Appendix

There is significant driver churn across teams from season to season. As an example, Table A.1 shows the teams and teammates for Alain Prost. Prost drove for four different teams in his career: McLaren (two stints), Renault, Ferrari, and Williams. During his career, he had 10 different teammates. For the cumulative teammate experience variables, we sum the experiences by teammates when the driver was on the same team as that teammate. For example, to calculate *Cumulative Teammate Wins*$_{dr}$ for Prost for race 10 in 1984, we sum the wins by Watson (on McLaren) in 1980, the wins by Arnoux (on Renault) in 1981 and 1982, the wins by Cheever (on Renault) in 1983, and the wins by Lauda (on McLaren) in 1984 in races 1 through 9.

Table A.1. Teams and Teammates for Alain Prost

Year	Team	Teammate	Teammate
1980	McLaren	John Watson	
1981	Renault	René Arnoux	
1982	Renault	René Arnoux	
1983	Renault	Eddie Cheever	
1984	McLaren	Niki Lauda	
1985	McLaren	Niki Lauda (races 1–13, 15–16)	John Watson (race 14)
1986	McLaren	Keke Rosberg	
1987	McLaren	Stefan Johansson	
1988	McLaren	Ayrton Senna	
1989	McLaren	Ayrton Senna	
1990	Ferrari	Nigel Mansell	
1991	Ferrari	Jean Alesi	
1992	—	—	
1993	Williams	Damon Hill	

Notes. In 1985, in qualifying for race 13, Lauda injured his wrist. In race 14, John Watson drove for the injured Lauda. It was Watson's only F1 race after retiring from F1 at the end of 1983.

Table A.2. Teams with at Least Five Wins: 1950–2017

Team	Seasons	Wins
Alfa Romeo*	1950–1951	10
Benetton*	1986–2001	27
Brabham*	1962–1992	35
Brawn*	2009	8
BRM*	1951–1977	17
Cooper*	1953–1968	12
Ferrari*	1950–2017	228
Ligier*	1976–1996	9
Lotus*	1958–1994	74
Maserati*	1950–1957	9
McLaren*	1966–2017	182
Mercedes* (1950s)	1954–1955	9
Mercedes* (2010s)	2010–2017	67
Renault* (1970s–1980s)	1977–1985	15
Renault* (2000s)	2002–2011	20
Red Bull*	2005–2017	55
Tyrrell**	1968–1998	33
Vanwall*	1954–1960	9
Walker***	1953–1970	9
Williams*	1977–2017	114

Note. The 20 teams listed combine for 942 wins in 965 races, i.e., 97.6% of all wins; 15 other teams combine for the remaining 23 wins.

* denotes a team who built their own chassis; ** Tyrrell had nine wins with a Matra (1968–1969), one win with a March (1970), and 23 wins as a works team with a Tyrrell (1971–1983); *** Privateer team Walker had four wins with a Cooper (1958–1959) and five wins with a Lotus (1960–1968).

Table A.3. Highest-Performing Drivers Driving for High-Performing Teams

Driver	Team	Stint	World champion
Jackie Stewart	Tyrrell	1968–1973	1969, 1971, 1973
Alain Prost	McLaren	1984–1989	1985, 1986, 1989
Ayrton Senna	McLaren	1988–1993	1988, 1990, 1991
Michael Schumacher	Ferrari	1996–2006	2000, 2001, 2002, 2003, 2004
Sebastian Vettel	Red Bull	2009–2014	2010, 2011, 2012, 2013
Lewis Hamilton	Mercedes	2013–2017*	2014, 2015, 2017*

*Our data set ends in 2017. Lewis Hamilton also became world champion with Mercedes in 2018, 2019, and 2020.

Table A.2 lists the 20 teams with at least five wins. These 20 teams include all 17 teams who have won the constructors' championship or would have won in the 1950–1957 period when the constructors' championship did not yet exist. The other three teams are Walker, the most successful privateer team, Renault (1977–1985), who pioneered the turbo engine and thereby revolutionized the sport, and Ligier (1976–1996), a runner-up for the constructors' championship.

References

Argote L (2013) *Organizational Learning: Creating, Retaining and Transferring Knowledge* 2nd ed. (Springer, New York).

Aversa P, Furnari S, Haefliger S (2015) Business model configurations and performance: A qualitative comparative analysis in Formula One racing, 2005–2013. *Indust. Corporate Change* 24(3):655–676.

Baum JAC, Dahlin KB (2007) Aspiration performance and railroads' patterns of learning from train wrecks and crashes. *Organ. Sci.* 18(3):368–385.

Baum JAC, Ingram P (1998) Survival-enhancing learning in the Manhattan hotel industry, 1898–1980. *Management Sci.* 44(7):996–1016.

Bell A, Smith J, Sabel CE, Jones K (2016) Formula for success: Multilevel modelling of Formula One driver and constructor performance, 1950–2014. *J. Quant. Anal. Sports* 12(2):99–112.

Bohn RE (1995) Noise and learning in semiconductor manufacturing. *Management Sci.* 41(1):31–42.

Bohn RE, Lapré MA (2011) Accelerated learning by experimentation. Jaber MY, ed. *Learning Curves: Theory, Models, and Applications* (CRC Press, Boca Raton, FL), 191–209.

