Essays in Behavioral Economics

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Weiwen Leung

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Aldo Rustichini, Adviser

August, 2018

Acknowledgements

I am thankful to my advisor and committee members (Aldo Rustichini, Thomas Holmes, David Rahman, and Paul Schrater) for their guidance during my PhD. I am also indebted to Professors Haiyi Zhu, Joseph Konstan, and Jerry Zhao for their support. In particular, the first chapter of my dissertation was joint work with Professors Zhu and Konstan, and is forthcoming at Computer-Supported Cooperative Work.

I am especially thankful to many graduate student colleagues such as Vegard Nygaard, Elena Falcettoni, Hao Fei Cheng, Zachary Levonian, Sarah McRoberts, Colleen Estelle Smith, Bowen Yu, Raghav Karumur, who always made time for me whenever I needed it.

Dedication

To my family and friends, but more importantly to God, who sustained me through the five years of my PhD, and through whom all things are possible.

Abstract

Behavioral economics has grown immensely since Richard Thaler founded the field. In this thesis, I explore three different areas of behavioral economics. The first chapter explores how mood affects contributions to Wikipedia, an online public good. The second chapter explores whether self-restraint can be a driver of crime. The third chapter explores whether rational inattention in valuations can be a source of risk aversion.

Contents

List of Tables

List of Figures

Chapter 1

Mood and Public Goods Contributions

The emergence of the internet has created a large number of public goods that rely entirely on user contributions. Examples include Wikipedia and Stackoverflow, which have been the subject of extensive study. These communities are often seen as shining examples of how public goods can be privately provided. However, the majority of online communities fail due to nonparticipation (Butler, 2001). Therefore, it is important from a policy and design perspective to examine factors that affect contributions to such public goods.

Such an endeavor also has implications for offline communities, for the public goods literature is well developed within experimental economics. Some key results dating from before the popularization of the internet: Bergstrom, Blume and Varian (1985) show that very weak conditions ensure the existence of a unique Nash equilibrium. On the experimental side, Fischbacher, Gachter and Fehr (2001) show that in the lab, people are conditionally cooperative in providing public goods. For reviews, one can see Chaudhuri (2001) as well as Ledyard (1995).

Here we study one possible determinant of contributions to public goods: mood. While emotions have been neglected by economists for a long time, there is an emerging literature on the subject. For example, Ifcher and Zarghamee (2011) show that mood can affect one's time preferences, while Oswald, Proto and Sgroi (2015) show that people in better moods are more productive.

The key contribution of this paper is to show that mood can affect contributions to public goods in a natural setting.

Specifically, we make use of the fact that around 1500 Wikipedia editors identify as fans of a specific American National Football League (NFL) team. Sport outcomes have clear emotional effects. Wann et al. (1994) studies a variety of sports, and finds that fans experience negative emotions after losses, but positive emotions after wins. Other studies also reach similar conclusions (Wann, Royalty and Rochelle, 2002; Crisp et al., 2007; Jones et al., 2010). Therefore, although we do not observe the exact emotions of the editors in our dataset, sports outcomes have been extensively used across several disciplines as emotional cues for sports fans. Indeed, the impact of emotional cues on sports fans is not negligible and can sometimes be detected among the general population (not just sports fans). For example, at-home male-on-female domestic violence increases after a shock NFL loss (Card and Dahl, 2011). Shock football losses even cause judges to hand out longer sentences (Eren and Mocan, 2016). Soccer even has macroeconomic effects: stock markets respond negatively to soccer losses (Edmans, Garcia and Norli, 2007).

We draw several conclusions from analyzing NFL game outcomes from the 2006 season to the 2016 season. First, after a loss, Wikipedia editors that identify as fans of a particular NFL team decrease their contributions towards football-related pages (relative to after a win). Second, the greater the losing margin, the greater the decrease in contributions. Also, our preferred estimates indicate that for a given losing margin, the effect is four times as large if the loss was a shock loss. Quantitatively, a shock loss of 10 points in our sample translates into a 10% reduction in contributions to football-related pages for each editor. Third, losing to a rival team has a greater impact on contributions than losing to a non-rival team.

In contrast, if a team wins, both whether the win was expected and the size of the victory margin do not affect the volume of contributions. Moreover, edits to pages unrelated to football are unaffected by NFL game results.

1.1 Theory and Hypotheses

How might mood affect editing behavior? Isen and Geva (1987) and Isen (2000) find that the better one's mood is, the more altruistic one is, and the more one elaborates in response to most stimuli.

A simple model will help us generate predictions. Suppose that when deciding whether and how much to edit, users weigh the benefits against the costs. Suppose that the total benefit to writing x characters is γbx, where γ is the weight a user places on other's welfare and b is the (average) benefit to other people from writing one character.

Suppose that $\gamma = \gamma^0 + \mu(y - s, p)$, where y is the team score, s is opponent score, and p is the ex-ante probability of winning. Adapting Koszegi and Rabin (2006), suppose that $\mu(y - s, p) = \alpha(1 - p)(y - s)$ if $(y - s) \geq 0$, and $\beta p(y - s)$ if $(y - s) < 0$. This piecewise definition gives us the classic "kink" of prospect theory. By loss aversion, $\beta > \alpha$.

On the other hand, suppose that costs of editing is superlinear in the number of characters written. For example, assume that the average cost of each character is ax for an edit of x characters; in other words, the total cost of writing x characters is ax^2. (Although we have assumed a quadratic cost function, the results generalize to many other functional forms as well.) Suppose that $a = a^0 + \mu(y - s, p)$, with $\mu(\cdot)$ defined as previously.

Hence $u(x) = \gamma bx - ax^2$. Taking the first order condition, $x^\star = \frac{\gamma b}{2a}$. It is clear that editing increases if either γ increases or a decreases. Both of these happen when the team wins. Also, holding the actual margin constant, the model predicts that the impact of a shock loss is bigger than a non-shock loss, and the impact of a shock win is bigger than a non-shock win.

1.2 Data

1.2.1 NFL data

We gather data on both predicted and actual outcomes of NFL games. Las Vegas bookmakers organize bets for NFL games; the predicted game outcome depends on both supply and demand for such bets.

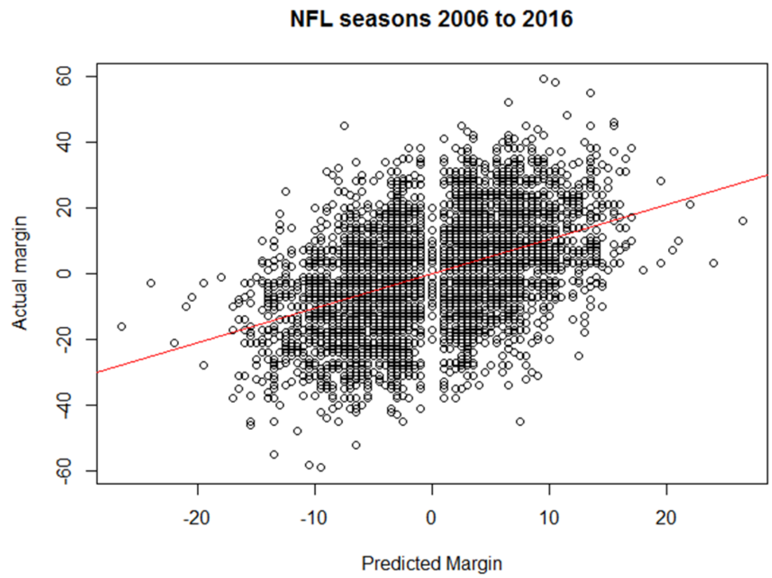

Figure 1.1: Actual versus predicted margins, 2006 to 2016.

We plot the actual margin against the predicted margin for all games from the 2006 season to 2016 season (see Figure 1.1). Note that by convention, more negative scores are better. For example, if a team won by 10 points, then the actual margin is -10. If the predicted margin is -5, then a team must win by 5 points or more for a bet to pay off.

There are a total of 5870 team-games, of which 2656 are predicted to be close games (defined as a predicted margin of between -3 and 3, inclusive of endpoints). Of the non-close games, one-sixth (16%) resulted in upsets.

A regression of actual score against predicted score gives an estimated intercept of -0.005 (s.e. 0.18) and an estimated slope of 1.04 (s.e. 0.03). The 95% confidence intervals for the intercept and slope contain 0 and 1 respectively. Hence, we conclude that the predicted margin does not systematically over- or under-estimate the actual margin. For example, if the predicted margin is 7, then the actual margin is 7 in expectation. Bets normally close a few minutes before the game starts. As such, the evidence suggests that conditional on the predicted margin, the outcome of a game is as good as random at the game's start. We can indirectly test for the impact of emotions by comparing, for example, the impact of shock losses to non-shock losses. (We subsequently explain how we deal with possible confounds).

1.2.2 Wikipedia data

Around 1,500 Wikipedia editors identify as fans of an NFL team in that they have a userbox of an NFL team on their userpage and thus belong to the category "Wikipedian National Football League fans".

We download the contribution history of each user from teams that have 10 or more fans. Around 10% of users identify as a fan of more than one team, and we exclude them from the analysis (since, for example, we are not sure how they will be affected if two of their favorite teams play each other). Thus, we are left with 1,279 users.

We primarily conduct analysis by aggregating contributions to the user-day level, though we also consider the user-hour level, the team-hour level, and the team-day level. Since game times are given in U.S. Eastern Time, we adjust Wikipedia edit timestamps from Greenwich Time to Eastern Time. User-day analysis aggregates all edits by a given user in a given day, and the process is analogous for the other levels of analysis.

We measure contributions in three ways: additions, deletions, and total contributions.

A user's additions during a time period is defined as the total size of all edits which result in a net lengthening of an article in that period (ignoring all edits which resulted in a net shortening of articles). A user's deletions during a time period is the total size of all edits which result in a net shortening of articles in that period (ignoring all edits which resulted in a net lengthening of articles). Total contribution is the sum of additions and deletions. To illustrate, suppose that on a particular day, a user made two edits which lengthened an article by 10 characters and 3 characters respectively. She also made an edit which shortened an article by 1 character. Then her additions for that day is 13, her deletions is 1, and her total contributions is 14.

