The Forgone-Option Fallacy*

Etan A. Green Joshua Lewis

University of Pennsylvania

June 12, 2017

Abstract

Forgoing an option motivates decision makers to achieve outcomes that "justify" the initial choice. Analyzing a natural experiment in the National Football League, we find that forgoing an option motivates professional athletes to do no worse than they would have done had they taken the option in the first place. We complement this result with evidence from survey experiments.

*We thank seminar participants at Stanford University, Microsoft Research, and the Society for Judgment and Decision Making conference, as well as Ashton Anderson, Isa Chaves, Luke Coffman, David Daniels, Thomas Dudley, Dan Goldstein, Jake Hofman, Dorothy Kronick, Jonathan Levav, Mark Machina, Alex Rees-Jones, Peter Reiss, Joe Simmons, Uri Simonsohn, Sid Suri, and Joel Watson for helpful thoughts and comments. We are also grateful to Laura Kuder, Johanna Matt-Navarro, Catherine O'Donnell, Noah Wilson, and Sophia Yang for research assistance.

Research in psychology, neuroscience, and economics shows that people evaluate outcomes in comparison to reference points, which separate psychological gains from psychological losses. These evaluations are characterized by loss aversion, in which outcomes worse than the reference point are felt more strongly than commensurate gains (Kahneman and Tversky, 1979; Tom et al., 2007; De Martino et al., 2010). A broad class of reference points comprises salient counterfactuals (Kahneman and Miller, 1986; Thaler and Johnson, 1990; Medvec et al., 1995), and loss aversion with respect to counterfactual outcomes can lead to anomalous behavior. For instance, investors may define gains and losses in relation to the counterfactual of not having invested. This reference point motivates a sunk-cost fallacy: investors refrain from selling for less than the prices they paid, as has been documented for stocks (Odean, 1998; Frazzini, 2006) and real estate (Genesove and Mayer, 2001).

We propose that individuals also define gains and losses in relation to the counterfactual outcomes of forgone-options, and that this reference point motivates a forgone-options fallacy. Just as a sunk-cost fallacy predicts that investors will refrain from selling for less than the prices they paid, a forgone-option fallacy predicts that investors will refrain from selling for less than the offers they declined. Similarly, employees who forgo a fixed wage to work on commission will work until their piecemeal earnings surpass the forgone fixed wage, and diners who decline the option of an all-you-can-eat buffet will stop eating before the a-la-carte cost of their consumption exceeds the buffet price. Forgoing an option renders its outcome obsolete. Yet in each of these examples, the counterfactual of having chosen differently frames how outcomes are valued, prompting decision makers to anticipate the regret of having decided poorly and motivating them to attain outcomes that "justify" their initial choices (Festinger, 1962; Bell, 1982; Loomes and Sugden, 1982; Arkes, 1996; Zeelenberg, 1999; Coricelli et al., 2005).

We measure the effect of forgoing an option using a natural experiment in the National Football League. This setting allows us to compare outcomes between players who forgo an option and otherwise equivalent players for whom the option was, by chance, not available. Importantly, players assigned the option almost always decline it (98% of the time), mitigating selection concerns. We find evidence in support of the prediction that forgone options provide reference points. Compared to players who are not assigned the option, those who forgo the option are 10 percentage points, or 14%, more likely to achieve an outcome at least as desirable as the option's counterfactual. Consistent with loss aversion, those who decline the option appear to work harder than those who never had the choice—merely to avoid the regret of having declined the option in the first place. We complement this result with evidence from survey experiments.

Study 1: NFL kickoff returns

We first examine forgone options in a natural experiment: kickoff returns in the National Football League. During the kickoff, the kicking team kicks the ball downfield toward the returning team's goal line.[1] The returner then tries to return the kickoff by running as far upfield as possible while the kicking team tries to limit the distance of the return by tackling the returner. If the returner fields the kickoff (i.e., first touches the ball) behind the goal line, he has a choice. He may return the kickoff, thus gambling on a lottery of potential outcomes by running until he is tackled. Alternatively, he may choose to take a "touchback", in which case his team receives the ball at the 20-yard line with certainty—i.e., 20 yards in front of the returning team's goal line.[2] If instead, the returner fields the kickoff in front of the goal line, he does not have the touchback option. Hence, the goal line demarcates the option space: behind it, returners may opt for a certain outcome; in front of it, they may not. We

[1] A football field is 100 yards long, with an additional region at each end—the end zone—that is 10 yards deep. The 50-yard line splits the field. On either side, yard markers count down to the goal lines, which separate the field of play from the end zones.

[2] A recent rule change moved this reward to the 25-yard line. During the observation window we analyze, the touchback reward is the 20-yard line.

describe the data in Appendix A.

Analysis

In our setting, the objective benefit of returning the kickoff an extra yard is the same regardless of where the extra yard begins and ends, as we show in Figure S1. For a loss-averse returner, however, the psychological benefit of an extra yard is greater when the returner is short of the reference point than when he has already eclipsed it. In Appendix B, we sketch a simple model in which the returner may exert effort to gain an extra yard. We assume that declining the touchback option makes the 20-yard line salient as a reference point. We show that loss aversion prompts the returner to exert extra effort when he would otherwise be tackled short of the reference point, and that this behavior deforms the distribution of return outcomes, shifting probability mass forward until it heaps at the 20-yard line.

We test this prediction by comparing kickoff returns from either side of the goal line. In doing so, we assume that returns from either side of the goal line vary solely in the returner's options. As we show in Appendix C, kickoffs fielded across the goal line do not vary on a set of observables, including game and situational attributes. However, returns may differ across the goal line because of selection. If kickoff returners take the touchback option when they are pessimistic about their chances of reaching the 20-yard line, the distribution of outcomes among those who forgo the option will appear to have missing mass short of the 20-yard line—merely because those who expect worse outcomes will opt for touchbacks.

We sidestep this concern by comparing kickoffs fielded from one yard in front of the goal line, for which the touchback is not an option, to those fielded one yard behind the goal line, for which the touchback is rarely taken. For kickoffs fielded one yard behind the goal line, 79% of returns reach at least the 20-yard line. Accordingly, 97.7% of kickoffs fielded one yard behind the goal line are returned, minimizing selection concerns.

Forgoing the touchback option precipitates a discontinuity in the distribution of return

Figure 1: Histograms of the yard line where the returner was tackled or ran out of bounds, for returns from one yard behind (a) and in front of (b) the goal line.

