Sticking with What (Barely) Worked:

A Test of Outcome Bias

Lars Lefgren
Brennan Platt
Joseph Price

Department of Economics
Brigham Young University

January 17, 2012

Abstract

Outcome bias occurs when an evaluator considers ex-post outcomes when judging whether a choice was correct, ex-ante. We formalize this cognitive bias in a simple model of distorted Bayesian updating. We then examine strategy changes made by professional football coaches. We find they are more likely to revise their strategy after a loss than a win — even for narrow losses, which are uninformative about future success. This increased revision following a loss occurs even when a loss was expected, and the offensive strategy is revised even when failure is attributable to the defense. These results are consistent with our model's predictions.

JEL Codes: C11, D03, D81, L83
Keywords: outcome bias, hindsight bias, Bayesian updating, strategy revision

1. Introduction

In a broad variety of settings, economic actors must regularly evaluate whether their current strategy is still optimal in a constantly shifting environment. Firms adjust product and pricing decisions as technology and consumer preferences change. University faculty update their research strategy in response to changes in professional norms and realized successes. Parents revise rules and incentives as their child's needs and circumstances change.

These revisions of strategy crucially rely on the actor's ability to process information about which parts of their strategy are working well and which need adjustment; yet in practice, these evaluations are not necessarily objective and dispassionate. When asking individuals to assess the appropriateness of an action, psychologists have documented an *outcome bias* in their evaluation: even when ex-ante information is identical, the action is considered more justified if the ex-post outcome was favorable. Baron and Hershey (1988) show that individuals were more likely to be critical of a medical decision when the outcome was poor, even though the objective risk of a poor outcome was the same. Thus, an outcome-biased actor is likely to revise strategies sub-optimally — failing to make needed adjustments after fortuitous successes and changing excessively after unlucky failures.

Prior empirical research has documented an outcome bias in various experimental laboratory settings. Relative to these experiments, we examine outcome bias in a high-stakes environment with subjects who evaluate their own decisions. In this setting, outcome bias can have a dramatic effect on an individual's career and earnings, and yet these individuals hold themselves accountable for success or failure outside their control. We also provide a simple theoretical model of outcome bias, which produces behavior consistent with our empirical findings. In our model actors primarily follow standard Bayesian updating, but place inordinately more weight on success/failure than other available information about future success.

Empirical examinations of how economic actors adjust strategies are complicated by the fact that it can be difficult to describe a strategy in a parsimonious fashion. Additionally, this requires panel data on the strategy and its consequences for multiple periods. We overcome these challenges by examining how NFL coaches adjust their play calling in response to past success or failure. This is a high-stakes setting in which coaches have a strong career incentive to implement a strategy each week that maximizes the probability of victory.[1] Additionally, we have detailed data from 5,661 games over 25 seasons.

The strategic decision we focus on is the offensive strategy of how frequently to pass versus run the ball. We show that teams' passing frequency is correlated from week to week. Furthermore, the team's offensive strategy is more than twice as persistent when the team wins relative to when the team loses. This simple finding can be readily explained by the fact that when the team wins, the coach receives a positive signal regarding the efficacy of his strategy. However, a closer examination of the empirical findings suggests the existence of outcome bias.

First, the persistence in pass frequency remains higher after narrow victories versus narrow losses, even though a narrow victory provides little information regarding future success. Second, teams with better past performance show greater persistence in their offensive strategy and exhibit less evidence of outcome bias. Third, coaches' decision to change strategies is equally responsive to expected and unexpected success. If coaches were optimally incorporating information they should respond only to unexpected performance. Fourth, coaches adjust their offensive strategy similarly whether their own team scores fewer points or their opponent scores

[1] Lefgren and Platt (2011) show that the probability an NFL coach is fired is strongly related to team performance in the current season, as measured by the number of games won. Romer (2006) also argues that NFL teams are useful subjects in testing firm behavior.

additional points, even though the opponent's score is more strongly affected by the coach's defensive strategy than his offensive strategy.

To provide a clearer definition of outcome bias, we develop a simple model of an individual evaluating whether to switch between two strategies. After choosing one, he observes a noisy measure (on the real line) of the quality of his choice. His goal is to maximize the probability of success, which occurs if the measure lies above zero.[2] The individual observes the success/failure (the *outcome*) as well as the continuous measure which generates it (overall *performance*); he then uses Bayes' rule to estimate the expected future success of both options, switching if the other strategy is more likely to succeed.

We incorporate outcome bias by assuming that the individual inflates the ex-ante likelihood of the outcome that actually occurred; that is, they act as if success was more likely than it really was following a successful outcome. This creates a discontinuous jump in the probability of switching strategies when comparing performance just below or just above zero, while decisions to switch will appear Bayesian for performance further from the threshold. This outcome bias can also induce an individual to switch strategies after failure due to events outside his control.

Our results suggest that NFL coaches[3] exhibit outcome bias in that they attribute excess importance to the role of their strategy in determining whether they won or lost a game. Consequently, they may switch strategies excessively after losses and not enough after wins.

---

[2] This will be natural in our empirical setting. The quality of a team's performance may be best measured by the difference between their score and the opponents, but the most salient measure of success is whether the team won or lost. Even so, this setup is relevant for a much broader variety of settings. For example, students are categorized as having failed a course based on their inability to cross a threshold level of performance. Managers may be considered unsuccessful if they do not hit predetermined targets for sales or profits. Doctors may be particularly concerned about whether a sick patient dies, despite other measures such as prolonging life or reducing pain.

[3]Our analysis cannot identify whether coaches themselves suffer this bias, or if they are responding to pressures from biased team owners or fans. Additionally, the quarterback (or other players) could cause this bias through in-game decisions despite the efforts of a Bayesian coach. Regardless of the source, our results indicate that the NFL team suffers from this bias; for ease of discussion, throughout the paper we refer to the coach as the decision maker.

Our findings suggest that outcome bias may make it difficult for economic agents to make optimal strategic choices in a variety of settings, most noticeably for decisions where the strategy barely worked (or barely failed).

2. Prior Literature

Evaluating decisions made under uncertainty is not an easy task; psychologists have documented a number of cognitive biases that can distort the evaluation.[4] For our current setting, the two most relevant are outcome bias and hindsight bias.

When judging the correctness of a decision, one typically evaluates it from the *ex-ante* position of the decision maker, asking whether it was the best choice given the information available at the time. *Outcome bias* occurs when the evaluator considers the *ex-post* outcome as well.[5] This bias was first labeled by Baron and Hershey (1988). In their study, students were given objective data on the risks of a medical procedure, and were asked to rate the correctness of a decision. The students consistently felt the decision was more justified when the outcome was successful than when it failed, even though all other information was unchanged. Positive outcomes also produced a more favorable view of gambling decisions.

Similar outcome bias has been shown in a variety of laboratory settings, rating ethically-questionable choices (Gino, Moore, and Bazerman, 2008), decisions that benefit one while causing greater harm to another (Gino, Shu, and Bazerman, 2010), decisions by a fictitious salesperson's to pursue one client over another (Marshall and Mowen, 1993), and hypothetical military decisions (Lipshitz 1989).

---

[4] Earl (1990) and Rabin (1998) survey a much larger set of cognitive limitations, describing their potential importance in economic settings.

[5] If the outcome could reveal additional information held by the decision maker, the evaluator is not considered biased for considering it. For instance, if a home buyer discovers a massive mold problem shortly after closing, he may reasonably suspect the seller had more information than she disclosed.

Ratner and Herbst (2005) take this a step further to consider how outcome bias affects future decisions. Students were given two investment options, one of which had a clearly higher expected return and was thus the initial choice of most students. After learning the realized returns, students showed outcome bias, rating their own decision more favorably when the outcome was positive. When revisiting the two options for a second round, 23% of those with bad outcomes switched to the lower average return option, while only 2% of those with good outcomes switched. Significantly, this demonstrates that outcome bias can occur in evaluating one's *own* decisions, and can distort one's *future* decisions.