Our main analysis removes all edits to talk pages (though we also consider what happens to talk pages)[1] . Contributions to all other pages are included.

We first analyze football-related page edits as they are most likely to be affected by emotions. Subsequently, we analyze non-football related page edits, i.e. edits to all

[1] In Wikipedia, talk pages are used by editors to discuss changes to pages

non-talk pages that were not classified as football-related.

1.3 Empirical Framework

1.3.1 Main Analysis

$$\begin{aligned}\log(C_{jdks}) = \beta_0\text{lossmargin} + \beta_1\text{lossmargin} \times 1(\text{predicted win}) + \\ \beta_2\text{winmargin} + \beta_3\text{winmargin} \times 1(\text{predicted loss}) + \\ \eta_j + \gamma_d + \delta_k + \theta_s + \epsilon_{jdks}\end{aligned} \tag{1.1}$$

C_{jdks} is the measure of contributions by user j on day d of week k in season s. The three different measures of contributions (additions, deletions, and total), as defined in the previous section, are used in different regressions. "Loss margin" is defined as max(actualmargin, 0), and "Win margin" is defined as -min(actualmargin, 0). So a winning margin of 3 would indicate the team won by 3 points, and a losing margin of 5 would indicate that a team lost by 5 points. $1(\cdot)$ denotes an indicator variable, with "predicted win" indicating that the team was predicted to win by at least 4 points, and "predicted loss" indicating that the team was predicted to lose by at least 4 points[2] . A set of fixed effects controls for time-invariant differences across users (η_j), as well as day-of-week (γ_d), week (δ_k), and NFL season effects that do not differentially impact users (θ_s). These collectively allow us to focus on differences in Wikipedia edits that are due to differences in game outcomes. Standard errors are clustered by user (Cameron and Miller, 2015).

Our coefficients β_0 to β_3 are estimated using OLS. They can be interpreted as follows:

- If a team was not predicted to win by at least four points, then a 1 point increase in the losing margin will change edits by $100\beta_0\%$, all else equal.
- A shock loss of X points changes contributions by $100\beta_1 X$ percentage points more compared to a non-shock loss[3] of the same margin.

[2] We take a predicted winning margin of 4 points as the threshold at (and above) which people are reasonably certain that a team will win. Our main results are robust to increasing this threshold.

[3] This is a loss when the team was expected to lose, or win by less than four points.

- If a team was not predicted to lose by at least four points, then a 1 point increase in the winning margin will change edits by $100\beta_2\%$, all else equal.

- A shock win of X points changes contributions by $100\beta_3 X$ percentage points more compared to an non-shock win[4] of the same margin.

Prospect theory (in particular, the Koszegi-Rabin framework) makes several predictions. First, $\beta_0 \leq 0, \beta_1 < 0$ (losses are painful and more so if unexpected). In addition, we would also expect that $|\beta_0| > |\beta_2|, |\beta_1| > |\beta_3|$, which intuitively means two things. First, changes in the losing margin have more impact than changes in the winning margin. Second, the impact of a shock loss relative to a non-shock loss is bigger than the impact of a shock win relative to a non-shock win. Finally, we would also expect that $\beta_2 \geq 0, \beta_3 \geq 0$.

1.3.2 Hourly Analysis (Football related pages)

The results from our analysis in the previous section are generally consistent with our hypotheses (see Results section), so we next perform analysis using three-hour time windows to further examine if our data are consistent with a causal effect. We consider several intervals on game day itself: noon to 3pm, 3pm to 6pm, 6pm to 9pm, and 9pm to midnight.

Specifically, we are interested in whether the changes in edit volume occur after the game, rather than during or before the game. We make use of the fact that around two-thirds of NFL games start at 1pm, and around a quarter start at 4pm. (Recall that all time windows are U.S. Eastern time.) Since 1pm games do not end before 3pm, we should not detect an effect before the 3pm window. Likewise, because 4pm games do not end before 6pm, we should not detect any effect before the 6pm window.

Our analysis extends the regression model we previously used to allow for separate coefficients for 1pm games and 4pm games[5] (We ignore all games that do not start at these times, because there are so few of them). We fit separate models on edits that occur in each of the three-hour intervals stated above. Our model is as follows:

[4] This is a win when the team was expected to win, or lose by less than four points.

[5] We could estimate the effects of the game relative to start time. However, our approach allows us to exploit the differential start times of games, thus giving us greater confidence for a causal interpretation.

$$
\begin{aligned}
\log(C_{jdks}) = \beta_0 \text{LM} \times \text{1pm} + \beta_1 \text{LM} \times 1(\text{predicted win}) \times \text{1pm}+ \\
\beta_2 \text{WM} \times \text{1pm} + \beta_3 \text{WM} \times 1(\text{predicted loss}) \times \text{1pm}+ \\
\beta_4 \text{LM} \times \text{4pm} + \beta_5 \text{LM} \times 1(\text{predicted win}) \times \text{4pm}+ \\
\beta_6 \text{WM} \times \text{4pm} + \beta_7 \text{WM} \times 1(\text{predicted loss}) \times \text{4pm}+ \\
\eta_j + \gamma_d + \delta_k + \theta_s + \epsilon_{jdks}
\end{aligned}
\tag{1.2}
$$

where LM refers to loss margin and WM refers to win margin. We then adapt our strategy to analyze what happens from 12:01am of the day after the NFL game to 9am, 9am to noon, noon to 3pm, and three hour intervals thereafter until midnight. We use all games regardless of start time.

In other words, we run the following regression six times:

$$
\begin{aligned}
\log(C_{jdks}) = \beta_0 \text{LM} + \beta_1 \text{LM} \times 1(\text{predicted win})+ \\
\beta_2 \text{WM} + \beta_3 \text{WM} \times 1(\text{predicted loss})+ \\
\eta_j + \gamma_d + \delta_k + \theta_s + \epsilon_{jdks}
\end{aligned}
\tag{1.3}
$$

Each time we run the regression, our dependent variable only uses contributions from the specified time window.

1.3.3 Games between rivals

Another sub-analysis exploits the idea that if the link between NFL outcomes and Wikipedia edits is driven by emotion, one would expect that games that are more emotionally charged have larger effects. Games that are played against traditional rivals are likely to induce stronger emotions than games that do not. We determine whether two teams are traditional rivals through the Wikipedia page titled "National Football League rivalries". We examine whether such emotionally salient games induce bigger reactions through the following model:

$$
\begin{aligned}
\log(C_{jdks}) = \beta_0 \mathrm{LM} + \beta_1 \mathrm{LM} \times 1(\text{predicted win}) + \\
\beta_2 \mathrm{WM} + \beta_3 \mathrm{WM} \times 1(\text{predicted loss}) + \\
\beta_4 \mathrm{LM} \times \mathrm{RT} + \beta_5 \mathrm{LM} \times 1(\text{predicted win}) \times \mathrm{RT} + \\
\beta_6 \mathrm{WM} \times \mathrm{RT} + \beta_7 \mathrm{WM} \times 1(\text{predicted loss}) \times \mathrm{RT} + \\
\eta_j + \gamma_d + \delta_k + \theta_s + \epsilon_{jdks}
\end{aligned}
\tag{1.4}
$$

where RT ("rival teams") is an indicator variable equal to 1 if both teams were traditional rivals.

Notice that $100\beta_0 X\%$ is the impact of a non-shock loss by X points in a "typical" (i.e. non-rival teams) game, while $100\beta_4 X\%$ captures the additional effect if the game was between rival teams. Hence the total impact of a loss in a game between rivals is $(100\beta_0 X + 100\beta_4 X)\%$. Other impacts can be calculated through an analogous process.

1.3.4 Non-football related pages

We also analyze whether edits to non-football-related pages are affected. Recall that non-football-related pages are all non-talk pages that were not classified as football-related. The analysis is essentially identical to our main analysis for football-related pages, except that the dependent variable now measures edits to non-football-related pages.

1.4 Results

As mentioned, we start by considering the impact of game results on edits to football related pages.

Tables 1.1, 1.2, 1.3, illustrates what happens to NFL pages on the day of the game, the day after the game, and two days after the game. Note that on the day of the game, we only use a twelve-hour time window (between noon and midnight Eastern time; all timings are normalized to Eastern time), since games only start in the afternoon. For the day after and two days after, however, we use the entire 24 hours of the day.

There appears to be an impact on the day of the game as well as the day after the game, but not two days after.

Table 1.1: Day of game: Football-related pages

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.00222**	-0.00209**	-0.00107
	(0.00108)	(0.00101)	(0.000870)
lossmargin×predictedwin	-0.00803**	-0.00596*	0.000102
	(0.00391)	(0.00312)	(0.00219)
winmargin	-0.000322	-0.000964	0.000233
	(0.00112)	(0.00115)	(0.000599)
winmargin×predictedloss	0.00127	0.00133	0.00169
	(0.00216)	(0.00231)	(0.00214)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses, clustered by user

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

When examining the coefficients in Tables 1.1 and 1.2, we note that they are generally consistent with our hypotheses. First, shock losses have a bigger impact than non-shock losses (since $\beta_1<0$), and losses decrease contributions (since $\beta_0<0$). This is consistent with our initial predictions that were based on the Koszegi-Rabin model (For a discussion of effect size, refer to the next section).

Second, wins do not have any statistically significant effect whether or not they were expected. This is not consistent with the psychology studies cited earlier: that people in better moods elaborate more. We can only speculate as to the reasons for this result. It could be that loss aversion causes an asymmetric effect of wins and losses; hence, people react much less strongly to wins than losses. If the asymmetry is strong, the effect of wins could be very small. Hence, it may be difficult or impossible to detect the effect of wins. Indeed, studies on the impact of sporting outcomes on the stock market (Edmans, Garcia and Norli, 2007), judge decisions (Eren and Mocan, 2016), and domestic violence (Card and Dahl, 2011) all failed to find an impact of wins.