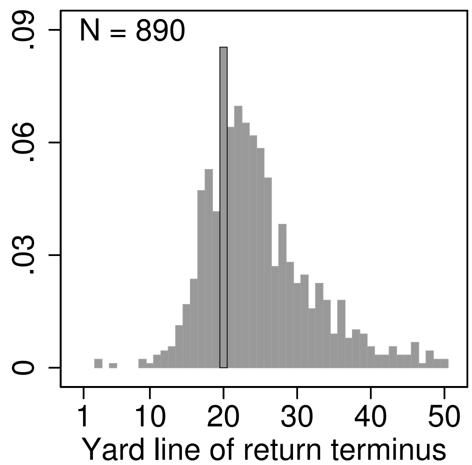

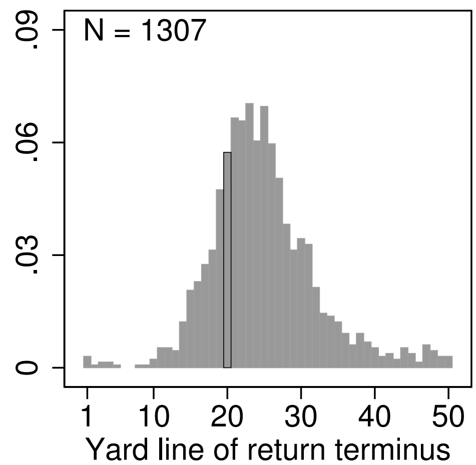

distance at the 20-yard line—i.e., the counterfactual outcome of having taken the touchback. Figure 1 shows the empirical probability of a return to each yard line, separately for returns from one yard behind the goal line (1a), and for returns from the one yard in front of the goal line (1b). For returns from just behind the goal line, the distribution of return distance jumps at the 20-yard line with probability mass apparently displaced from left of that marker, and from the 19-yard line in particular. By contrast, the distribution of return distance appears smooth across the 20-yard line for returns from just in front of the goal line.

Explanations which neglect the declined option are inconsistent with these distributions. Reference dependence around round numbers (Pope and Simonsohn, 2011; Allen et al., 2016) predicts a displacement of probability mass to the 20-yard line for all returns, regardless of whether the touchback option was available. A tendency by game officials to award generous placements to returners who fall short of the 20-yard line makes an identical prediction. Consistent with these explanations, probability mass concentrates at the 20-yard line after the returner forgoes the touchback option. Yet when the touchback option is unavailable, the 20-yard line is no longer alluring and referees are no longer generous.

We first estimate the effect of declining the touchback option by comparing treated returns, or those from one yard behind the goal line, to a control group of returns from one yard in front of the goal line. By construction, returns from either side of the goal line differ not only in the availability of the touchback option but also in the yard line from which the return begins. To isolate the effect of the declining the touchback option, we amend returns in the control group by shifting the terminus of the return back two yards.[3]

Figure 2: The effect of declining the touchback option on the probability of returning the kickoff to the given yard line, measured as the difference between the distribution of the return terminus for returns from one yard behind the goal line and the distribution of the return terminus for returns from one yard in front of the goal line, shifted two yards back. 95% confidence intervals from 10,000 bootstrap samples of the underlying distributions.

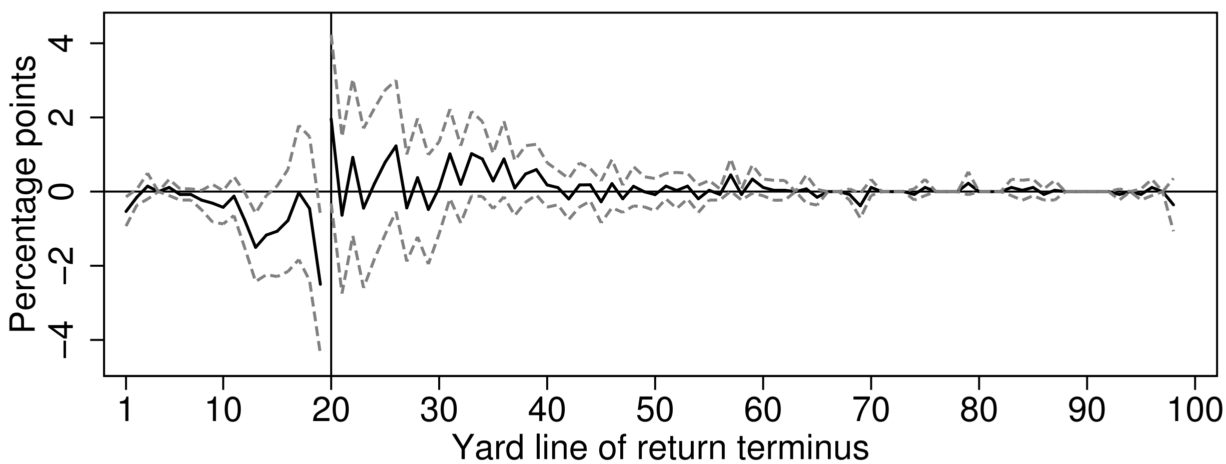

Declining the touchback option displaces probability mass from short of the 20-yard line to the 20-yard line and just beyond. Figure 2 shows the effect of declining the touchback option on the probability of returning the kickoff to each yard line, measured as the difference between the treatment and control distributions. Missing mass accumulates short of the 20-yard line, and excess mass is concentrated at the 20-yard line and just beyond, implying

[3]This assumes that absent an effect of declining the touchback option, a kickoff returned from the 1-yard line to the 25-yard line, say, would have been returned to the 23-yard line had it been fielded one yard deep in the end zone. While we cannot verify this assumption directly, we note that returns from the 1- and 2-yard lines are on average returned 25.4 and 25.3 yards, respectively ($p = 0.83$).

a shift of probability mass from outcomes worse than a touchback to outcomes that are at least as desirable. Declining the touchback option shifts 4.5% of returns from the 19-yard line to the 20-yard line exactly (se: 1.5pp from 10,000 bootstrap samples of both distributions; $p = 0.004$) and increases the probability of a return to at least the 20-yard line by 9.9 percentage points (se: 1.9pp; $p < 10^{-4}$), or 14%. Selection into touchbacks can only account for a small portion of this increase. Under a worst-case scenario in which we assume that the 22 observed touchbacks on kickoffs fielded one yard behind the goal line would all have been returned short of the 20-yard line, returns to at least the 20-yard line are 8.1 percentage points (se: 1.9pp; $p < 10^{-4}$), or 12%, more likely after declining the touchback option.

We also conduct a more conservative version of this test by comparing returns from the goal line to a control distribution of returns from the 1-yard line, now shifted one yard back. Whether a kickoff fielded on the goal line may be taken for a touchback is often ambiguous, and just 1.5% of these kickoffs result in touchbacks. For returns from the goal line, the distribution of outcomes, shown in Figure S5, appears to be missing mass short of the 20-yard line, though the deformation is less severe than for returns from one yard behind the goal line (1a). The difference between the treatment and control distributions reveals a shift of mass from short of the 20-yard line to the 20-yard line and beyond, as shown in Figure S6. Declining the touchback option increases the probability of a return to at least the 20-yard line by 4.6 percentage points (se: 1.6pp; $p = 0.004$), or 6%.[4] Under our worst-case assumption, returns to at least the 20-yard line are 3.3 percentage points (se: 1.6pp; $p = 0.041$), or 4%, more likely after declining the touchback option. Finally, in Appendix D, we adjust for small differences in the origin of the return using a difference-in-difference approach, rather than a shift of the control distribution.

[4]The displacement is also less pronounced near the threshold: declining the touchback option shifts 1.3% of returns from the 19-yard line to the 20-yard line exactly (se: 1.3pp; $p = 0.316$).

Discussion

Returners who decline the touchback option reach its counterfactual outcome more frequently than comparable returners who never have the option. A likely interpretation is that declining the touchback option motivates returners to try harder than they otherwise would. Although increased exertion improves return outcomes, it is not optimal. If this marginal effort is sufficiently costly—if it makes injury likely, for instance—it is inadvisable. If not, the returner should exert himself similarly when returning kickoffs from in front of the goal line.