A related literature considers *hindsight bias* (Fischhoff 1975), which occurs when people with knowledge of an outcome falsely believe they would have predicted that outcome. Hindsight bias has been studied extensively in the laboratory as well as in political polling (surveyed in Hawkins and Hastie, 1990). Of course, hindsight bias could easily lead to outcome bias: if the evaluator believes the outcome was inevitable, he will condemn the decision maker for not acting accordingly. For example, LaBine and LaBine (1996) asked participants to act as jurors in a hypothetical malpractice suit for a therapist of a potentially violent patient. They assessed what the therapist should have known (hindsight bias was found) and should have done (outcome bias was found).

Outcome and hindsight bias have received some limited attention from economists. Camerer, Loewenstein and Weber (1989) study whether more-informed participants in a market game can reproduce the judgments of participants with a given subset of information. Consistent with hindsight bias, the informed are swayed by their added knowledge of the outcome and tend to make worse decisions as a consequence. Outcome bias could also contribute to the *hot-hand effect* in sports, in which bettors overestimate the autocorrelation in performance; Offerman and

Sonnemans (2004) introduce two distortions to standard Bayesian updating and determine that the hot-hand distortion can effectively explain their data, while the other distortion (recency) has no significant impact.[6]

The current study offers several advantages over the preceding literature on outcome and hindsight bias. First, we study real-life decisions by experts in a high-stakes setting; our results confirm that outcome bias occurs outside the laboratory. Second, in our setting, coaches are evaluating their own decisions, rather than the decisions of others (as is common in most of the preceding literature).[7] It is easier to be critical or dismissive of the choices of others, and hence more significant to find outcome bias in self-evaluation. Third, we examine how these biases distort future decisions, which ultimately determines whether these biases actually matter. This issue is directly addressed only in Camerer, Loewenstein and Weber (1989) and Ratner and Herbst (2005). Fourth, we construct a model of decision making in which otherwise Bayesian agents exhibit outcome bias. In doing so, we translate a specific cognitive limitation into a set of theoretical predictions that are highly consistent with our empirical analysis.

Our project adds to a body of prior work in which sports competitions have provided a fertile setting for testing economic theories. For example, Pope and Schweitzer (2011) test for players' loss aversion in golf tournaments, and Card and Dahl (2011) test for loss aversion among NFL fans. Romer (2006) shows that football coaches fail to optimize when deciding whether to "go for it" on fourth down. Chiappori, Groseclose, and Levitt (2002) test whether soccer players behave as predicted by a simple "matching pennies" game when making penalty kicks. Gray and Gray (1997) test the efficiency of betting markets in the NFL. Our work is

[6] In a purely theoretical study, Dobbs (1991) shows that agents will appear to have hindsight bias if they are initially unsure of which probability distribution governs a random event. In fact, they are just updating their belief about the distribution itself.

[7] Jones, Yurak, and Frisch (1997), Ratner and Herbst (2005), Gray, Beilock, and Carr (2007) are notable exceptions.

similar in spirit to these earlier papers.

Our paper also relates to a growing literature that examines whether economic agents are held accountable for factors outside of their control. For example, Bertrand and Mullainathan (2001) show that CEOs are rewarded on the basis of share price movements that are due to factors outside of the manager's control. Wolfers (2002) presents evidence that voters are less likely to vote for incumbent politicians during poor economic circumstances, even when the difficulties cannot be attributed to the politician. Lefgren and Platt (2011) show that football franchises fire and retain coaches on the basis of team performance outside of the coach's control. Collectively, these and other papers show that individuals have difficulty assessing other people's contribution to a successful or unsuccessful outcome. Our paper builds on this literature by demonstrating how people hold *themselves* accountable for success or failure unrelated to the quality of their own decision.

3. A Model of Outcome Bias

We now present a simple model in which an agent evaluates the effectiveness of a particular strategy via Bayesian updating. We then expand the model to allow for an outcome bias. Our model has particular importance since beliefs are not empirically observable. The model allows us to explicitly identify how the underlying updating method translates into observed performance and strategic choices.

Bayesian Updating

Consider an environment in which a single decision maker (*i.e.* the coach) must select either game plan a or b. For instance, plan a could implement a pass-oriented offense while plan b implements a run-oriented offense.[8] The team then plays the game and realizes performance P, which we interpret as the difference between the final scores of a team and its opponent. Thus if $P > 0$, the team wins the game.

Performance is a random variable whose distribution depends on both the game plan and the team's (unknown) state, denoted either A or B. In state A, the composition of team personnel and opponent weaknesses makes passing more productive; in state B, running leads to a better expected outcome. We formalize this by assuming that if a team in state A uses plan a, performance P is normally distributed with mean h and variance σ^2. Using plan b in state A produces the same variance but with a mean of ℓ, where $\ell < h$.[9] In state B, the means of the performance are reversed.

Between games, the coach decides whether to change strategies; that is, he must assess whether state A or B is more likely. Suppose that prior to the latest game, the coach believed with probability $\alpha \in [0,1]$ that his team was in state A. Thus, for α near 1, the coach was certain that implementing a pass-oriented offense would yield the best expected outcome. The opposite is true for α near 0. After observing performance P, the coach uses Bayesian updating to revise this estimated probability. Also, he must account for the possibility that injuries, player development, or opponent characteristics can change the optimal strategy. We capture this by assuming that the state changes between games with probability $\rho \in [0, 1/2)$.

[8] Here, the game plan does not specify each play, which needs to be randomized to avoid exploitation by the opponent. Rather, it should be seen as the broad strategy (such as how frequently to call passing plays). The effectiveness of the game plan is mostly a question of how well it draws on the particular strengths of the team — a good fit will lead to a higher average performance. Strategic decisions (such as choosing a specific play that the opponent does not anticipate) would be reflected in the variance of performance.

[9] The assumptions of a discrete strategy choice and normally-distributed performance are made solely for expositional clarity. When we used a continuous strategy variable or alternative distributions, belief updating could only be numerically calculated, and yet it displayed nearly the same behavior as in our predictions.

Suppose the coach entered a game with prior α, used plan a, and observed performance P. The likelihood of this performance in state A and B, respectively, are:

$$\Pr(P \mid A) = \frac{e^{-\frac{1}{2\sigma^2}(P-h)^2}}{\sqrt{2\pi}\cdot\sigma} \quad \text{and} \quad \Pr(P \mid B) = \frac{e^{-\frac{1}{2\sigma^2}(P-\ell)^2}}{\sqrt{2\pi}\cdot\sigma}.$$

Thus, the estimated posterior of the *Bayesian coach* would be:

$$\begin{aligned}\hat{\alpha} &= (1-\rho)\cdot\frac{\alpha\cdot\Pr(P \mid A)}{\alpha\cdot\Pr(P \mid A)+(1-\alpha)\cdot\Pr(P \mid B)} + \rho\cdot\frac{(1-\alpha)\cdot\Pr(P \mid B)}{\alpha\cdot\Pr(P \mid A)+(1-\alpha)\cdot\Pr(P \mid B)} \\ &= \rho + \frac{\alpha\cdot(1-2\rho)}{\alpha+(1-\alpha)\cdot e^{\frac{1}{2\sigma^2}(h-l)(h+\ell-2P)}}\end{aligned}.$$

If the coach's objective is to win as frequently as possible, he should continue to employ plan a as long as $\hat{\alpha} \geq 1/2$, and switch to plan b otherwise. This simply means that a coach will employ the strategy that he believes has the highest expected performance and hence a higher probability of winning. Similar calculations arise in evaluating whether to switch from plan b.

By solving for the P where $\hat{\alpha} = 1/2$ for a given α, one obtains a performance threshold $\hat{P}(\alpha) = \frac{h+\ell}{2} + \frac{\sigma^2}{h-\ell}\ln\left(\frac{1-\alpha}{\alpha}\right)$. If performance falls below this threshold, the coach will switch plans before the next game. It is noteworthy that generally $\hat{P}(\alpha) \neq 0$; that is, changing plans will not typically hinge on whether the team wins or loses. For example, if a coach has a strong prior that he is using the right strategy, the last term will be strongly negative. He could suffer a large loss and still remain convinced that he has the right strategy. Indeed, $\hat{P}(\alpha) = 0$ only when $\alpha = 1\Big/\left(1+e^{-\left(h^2-\ell^2\right)/2\sigma^2}\right)$.