In addition, there are other possible explanations for why we do not observe an effect of win margins, or unexpected wins relative to wins that were not unexpected. For example, increased productivity (or altruism) could be counterbalanced by an increased

Table 1.2: Next day: Football-related pages

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.00252*	-0.00328**	0.0000600
	(0.00146)	(0.00143)	(0.000800)
lossmargin×predictedwin	-0.00751**	-0.00557*	-0.00289
	(0.00350)	(0.00331)	(0.00218)
winmargin	0.00119	0.00131	-0.00000640
	(0.00122)	(0.00114)	(0.000726)
winmargin×predictedloss	0.00142	0.00160	-0.000849
	(0.00421)	(0.00394)	(0.00303)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses, clustered by user

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3: Day after: Football-related pages

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.00122	-0.00272	-0.000111
	(0.00119)	(0.00207)	(0.000868)
lossmargin×predictedwin	0.00220	0.00195	0.00121
	(0.00355)	(0.00350)	(0.00231)
winmargin	-0.000880	-0.000897	0.000159
	(0.00109)	(0.00101)	(0.000815)
winmargin×predictedloss	-0.00318	-0.00271	0.000850
	(0.00330)	(0.00544)	(0.00222)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses, clustered by user

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

tendency to go outdoors and celebrate. Future research could attempt to identify the mechanism. However, our results are still consistent with the behavioral economics theories we cited earlier, specifically loss aversion: people react more strongly to losses than gains. That in itself has design implications (see the Discussion section).

Third, consistent with the psychological theories cited earlier (Isen, 2000, 2007), which predict a change to additions and not deletions, the statistically significant coefficients of losses of total contributions (i.e. column 1 of these tables) are driven by reductions to additions; there are no statistically significant changes to deletions.

1.4.1 Results: Hourly Analysis (Football related pages)

For our hourly analysis, we indeed start detecting effects from the 3pm window for the 1pm games, and from the 6pm window for the 4pm games (though the estimates are noisier for 4pm games due to the smaller sample size). For readers' convenience, we use our regression results to evaluate the predicted percentage reduction in total contributions due to a 10 point shock loss and display it in Figures 1.2 and 1.3.

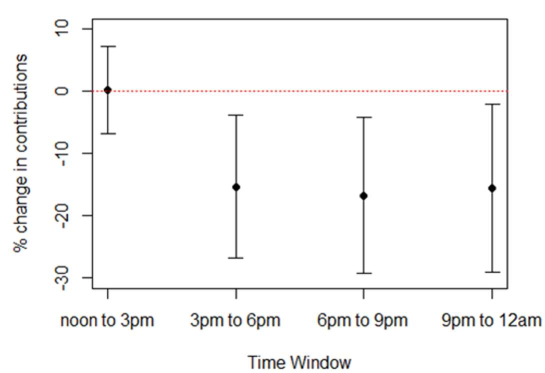

Figure 1.2: Effect of a 10 point shock loss in 1pm games on edits that day.

Figure 1.2 shows the impact of shock losses in games that start at 1pm (for which the results are only known after 3pm). Notice that contributions from noon to 3pm are unaffected by a future shock loss that occurs after 3pm. Indeed, the point estimate of the leftmost confidence interval in Figure 1.2 is close to zero and relatively precisely estimated. In contrast, the other three confidence intervals in Figure 1.2, which study edits after the shock loss occurs, all show negative impacts of the shock loss. Not only are all the point estimates negative, the upper bounds of the confidence intervals all lie below zero.

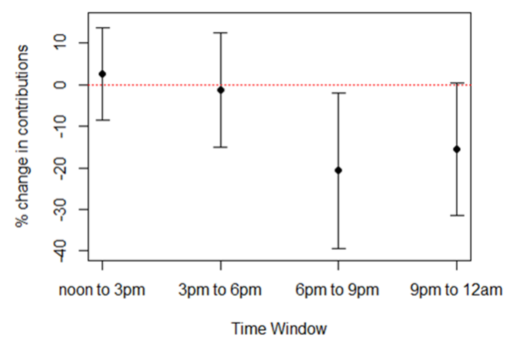

Figure 1.3: Effect of a 10 point shock loss in 4pm games on edits that day.

Figure 1.3, which shows the impact of shock losses in games that start at 4pm, shows that an effect of a shock loss can only be detected after 6pm, as the two confidence intervals on the left comfortably encompass zero. In contrast, the two confidence intervals on the right side of the figure either do not include zero, or barely include it. Hence, these results lend support to the idea that it is the game outcome that matters.

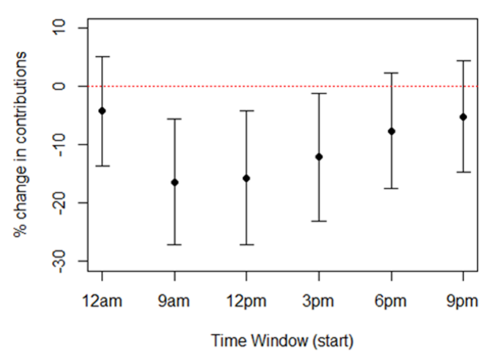

Figure 1.4: Effect of 10 point shock loss (all games) on next-day edits.

For contributions on the day following the loss, the pattern of decline is largely consistent with a short-term emotional cue in that it decays from the 9am window until midnight; Figure 1.4 summarizes our results; recall from the Methods section that the time windows are 12am to 9am, 9am to 12pm, 12pm to 3pm, 3pm to 6pm, 6pm to 9pm, and 9pm to 12am. Except for the first window, notice that the confidence intervals gradually get closer and closer to zero as the day progresses. (We may not have been able to detect an effect in the 12:01am to 9am window as most people may have been asleep.)

1.4.2 Results: Games between rivals (Football related pages)

Consistent with our interpretation that the underlying mechanism is mood, we find that losses to a rival indeed result in larger effects that losses to a non-rival. Table 1.4 shows the coefficient estimates when we use the day after the NFL game (as an example). Notice that the coefficients in the fifth and sixth rows (lossmargin$\times$ES and lossmargin$\times$1(predicted win)$\times$ES) are negative. In fact, both coefficients are larger than their counterparts in the first two rows, suggesting that a loss to a rival has at least twice as much impact compared to non-rival.

1.4.3 Results: Non-football-related pages

We find that NFL outcomes have no significant effect on edits to non-football-related pages. (Recall that non-football-related pages are all non-talk pages that were not in the "National Football League" category, or related categories.) Tables 1.5 to 1.7 do not give any evidence of any effect at the daily level. Nor do we find any effects when using three-hour windows, or when we use emotionally salient games. These results are unlikely to be due to a lack of statistical power as the typical editor in our sample makes around 10 edits to non-football related pages for every edit to a football-related page.

One may be concerned that some users may not follow NFL games and their results, even though they have a userbox of an NFL team on their userpage. To examine if this could explain the failure to detect a significant result, we remove the subset of editors that have never contributed to a football-related page (even though they identify as a football fan). Our estimates remain statistically insignificant. We conclude that NFL game outcomes have no effect on the volume of contributions to non-football related

Table 1.4: Games between rivals

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.00198**	-0.00253**	0.0000721
	(0.000901)	(0.00122)	(0.000912)
lossmargin×predictedwin	-0.00372**	-0.00429*	-0.00311
	(0.00156)	(0.00220)	(0.00268)
winmargin	0.00123	0.00144	-0.0000104
	(0.00172)	(0.00126)	(0.00899)
winmargin×predictedloss	0.00116	0.00150	-0.000249
	(0.00454)	(0.00421)	(0.00134)
lossmargin×ES	-0.00203**	-0.00331*	-0.0000326
	(0.00100)	(0.00167)	(0.000311)
lossmargin×predictedwin×ES	-0.00724*	-0.00461*	0.00178
	(0.00360)	(0.00242)	(0.00209)
winmargin×ES	-0.000923	0.000873	0.0000113
	(0.00157)	(0.00102)	(0.0000294)
winmargin×predictedloss×ES	0.00107	-0.000649	-0.000179
	(0.000982)	(0.00137)	(0.00123)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses; *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 1.5: Day of game: non-football-related pages

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.00137	-0.00115	-0.000949
	(0.00158)	(0.00151)	(0.00102)
lossmargin×predictedwin	-0.00143	0.000480	-0.00110
	(0.00363)	(0.00348)	(0.00255)
winmargin	-0.00167	-0.00178	-0.000193
	(0.00133)	(0.00133)	(0.000896)
winmargin×predictedloss	0.00115	0.00212	-0.000818
	(0.00493)	(0.00478)	(0.00344)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Next day: non-football-related pages

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.00147	-0.00148	-5.51×10^{-6}
	(0.00167)	(0.00154)	(0.00112)
lossmargin×predictedwin	-0.00248	-0.00324	0.000429
	(0.00413)	(0.00376)	(0.00291)
winmargin	-0.00171	-0.00134	-0.00195
	(0.00149)	(0.00142)	(0.00917)
winmargin×predictedloss	-0.00135	0.00220	0.00423
	(0.00535)	(0.00505)	(0.00367)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.7: Day after: non-football-related pages

	(1) Total	(2) Additions	(3) Deletions
lossmargin	-0.000509	-0.000827	0.000650
	(0.00170)	(0.00164)	(0.00112)
lossmargin×predictedwin	-0.000401	-0.000294	-0.000535
	(0.00387)	(0.00365)	(0.00288)
winmargin	-0.00183	-0.00142	-0.00243
	(0.00241)	(0.00136)	(0.00940)
winmargin×predictedloss	-0.000584	0.000988	-0.00368
	(0.00527)	(0.00503)	(0.00358)
User Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

pages.

We also do not find any statistically significant effect on edits to talk pages, whether football-related or not. However, given that the football-related talk page edits is only 4% of the size of the football-related page edits, and the effects we detect for football-related page edits have p-values between 0.02 to 0.09, it is unlikely that we have sufficient statistical power to examine whether there are changes in football-related talk page edits[6] .