The principal limitation of our design is that we cannot identify whom the forgone option motivates. One alternative interpretation is that the returner's motivation is extrinsic and comes from incentives offered by his coaches. Another possibility is that game officials favor returners only when they decline the touchback option and then fall short of its counterfactual outcome. Since our data does not allow us to parse these explanations, we conduct an experiment in which the decision maker acts alone.

Study 2: Vignettes

We asked 1,974 participants from Mechanical Turk to evaluate vignettes in exchange for 50 cents.[5] Subjects were randomly presented with one of 2 vignettes in one of 2 conditions (forgone option or no option). The vignettes were as follows. Parentheses surround text specific to the forgone-option condition, whereas brackets surround text specific to the no-option condition. The protagonist's name was randomly sampled from 6 male and 6 female

[5] We pre-registered a goal of 2000 responses and committed to removing responses from participants who had taken the survey previously; our pre-registration document can be viewed at `https://aspredicted.org/me5t3.pdf`. The final sample reflects the removal of 35 responses from respondents who had participated previously.

names.[6]

1. For dinner, Amy goes to a restaurant that has recently been offering diners the option to order as many dishes from the menu as they can eat, including desserts, for $24. Amy knows that once she starts ordering a la carte, she will no longer be allowed to get the all-you-can-eat deal. However, she starts ordering a la carte anyway, as she thinks it will be cheaper. After her main course, Amy finds out that the all-you-can-eat offer (was ongoing) [had expired], and so it would [not] have been available [even] if she had wanted it. Having already spent $21, Amy eyes a dessert that will bring her bill to $27. How likely is she to order dessert?

2. Amy's cable provider has been advertising a one-time special offer: one month of unlimited movies for $14. She knows that once she buys a pay-per-view movie, she will no longer be eligible for the unlimited movie offer. However, she decides to buy her first movie of the month with pay-per-view anyway, as she thinks it will work out cheaper than the unlimited offer. Nearing the end of the month, Amy learns that the unlimited offer was for (this month) [last month], so it would [not] have been available [even] if she had wanted it. Having so far spent $12 on pay-per-view movies, Amy eyes a movie that would bring her pay-per-view movie costs to $16 for the month. How likely is Amy to watch the movie?

Participants responded using a 1-7 scale from *very unlikely* to *very likely*. Thereafter, we required participants to justify their choice.

Each vignette describes a protagonist weighing consumption against expenditure. In both conditions, the protagonist knows about an option for unlimited consumption with a fixed cost but expects to spend less paying piecemeal, and hence ignores the option. In the forgone-option condition, the protagonist finds out that the option would have been available

[6]These names were Amy, Meredith, Claire, Keisha, Mei, Gabriella, Bob, Steven, Brad, Lamar, Chang, Carlos (c.f. Milkman et al., 2012).

if he or she had wanted it, and in the no-option condition, the protagonist finds out that the option would *not* have been available even if he or she had wanted it. The protagonist then decides whether to consume an additional unit that would push total expenditures beyond the fixed-cost counterfactual.

Figure 3: Results by condition, with 95% confidence intervals.

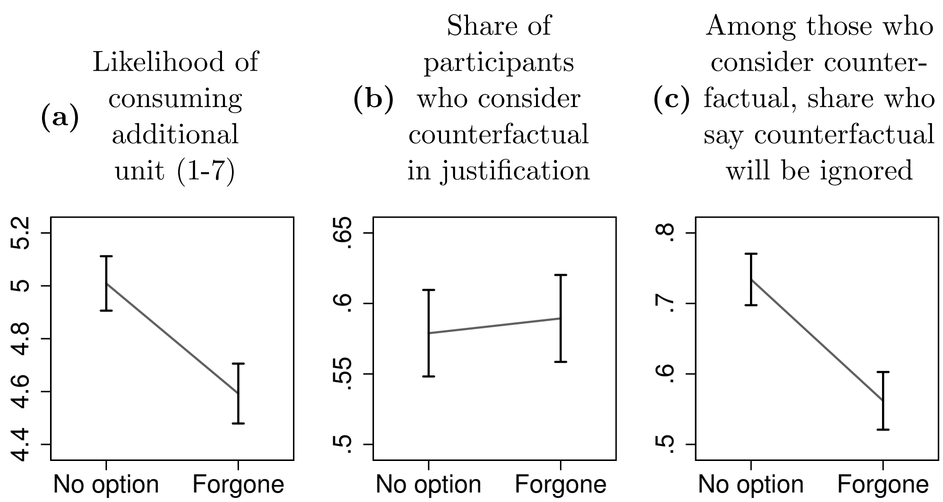

We hypothesized that participants would rate the protagonist as less likely to consume the additional unit in the forgone-option condition than in the no-option condition. Results are similar for each vignette, as we show in Appendix E; hence, we pool vignettes here. Figure 3a shows average responses by condition, with 95% confidence intervals. Protagonists are seen as less likely to consume the additional unit in the forgone-option condition than in the no-option condition (Cohen's $d = 0.24$, $t(1972) = 5.34$, $p < 10^{-4}$).

We also analyzed participants' text justifications. Research assistants coded whether the justification considered the counterfactual of unlimited consumption, and if so, whether it argued that the counterfactual would be ignored.[7] Participants considered the counterfactual at similar rates in both conditions (Cohen's $d = -0.02$, $t(1972) = -0.47$, $p = 0.637$). Among

[7]Two research assistants independently coded each response. When they disagreed, a third coder made the final decision. The research assistants were blinded to our hypothesis, as well as to the condition and numerical response associated with each justification.

those who considered the counterfactual, however, proportionally more participants said that the counterfactual would be ignored in the no-option condition than in the forgone-option condition (Cohen's $d = 0.37$, $t(1128) = 6.15$, $p < 10^{-4}$). Declining an option makes its counterfactual outcome difficult to ignore.

Conclusion

Using a natural experiment and survey experiments, we find that the counterfactual outcomes of forgone options provide reference points. Declining the touchback option motivates NFL kickoff returners to reach the 20-yard line—i.e., the counterfactual outcome of having taken the touchback. And when a-la-carte consumption costs approach the price of an unlimited option, experimental subjects attribute greater self control to those who declined the option than to otherwise identical individuals for whom that option was never available. Though it has no bearing on future payoffs, the counterfactual outcome of a forgone option may pose as a reference point for subsequent decisions, motivating decision makers to exert effort or exercise self control in order to avoid psychological losses.

In the contexts we study, the counterfactual outcome is always certain—the 20-yard line or a fixed cost. However, not every forgone option offers a certain counterfactual. Many assets have uncertain valuations. Hence, the investor who opts to buy one asset over another may never know precisely what would have happened had she made a different decision. When counterfactual outcomes are uncertain, we expect that the reference point will be less salient, and as a result, motivation to avoid psychological losses will be diminished.