Outcome Bias

To incorporate an outcome bias, we assume that the coach over-weights the likelihood of the outcome that actually occurred by a factor $\beta \geq 1$. That is, he acts as though the realized outcome was more likely than it actually was ex-ante. If he wins after using plan a, he attributes much of this success to having correctly deduced A as the current state. After a loss, he second-guesses himself, thinking he should have known B was more likely.

Formally, a *biased coach* using plan a and observing outcome $P > 0$ obtains a posterior:

$$\begin{aligned}\tilde{\alpha} &= (1-\rho)\cdot\frac{\alpha\cdot\beta\cdot\Pr(P\mid A)}{\alpha\cdot\beta\cdot\Pr(P\mid A)+(1-\alpha)\cdot\Pr(P\mid B)}+\rho\cdot\frac{(1-\alpha)\cdot\Pr(P\mid B)}{\alpha\cdot\beta\cdot\Pr(P\mid A)+(1-\alpha)\cdot\Pr(P\mid B)}\\ &= \rho+\frac{\alpha\cdot\beta\cdot(1-2\rho)}{\alpha\cdot\beta+(1-\alpha)\cdot e^{\frac{1}{2\sigma^2}(h-l)(h+\ell-2P)}}\end{aligned}.$$

Using plan a and observing outcome $P < 0$, the coach reaches a posterior:[10]

$$\begin{aligned}\tilde{\alpha} &= (1-\rho)\cdot\frac{\alpha\cdot\Pr(P\mid A)}{\alpha\cdot\Pr(P\mid A)+(1-\alpha)\cdot\beta\cdot\Pr(P\mid B)}+\rho\cdot\frac{(1-\alpha)\cdot\beta\cdot\Pr(P\mid B)}{\alpha\cdot\Pr(P\mid A)+(1-\alpha)\cdot\beta\cdot\Pr(P\mid B)}\\ &= \rho+\frac{\alpha\cdot(1-2\rho)}{\alpha+(1-\alpha)\cdot\beta\cdot e^{\frac{1}{2\sigma^2}(h-l)(h+\ell-2P)}}\end{aligned}.$$

In this simple model, performance $P = 0$ almost never occurs, and hence can be resolved either way without loss of generality. Note, however, that as P increases, $\tilde{\alpha}$ jumps upward discontinuously at $P = 0$. This reflects the additional impact that winning has on validating the coach's strategic choice. This affects performance threshold $\tilde{P}(\alpha)$, below which the coach switches plans:

---

[10] The hot-hand distortion in Offerman and Sonnemans (2004) has some similarity to our model, in that agents overweight the ex-ante probability that the coin is unfair when using Bayes' rule. The key difference is that our coaches overweight the ex-ante probability of whichever outcome (win/loss) actually occurred, rather than always overweighting the same event.

$$\tilde{P}(\alpha) = \begin{cases} \frac{h+\ell}{2} + \frac{\sigma^2}{h-\ell}\ln\left(\frac{1-\alpha}{\alpha\beta}\right) & \text{if } \alpha > \frac{\beta}{\beta+e^{-\frac{h^2-\ell^2}{2\sigma^2}}} \\ \frac{h+\ell}{2} + \frac{\sigma^2}{h-\ell}\ln\left(\frac{(1-\alpha)\beta}{\alpha}\right) & \text{if } \alpha < \frac{1}{1+\beta e^{-\frac{h^2-\ell^2}{2\sigma^2}}} \\ 0 & \text{otherwise} \end{cases}.$$

With outcome bias, $\tilde{P}(\alpha)$ equals 0 for an interval of priors, which we call S:

$$S \equiv \left[1\Big/\left(1+\beta e^{-\frac{h^2-\ell^2}{2\sigma^2}}\right), \beta\Big/\left(\beta+e^{-\frac{h^2-\ell^2}{2\sigma^2}}\right)\right].$$

Within this subset of priors, the coach will reach a posterior $\tilde{\alpha} > 1/2$ for any $P > 0$, and reach a posterior $\tilde{\alpha} < 1/2$ for any $P < 0$. Thus, if $\alpha \in S$ then the coach switches strategies if and only if he loses the game. As β increases, the posterior increases for $P > 0$ and decreases it for $P < 0$, making the size of the discontinuity (and the interval S) larger. Intuitively, for biased coaches who are somewhat uncertain regarding the appropriate strategy, the decision to switch strategies hinges entirely on whether they won or lost the prior game.

Consequences

We now present several predictions of this model that shed light on our empirical work. Each prediction indicates how to distinguish the unobservable updating process by observing performance and strategy choices.

1. Under either Bayesian or biased updating, a coach is more likely to change strategy after worse performance. Yet under biased updating, a coach is more likely to change strategy even when comparing a narrow loss to a narrow victory.

The first claim simply comes from $\partial\hat{\alpha}/\partial P > 0$ and $\partial\tilde{\alpha}/\partial P > 0$. A worse performance result in a lower posterior, and thus a coach is more likely to switch strategies. The second claim comes from the discontinuity at $P = 0$. Any coaches who entered the game with $\alpha \in S$ will switch strategies if and only if $P < 0$. Since the interval S likely contains many coaches if β is moderately large, we would see that bias leads coaches (on average) to switch more often after losing than after winning, even when the margin of victory (or loss) was small. This would not show up for Bayesian coaches because a narrow loss only has a little more negative information (about future performance) than a narrow win, and would thus only cause the tiny fraction of coaches with priors near $\alpha = 1\Big/\left(1 + \mathrm{e}^{-\left(h^2-\ell^2\right)/2\sigma^2}\right)$ to change strategy.

It is worth noting that biased coaches are mistaken on both sides of $P = 0$. They are too complacent after a narrow victory, over-estimating $\tilde{\alpha}$, and too worrisome after a narrow loss, under-estimating $\tilde{\alpha}$. These mistakes lessen as the magnitude of performance increases (such as decisive victories or defeats).

2. Under either Bayesian or biased updating, a team with stronger past performance is less likely to change strategy after worse performance. Under biased updating, a team with stronger past performance will appear to be less biased than one with weak past performance.

Intuitively, coaches with a long history of success with a particular strategy are more confident in their strategic approach and less likely to switch strategies even in the event of a loss. More formally, past performance is encapsulated in the prior belief α, which is higher after repeated success. This has two effects on updating. First, $\partial\hat{\alpha}/\partial\alpha > 0$ and $\partial\tilde{\alpha}/\partial\alpha > 0$; which is to say that with a stronger prior and the same realized performance, the posterior will also be stronger. Because of the accumulated positive evidence, the coach's confidence in its current strategy is less shaken by a given event. Thus, the posterior is less likely to fall below $1/2$ and lead to a change of game plan.

A strong prior will also make outcome bias less visible, though. For instance, if $\alpha > \beta \Big/ \left(\beta + \mathrm{e}^{-\left(h^2-\ell^2\right)/2\sigma^2}\right)$, then $\tilde{P}(\alpha) < 0$; that is, barely losing is not enough to warrant a change in game plan. Indeed, the higher α is, the worse performance that will be tolerated without switching game plans. Yet even then, β still distorts the calculation of $\tilde{\alpha}$, as the posterior still drops discontinuously as performance falls below $P = 0$. Our point is merely that it will not fall below $1/2$ for those with high α. Empirically, coaches with a strong record will appear to use standard Bayesian updating, in contrast to those with a weak record, even if they use the same biased updating process.

3. Under Bayesian updating, only unexpected performance will affect the likelihood of changing game plans; higher expected performance has no effect. Under biased updating, unexpected performance also matters, but even expected losses can lead to game plan changes.