[6] To examine how the statistical power we have for football-related talk page edits compares to that of football-related page edits, we randomly take a subsample of 4% of edits from our football-related page edits dataset, conduct the same analysis using our subsample, note whether the results are statistically significant, and repeat this procedure ten times. We only find a statistically significant effect for shock losses for two out of ten times, suggesting that we have low statistical power for football-related talk page edits.

1.5 Effect Size, Duration, and Scope

1.5.1 Effect Size

The average shock loss in our sample is 8.5 points[7]. Using our preferred estimates in Table 2, this translates into an 8.5% reduction in editing activity to football-related pages (relative to the case of a win)[8].

This effect size is comparable to previous studies that have examined other effects of football game results. Eren and Mocan (2016) find that a shock loss increases sentence severity for black defendants by almost 9 percent (white defendants are unaffected). Card and Dahl (2011) find that at-home male-on-female violence increases by 10 percent (violence against other family members, as well as away-from-home male-on-female violence, are unaffected)[9].

How does this translate into characters? Conditional on there being edits to football-related pages, the length of all edits to such pages averages 1795 characters at the individual-day level, and the corresponding median is 227 characters. This translates into an impact of 153 characters for the mean individual-day, and 19 characters for the median individual-day. This suggests that the loss does not merely impact whether game results are updated; other substantive contributions may also be affected.

Though this may not seem large, small changes add up: In May 2016, the longest featured Wikipedia article (on Elvis Presley) was 17659 words[10], or around 80000 letters. That is equivalent to 524 installments of 153 characters each.

Moreover, our regressions may underestimate the conditional mean effect, for as discussed later, our estimates for β_0 and β_1 are biased towards zero.

---

[7] The average non-shock loss is 11.8 points, which translates into a 3% reduction in activity using the method explained in the next footnote.

[8] The 8.5% figure is calculated as follows. We add the estimated coefficients of β_0 and β_1 together (-0.00252 - 0.00751 $\approx$ -0.01). Next, we multiply it by 100 since it is a log-linear model. Finally, we multiply it by 8.5, since the average shock loss is 8.5 points.

[9] Our results are not directly comparable to studies which use other emotional cues and do not look at the differential effect of shock losses, but for convenience, Edmans, Garcia and Norli (2007) find that soccer World Cup losses are associated with an abnormal stock return of -49 basis points, while the standard deviation of daily returns is 144 basis points.

[10] https://blog.wikimedia.org/2016/05/12/rock-n-scroll-english-wikipedias-longest-featured-articles/

1.5.2 Effect Duration

The effect duration we have found (around 24 hours) is longer than that of Card and Dahl (2011), who found that domestic violence increases in a three-hour window following an NFL game, but shorter than Eren and Mocan (2016), who find that judge decisions are affected up to a week after a game.

1.5.3 Effect Scope

One very important question is why only football-related page edits are affected. Here, we offer several candidate explanations. First, it could be due to "escapism": editors avoid editing NFL pages because they dont want to be reminded of the recent unpleasant experience. Escapism is well documented in psychology, (e.g. Hirschman, 1983; Addis and Holbroock, 2007). Boen, Vanbeselaere and Feys (2002) give evidence consistent with this interpretation by showing that soccer teams had fewer visitors after losses compared to wins. Bizman and Yinon (2002) show that basketball fans engage in distancing tactics after their team loses. Second, even if editors still edit football-related pages, it could be that editors think of the recent loss when they edit football-related pages, but not when they edit unrelated pages. Third, due to space constraints, we have only studied the effect on contribution volume, and not other aspects such as contribution quality. It could be that contribution quality to non-football-related pages are affected. Fourth, the two other studies on domestic violence and judge sentencing have also found domain-specific effects of NFL matches, so it could be that the effect of NFL emotional cues are arbitrarily domain-specific.

1.5.4 Possible confound: news

One factor that can threaten the validity of the analysis is that the actual score may be correlated with the amount of Wikipedia-worthy news. For example, huge losses or victories may provide material for Wikipedia edits.

Given that our results indicate that fans of teams which suffer huge losses edit less compared to the counterfactual where their team had experienced a huge victory, this does not seem very likely. However, we conduct two further checks to determine whether Wikipedia-worthy news could be a confounding factor.

First, we restrict our dataset to games where the actual margin was between -10 and 10. Conditional on the actual margin being within this interval, and keeping the actual game result constant (i.e. not allowing wins to become losses and vice-versa), it seems unlikely that changes in the winning margin would translate into additional Wikipedia worthy news. When we perform this check, we find that our estimated coefficients are essentially unchanged. These suggest that our results are unlikely to be driven by additional news associated with NFL game scores.

Second, we analyze theoretically how news associated with NFL games might bias our coefficients, and we find that the bias often has the opposite sign of our estimated coefficients. To elaborate, since a larger absolute winning/losing margin is associated with additional Wikipedia news (and shock losses likely generate more news relative to non-shock losses), this should bias our estimates for β_0 and β_1 upwards. In other words, the news associated with NFL scores could prevent us from finding that losses have an effect if there is indeed an effect. However, given that we find that losses have an effect (there is sufficient evidence to conclude that $\beta_0, \beta_1 < 0$), we conclude that the effect of losses is at least as strong as what our estimates suggest.

Also, β_2 and β_3 will be biased upwards. This means that if we observe that $\beta_2 > 0$ and/or $\beta_3 > 0$, it could be spurious (i.e. due to extra news and emotions). Since we find that wins do not have an effect ($\beta_2, \beta_3 = 0$ is not rejected), and from a theoretical standpoint it is not as easy to argue that wins should be associated with negative emotions as compared to losses, we are not too concerned with the possibility of bias. Finally, under the assumption that a loss by a certain margin generates as much news as a win by a certain margin, the mere fact that $|\beta_0| > |\beta_2|$ and $|\beta_1| > |\beta_3|$ is suggestive of emotions playing a role.

Therefore, we conclude that our result is unlikely to be driven by news associated with the game outcome, or by the margin associated with that outcome.

1.5.5 Other confounds

Reverse causality is not an issue as the changes in editing volume occur after the game finishes. We now consider whether the link between game outcomes and Wikipedia editing volume can be driven by a third variable. Because bets close a few minutes before the game starts, the predicted margin essentially takes into account all events leading up to the game. It remains for us to consider whether events during the game can drive both game outcomes and Wikipedia editing volume (independent of emotions). Events such as unexpected injuries to players of a team one is a fan of could lead to shock losses, be worthy Wikipedia news, and may not necessarily affect ones emotions (if that player is not liked). However, that would actually lead to an effect that is the opposite of what we observe. To be sure, some events could conceivably explain the links we observe between shock losses and decreases in edits, but they would have to regularly occur, and we do not feel it is very likely that this is the case.

1.5.6 Summary of results

Analyzing the edits of Wikipedia editors who identify as football fans, we find that:

- Losses are associated with fewer edits to football-related pages as compared to wins
- Shock losses have a bigger impact on football-related pages as compared to non-shock losses
- Losses to a rival, which are likely to be more emotionally charged, induce bigger effects to football-related pages.
- Changes in edit volume occur after the game ends (and not before)
- Conditional on winning, neither the margin of a win, nor whether the win was expected, affect football-related edits.
- Edits to non-football-related pages are unaffected

1.6 Discussion

Our study has several implications. First, that mood can affect contributions to public goods in a natural setting. In doing so, we contribute both to the literature on the effects of mood, as well as the public goods literature. Mood has been shown to affect a wide variety of outcomes, such as productivity, risk aversion, and time preference. Here we show that contributions to one of the most important sites on the internet can be affected by mood.

Second, we also find field evidence consistent with the Koszegi-Rabin model of "expectations-based reference-dependence preferences", suggesting that reference points ought to be taken into consideration when analyzing observed behavior, or even theorizing about behavior.

Third, our study also has implications for online community designers. Designers can consider adopting mood detection software, such as those trained to detect emotion through keystrokes (Kolakowska, 2013). Administrators of peer production communities can consider employing such software to detect users' mood to (i) determine what affects their user's moods, (ii) to study how emotions affect their users interaction with their platforms, and (iii) study the effectiveness of proposed interventions, such as those we speculate on. Our study is also a clarion call for the further development of the software in Kolakowska (2013) and related software. Moreover, while designers may consider both increasing positive affect and decreasing negative affect, our results suggest that the latter could be more effective in spurring contributions. As such, designers of peer production communities ought to consider having "nudges" that can reduce negative affect. One reliable method for mood inducement is to give feedback (Cheng, Bernstein and Leskovec, 2017), or to show video clips e.g. of comedies (Ifcher and Zarghamee, 2011).

1.6.1 Limitations and Future Research

One key limitation of our study relates to scope: we have only studied the volume of edits, primarily due to space constraints. It could be that edit impact and quality Halfaker et al. (2009) are affected. Future research could examine whether, for example, edits that are made after losses are more likely to be reverted because they are of lower quality. One can even imagine that "edit wars" (battles between editors over what phrases should be used in an article) may more likely after one is exposed to certain kinds of negative affect. Conceivably, one could also employ sentiment analysis to determine whether the tone of words changes. Even though Wikipedia is ostensibly a neutral-point-of-view encyclopedia, one can imagine that editors have some latitude to choose adjectives when editing, and some adjectives are stronger than others.

Two other limitations relate to the nature of our sample. Consider "false negatives": some Wikipedia editors may not report that they are NFL fans (even though they are). Our estimates capture the effect of NFL game outcomes of self-identified NFL fans, and are internally valid as such. However, if the average "false negative" is less affected by NFL game outcomes than the average self-identified fan, then our estimates may not be readily interpretable as applying to all Wikipedian NFL fans (only self-identified ones).

A second limitation arising from the nature of our sample is external validity. While we believe that our findings apply to other emotional cues (e.g. weather or disaster-like events), and have discussed design implications, our claims should be taken with caution. Future research could examine the degree to which our findings are externally valid.