Nonetheless, forgone options with certain outcomes are ubiquitous, and as such, their use as reference points has broad implications. Consider the disposition effect, in which individuals avoid selling investments for less than the sunk cost of the price paid (Shefrin and Statman, 1985; Barberis and Xiong, 2009). Forgone options suggest a different disposition effect, in which obsolete counterfactuals, rather than sunk costs, provide reference points. An individual who refuses a buyer's offer for her house or declines to cash a stock option when it is in the money may refrain from selling the house or exercising the option until the value of the asset "justifies" her past decision. This psychology can also explain an anomalous pattern in financial markets wherein divestment rates spike when prices eclipse recent peaks (Baker et al., 2012). Every moment that an asset is held is a moment in which the option to sell is forgone. Under the illogic of forgone options, reference points will coincide with past peak valuations, and investors will hold onto assets until prices return to past peaks.

Given the pervasiveness of forgone options, future work might test interventions that promise to enhance motivation or self control by providing decision makers with options that they are likely to forgo. Alternatively, future work might explore ways to help individuals ignore obsolete counterfactuals.

References

Allen, Eric J, Patricia M Dechow, Devin G Pope, and George Wu (2016) "Reference-dependent preferences: Evidence from marathon runners," *Management Science*.

Arkes, Hal R (1996) "The psychology of waste," *Journal of Behavioral Decision Making*, Vol. 9, pp. 213–224.

Baker, Malcolm, Xin Pan, and Jeffrey Wurgler (2012) "The effect of reference point prices on mergers and acquisitions," *Journal of Financial Economics*, Vol. 106, pp. 49–71.

Barberis, Nicholas and Wei Xiong (2009) "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation," *the Journal of Finance*, Vol. 64, pp. 751–784.

Bell, David E (1982) "Regret in Decision Making under Uncertainty," *Operations Research*, pp. 961–981.

Coricelli, Giorgio, Hugo D Critchley, Mateus Joffily, John P O'Doherty, Angela Sirigu, and Raymond J Dolan (2005) "Regret and its avoidance: a neuroimaging study of choice behavior.," *Nature Neuroscience*, Vol. 8.

De Martino, Benedetto, Colin F Camerer, and Ralph Adolphs (2010) "Amygdala damage eliminates monetary loss aversion," *Proceedings of the National Academy of Sciences*, Vol. 107, pp. 3788–3792.

Festinger, Leon (1962) *A theory of cognitive dissonance*, Vol. 2: Stanford university press.

Frazzini, Andrea (2006) "The disposition effect and underreaction to news," *The Journal of Finance*, Vol. 61, pp. 2017–2046.

Genesove, David and Christopher Mayer (2001) "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *The Quarterly Journal of Economics*, Vol. 116, pp. 1233–1260.

Kahneman, Daniel and Dale T Miller (1986) "Norm theory: Comparing reality to its alternatives.," *Psychological review*, Vol. 93, p. 136.

Kahneman, Daniel and Amos Tversky (1979) "Prospect theory: An analysis of decision under risk," *Econometrica: Journal of the Econometric Society*, pp. 263–291.

Loomes, Graham and Robert Sugden (1982) "Regret theory: An alternative theory of rational choice under uncertainty," *The Economic Journal*, pp. 805–824.

Medvec, Victoria Husted, Scott F Madey, and Thomas Gilovich (1995) "When Less Is More: Counterfactual Thinking and Satisfaction Among Olympic Medalists," *Journal of Personality and Social Psychology*, Vol. 69, pp. 603–610.

Milkman, Katherine L, Modupe Akinola, and Dolly Chugh (2012) "Temporal distance and discrimination: An audit study in academia," *Psychological Science*, Vol. 23, pp. 710–717.

Odean, Terrance (1998) "Are Investors Reluctant to Realize Their Losses?" *The Journal of Finance*, Vol. 53, pp. 1775–1798.

Pope, Devin and Uri Simonsohn (2011) "Round Numbers as Goals: Evidence From Baseball, SAT Takers, and the Lab," *Psychological Science*, Vol. 22, pp. 71–79.

Romer, David (2006) "Do firms maximize? Evidence from professional football," *Journal of Political Economy*, Vol. 114, pp. 340–365.

Shefrin, Hersh and Meir Statman (1985) "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, pp. 777–790.

Thaler, Richard H and Eric J Johnson (1990) "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice," *Management Science*, Vol. 36, pp. 643–660.

Tom, Sabrina M, Craig R Fox, Christopher Trepel, and Russell A Poldrack (2007) "The neural basis of loss aversion in decision-making under risk," *Science*, Vol. 315, pp. 515–518.

Zeelenberg, M (1999) "Anticipated regret, expected feedback and behavioral decision-making," *Journal of Behavioural Decision Making*, Vol. 12, pp. 93–106.

A Data

We compile the dataset from NFL play-by-play records, which list the yard line from where the ball was kicked, the yard line where it was fielded, whether the kickoff resulted in a touchback, and if not, where the returner was tackled or ran out-of-bounds.[8] The data cover the 2000-10 NFL seasons, an 11-year period during which the NFL employed the same kickoff rules: the ball is kicked 70 yards from the returning team's goal line, and the reward for a touchback is the 20-yard line. To ensure consistency between kickoffs, we exclude the 2.6% of kickoffs for which the location of the kickoff was moved because of a prior penalty. We also exclude the 1.5% of remaining kickoffs that were kicked out of bounds, which result in a penalty on the kicking team and which are generally unreturnable. This leaves a dataset of 26,324 kickoffs. When a penalty occurs during the return, we record the end of the return as the yard line where the returner was tackled or ran out-of-bounds, not the yard line awarded to the returning team after the penalty.

Our main analysis compares kickoff returns from one yard behind the goal line to kickoff returns from one yard in front of the goal line. We exclude the 3.5% of returns from these depths in which the returner fumbles the ball or a penalty on the return censors the yard at which the returner was tackled or ran out of bounds. This leaves 890 returns from one yard behind the goal line and 1,307 returns from one yard in front of the goal line.

B Theoretical framework

Loss aversion around the certain outcome of the declined touchback option—the 20-yard line—generates a discontinuity in the distribution of kickoff return distance at that reference point. During the return, the returner faces some probability of being contacted by a member

[8]Unfortunately, the play-by-play data for college football censor the yard line where the kickoff is fielded at the goal line—i.e., marking all kickoffs fielded in the end zone as having been fielded at the goal line—precluding use of data from college football for our purposes.

of the kicking team at each yard line y. We assume that if the returner is contacted at y, either the returner is tackled at the current yard line, y, or he falls forward to the next yard line, $y+1$. This assumption gives a simple form for the distribution of return distance, $f(y)$:

$$f(y) = P(\text{tackle at } y) = P(\text{contact at } y) \cdot P(\text{tackled}|\text{contact}, e) + P(\text{contact at } y-1) \cdot \big(1 - P(\text{tackled}|\text{contact}, e)\big) \quad (1)$$

where e is the level of effort exerted by the returner. The return ends at y if the returner is contacted at y and tackled there, or if he is contacted at $y-1$ and falls forward for an extra yard.[9]

The likelihood of contact at any yard line stems from a chaotic process involving 11 kicking team members trying to evade 10 blockers and tackle one returner. We abstract over this physical complexity by letting $P(\text{contact at } y)$ follow a smooth distribution function G.[10] Upon contact at y, the returner chooses a binary effort level $e \in \{e_L, e_H\}$. Under low effort, $P(\text{tackled}|\text{contact}, e_L) = p$, and under high effort $P(\text{tackled}|\text{contact}, e_H) = q < p$. High effort decreases the probability of being tackled at the point of contact and thus increases the probability of falling forward for an extra yard.