Coaches receive information only from unexpected events. Consequently, a Bayesian coach should only switch strategies when an approach works less well than expected. A biased coach reacts even to expected performance because it affects the probability of crossing the win threshold.

In the context of our model, we incorporate higher expected performance by assuming that both h and ℓ increase by equal amounts. That is, average team performance increases whether the team is using the right game plan or the wrong one, with the difference between them remaining the same.

When a Bayesian coach has higher expectations, he will hold his team to a higher standard: for a given performance P, the posterior $\hat{\alpha}$ is lower as h and ℓ both increase. Since the team with higher expectations should have been able to accomplish more, an unchanged performance gives the coach less confidence that he has the right game plan. In fact, it is only the *difference* between expected and realized performance that matters in the Bayesian updating. If h, ℓ, and P each increase by ε, there is literally no change in $\hat{\alpha}$. Thus, when both expected and unexpected performance are included in the same regression, the former should have no impact on the likelihood of changing strategy, while the latter will be negatively correlated with strategy changes.[11]

For a biased coach, higher expected performance operates identically on $\tilde{\alpha}$. However, this does not mean they only react to unexpected performance, as Bayesian coaches do. This is because $\partial\tilde{\alpha}/\partial\beta > 0$ if and only if $P > 0$; bias increases the posterior for wins and lowers it for losses, regardless of expectations in h and ℓ. Indeed, in the extreme case (as β becomes very large), the only factor that coaches consider in retaining their game plan is whether they won or

[11] Alternative interpretations of increased expectations (*e.g.* h increases by more than ℓ) could potentially introduce minor correlation between higher expectation and strategy changes. As the gap widens between the right and the wrong plan, the same amount of sub-par (*i.e.* unexpected) performance is interpreted more harshly. Even then, this indirect effect of expected performance would be second-order compared to the direct effect of unexpected performance.

lost, regardless of whether the outcome was expected or not. Thus, even with moderate bias, coaches will be more likely to adjust their game plan when performance is lower, even (to some extent) when that performance was expected.

4. Under Bayesian updating, a coach should not switch game plans due to events that influence the outcome but are unrelated to the plan. Under biased updating, unrelated event may induce switching, particularly in close outcomes.

Many random events contribute the final outcome of the game. In particular, a coach's offensive strategy may be effective; yet they will still lose the game if their defense cannot stop opposing drives. A Bayesian coach would not allow defensive performance to taint the evaluation of offensive strategy. Rather, performance P_o would be limited to offensive production (*i.e.* points scored), with h and ℓ reflecting the expected production under the right or wrong plan. Thus, the coach still responds to the performance measure as before, but completely ignore irrelevant factors.

A coach with outcome bias can also attempt to focus on offensive performance, but is still biased by the actual win/loss outcome. If the team only lost due to defensive errors that resulted in a high opponent score $P_d > P_o$, the coach places extra weight on the ex-ante likelihood of state B. The resulting posterior $\tilde{\alpha}$ will be lower than what a Bayesian coach would conclude, and hence the coach could switch strategies even when the offense-specific evidence is favorable. This bias will be most noticeable in close games, since factors that are irrelevant to the offensive performance are likely to still be relevant to the actual outcome.

4. Data

Before we move forward to our empirical analyses, it is crucial to identify our empirical measure of a coach's strategy. Football is a complicated sport. Indeed, the official rulebook for the NFL is 289 pages. Coaches choose between many offensive and defensive plays. Additionally, a single play can include contingencies, in which a player's actions depend on the behavior of players on the other team. In order to make empirical headway it is necessary to characterize a uni-dimensional index of strategy that is both measurable and relevant.

The measure we construct is the fraction of offensive plays (not including punts and field goal attempts) in which the team attempts to pass the ball, which we refer to as *fraction pass*. Passing is generally considered a high risk / high reward offensive approach. On average, when a team passes the ball they advance the ball further down the field than would be the case if they ran the ball (about 2.6 yards more). However, there is a higher probability of a negative event including an interception (where the team loses possession of the ball), a loss of yards if the quarterback is tackled prior to the pass, or an incomplete pass in which the ball is not advanced. Due to differences in personnel and coaching philosophy, teams differ systematically in their propensity to pass the ball. Additionally, coaches may adjust their game plan, including the planned mix of runs and passes, from week to week to adjust to their opponent or improve their offensive efficiency.

We use data on NFL team performance from the 1985 to 2009 seasons, provided by NFLData.com. Our complete sample of data includes 11,322 team-game observations, each of which indicates the fraction passing plays for the team in the game and the final score. The average team runs a passing play about 53.7% of the time though this passing rate varies from

38.5% at the 10th percentile to 68.6% at the 90th percentile and overall has a standard deviation of 11.6 percentage points.[12]

From the final score, we compute the score differential between the team and its opponent (which also indicates by its sign whether the team won or lost). In football, each touchdown is worth six points (with the opportunity to achieve an extra point, which happens most of the time). The average team scores 20.7 points per game with a standard deviation of 10.3 points. The absolute value of the score differential has a mean of 11.6 points with a standard deviation of 9.2.

To measure expected performance, we use the gambling spread for the game. The gambling spread is empirically an unbiased measure of the expected score differential (Wolfers 2006; Sauers 1998), and is constructed by gambling establishments to balance the number of people betting for or against a particular team. We define unexpected performance as the difference between the actual score differential and the gambling spread. The standard deviation of the measure of expected performance (the gambling spread) is 6.6 points, while unexpected performance has a standard deviation of 13.3 points, indicating that the majority of variation in team performance is unexpected. We also measure past team performance using the fraction of current-season games won by a team prior to the current game, though, in the cases where we use this measure in our analysis, we restrict our sample to games that occur during the second half of the season.

[12] There is likely to be some amount of measurement error in this variable since some passes may end up as sacks or as runs by the quarterback, both of which are coded in the data we use as runs. Other situations that can create measurement problems include fumbles or penalties.

5. Empirical Evidence

Sticking with What Barely Worked

Our theoretical framework suggests that when coaches suffer from biased updating, their strategic decisions will be overly sensitive to the outcome of their last game. One manifestation of this is that even though team strategy affects a continuous performance measure (the score differential), coaches will ascribe additional importance to whether this performance results in a positive outcome (victory). Hence, coaches will be more likely to adjust their strategies after a narrow loss than after a narrow victory.

In Figure 1a, we show the probability of winning the next week's game[13] as a function of the score differential in the current week. The y-axis represents the difference in probability of winning next week's game between the winner and loser of the current game. The x-axis represents the closeness of the game. For example, restricted to games determined by ten or fewer points, the winner has about a three percentage point higher probability of winning the next game than the loser. This difference is statistically significant, as shown by the dashed lines indicating the 95 percent confidence interval (based on standard errors cluster-corrected at the team-year level). For games decided by six or fewer points, the difference in the probability of winning the next game between winners and losers of the current game is only 1.8 percentage points and is statistically insignificant (p-value of 0.297). This difference is estimated using 3,368 observations.[14] This figure suggests that for games decided by less than a touchdown, winning (or losing) the game provides only weak inference regarding the team's future success.

[13] We only consider the relationship between games within the same season.

[14] Curiously, winners of games decided by one or two points are significantly more likely to win their next game. This small anomaly does not match up with the pattern for the rest of the figure, and loses significance after accounting for team strength (discussed in this subsection and illustrated in Figure 2a).

We now examine whether even for these close games, winning teams persist with their strategy at a higher rate than losing teams. To test this, we estimate the following regression:

$$frac\_pass_{i,t+1} = \beta_0 + \beta_1 frac\_pass_{i,t} * win_{i,t} + \beta_2 frac\_pass_{i,t} + \beta_3 win_{i,t} + \varepsilon_{i,t+1},$$

where $frac\_pass_{i,t}$ represents team *i*'s fraction of passing plays in game *t*, and $win_{i,t}$ is a binary variable measuring whether team *i* won game *t*. In this specification, β_2 measures the persistence of team *i*'s strategy from week *t*-1 to week *t* if the team lost, as in an AR1 process. This coefficient will be greater than zero as long as there is some persistence. β_1 shows how the persistence of the team's strategy increases after a win. Our model predicts that even for close games, where winning provides little information, β_1 will be positive if coaches exhibit outcome bias.