Two other limitations arise from our study being a natural experiment. First, while our study is well designed to identify the effect of emotional cues, it does not directly inform us what can be done to mitigate negative affect, including the negative affect generated by NFL shock losses. Hence, the design implications we previously gave are speculative. Second, while we have discovered a link between negative affect and contribution volume, and studies examining sport and mood suggest this negative affect is likely to be a combination of anger and sadness Crisp et al. (e.g. 2007), we do not whether anger or sadness is more important (if any). Future research, possibly involving sentiment analysis, should examine this.

Chapter 2

Self-Restraint and Crime

Crime imposes huge costs on society. For example, the social costs of a murder have been estimated to be at least \$1 million (McCollister, French and Fang, 2010). Under the rational model of crime (Becker, 1968), potential criminals commit crimes if the expected benefits to committing the crime exceed the costs. Indeed, much of the literature has focused on criminals' response to the costs (e.g. Drago, Galbiati and Vertova, 2009; Levitt, 1997; Doleac and Sanders, 2015) and benefits (e.g. Shoukry, 2016; Levitt, 2006) of committing a crime.

However, in economics at least, most existing research has abstracted away from the fact that behavioral factors such as self-restraint can play an important role in the decision to commit crime[1] . At some level, it seems intuitive that some crime could be carried out when, for example, one is caught in a situation of low self-restraint (such as when in a fit of rage), and experiments indeed suggest that lowered self-control makes one respond more aggressively to provocations (DeWall et al., 2007). However, evidence is scarce that self-restraint could lead to crime at the aggregate level.

On the other hand, criminology has long recognized the importance of self-restraint, starting from Gottfredson and Hirschi (1990)'s famous theory that individuals with low self-control are more prone to commit crime. While there have been many empirical tests of this theory (e.g. Vazsonyi et al., 2001; Pratt and Cullen, 2000), most (if not

[1] There has been some literature linking environmental factors to crime, but factors like temperature and rainfall can very plausibly affect income. Perhaps the notable exception is the working paper of Herrnstadt and Muehlegger (2015), who show that battery arrests in Chicago increase when air pollution increases

all) of them are correlational in nature and arguably use measures of self-control that are poorly defined and measured (Marcus, 2004). Therefore, our use of Ramadan as a natural experiment can also be viewed as an attempt to examine whether the correlation is actually due to a causal relationship.

In this paper, we use a natural experiment to examine if self-restraint can play an important role. In particular, we analyze UK data, focusing on the Islamic fasting month of Ramadan from 2011 to 2015. During Ramadan, Muslims are not supposed to eat, drink, smoke, or engage in sexual activity from dawn to dusk during that month. It is well known that self-control operates like a muscle (Muraven and Baumeister, 2000): people have a limited capacity for self-control, and engaging in acts of restraint deplete this resource. Fasting requirements are particularly demanding in the country and time period we are studying: since Ramadan occurs during the summer, fasting hours average over 16 hours each day. Moreover, survey evidence indicates that while 76% of Muslims in Islamic countries have reduced Ramadan working hours, only 25% of Muslims in non-Islamic countries have such working hours (Dinar Standard, 2011).

We find that areas in the UK with more Muslims see a greater crime increase during Ramadan. This is mainly driven by categories of crime which involve a loss of self-restraint, such as violence and sexual offences. A conservative estimate is that an area which is almost completely Muslim sees an increase of around 1 incident per 1000 people, which translates into an increase of around 12.5% in an area with the median crime rate. Since some types of crimes that increase are serious crimes (e.g. violent crimes) that are unlikely to suffer from reporting bias, it is unlikely that the crime increase is solely due to an increase in reporting.

However, while we believe that fasting is a very plausible explanation for our results, we are unable to rule out that other factors associated with Ramadan are driving the result. We attempt to show that fasting (defined broadly as abstinence from food, water, sleep, and smoking) drives the link between Ramadan and crime, but are unable to rule out completely the role of non-Muslims, Ramadan related gatherings, and non-fasting Muslims[2] . As such, our findings can be cautiously viewed as the reduced-form effect of Ramadan, and we urge future studies to corroborate our findings. We note that this

[2] However, we do give evidence our findings are not driven by Islamophobia or gatherings near mosques. Also, access to administrative data, which is not available to researchers outside of the EU, would only solve a few of these problems.

problem is not unique to us: For example, Almond and Mazumder (2011) attempt to show that food deprivation is mechanism by which Ramadan affects childbirth outcomes, but are unable to rule out fluid restriction and sleep deprivation, among others. They also do not know whether children born with disabilities had mothers who fasted.

Our study contributes to various strands of the literature. First, we contribute to the behavioral economics literature by illustrating the effect of emotional factors in decision making, adding to studies such as Bertrand et al. (2010) (though the emotional factor here is different: that of irritability/self-restraint). Second, we also build on the literature on studying the determinants of crime through natural experiments (e.g. Levitt, 1997; Drago, Galbiati and Vertova, 2009). Finally, we build on a small but growing literature (previously cited) that uses Ramadan as a natural experiment. Our focus, however, is slightly different: most of the existing Ramadan literature has focused on child health outcomes.

Last but not least, our study potentially has implications for criminology, as it sheds light on the possibility that emotions can systematically increase the propensity to commit crime, with effects visible at the aggregate level. At the very least, our paper is a call for more studies and field experiments aimed at reducing crime to be conducted, particularly those relating to self-restraint.

The paper continues as follows: In Section 2.1, we explain our natural experimental setting. Section 2.2 describes the data and identification strategy. Section 2.3 describes the results, and attempts to uncover the underlying mechanisms behind the crime increase as thoroughly as possible. Section 2.4 concludes.

2.1 The natural experimental setting

The holy month of Ramadan is the ninth month of the Islamic calendar. It generally lasts for thirty days; however, because the Islamic year only has 354 days, Ramadan gradually moves forward through the Western calendar (by around 11 days each year).

Muslims are supposed to abstain from eating, drinking, smoking, and sexual activities from dawn to dusk for the entire month, unless exempted for reasons such as pregnancy. A meta-analysis by Leiper, Molla and Molla (2003) found that Muslims suffer from increased irritability during the fast. Both smokers and non-smokers experience a continuous increase in irritability, but smokers experience a greater increase in irritability than non-smokers (Kadri et al., 2000), suggesting that several elements of the fast[3] contribute to increased irritability.

Our analysis studies the UK during 2011 to 2015. During those years, although Ramadan moves forward from late summer to early summer, it always occurs during the summer. As mentioned, this results in long fasting hours due to the UK's northerly latitude. For example, in 2013, Ramadan was from 8 July to 7 August. In mid-July, the sun rose in London at 5am and set at 9pm, resulting in 16 hours of fasting. Considering also that special Ramadan working hours are not common in non-Islamic countries (Dinar Standard, 2011), the effects of increased irritability are likely to be relatively large.

One popular (if oversimplified) framework of a possible mechanism by which irritability could drive crime is the dual process theory, or that of "System 1" and "System 2" (Kahneman, 2003). System 1 corresponds to an "emotional self" while System 2 corresponds to a "rational self". Under this theory, decisions are sometimes largely the product of System 1, and sometimes largely the product of System 2[4] . Because System 1 processes are often thought to require less energy and time than System 2 processes, System 1 often dominates System 2 when the decision maker is tired. It is plausible that Ramadan fasting induces greater reliance on System 1, due to increased tiredness and irritability. *A priori*, an increase in irritability is more likely to cause non-economically motivated crime (such as violence) to increase, compared to economically motivated crime (such as vehicle theft). Additionally, crimes that do not require careful planning are more likely to increase than crimes that require careful planning.

There is also substantial evidence that many Muslims fast; 98% of all respondents of a Muslim survey of both developed and developing countries indicated they will fast (Dinar Standard, 2011). Even though pregnant women are exempt from fasting, a majority of them still fast on some days of Ramadan (Almond and Mazumder, 2011). However, the exact percentage of Muslims that fast is not especially important for our

[3] Fasting can be narrowly defined as abstinence from food and drink, or be more broadly defined to also include abstinence from smoking and sexual relations. We use the broader definition here.

[4] The presence of multiple decision making systems can account for many anomalies not seen in expected utility theory, such as hyperbolic discounting (Fudenberg and Levine, 2006)

purposes; we simply note that to the extent that Ramadan is not universally observed, our estimates are downward biased. In other words, we estimate the ITT (intent to treat) effect and note that it forms the lower bound for LATE (the local average treatment effect).

Note that Ramadan does not solely entail fasting; there are some other factors associated with Ramadan which may increase or decrease crime simultaneously. For example, the daily routine of Muslims often incorporate major pre-dawn (suhur) and fast-breaking (iftar) meals, which are social events involving family, friends and acquaintances, and co-workers. Both events, particularly the latter, are opportunities for socializing. It is possible that people (in general) may feel more emboldened to commit crimes when they are in large groups. If so, we would expect to see crime increase around gathering areas such as mosques.

Another mechanism through which Ramadan could potentially affect crime is that since Ramadan involves greater displays of religiosity, xenophobic non-Muslims may commit crimes against Muslims during Ramadan. In Section 1.3, we use data on Islamophobic reports to examine this claim. We also consider the possibility that Muslims may be targeted by non-Muslims for economically motivated crime (such as robbery) because the fast has weakened them.

We address the possibilities mentioned in the previous two paragraphs (among other possibilities) in the falsification section. While the available evidence cannot show that no crime is generated under these mechanisms, these mechanisms are not driving our result.

2.2 Empirical Framework

Our analysis centers around the United Kingdom[5] . Several aspects jointly make the UK a good place to study.

First, the UK has month-by-month data on individual crimes from December 2010 onwards. Second, the UK also publically releases census data down to areas of less than 1500 people. These contain information on religion and other socioeconomic variables.

[5] The analysis actually covers only England and Wales. However, for simplicity, and since England and Wales contain the vast majority of the UK's population, we shall refer to our study as analyzing the UK.

Finally, and perhaps most importantly, there is significant variation in concentration of Muslims across UK areas. Some areas of the UK have substantial Muslim majorities (e.g. parts of Birmingham), but many areas of the UK have almost no Muslims too. This allows for the treatment effect of Ramadan to vary across areas.