Let the returner's utility for a given effort level, $U(e)$, be the difference between the benefit of achieving a particular yard line, $b(y)$, and cost of exerting the associated effort,

[9] We assume that $P(\text{tackled}|\text{contact}, e)$ is independent of y. In other words, whether contact leads to a tackle depends on the returner's effort but not his position on the field.

[10] In practice, G represents the discretization of a continuous distribution H, where $g(y) = G(y) - G(y-1) = \int_{y-1}^{y} dH(u)$. We assume that the second derivative of H is small in magnitude, which implies that g is "smooth"—the jump in the probability of contact at y is similar to the jump in the probability of contact at $y-1$, or equivalently, that $g(y) - g(y-1) \approx g(y-1) - g(y-2)$ for all y.

$c(e)$. Then,

$$\mathbb{E}\big[U(e|\text{contact at } y)\big] = b(y) \cdot P(\text{tackled}|\text{contact}, e) + b(y+1) \cdot \big(1 - P(\text{tackled}|\text{contact}, e)\big) - c(e)$$

We normalize $c(e_L) = 0$ and $c(e_H) = C > 0$.

The empirical setting provides a clear normative benchmark for the shape of $b(y)$. As Figure S1 shows, the expected number of points that the returner's team can expect to score in the ensuing drive is linearly increasing in the terminus of the return through the 20-yard line. Hence, the returner should view an advance from the 19- to the 20-yard lines as providing the same marginal benefit as any other one-yard advance. Hence, the returner's benefit function should be reference independent (RI) and should be linearly increasing in y, as for $b_{RI}(y) = m(y - 20)$, for $m > 0$.[11]

For a loss-averse returner, however, outcomes just above and below the 20-yard line will feel more different than outcomes just above and below any other marker. In particular, a loss-averse returner may feel disappointed if he falls short of the 20-yard line after declining the touchback option. Accordingly, we consider a benefit function with a kink at the 20-yard line, reflecting loss aversion (LA) around that reference point:

$$b_{LA}(y) = \begin{cases} m(y-20) & y \geq 20 \\ \lambda m(y-20) & y < 20 \end{cases}$$

Here, $\lambda > 1$, implying that losses, or returns that fall short of the 20-yard line, are felt more

[11] Another possible form for $b(y)$ is strict concavity in y, or risk aversion in the yard line of the return. Whereas the expected extrinsic benefit function in Figure S1 is linear in y, players may nonetheless exhibit risk aversion in y. Indeed, decisions by NFL coaches about whether to "go for it" on fourth down reflect risk aversion over the in-game probability of winning (Romer, 2006). Coaches typically choose the relative certainty of punting on fourth down over the comparatively risky option of going for the first down, even when the riskier option has a better expected outcome. While it is possible that returners are also risk averse, We do not consider this possibility because it does not reflect their normative goal.

strongly than commensurate gains. For the returner, gaining an extra yard short of the 20-yard line is seen as more beneficial than gaining an extra yard beyond the 20-yard line.

If the cost of effort falls in an intermediate range, then the loss-averse returner exerts high effort only when initial contact occurs before the 20-yard line, whereas the reference-independent returner exerts low effort everywhere. This intermediate range is defined by $C \in \big[m(p-q), \lambda m(p-q)\big)$. If $C < m(p-q)$, then the expected benefit of high effort is always greater than the cost, and both returners exert high effort everywhere. If $C \geq \lambda m(p-q)$, then the expected benefit of high effort for the loss-averse returner when short of the 20-yard line is less than the cost, and neither returner ever exerts high effort.[12] In between, only the loss-averse returner exerts high effort, and only when initial contact occurs short of the 20-yard line.

The distribution of return distance, $f(y)$, follows from Equation 1. Let $f_0(y)$ be the distribution of return distance when the returner is reference independent. For all y, $f_0(y) = (1-p) \cdot g(y-1) + p \cdot g(y)$, a weighted average of the probabilities of contact at $y-1$ and y. Because g is smooth around the 20-yard line:

Proposition 1. *When the returner is reference independent, the distribution of return distance will be smooth around the 20-yard line.*

When the returner has reference-dependent preferences, probability mass is displaced a yard forward when contact occurs short of 20-yard line. Define the excess mass at y relative to the reference-independent baseline as $f(y) - f_0(y)$. Short of the 20-yard line, the excess mass is $(p-q) \cdot \big(g(y-1) - g(y)\big)$, which is negative so long as g is increasing in y. At the 20-yard line, the excess mass is $(p-q) \cdot g(y-1)$, which is positive. Beyond the 20-yard line, the excess mass is zero.

Proposition 2. *Compared to reference independence, the distribution of return distance un-*

[12] We resolve indifference in favor of low effort, yielding this weak inequality.

der loss aversion will be characterized by 1) short of the 20-yard line, displacement probability mass to higher yard lines; and 2) heaping of probability mass at the 20-yard line.

One concern is that $g(y)$, the probability of contact, may not be increasing in y short of the 20-yard line, as assumed. To verify this assumption, a research assistant (who was not informed of the purpose) watched video replays of the 1,122 kickoff returns fielded beyond the goal line during the 2011-14 seasons and marked the yard line of initial contact.[13] Figure S2 shows the distribution of the yard line of initial contact. Short of the 20-yard line, the probability of contact is generally greater at higher yard lines, in line with our assumption.

C Kickoffs

Whether the returner fields the kickoff just behind the goal line or just in front of it appears to be random. Figure S3 shows the distribution of the yard line at which the kickoff is fielded, measured in two-yard increments (marking the depth of an increment by the greater of the yard lines). The distribution appears smooth across the goal line, suggesting that kickers and returners do not manipulate whether the kickoff is fielded just behind the goal line or just in front of it.

We also observe balance on observables for kickoffs fielded on either side of the goal line. Figure S4 shows regression discontinuity estimates for various events as a function of kickoff depth. None of these variables appears to be discontinuous at the goal line. In addition, kickoff depth does not meaningfully vary with the identity of the returner. Restricting the sample to kickoffs fielded one yard on either side of the goal line, a regression of an indicator for kickoffs fielded one yard behind the goal line on returner fixed effects shows no significant differences among returners ($F(407, 1854) = 0.953, p = 0.728$).

[13] Video replays are only available after our window of observation for the main analyses and during a shorter period in which the ball was kicked 65 yards from the goal line, rather than 70 yards.

D Difference-in-difference

We consider an alternative means of adjusting for small differences in the origin of the return. Instead of shifting the control distribution back two yards, we isolate the effect of a two-yard difference absent the touchback option by comparing returns from the 1-yard line to returns from the 3-yard line. We measure the effect of declining the touchback option as a difference-in-difference: the difference between the distributions of returns from one yard behind and one yard in front of the goal line, which confounds a difference in options with a difference in return origin (Figure S7); and the difference between the distributions of returns from one yard and three yards in front of the goal line, which isolates the difference in return origin (Figure S8). This approach assumes that the two-yard difference between the 1- and 3-yard lines has the same effect as a two-yard difference across the goal line. By this estimate (Figure S9), declining the touchback option shifts 4.3% of returns from the 19-yard line to the 20-yard line exactly (se: 2.3pp from 10,000 bootstrap samples of the underlying distributions; $p = 0.062$) but increases the probability of a return to at least the 20-yard line by only 3.2 percentage points (se: 2.7pp; $p = 0.237$).