Figure 1b maps out β_1 for a set of games decided by no more than a particular number of points. We see that β_1 is consistently positive and relatively stable; the slight upward slope is consistent with either Bayesian or biased updating, since larger victories are taken as stronger signals that the right strategy is in use. However, under Bayesian updating, this coefficient should be statistically insignificant once winning is uninformative; but this does not hold. The only exception is as the sample size gets small and standard errors get large for games decided only by one or two points. Examining Figures 1a and Figures 1b together, we see that even for close games in which winning has no significant relationship with the next game's outcome, the impact of winning on the persistence of offensive strategy is still large and significant.

Table 2 helps us better describe the magnitude of the effect. Column 1 reports the effect for the entire sample (as a baseline), while column 2 provides estimates for just those games decided by six or fewer points. Focusing on these close games, we see that the losing team's persistence is 0.145, which means a losing team that passes ten percentage points more than average in week $t-1$ will pass 1.45 percentage points more often in week t on average. A winning team's offensive strategy persists at a rate 0.107 higher than losing teams (a 74 percent increase in persistence). Consequently, a winning team that passes ten percentage points more than average in week $t-1$, on average, passes 2.5 percentage points more often during the following week.

It may be the case that even for games decided by six or fewer points, winning still contains sufficient information to rationalize the observed increase in persistence. To address this concern, we examine games in the second half of the season in which the coaches already have a large amount of information regarding their team's strengths and weaknesses. For such coaches, winning instead of losing one additional game by a small margin should have even less information. To see that this is the case, we estimate a linear probability model in which the dependent variable is whether the team won its game the next week. The independent variables include whether the team won in the current week as well as the fraction of games won in the weeks leading up to the current week. Figure 2a shows the coefficient on the variable indicating a current week win. Consistent with our conjecture, winning a marginal game has very little information about future success (even less than in Figure 1a). For example, winning a game by six or fewer points has a slight negative, though insignificant, relationship with the probability of winning the next game.

While winning a close game in these circumstances provides very little information to coaches, they still update their strategy based on whether they win. To show this we estimate the following regression model:

$$\begin{aligned} frac\_pass_{i,t+1} = \beta_0 + \beta_1 frac\_pass_{i,t} * win_{i,t} + \beta_2 frac\_pass_{i,t} + \beta_3 win_{i,t} \\ + \beta_4 frac\_pass_{i,t} * win\_fraction_{i,t} + \beta_5 win\_fraction_{i,t} + \varepsilon_{i,t+1} \end{aligned}.$$

For ease of interpretation, $frac\_pass_{i,t}$ and $win\_fraction_{i,t}$ are centered about their means. Figure 2b shows the value of β_1 estimated on different samples. Because of the centering, this represents the mean increase in persistence associated with winning the current game. We observe a large increase in persistence that is very stable regardless of our choice of samples, again indicating that when teams win, they stick with what barely worked even if the success contains very little information. Indeed, the point estimates are nearly constant; they react almost identically to a win regardless of the margin of victory. The third and fourth columns of Table 2 quantify the magnitude of the impact. Focusing on column 4 we see that winning increases persistence by 0.163 off of a baseline of 0.143 (a 112% increase).

Another potential concern with our approach is that end-game dynamics have an effect on the observed strategy played by winning and losing teams. Teams that are behind attempt to pass more frequently so as to quickly score in a short period of time. Teams that are ahead run the ball more often as this uses more game time. Empirically, the winning team passes about 13 percentage points less frequently than the losing team. The difference is only six percentage points for games decided by six or fewer points.

Such end-game patterns do not necessarily invalidate our model or empirical approach. A team that passes more than usual to try to win from behind may update their information regarding the efficacy of passing based on the whether their comeback is successful. Indeed, the success of the end of game strategy may be the most salient information of the entire game. The only concern arises if the strategy of the winning team is systematically more informative of the team's usual passing strategy. This would be true if a losing team systematically deviates from their usual approach more than the winning team. If so, the interaction of the fraction pass with winning just indicates that the team which won didn't have to change as much as the team that lost.

If this mechanical correlation explained our results, however, we should see that the fraction pass in a win is also differentially informative regarding the strategy of the *prior* week. That is, a team that won this week also would not have changed much relative to the prior week either. We test this in Table 3. Columns (1) repeats our result from Table 2. In column (2) we estimate the following regression:

$$frac\_pass_{i,t-1} = \beta_0 + \beta_1 frac\_pass_{i,t} * win_{i,t} + \beta_2 frac\_pass_{i,t} + \beta_3 win_{i,t} + \varepsilon_{i,t-1}\,.$$

Under our updating hypothesis, β_1 should be zero. However, if our results reflect something mechanical, the coefficient would be positive and similar to our baseline results. Consistent with the updating hypothesis, the coefficient is very close to zero and statistically insignificant. Indeed, the point estimate is actually negative. This provides strong evidence that our results are not driven by a mechanical correlation induced by end-game strategic deviations but rather reflects updating on the part of coaches.

In Table 3, we also examine two other threats to validity. First, a non-linear impact of week *t*'s fraction on week *t*+1's fraction pass could induce a bias in the interaction term. In column (3), we see that this is not the case. The interaction term is very similar while squared and cubed terms in the lagged fraction pass are insignificant. Second, our results may reflect differences in strategies over time or across teams. In column (4), we include a specification with team and season fixed effects. Again, our estimates are substantively unchanged.

Seemingly Rational Winners

The second prediction of our model is that stronger teams exhibit more persistence in their strategic choice and switch less based on the outcome of a particular game. To test this hypothesis, we focus on teams in the last half of the season, for which we have a good measure of team strength (namely, their record from the first half). We run the following regression:

$$\begin{aligned} frac\_pass_{i,t+1} = {} & \beta_0 + \beta_1 frac\_pass_{i,t} * win_{i,t} * win\_fraction_{i,t} + \beta_2 frac\_pass_{i,t} * win\_fraction_{i,t} \\ & + \beta_3 frac\_pass_{i,t} * win_{i,t} + \beta_4 win_{i,t} * win\_fraction_{i,t} + \beta_5 frac\_pass_{i,t} + \beta_6 win_{i,t} \\ & + \beta_7 win\_fraction_{i,t} + \varepsilon_{i,t+1} \end{aligned}$$

.

Again, $frac\_pass_{i,t}$ and $win\_fraction_{i,t}$, both measured between 0 and 1, are centered around their respective means. In the context of this regression specification, our model predicts that β_1 will be negative. This means that the persistence of the offensive strategy for a strong team depends less on winning and losing than for a weak team. Our model also predicts that β_2 will be positive, indicating that stronger teams will exhibit a more persistent offensive strategy, regardless of wins or losses.

In Table 4, we see that these predictions hold both for the full sample as well as for games decided by six or fewer points. Focusing on close games, we see that β_1 is negative while β_2 is positive; both coefficients are statistically significant.

To put the magnitudes into perspective, we predict the persistence of offensive strategy for winning and losing teams based on the fraction of games they won prior to the game. Table 5 shows these predicted persistence measures. Note that because we demeaned fraction pass and the win fraction variables, the figures in the table do not simply reflect the simple sum of coefficients in Table 4. We see that for teams that won all of their games in the first half of the season, the persistence of offensive strategy is roughly 0.3, regardless of whether the team wins or loses. Conversely, teams which lost all games in the first half of the season have predicted persistence measures of about 0.3 if they win but the persistence is actually slightly negative when they lose. Furthermore, because weak teams lose the majority of their games, the average persistence of offensive strategy for such teams will be much lower than for strong teams.