2.2.1 Data

Crime data is taken from the UK police (data.police.uk). For each crime reported, we have the month in which the crime was committed, the type of crime (roughly speaking, divided into eleven categories before 2012 and fourteen categories after that[6]), as well as the Lower Layer Super Output Area (LSOA) the crime was committed in. In the UK, an LSOA is a geographic area with around 1,500 people. These LSOAs can be aggregated into Middle Layer Super Output Areas (MSOAs), which contain around 10,000 people. MSOAs then can be aggregated into local authorities (LAs)[7] . Local authorities are political in the sense that each LA has a government. However, both LSOAs and MSOAs are non-political, as they were formed by the UK's Office of National Statistics for statistical purposes. Furthermore, we have slightly anonymized GPS coordinates of each crime, which generally allow us to determine the location of a crime up to a catchment area of eight postal addresses.

Data on crime is available from December 2010 onwards. However, it is important to note that a government audit from December 2013 to November 2014 found that police were substantially underreporting crime. Nineteen percent of incidents that were supposed to be recorded as crimes were not recorded as a crime (Her Majesty's Inspectorate of Constabulary, 2014). This in itself is not surprising even for a developed country: In the US, underreporting on a similar scale has been found in the Los Angeles and Milwaukee police forces[8] , to cite just two examples. However, it is still important to think of the impact of misreporting on our results. Assuming that Ramadan leads to a crime increase (which is what we indeed find), such misreporting is likely to cause us to understate the true extent of the increase (or make the increase less statistically

[6] The Appendix contains more details.
[7] While MSOAs are relatively homogenous in population, LAs are not.
[8] See http://time.com/4074896/los-angeles-crime-rates-higher-assaults/ as well as http://www.jsonline.com/watchdog/watchdogreports/crimes-underreported-by-police-include-robbery-rape-e567cu0-167448105.html

significant), under the assumption that police officers try to suppress crime increases. Therefore, we analyze pre-audit data and post-audit data separately, to ensure that our results are as minimally affected by police misreporting as possible (among other reasons). In fact, we analyze data year-by-year most of the time.

Data on religion is taken from the 2011 UK Census provided by the UK's Office of National Statistics (ONS). We obtain the number of Muslims and non-Muslims in each LSOA, and then aggregate to the local authority level[9] . We also obtain Census data on the national origin, education level, population density, and length of residence in the UK. We also obtain some non-Census statistics from the ONS. These include yearly population and unemployment at the local authority level.

We also obtain the annual measures of police officer strength, as measured by the number of full-time equivalent police officers, for each of the UK's 43 police forces[10] .

Finally, we also obtain weather data from the UK's meteorological service (the Met Office). Our main measure of weather is mean temperature, since studies have shown a link between temperature and crime (Hsiang, Burke and Miguel, 2013).

Figure 2.1: Proportion of Muslims in England

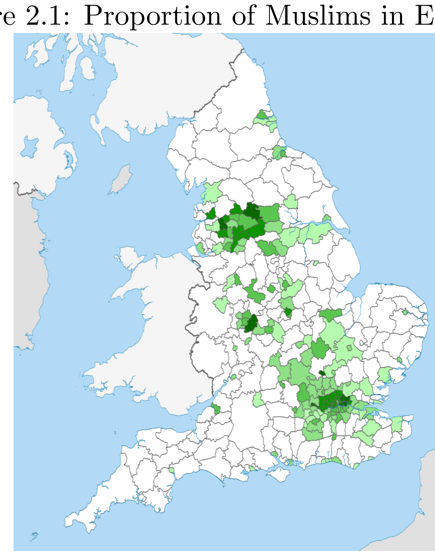

[9] Our main analysis assumes no population growth, though in reality areas with more Muslims see greater population growth. We relax this assumption in our robustness checks. Instead of using the 2011 population, we multiply 2011 percentage of Muslims with the mid-year population estimates of 2012, 2013, 2014, and 2015. Our results are unaffected.

[10] The UK does not have a national police force. Instead, Her Majesty's Inspectorates of Constabulary supervises all police forces; each individual police force serves an average of nine local authorities.

Table 2.1: Summary Statistics (Local Authority Level)

	Min	10th percentile	Median	90th percentile	Max
Monthly Crime	0	418	962	2887	13218
Monthly Crime Rate	0%	0.46%	0.76%	1.21%	11.9%
# of Muslims	3	212	1204	23665	234411
# of Residents	2203	75356	199954	302402	1073045
% Muslim	0.08%	0.23%	0.83%	9.47%	34.5%

As Figure 1.1 illustrates, Muslims mostly reside in urban areas, which tend to be densely populated and less affluent than average. This is something we have to take into consideration when devising an estimation strategy.

2.2.2 Identification

Our identification strategy relies on the fact that Ramadan can be viewed as a treatment, with the treatment strength dependent on the number of Muslims in an area. Here, the treatment is not binary; it would be inaccurate to say that an area is either impacted by Ramadan or not impacted by Ramadan. It would instead be more accurate to model the treatment as a continuous variable, with some areas more impacted by Ramadan than others. Therefore, we use a continuous version of a difference-in-differences estimator (similar to Nunn and Qian (2011)).

$$CrimeRate_{it} = \beta PctMuslims_i R_t + \sum_a \gamma_a I_i^a + \sum_j \rho_j I_t^j + \sum_j X_i' I_t^j \Phi_j + \epsilon_{it}$$

Here, $CrimeRate_{it}$ is the number of crimes in area i at time t divided by the total number of residents, and $PctMuslims_i$ is the percentage of Muslims in area i. Note that in a standard difference-in-differences analysis, where an area either receives a shock or does not, the treatment is binary. Here it is non-binary, and for all practical purposes continuous. We also weight each observation by the total number of residents, which makes the estimator more efficient by reducing heteroskedasticity.

Recall that Ramadan is the ninth month of the Islamic calendar, and the months of the Western and Islamic calendar rarely coincide perfectly. With this in mind, R_t is the proportion of the month which falls into Ramadan. For example, if 15 out of 30 days were Ramadan days, then R_t would equal 0.5. β therefore represents the treatment effect, which is our parameter of interest[11] .

γ_a, I_i^a are the area fixed effects and their indicator variables, while ρ_j, I_t^j are time fixed effects and their indicator variables. The former helps to control for all time invariant factors that vary across areas of the UK, such as geographical factors. The latter controls for factors that change over time but are the same across areas, such as business cycles.

Our strategy relies on there being no shocks occurring that are correlated with the strength of the Ramadan treatment. In other words, the identifying assumption is $Cov(PctMuslims_i R_t, \epsilon_{it}) = 0$. To fix ideas, suppose that a non-Ramadan crime shock occurs in August 2011, one of the Ramadan months in our sample. Suppose further that the size of this shock differed across areas, and was proportional to the number of Muslims in that area. Then our estimate of the treatment effect will be biased.

We view this as the greatest threat to our estimation strategy, and therefore take steps to address this identification concern. First, we estimate the treatment effect for each year separately[12] . If the coefficient estimates for all (or most) years are significant, then this would be evidence of a crime trend moving synchronously with Ramadan. Since Ramadan progressively moves forward through the Western calendar, it would not be easy to argue that our results are spurious. On the other hand, if we only observe a statistically significant relationship in one or two years, then we might be concerned that the observed relationship may be spurious. For example, it could be due to a Type I error, or due to event(s) that coincided with Ramadan in those years. A second measure we can take is to control for as many time- and area-varying variables as possible. Finally, we control for variables that are constant over time, but allow for their impact to vary across months.

We do so by including $\sum_j X_i' I_t^j \Phi_j$, which represents area-specific characteristics interacted with time-period fixed effects. In our most flexible specifications, we control for four time varying characteristics, and also interact time invariant characteristics with time dummies.

---

[11] Recall that this is the intent to treat estimate

[12] This also helps to ensure that crime data before 2014 (during the regime of misreporting) is not analyzed together with data from 2014 onwards, when an audit revealed misreporting.

2.2.3 Events that were correlated with Ramadan shocks

On top of the previously mentioned strategies, we examine major events reported in British newspapers that could influence our results. We identified three events that were likely large enough to have a significant impact: the 2011 nationwide riots, the 2012 London Olympics, and the 2013 religiously-motivated murder of a white soldier by Islamic extremists. The first and third events are likely to downward bias the results. The impact of the second event on coefficient estimates is not clear, but can be accounted for. We elaborate on the events and their likely impact on our coefficient estimates in the below paragraphs.

First, in August 2011 (coincident with Ramadan), there was a wave of riots in many areas of England following the police shooting of Mark Duggan, a Londoner. Cities which were affected included London, Birmingham, Manchester and Bristol. In total, more than 3,443 crimes were recorded in London alone (Mirror, 2011). While almost all the cities involved had substantial Muslim populations, news reports indicate that most of the rioting did not occur in the Muslim areas of the cities (Guardian, 2011). If true, this effect downward biases estimates - an assertion we subsequently verify.

A second plausible "crime shock" that coincided with Ramadan was the 2012 Olympics, which saw an influx of tourists into the British capital. The direction of the bias is not clear. While some Olympic venues were located in areas which had many Muslims, others were located in areas with few Muslims. Therefore, it is not clear whether the Olympics caused crime to rise more in high-Muslim areas or low-Muslim areas. However, the Olympics should be little cause for concern for our estimation strategy. First, the Olympics only affected London while 2011's riots were nationwide, so one can easily remove London from the analysis of 2012 data. Moreover, even if London were included, the size of the bias due to the Olympics is probably not as great as the bias due to the 2011 riots, given that London police reported that crime during the Olympics was actually 6% lower than the year before.

A third event that could bias our estimates was the brutal murder of British soldier Lee Rigby by two Muslim converts on May 22, 2013. This widely publicized event sparked a wave of revenge attacks in the form of Islamophobic crime across the UK. Under the assumption that the amount of revenge attacks is likely to be positively correlated in the number of Muslims in an area, regressions using 2013 data that do not control for these revenge attacks will be affected by downward bias. (To see this clearly, notice that the revenge attacks can be viewed as an omitted variable that is negatively correlated with the treatment effect.) Indeed, we illustrate this in the Appendix[13] .