E Results by vignette

Similar results pertain for each vignette. Figure S10 shows average responses by vignette and condition. For both vignettes, protagonists are seen as less likely to consume the additional unit in the forgone-option condition than in the no-option condition (dessert: Cohen's $d = 0.29$, $t(981) = 4.51$, $p < 10^{-4}$; movie: Cohen's $d = 0.20$, $t(989) = 3.19$, $p = 0.001$). Pooling the vignettes, a regression of the rating on vignette fixed effects and an indicator for the forgone-option condition shows that participants rate the protagonist as significantly less likely to pay more than the counterfactual cost of an option forgone than the counterfactual cost of an option that was never available ($b = -0.42$, se: 0.08, $p < 10^{-4}$).

F Supplemental figures

Figure S1: The expected number of points scored on drives that start with a kickoff returned to the given yard line, excluding touchbacks, for 2000-14 NFL kickoffs. We estimate the expectation as two lines, one for returns that end short of the 20-yard line, and another for returns that reach the 20-yard line and beyond. 95% confidence intervals from Huber-White standard errors.

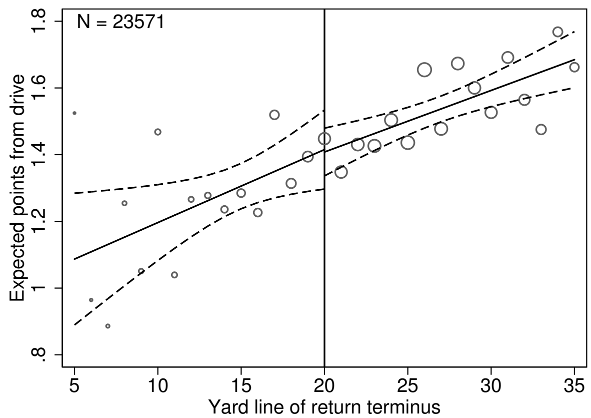

Note: The expectation has a slope of 0.022 left of the 20-yard line—i.e., each extra yard on the return adds 0.022 points in expectation—and a slope of 0.018 right of the 20-yard line; the p-value of this difference is 0.75. Though we allow for a discontinuous jump in the expectation at the 20-yard line, our estimate of this jump is indistinguishable from zero ($p = 0.75$). The marginal benefit of an extra yard is smooth across the 20-yard line. Hence, the kickoff returner's preferences should also be smooth across the 20-yard line. Even if there were a discontinuity in the marginal benefits at the 20-yard line, this jump would presumably be the same for returns from the end zone as for returns from in front of the goal line.

Figure S2: Probability of initial contact at each yard line, for returns from one yard deep in the end zone, the goal line, and the 1- and 2-yard lines. Coded from video replays from the 2011-14 seasons.

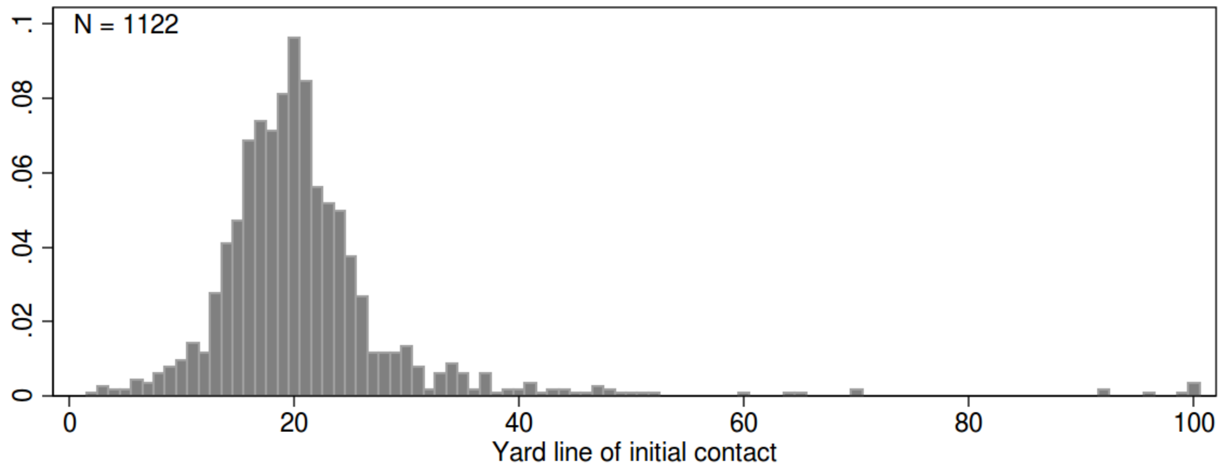

Figure S3: The empirical distribution of kickoff depth.

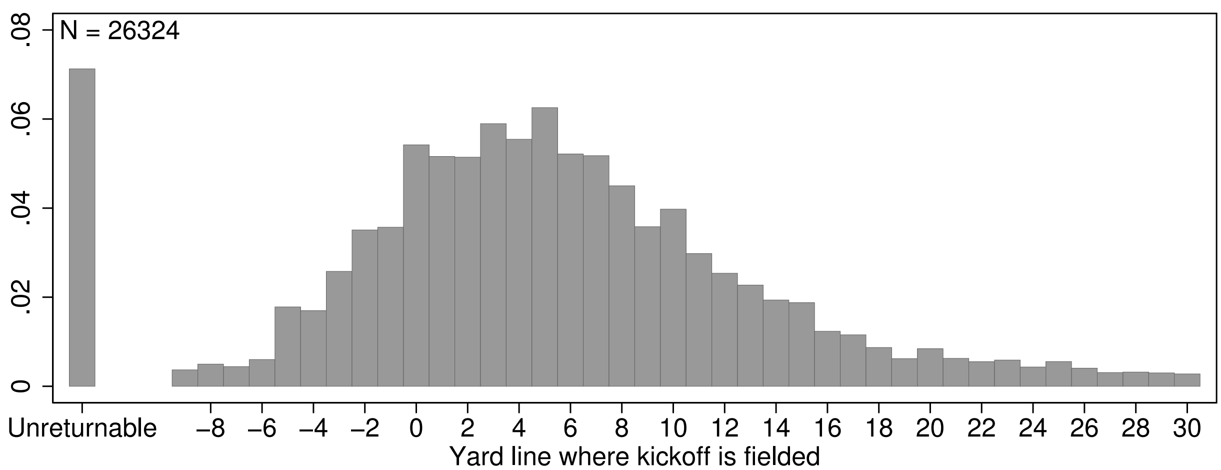

Figure S4: Probability of listed events for kickoffs fielded at given yard lines. We estimate the expectation as two lines, one for kickoffs fielded in the end zone (including the goal line), and another for kickoffs fielded in front of the goal line. 95% confidence intervals from Huber-White standard errors. p-values correspond to tests of the null hypothesis that the two expectations are coincident at the goal line.