Reacting to the Expected

The third prediction of our model is that Bayesian coaches should update their strategy based only on performance which deviates from what was expected. For instance, when teams unexpectedly lose to a weak team they should be more likely to update their strategy than when they lose to a strong team. However, both expected and unexpected performance will influence whether the team ultimately wins or loses. Consequently, coaches suffering from outcome bias will update their strategy both on the basis of expected and unexpected performance. We quantify expected performance by looking at the gambling spread, which is empirically an unbiased estimate of the expected point differential. Unexpected performance is the difference between the actual point differential and the gambling spread.

More concretely, we estimate the following regression equation:

$$\begin{aligned} frac\_pass_{i,t+1} = \beta_0 + \beta_1 frac\_pass_{i,t} * expected\_perf_{i,t} + \beta_2 frac\_pass_{i,t} * unexpected\_perf_{i,t} \\ + \beta_3 frac\_pass_{i,t} + \beta_4 expected\_perf_{i,t} + \beta_5 unexpected\_perf_{i,t} + \varepsilon_{i,t+1} \end{aligned},$$

where $expected\_perf_{i,t}$ is the expected performance of team i in week t, and $unexpected\_perf_{i,t}$ represents the unexpected performance. If coaches use standard Bayesian updating, then only β_1 should be positive and significant. With outcome bias, however, we might expect both β_1 and β_2 to be positive and significant. In Table 6, we find that expected and unexpected performance affect the persistence of the coach's strategy in virtually identical ways. This is particularly true for games decided by six or fewer points. This suggests that coaches do not take into account the relative strength of their opponent when evaluating whether a strategy is likely to be effective going forward.[15]

The Relevance of Irrelevant Information

Factors that are irrelevant to a team's offensive strategy may also affect whether a team ultimately wins; our fourth prediction says a biased coach may allow such events to taint the evaluation of his offensive strategy. For example, winning a football game depends both on how many points a team scores as well as the number of points scored by their opponent. To the extent that the offensive strategy should have a larger effect on own points scored than on the points scored by the other team, the coach should place more weight on offensive production when evaluating the efficacy of his offensive strategy.

We test this prediction by estimating the following regression model:

$$\begin{aligned} frac\_pass_{i,t+1} = \beta_0 + \beta_1 frac\_pass_{i,t} * own\_points_{i,t} + \beta_2 frac\_pass_{i,t} * opp\_points_{i,t} \\ + \beta_3 frac\_pass_{i,t} + \beta_4 own\_points_{i,t} + \beta_5 opp\_points_{i,t} + \varepsilon_{i,t+1} \end{aligned},$$

where $own\_points_{i,t}$ is the points scored by the reference team and $opp\_points_{i,t}$ is the points scored by the opposing team. We would expect a Bayesian coach to update the fraction of passing plays based on offensive production; hence, β_1 should be positive and significant. However, if coaches suffer from outcome bias, we would also expect β_2 to be statistically significant, though negative.

[15] One might be concerned that the expected outcome of the game affects coach strategic decisions in a way that drives our results. In a regression of fraction pass on expected score differential, the r-squared is only 0.01. For games decided by six or fewer points this is even smaller at 0.002, suggesting that this concern is likely to be of minimal importance.

Table 7 shows the results from estimating this model. As expected, we find that if a team scored more points in the previous game, they will have more persistence in their offensive strategy. However, we also find that the number of points scored by the opposing team decreases the amount of persistence of the offensive strategy. In fact, when we look at just those situations where the previous game was decided by six or fewer points, we find that the effect of a team's own score on persistence has nearly the same magnitude as the opponent's score (with the opposite sign). For these games, one cannot reject the hypothesis that the coefficients are the same in absolute value. Taken literally, this would mean that a coach is as likely to revise his offensive strategy when his offense scores one less touchdown as when his defense allows one more touchdown.

It is, of course, not the case that a team's offensive strategy has no effect on the number of points scored by their opponent. For instance, passing plays require less game time on average to execute. Consequently, teams that pass more provide more opportunities for their opponent to score. Teams that give up many points may want to pass less in subsequent games. This, however, does not affect our prediction about β_1 and β_2 but rather β_5, which is the direct effect of opponent points on next week's fraction pass. This is born out in our empirical analysis.

Another reason that opponent score may drive the persistence of offensive strategy is because it contains subtle information regarding the efficacy of the offensive strategy. For example, a strategy which leads to excessive turnovers could put an opponent in a favorable scoring position. Indeed, for each additional point an opponent scores the reference team scores about .05 points less in contemporaneous games. The r-squared, however, is only 0.002. Across weeks, the correlation is even lower. Giving up an extra point in week *t* predicts scoring only 0.025 fewer points the following week, with an r-squared of 0.0007. While the coefficients are statistically significant, the amount of information in points allowed about offensive efficacy is trivial in absolute terms and relative to the information in the team's own offensive output.

6. Conclusion

Decision makers have difficulty evaluating the efficacy of a strategy when random events also influence the final outcome. One can easily misinterpret a favorable outcome as justification for a given strategy, overriding more subtle evidence to the contrary. In this paper, we provide a theoretical definition of outcome bias, which distorts Bayesian updating by over-weighting the ex-ante likelihood of the outcome that ex-post occurred. This theory provides four clear predictions, all of which are borne out in our empirical application to strategy selection by NFL coaches.

In particular, coaches tend to change their strategy more frequently after losing a game compared to winning the game. This occurs even when comparing narrow losses or victories, where winning has no predictive power about future success. Teams that win frequently are much less sensitive to a single loss; but this higher persistence in strategy is also consistent with our theory of outcome bias. In addition, coaches react equally to expected and unexpected performance. Moreover, their offensive strategy is as likely to be revised whether responsibility for a loss lies with the offense or the defense.

It is not surprising that coaches would focus on the binary outcome of winning or losing a game; after all, this is what matters most to the team owners and fans, and will largely determine whether the coach retains his job. But to maximize the chance of future wins, a coach ought to base his strategy revisions on information that most accurately predicts future success. Outcome bias reduces the accuracy of his judgments, leading to complacency after narrow wins and excessive switching after narrow losses.

We chose to document outcome bias in this sports setting due to the ease of quantifying strategies, the availability of uniform data, and the high incentives for effective evaluation. However, we anticipate that decision makers in many other settings are equally susceptible to outcome bias. For instance, sales personnel are often judged relative to sales goals. A manager could easily place excessive weight on whether the salesperson cleared the mark, even when those just above the threshold may differ only in good fortune from those just below. If so, the manager would make inefficient decisions in revamping sales incentives or retaining employees. Good programs or workers (with a modest amount of bad luck) would be scrapped, while those that barely cleared the threshold would be given too much credit for their success.

References

Baron, Jonathan, and John C. Hershey. 1988. "Outcome Bias in Decision Evaluation." *Journal of Personality and Social Psychology*, 54(4), 569-579.

Bertrand, Marianne, and Sendhil Mullainathan. 2001. "Are CEOs Rewarded for Luck? The Ones without Principals Are." *The Quarterly Journal of Economics*, 116(3), 901-932.

Camerer, Colin, George Loewenstein, and Martin Weber. 1989. "The Curse of Knowledge in Economic Settings: An Experimental Analysis." *Journal of Political Economy*, 97(5), 1232-1254.

Card, David, and Gordon B. Dahl. 2011. "Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior." *The Quarterly Journal of Economics,* 126(1), 103-143.

Chiappori, P.-A., S. Levitt, and T. Groseclose. 2002. "Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer." *The American Economic Review*, 92(4), 1138-1151.

Dobbs, Ian M. 1991. "A Bayesian Approach to Decision-Making under Ambiguity." *Economica*, 58(232), 417-440.

Earl, Peter E. 1990. "Economics and Psychology: A Survey." *The Economic Journal*, 100(402), 718-755.

Fischhoff, Baruch. 1975. "Hindsight ≠ Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty." *Journal of Experimental Psychology: Human Perception and Performance,* 1(3), 288-299.

Gino, Francesca, Don A. Moore, and Max H. Bazerman. 2008. "No Harm, No Foul: The Outcome Bias in Ethical Judgments." Harvard Business School NOM Working Paper.