We do not review the entire series of events that could bias our estimates (no empirical study can do so). However, all other events appear to be orders of magnitude smaller in impact compared to the previously mentioned events.

2.3 Results

Because most of our time- and area- varying covariates are only available at the local authority level, we present results at the LA level[14] . For each year studied except 2011, the estimated coefficient of the treatment effect is clearly positive. In other words, areas with more Muslims see a greater increase in crime. The results therefore indicate that $\beta > 0$. A naive estimate of the average crime increase, obtained simply by averaging the coefficients across all five regressions, reveals that on average there will be an increase of 2.5 crimes for every 1000 people.

Table 2.2: Local authority level results

	(1)	(2)	(3)	(4)	(5)
	2011	2012	2013	2014	2015
$PctMuslim_i * R_t/10^3$	-0.209	2.46^{**}	4.62^{***}	2.75^{***}	3.05^{***}
	(0.832)	(1.12)	(1.09)	(0.644)	(0.903)
Observations/Month	348	348	348	348	348

Standard errors in parentheses. All regressions include time and area fixed effects.
Sample is from May to September of each year.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The results for 2011 could be anomalous due to the nationwide riots. Indeed, when

[13] Alternatively, consider the standard difference-in-differences estimator $DD = W - X - (Y - Z)$, where W and X denote crime in an area with a substantial Muslim population during Ramadan month and non-Ramadan respectively, and Y and Z denote the same for an area with no Muslims. Because Lee Rigby's murder occurred well before Ramadan, it pushed up X. As Section III illustrates, there is no evidence that it pushed up W to a comparable extent.

[14] See the Appendix for LSOA level results, which are similar but smaller in magnitude, possibly because LSOAs are small areas of around 1500 people. As such, people may commit crimes outside the LSOA they live, which introduces measurement error. That said, the LSOA results are much more precise due to the larger sample size, and they indicate a significant increase even for 2011.

we remove London, Birmingham and Manchester (three cities that were hit especially heavily by the riots) from the sample when analyzing 2011 data our coefficient becomes positive, though still insignificant. Removing these three cities in other years, however, results in a *decrease* in the estimated coefficient. As such, we exclude 2011 from our analysis from now on[15] .

One concern is that we may be actually picking up other effects. For example, UK Muslims tend to be poorer than average, and also stay in densely populated areas. Although Ramadan occurs at different times throughout our sample, it always occurs during the summer. Accordingly, one may be concerned that the summer heat could affect poorer areas differently from richer areas, or densely populated areas differently from less dense areas[16] .

We take three steps to address these legitimate concerns. First, our regressions above analyzed only data from May to September, thus minimizing any interaction effects that temperature may have on crime[17] .

The second step we take is to interact some variables that are effectively time invariant with monthly dummies. These variables include the population density of a local authority, as well as three measures from the English indices[18] of deprivation, namely income deprivation, young people deprivation, and adult skills deprivation. Deprivation refers to unmet needs caused by a lack of resources of all kinds. For example, factors that affect an area's score for the income deprivation index include the number of adults and children that are receiving income support from the government, as well as the number of asylum seekers that receive subsistence support. Factors that are included into an area's young people deprivation score include the secondary school absence rate and the

[15] Indeed, the nationwide riots not only impacted crime in August 2011, but had crime substitution effects that lasted for months and therefore impacted September as well, at least in London (Bell, Jaitman and Machin, 2014).

[16] Interestingly, the direction of the bias is not immediately obvious. Suppose that during the summer, crime increased more in poorer areas much more than richer areas, as one might intuitively expect. Then our coefficient estimate would suffer from upward bias because we have not controlled for the treatment effect of summer, but it would be counterbalanced by downward bias because Ramadan only stretches over a part of the summer. A similar argument holds if crime increases more in richer areas during the summer.

[17] Across England, the average temperature for the May and July (the coldest and warmest months in the sample) are 14 and 19 degrees Celsius respectively.

[18] The sample size is smaller in these regressions as the indices of deprivation for England and Wales are not comparable. Hence, we only analyze local authorities across England

Table 2.3: LA Results with interactions

Control (× Month Fixed Effect)	(1) 2012	(2) 2013	(3) 2014	(4) 2015
Income deprivation	0.79	2.385*	2.14***	3.55***
	(1.40)	(1.32)	(0.759)	(1.03)
Young people deprivation	2.44**	4.75***	2.77***	3.12***
	(1.13)	(1.09)	(0.650)	(0.920)
Adult skills deprivation	2.45**	4.64***	2.71***	3.19***
	(1.19)	(1.17)	(0.682)	(0.905)
Population density	1.19	3.15***	2.39***	1.65***
	(0.957)	(0.880)	(0.622)	(0.916)
All controls	-0.268	1.10	2.18***	2.89***
	(0.931)	(0.972)	(0.727)	(1.00)
Observations/Month	326	326	326	326

The coefficient on each box indicates the coefficient of $PctMuslims \times R_t/10^3$, for the year in the given column, and using the controls in the given row. Standard errors in parentheses. All regressions include time and area fixed effects. Observations are weighted by population.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

proportion of those aged 21 not entering higher education. Adult deprivation is based on the number of adults aged 25 to 54 with no or low qualifications[19] . If we were actually picking up the differential impact of summer on densely populated areas, or areas that are income deprived, or on areas where young people and adults are deprived of education and skills, then the coefficient of our parameter of interests should change. By and large our estimates still remain statistically significant (see Table 2.3, which represents the results of 16 regressions, one for each cell). Most notably, when we use all controls in 2015 data (after the crime audit had concluded and police forces had taken steps to correct misreporting), our point estimate is remarkably similar to the corresponding point estimate without any controls in Table 2.2.

To further examine whether the effect which we attribute to Ramadan could be spurious, we also include time- and area- varying controls into our regression. These include mean temperature, unemployment, and for good measure, the number of police officers, even though it is endogenously determined. None of these variables affect the

[19] See https://www.gov.uk/government/uploads/system/uploads/attachment˙data/file/6871/1871208.pdf for more details

significance or magnitude of our estimates significantly (see Table 2.4).

Table 2.4: LA Results with time varying controls

Control	(1) 2012	(2) 2013	(3) 2014	(4) 2015
Mean Temperature	2.48**	4.61***	2.64***	2.73***
	(1.13)	(1.10)	(0.647)	(0.868)
Unemployment	2.36**	4.53***	2.66***	3.03***
	(1.12)	(1.10)	(0.651)	(0.918)
Police Officers	3.33***	4.67***	2.75***	
	(1.23)	(1.11)	(0.644)	
All controls	3.20**	4.54***	2.54***	
	(1.25)	(1.14)	(0.655)	
Observations/Month	348	348	348	348

The coefficient on each box indicates the coefficient of $PctMuslims \times R_t/10^3$, for the year in the given column, and using the controls in the given row.
Standard errors in parentheses. All regressions include time and area fixed effects.
Observations are weighted by population.
2015 police officer data was unavailable at time of writing.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Therefore, it seems reasonable to conclude that either the summer heat does not differentially impact poorer areas compared to richer ones at a magnitude that is large enough for us to be concerned, or that the biases mentioned in footnote 15 cancel each other out.

Fourth, we perform placebo tests by moving Ramadan forward by two months[20] and restricting the sample to be from two months before our placebo Ramadan (i.e. March in 2012 and 2013, and February in 2014 and 2015) to the month before Ramadan starts (i.e. Ramadan months are deleted). None of the coefficients in Table 2.5 are significant, which gives us more confidence in our results.

We conduct further placebo shocks (see below table). For example, the placebo shock in May 2014 represents the estimate of β if the regression is run where $R_t = 1$ if and only if $t =$ May 2014. The table below therefore summarizes the results of 20 regressions.

Several points are clear from the table:

[20] Since Ramadan often overlaps with two months in the Western calendar, moving Ramadan forward by one month only would not work.

Table 2.5: LA level fake shock (move Ramadan two months forward, sample starts two months before placebo Ramadan, ends in the placebo Ramadan month)

	(1)	(2)	(3)	(4)
	2012	2013	2014	2015
$Muslims_i * R_t/10^3$	1.90	-0.344	-0.606	-0.725
	(1.48)	(0.898)	(5.03)	(1.14)
Observations/Month	348	348	348	348

Standard errors in parentheses. All regressions include time and area fixed effects.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.6: LA level fake shock table

	(1)	(2)	(3)	(4)
	2012	2013	2014	2015
Placebo May shock	1.31**	0.115	0.532	-0.683
	(0.627)	(0.396)	(0.492)	(0.505)
Placebo June shock	0.00182	-0.314	0.918***	0.651
	(0.5625)	(0.530)	(0.254)	(0.430)
Placebo July shock	1.24**	3.42***	2.33***	1.54**
	(0.498)	(0.712)	(0.565)	(0.597)
Placebo August shock	0.765	0.653	-2.02***	-0.560
	(0.576)	(0.518)	(0.475)	(0.529)
Placebo September shock	-3.33***	-2.56***	-1.76***	-0.683
	(0.921)	(0.635)	(0.438)	(0.504)

Shaded cells indicate Ramadan months. Actual Ramadan "shock" was not included, so some coefficients may be negative and significant.

- If a month overlaps with Ramadan, the corresponding fake shock has a positive coefficient. Moreover, the estimated coefficients are smaller than the actual coefficients from the "real" Ramadan shocks and are also not always statistically significant, reflecting imperfect overlap between the fake shocks and the actual Ramadan.

- If a month does not overlap with Ramadan, we do not get a positive and statistically significant coefficient. There is one sole exception (May 2012), which could be a Type I error.

We employ one last robustness check in an attempt to assess whether we are picking up something spurious: continuous triple-differences estimation. Recall that in 2011, Ramadan occurred in August, but in 2015, Ramadan occurred in June and July. (For this explanation, ignore the fact that 2011's riots were a confounding factor; 2011 is chosen solely for illustration because Ramadan overlapped with August nearly perfectly in that year. Also, because continuous triple-differences is not necessarily intuitive, we illustrate the intuition behind our strategy by describing discrete triple-differences first.)