(a) Home team kicking off ($p = 0.735$)

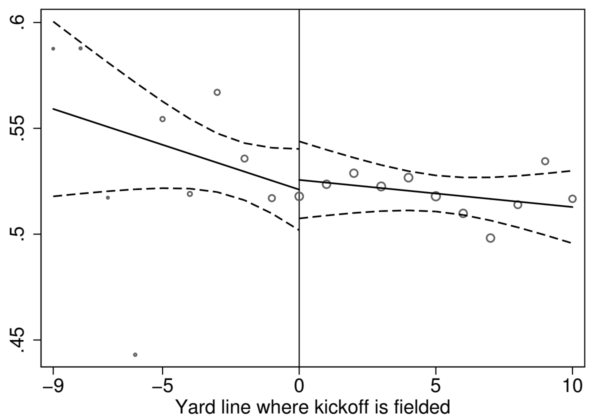

(b) Kicking team leading ($p = 0.358$)

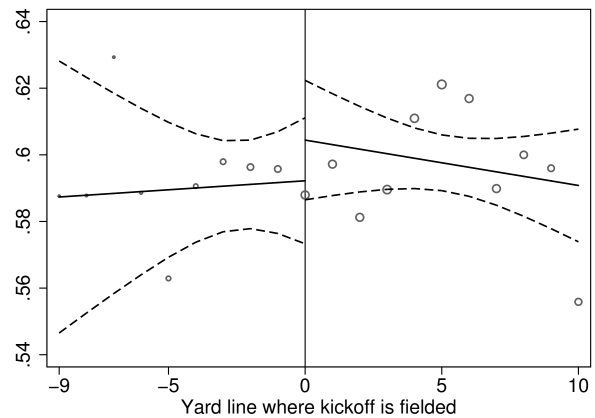

(c) Tie game ($p = 0.367$)

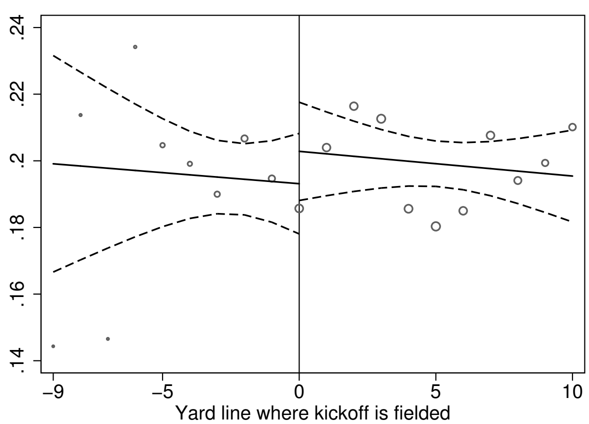

(d) Playoff game ($p = 0.333$)

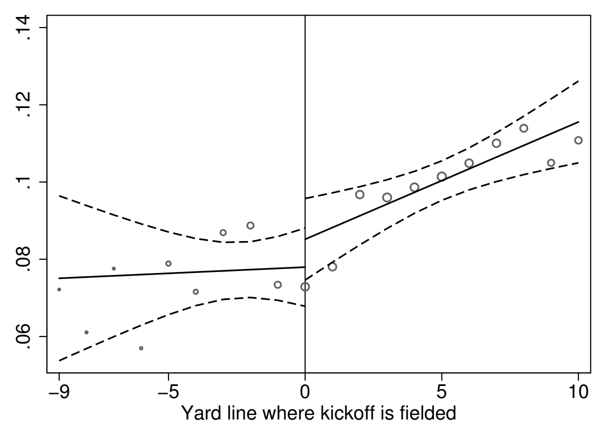

(e) First half ($p = 0.728$)

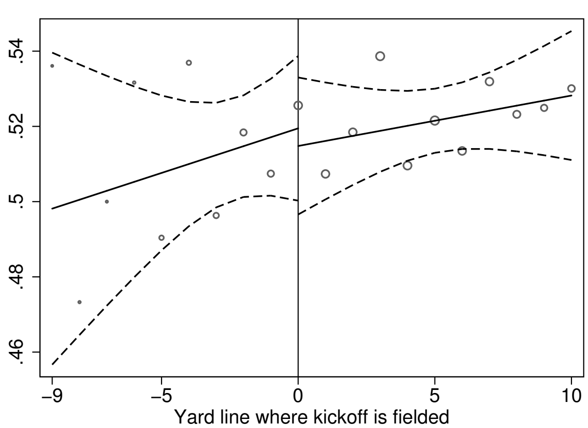

Figure S5: Histogram of the yard line where the returner was tackled or ran out of bounds, for returns from the goal line.

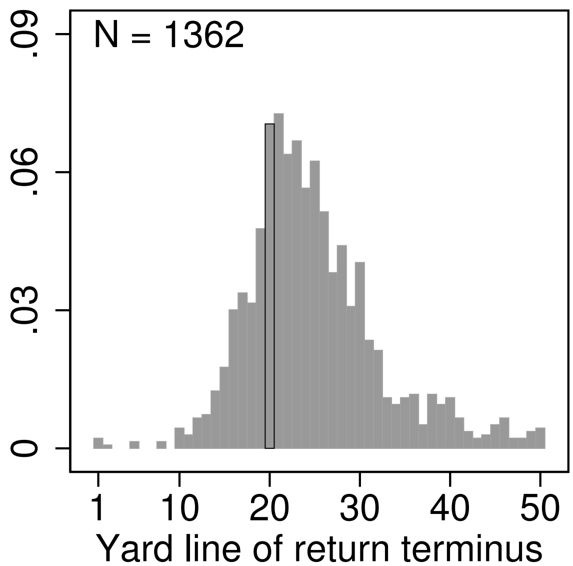

Figure S6: The effect of declining the touchback option on the probability of returning the kickoff to the given yard line, measured as the difference between the distribution of the return terminus for returns from the goal line and the distribution of the return terminus for returns from the 1-yard line, shifted one yard back. 95% confidence intervals from 10,000 bootstrap samples of the underlying distributions.

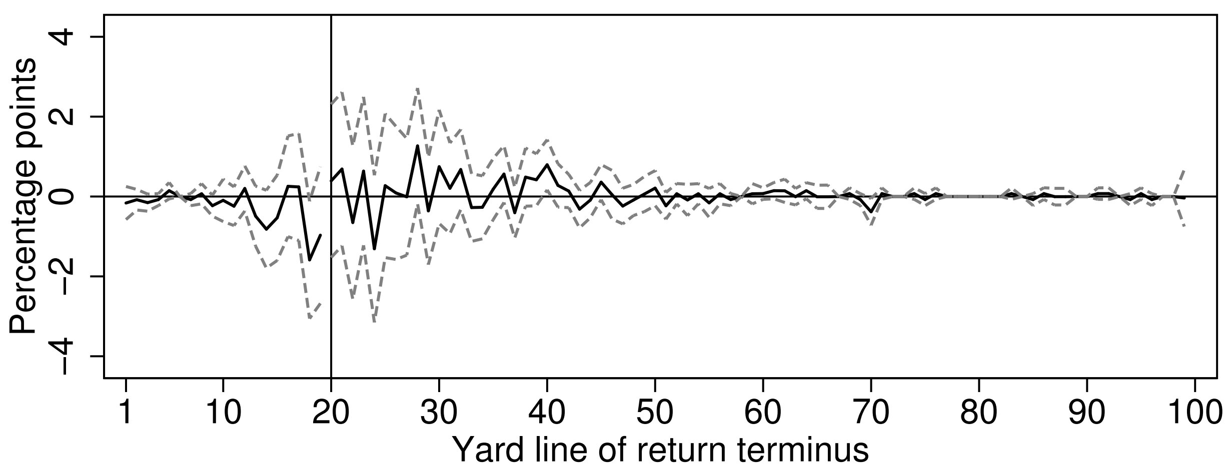

Figure S7: The effect of declining the touchback option on the probability of returning the kickoff to the given yard line, measured as the difference between the distributions of the return terminus for returns from one yard behind and one yard in front of the goal line. 95% confidence intervals from 10,000 bootstrap samples of the underlying distributions.