Gino, Francesca, Lisa L. Shu, and Max H. Bazerman. 2010. "Nameless + Harmless = Blameless; When Seemingly Irrelevant Factors Influence Judgment of (Un)ethical Behavior." *Organizational Behavior and Human Decision Processes*, 111, 93-101.

Gray, Rob, Sian L. Beilock, and Thomas H. Carr. 2007. "'As Soon as the Bat Met the Ball, I Knew It Was Gone': Outcome Predication, Hindsight Bias, and the Representation and Control of Action in Expert and Novice Baseball Players." *Psychonomic Bulletin & Review*, 14(4), 669-675.

Gray, Philip K., and Stephen F. Gray. 1997. "Testing Market Efficiency: Evidence from the NFL Sports Betting Market." *The Journal of Finance*, 52(4), 1725-1737.

Hawkins, Scott A., and Reid Hastie. 1990. "Hindsight: Biased Judgments of Past Events After the Outcomes are Known." *Psychological Bulletin*, 107(3), 311-327.

Jones, Steven K., Tricia J. Yurak, and Deborah Frisch. 1997. "The Effect of Outcome Information on the Evaluation and Recall of Individuals' Own Decisions." *Organizational Behavior and Human Decision Processes,* 71(1), 95-120.

LaBine, Susan J., and Gary LaBine. 1996. "Determinations of Negligence and the Hindsight Bias." *Law and Human Behavior*, 20(5), 501-516.

Lefgren, Lars, and Brennan C. Platt. 2010. "Coaches on the Hot Seat: Testing Models of Moral Hazard and Screening." Mimeo, Brigham Young University.

Lipshitz, Raanan. 1989. "'Either a Medal or a Corporal': The Effects of Success and Failure on the Evaluation of Decision Making and Decision Makers." *Organizational Behavior and Human Decision Processes*, 44(3), 380-395.

Marshall, Greg W., and John C. Mowen. 1993. "An Experimental Investigation of the Outcome Bias in Salesperson Performance Evaluations." *Journal of Personal Selling & Sales Management,* 13(3), 31-47.

Offerman, Theo, and Joep Sonnemans. 2004. "What's Causing Overreaction? An Experimental Investigation of Recency and the Hot-Hand Effect." *Scandinavian Journal of Economics*, 106(3), 533-553.

Pope, Devin G., and Maurice E. Schweitzer. 2011. "Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes." *American Economic Review*, 101(1), 129-157.

Rabin, Matthew. 1998. "Psychology and Economics." *Journal of Economic Literature,* 36(1), 11-46.

Ratner, Rebecca K., and Kenneth C. Herbst. 2005. "When Good Decisions have Bad Outcomes: The Impact of Affect on Switching Behavior." *Organizational Behavior and Human Decision Processes,* 96, 23-37.

Romer, David. 2006. "Do Firms Maximize? Evidence from Professional Football." *Journal of Political Economy,* 114(2), 340-365.

Sauer, Raymond D. 1998. "The Economics of Wagering Markets." *Journal of Economics Literature, 36*(4), 2021-2064.

Wolfers, Justin. 2002. "Are Voters Rational? Evidence from Gubernatorial Elections." Stanford GSB Working Paper No. 1730.

Wolfers, Justin. 2006. "Point Shaving: Corruption in NCAA Basketball." *The American Economic Review*, 96(2), 279-283

Table 1. Descriptive Statistics

	Mean	SD
Fraction Pass	0.535	0.116
Win	0.499	0.500
Points Scored	20.72	10.24
Score Differential	0	14.78
Abs(Score Differential)	11.59	9.17
Expected Performance	0	6.65
Unexpected Performance	0	13.25

Notes: The sample includes all regular season games during the 1985-2009 NFL seasons and includes 10,578 team-game observations. The data comes from NFLData.com.

Table 2: The Impact of Winning on Persistence of Offensive Strategy

	(1)	(2)	(3)	(4)
Fraction Pass*Win	0.166**	0.107**	0.173**	0.163**
	[0.023]	[0.038]	[0.032]	[0.049]
Fraction Pass	0.104**	0.145**	0.125**	0.145**
	[0.019]	[0.029]	[0.026]	[0.038]
Win	-0.065**	-0.044**	0.030**	0.019**
	[0.013]	[0.021]	[0.004]	[0.005]
Win Fraction*Fraction Pass			0.138*	0.291**
			[0.080]	[0.135]
Win Fraction			-0.047**	-0.039**
			[0.008]	[0.012]
Restrictions	Full Sample	Games within 6 points	Games after Week 8	Games after Week 8/ within 6 points
Observations	10,573	3,972	5,330	2,033
R-squared	0.029	0.031	0.045	0.050

Notes: Standard errors are in brackets and are cluster corrected at the team-year level. ** and * indicate statistical significance at the 5 and 10 percent level respectively. “Win” indicates whether the team won their previous game and “Win Fraction” is the fraction of current-season games prior to the current game that the team has won.

Table 3: Robustness Checks

	Baseline	*Persistence from t to t-1*	*Including 3rd Order Polynomial in Fraction Pass*	*Including Team and Year Fixed Effects*
Fraction Pass*Win	0.107**	-0.021	0.091**	0.105**
	[0.038]	[0.038]	(0.043)	(0.037)
Fraction Pass	0.145**	0.157**	0.136**	0.103**
	[0.029]	[0.028]	(0.036)	(0.028)
Win	-0.044**	0.015**	0.014**	0.012**
	[0.021]	[0.004]	(0.004)	(0.004)
Fraction Pass Squared			-0.135	
			(0.155)	
Fraction Pass Cubed			0.547	
			(0.809)	
Restrictions	Games within 6 points	Games within 6 points	Games within 6 points	Games within 6 points
Observations	3,972	3,972	3,972	3,972
R-squared	0.031	0.017	0.031	0.070

Notes: Standard errors are in brackets and are cluster corrected at the team-year level. ** and * indicate statistical significance at the 5 and 10 percent level respectively. "Win" indicates whether the team won their previous game and "Win Fraction" is the fraction of current-season games prior to the current game that the team has won.

Table 4: The Impact of Winning on the Persistence of Offensive Strategy by Team Strength

	(1)	(2)
Win Fraction*Win* Fraction Pass	-0.218	-0.480*
	[0.158]	[0.250]
Win Fraction *Fraction Pass	0.222*	0.440**
	[0.127]	[0.187]
Win*Fraction Pass	0.177**	0.170**
	[0.033]	[0.049]
Win Fraction *Win	-0.012	-0.063**
	[0.018]	[0.024]
Fraction Pass	0.127**	0.144**
	[0.026]	[0.039]
Win	-0.031**	0.020**
	[0.004]	[0.005]
Win Fraction	-0.048**	-0.016
	[0.012]	[0.016]
Restrictions	Games after Week 8	Games after Week 8/ within 6 points
Observations	5,330	2.033
R-squared	0.045	0.055

Notes: Standard errors are in brackets and are cluster corrected at the team-year level. ** and * indicate statistical significance at the 5 and 10 percent level respectively. "Win" indicates whether the team won their previous game and "Win Fraction" is the fraction of current-season games prior to the current game that the team has won.

Table 5: Predicted Persistence for Winners and Losers by Team Strength

	Won Current Game	Lost Current Game
Win Fraction=1	0.293	0.364
Win Fraction=0	0.334	-0.076

Table 6: The Impact of Expected and Unexpected Performance on Persistence of Offensive Strategy

	(1)	(2)
Expected Performance * Fraction Pass	0.004**	0.016**
	[0.002]	[0.006]
Unexpected Performance * Fraction Pass	0.005**	0.013**
	[0.001]	[0.005]
Fraction Pass	0.201**	0.208**
	[0.033]	[0.021]
Expected Performance	-0.002**	-0.006**
	[0.001]	[0.003]
Unexpected Performance	-0.002**	-0.004**
	[0.000]	[0.003]
<u>P-values</u>		
Ho: Expected Performance Interaction=Unexpected Performance Interaction	0.550	0.354
Restrictions	Full Sample	Games within 6 points
Observations	10,573	3,972
R-squared	0.032	0.033

Notes: Standard errors are in brackets and are cluster corrected at the team-year level. ** and * indicate statistical significance at the 5 and 10 percent level respectively.