Bothner MS, Kang J-H, Stuart TE (2007) Competitive crowding and risk taking in a tournament: Evidence from NASCAR racing. *Admin. Sci. Quart.* 52(2):208–247.

Brawn R, Parr A (2017) *Total Competition: Lessons in Strategy from Formula One* (Simon & Schuster, London).

Castellucci F, Ertug G (2010) What's in it for them? Advantages of higher-status partners in exchange relationships. *Acad. Management J.* 53(1):149–166.

Castellucci F, Padula M, Pica G (2011) The age-productivity gradient: Evidence from a sample of F1 drivers. *Labour Econom.* 18(4):464–473.

Chuang Y-T, Baum JAC (2003) It's all in the name: Failure-induced learning by multiunit chains. *Admin. Sci. Quart.* 48(1):33–59.

Clark JR, Kuppuswamy V, Staats BR (2018) Goal relatedness and learning: Evidence from hospitals. *Organ. Sci.* 29(1):100–117.

Dahlin KB, Chuang Y-T, Roulet TJ (2018) Opportunity, motivation, and ability to learn from failures and errors: Review, synthesis, and ways to move forward. *Acad. Management Ann.* 12(1):252–277.

Desai V (2015) Learning through the distribution of failures within an organization: Evidence from heart bypass surgery performance. *Acad. Management J.* 58(4):1032–1050.

Eichenberger R, Stadelmann D (2009) Who is the best Formula 1 driver? An economic approach to evaluating talent. *Econom. Anal. Policy* 39(3):389–406.

Ellis S, Mendel R, Nir M (2006) Learning from successful and failed experience: The moderating role of kind of after-event review. *J. Appl. Psych.* 91(3):669–680.

Gino F, Pisano GP (2011) Why leaders don't learn from success. *Harvard Bus. Rev.* 89(4):68–74.

Haltiwanger JC, Lane JI, Spletzer JR (1999) Productivity differences across employers: The roles of employer size, age, and human capital. *Amer. Econom. Rev.* 89(2):94–98.

Haunschild PR, Sullivan BN (2002) Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Admin. Sci. Quart.* 47(4):609–643.

Hosmer DW, Lemeshow S, Sturdivant RX (2013) *Applied Logistic Regression*, 3rd ed. (Wiley, Hoboken, NJ).

Ingram P, Baum JAC (1997) Opportunity and constraint: Organizations' learning from the operating and competitive experience of industries. *Strategic Management J.* 18(Summer Special Issue):75–98.

Ingram P, Simons T (2002) The transfer of experience in groups of organizations: Implications for performance and competition. *Management Sci.* 48(12):1517–1533.

KC D, Staats BR, Gino F (2013) Learning from my success and others' failure: Evidence from minimally invasive cardiac surgery. *Management Sci.* 59(11):2435–2449.

Lapré MA, Nembhard IM (2010) Inside the organizational learning curve: Understanding the organizational learning process. *Foundations Trends Tech. Inform. Oper. Management* 4(1):1–103.

Lapré MA, Tsikriktsis N (2006) Organizational learning curves for customer dissatisfaction: Heterogeneity across airlines. *Management Sci.* 52(30):352–366.

Lapré MA, Mukherjee AS, Van Wassenhove LN (2000) Behind the learning curve: Linking learning activities to waste reduction. *Management Sci.* 46(5):597–611.

Lauda N (1977) *The Art and Science of Grand Prix Driving* (Motorbooks, Osceola, WI).

Levitt F, March JG (1988) Organizational learning. *Annual Rev. Sociol.* 14:319–340.

Madsen PM, Desai V (2010) Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry. *Acad. Management J.* 53(3):451–476.

Ménard P, Cahier B, Roebuck N (2006) *The Great Encyclopedia of Formula 1*, 3rd ed. (Chronosports, St-Sulpice, Switzerland).

Mourão P (2017) *The Economics of Motorsports–The Case of Formula One* (Palgrave-Macmillan, London).

Mourão PR (2018) Surviving in the shadows—An economic and empirical discussion about the survival of the non-winning F1 drivers. *Econom. Anal. Policy* 59:54–68.

Muthulingam S, Agrawal A (2016) Does quality knowledge spillover at shared suppliers? An empirical investigation. *Manufacturing Service Oper. Management* 18(4):525–544.

Piezunka H, Lee W, Haynes R, Bothner MS (2018) Escalation of competition into conflict in competitive networks of Formula One drivers. *Proc. Natl. Acad. Sci. USA* 115(15):E3361–E3367.

Prost A (1990) *Life in the Fast Lane* (Stanley Paul, London).

Smith R (2016) *Formula 1: All the Races: The World Championship Story Race-by-Race 1950–2015* (Evro, Dorset, UK).

Stan M, Vermeulen F (2013) Selection at the gate: Difficult cases, spillovers, and organizational learning. *Organ. Sci.* 24(3):796–812.

Staw BM (1981) The escalation of commitment to a course of action. *Acad. Management Rev.* 6(4):577–587.

Stewart J (2009) *Winning Is Not Enough* (Headline, London).