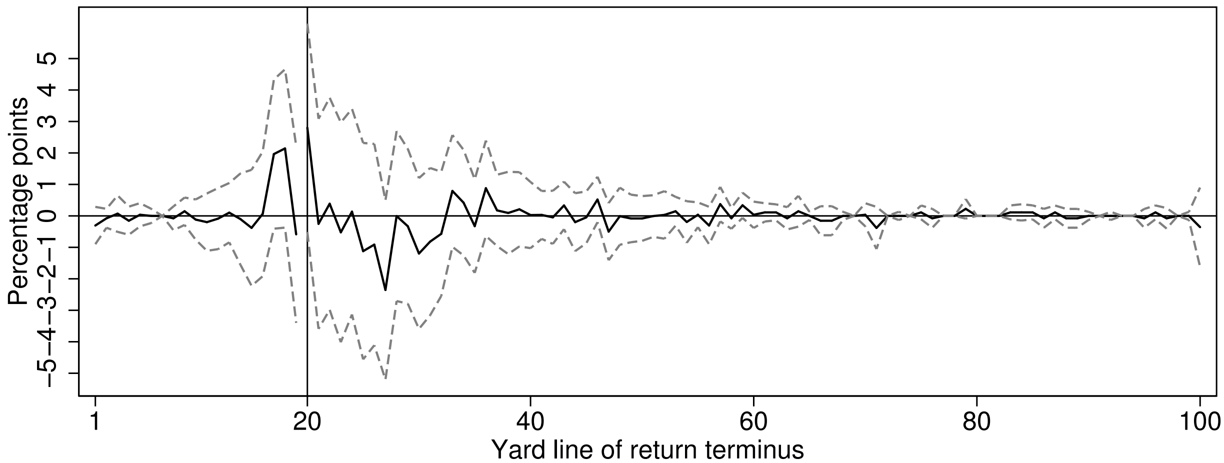

Figure S8: The effect of a 2-yard difference in field position on the probability of returning the kickoff to the given yard line, measured as the difference between the distributions of the return terminus for returns from the 1- and 3-yard lines. 95% confidence intervals from 10,000 bootstrap samples of the underlying distributions.

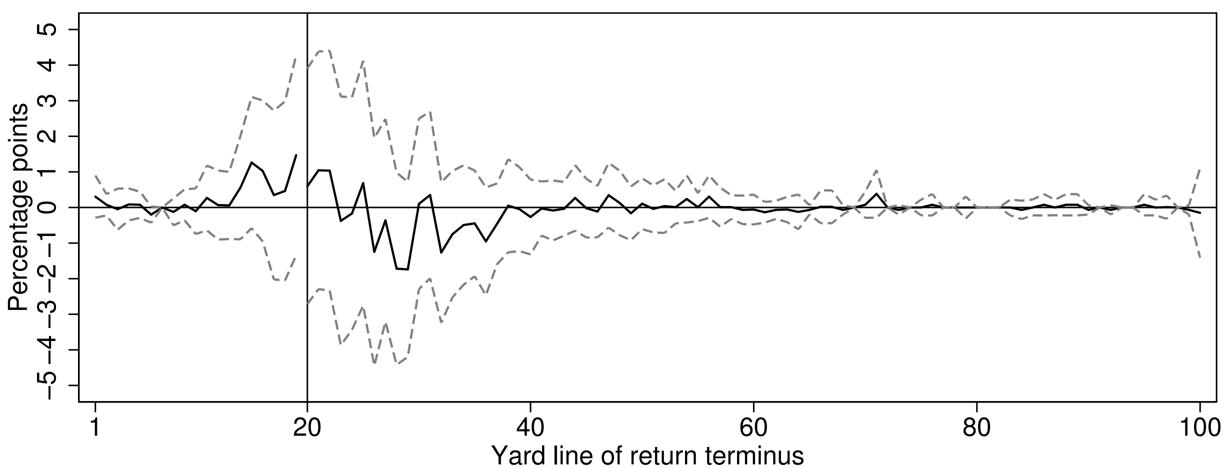

Figure S9: The effect of declining the touchback option on the distribution of return outcomes, measured as a difference in difference—i.e, the difference between 1) the difference between the distributions of the return terminus for returns from one yard behind and one yard in front of the goal line (S7) and 2) the difference between the distributions of the return terminus for returns from the 1- and 3-yard lines (S8). 95% confidence intervals from 10,000 bootstrap samples of the underlying distributions.

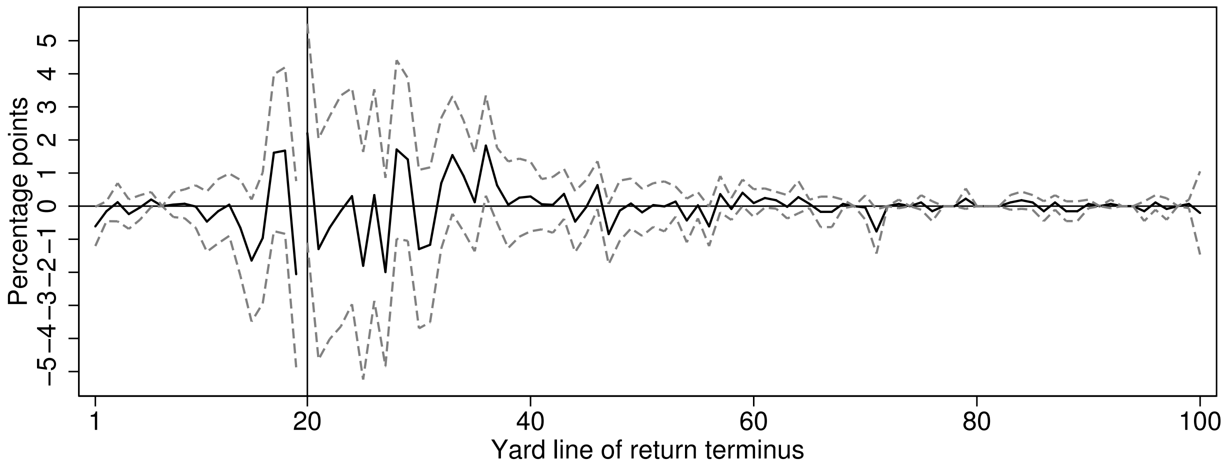

Figure S10: Average rated likelihood of ordering dessert or a pay-per-view movie when doing so would push total expenditures beyond the fixed-cost counterfactual. Error bars denote 95% confidence intervals.