Table 7: The Impact of Own and Opponent Scoring on Persistence of Offensive Strategy

	(1)	(2)
Own Score * Fraction Pass	0.007**	0.014**
	[0.001]	[0.005]
Opponent Score * Fraction Passed	-0.003**	-0.013**
	[0.001]	[0.005]
Fraction Passed	0.117**	0.186**
	[0.032]	[0.056]
Own Score	0.001**	0.002**
	[0.000]	[0.001]
Opponent Score	-0.001*	-0.002**
	[0.000]	[0.001]
P-values		
Ho: Own Score Interaction + Opponent Score Interaction =0	0.007	0.699
Restrictions	Full Sample	Games within 6 points
Observations	10,573	3,972
R-squared	0.032	0.033

Notes: Standard errors are in brackets and are cluster corrected at the team-year level. ** and * indicate statistical significance at the 5 and 10 percent level respectively.

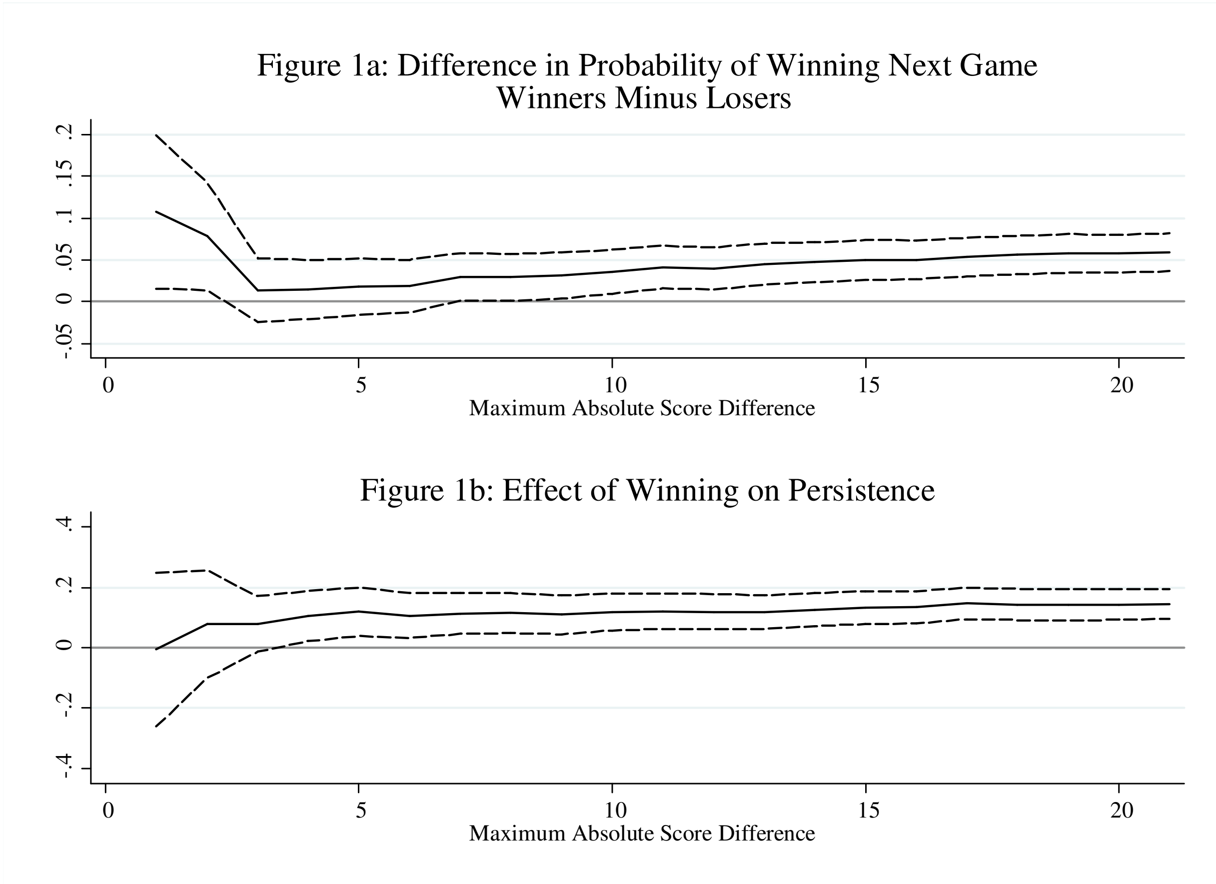

Notes: Each point in the figure represents a separate estimate. The X-axis indicates the maximum absolute score difference that is allowed for the observation to be included in the estimate, e.g. a value of 5 indicates that the previous game was decided by 5 or fewer points.

Figure 1a provides the difference in the probability of winning the next game between the team that won and the team that lost. Figure 1b provides the estimated persistence effect; that is, the effect that winning has on whether a team continues using a particular offensive strategy. The dotted lines provide the 95% confidence interval for each estimate.

Figure 2a: Difference in Probability of Winning Next Game
Winners Minus Losers-Controlling for Team Strength

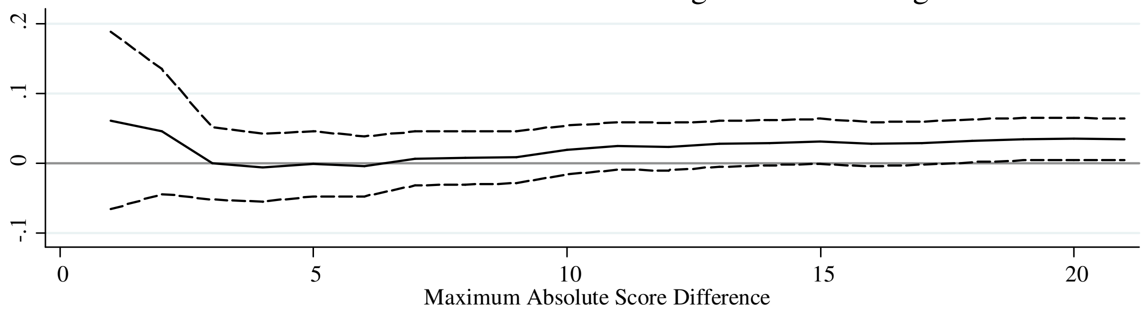

Figure 2b: Effect of Winning on Persistence-Controlling
for Team Strength

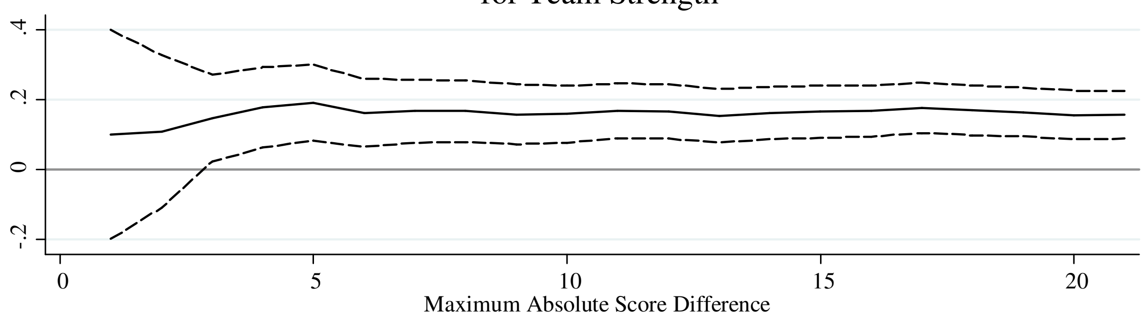

Notes: These figures are the same as those in Figure 1a and 1b except we only use games from the second half of the season and control for the team's winning percentage in the first half of the season. In Figure 2b, we also control for the interaction of the winning percentage and fraction pass. Because both variables are centered about their means, this does not affect the interpretation of the coefficient.