Table 2.7: Definitions of Variables used

	August	May
Muslim Area	W	X
Non-Muslim Area	Y	Z

Recall that our original difference-in-differences (DD) estimator involved calculating $DD_{2011} \equiv W_{2011} - X_{2011} - (Y_{2011} - Z_{2011})$ (see Table 7 for variable definitions).

Here, we exploit the fact that Ramadan cycles through the Western calendar to create a third difference. Specifically, the naive version of our triple-differences estimator is $DDD = DD_{2011} - DD_{2015} = (W_{2011} - X_{2011} - (Y_{2011} - Z_{2011})) - W_{2015} - X_{2015} - (Y_{2015} - Z_{2015})$. Because Ramadan did not occur in either August 2015 or May 2015, we can be more certain that the effect we are picking up is due to Ramadan.

The above discrete estimator can be generalized to become continuous:

$$\begin{aligned} CrimeRate_{it} = \alpha + \delta R_t \times PctMuslim + \sum \beta_m I_m + \sum \beta_y I_y + \sum \beta_a I_a \\ + \sum \gamma_1 I_m \times PctMuslim + \sum \gamma_2 I_m \times I_y + \sum \gamma_3 I_y \times PctMuslim \end{aligned} \tag{2.1}$$

where I_m is an indicator variable denoting month (e.g. July), and I_y is an indicator variable denoting year (e.g. 2014) and I_a is an area indicator variable. δ is the coefficient of interest. As before, R_t is the portion of the Western calendar month that overlaps with Ramadan.

The table below shows the DDD estimator. In column 1, notice that the DDD estimate of 2.95 deviates less than 10% from the average DD estimate from 2012 to 2015 of 3.22 taken from Table 2.2.

Furthermore, the positive coefficient observed in column 1 does not appear to be an artefact of misreporting in the earlier years (and not others). When the analysis is restricted to areas under the control of police forces which had a 90% reporting accuracy or higher (column 2), the estimated effect actually rises[21] .

Last but not least, suppose we were to include 2011 data to our analysis in column 1. As shown in column 3, if the analysis is expanded to include 2011 data as well, the coefficient is still positive and significant.

Table 2.8: Continuous DDD estimates

	(1)	(2)	(3)
	2012-2015	2012-2015	2011-2015
$Muslims_i * R_t$	2.95^{***}	5.90^{***}	1.47^{**}
	(0.640)	(1.29)	(0.702)
Observations/Month	348	73	348
Low misreporting areas	N	Y	N

Standard errors in parentheses. All regressions include time and area fixed effects.
"Observations" refers to observations per time period (month).
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In sum, we have detected a significant effect of Ramadan on crime.

One concern may be that people who declared themselves as Muslim in the 2011 census may be Muslim in name only. Therefore, we use the number of mosques in an area as a proxy for the number of devout Muslims in an area. It is worth noting that the correlation between the number of Muslims and number of mosques in an area is very high. Indeed, the correlation between the number of mosques and the number

[21] We are unsure as to whether the higher coefficient reflects greater reporting accuracy of those particular police areas, or a greater impact of Ramadan in those areas, but the reason why the coefficient increases is not particularly important for our purposes.

of Muslims in each local authority is over 0.95, and our results are still positive and significant when the number of Muslims is replaced by the number of mosques (see the Appendix).

2.3.1 Exploring potential mechanisms

Having identified the reduced-form effect of Ramadan on crime, we now examine what types of crime increase, in order to shed light on the mechanisms underlying the crime increase.

We again perform our analysis year-by-year, not only to avoid bias induced by any change in police reporting standards during and after the audit, but also because crime classifications changed in the middle of our sample.

One note is that the bulk of the crime increase is largely driven by "anti-social behavior" (ASB). ASB is not defined precisely by the UK police forces. However, ASB typically refers to crimes for which there is insufficient evidence to generate a crime report. For example, if citizens reported street drinking to the police, but there was not enough evidence to link it to a specific person, then it may be recorded as "Anti-social behavior" rather than "Public order". For these reasons, incidences of anti-social behavior are not calculated in the UK crime statistics. However, omitting anti-social behavior would restrict us to police-recorded crime, and it is well known that police recorded crime understates the true incidence of crime. Therefore, we feel that there is merit to counting ASB as crimes in our analysis. We also note that ASB is not the only crime category for which we see an increase.

What types of crime see an increase? We (somewhat arbitrarily) say that a certain type of crime increased if there was a statistically significant change in the same direction for at least two years during the sample, regardless of whether that crime classification was present for three or four years.

From Table 2.9, we see that according to this definition, a number of types of crime increased: anti-social behavior, bicycle theft, criminal damage and arson, public order, and violence and sexual offences. All of these offences could reasonably be linked to a decrease in self-restraint, and do not require careful planning. They are also not generally economically motivated. (The fact that bicycle theft increases initially may seem odd. Indeed, theft of locked bicycles is unlikely to be due to a loss in self-restraint.

However, two-thirds of all stolen bicycles are left unlocked (Crime Survey of England and Wales, 2013), so bicycle theft does not require careful planning. It could be that unlocked bicycles are stolen as a means of transportation, thus saving the thief energy. Most importantly, the increase in bicycle theft does not by itself show that increases in other crimes are caused by a loss in self-restraint; at most, it suggests that there may be an alternative mechanism occurring in parallel, such as reporting.) One may be concerned that increases in reporting may confound the estimates. For example, people that fast may be more irritable and hence more likely to make crime reports, all else equal. However, some of these categories are serious offences (such as violence) and reporting biases are likely to be minimal. Hence, the result is unlikely to be solely due to an increase in reporting.

The increase in crimes is economically significant for all five previously mentioned crime categories. For example, there may be an increase of less than one incident of "violence and sexual offences" per 1000 Muslim inhabitants, but since these serious offences occur rarely, this is actually an increase of at least 10%, and as much as 60%. In fact, in terms of percentages, the increase in all five categories is comparable.

Muslims or non-Muslims?

We now address the question: Is the increase in crime largely attributable to Muslims, or mostly attributable to non-Muslims?

After all, the increased religiosity during Ramadan could well spur Islamophobic attacks. However, data obtained from the London police and West Yorkshire police show no increase in Islamophobic attacks during Ramadan. Formally, we use the same DD model described previously, except that we replace crime by Islamophobic attacks per resident. In fact the point estimate is negative at -0.000312 (p = 0.794). The corresponding estimate for London is also small (0.000237) and insignificant ($p = 0.23$). It is important to note that areas policed by the London force and West Yorkshire force are very different from each other; the former is highly urban while the latter is much less urban. Moreover, West Yorkshire is in the north of England while London is in the south. The areas comprise around 20% of the population of England and Wales[22] .

[22] We requested data from some other UK police forces, but they did not classify Islamophobic hate crimes at a monthly level or finer. The remaining UK police forces covered areas which did not have

Table 2.9: Ramadan's effect on different crimes (Increases per 1000 Muslim inhabitants per Ramadan month)

	2012	2013	2014	2015
Antisocial behavior	1.01*	2.99***	1.84***	1.15**
	(0.533)	(0.672)	(0.546)	(0.486)
Bicycle theft		0.253***	0.137***	0.108**
		(0.0816)	(0.0469)	(0.0568)
Burglary	0.402**	-0.00928	-0.00519	-0.107
	(0.181)	(0.145)	(0.0652)	(0.154)
Criminal damage and arson	0.139	0.358***	0.259***	0.240
	(0.149)	(0.124)	(0.0784)	(0.150)
Drugs	0.303	0.161	-0.1614 *	0.00384
	(0.363)	(0.203)	(0.0852)	(0.121)
Public order		0.267***	0.0797*	0.112
		(0.0525)	(0.0452)	(0.0878)
Robbery	-0.0450	0.110*	-0.0756**	-0.0262
	(0.0938)	(0.0641)	(0.0297)	(0.0583)
Shoplifting	$< 10^{-3}$	-0.163*	0.0992	0.00624
	(0.117)	(0.0922)	(0.0614)	(0.163)
Theft from the person		0.198	-0.0388	0.0363
		(0.199)	(0.0725)	(0.0667)
Vehicle crime	-0.153	-0.0331	0.0681	0.104
	(0.200)	(0.147)	(0.0934)	(0.181)
Violence and sexual offences		0.270	0.426***	1.18***
		(0.198)	(0.111)	(0.197)
Weapon possession		0.0289	-0.00753	0.0306
		(0.0185)	(0.0178)	(0.0227)

Note there were changes in the categories of crimes over the years.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Hence, generalizability of these results is not a major concern.

Still, should there be residual concern about the generalizability of results from these areas, we analyze the number of calls made to TELL MAMA, to a non-government hotline for reporting anti-Muslim attacks throughout the UK (Littler and Feldman, 2014).

Table 2.10: Reports of Online and Offline Abuse to TELL MAMA

	Total	Online	Offline
May 2013	223	185	38
Jun 2013	131	96	35
Jul 2013	74	61	13
Aug 2013	58	46	12
Sep 2013	47	36	11
Oct 2013	36	31	5
Nov 2013	41	30	11
Dec 2013	38	33	5
Jan 2014	58	54	4
Feb 2014	28	27	1

Source: Littler and Feldman (2014)

Within this sample, Ramadan occurred from 8 July to 7 August. Since the number of calls during Ramadan is lower than the number of calls before Ramadan, it is likely that any possible impact of Ramadan on Muslim hate crime is orders of magnitude smaller than the hate crime caused by other events such as the brutal murder of British soldier Lee Rigby on 22 May 2013 (and as we mentioned, omitting this incident from our analysis downward biases estimates).

In fact, it is entirely reasonable to assert that Ramadan had no effect on reports to TELL MAMA, and the fact that reports are higher in July than the following months is due to the time it took for this incident to die off (indeed, media reports on the incident continued weeks after the murder). More tellingly, one might expect that if Ramadan were a major generator of anti-Islamic violence, this issue would be analyzed by TELL MAMA. However, none of the reports on the TELL MAMA website (as of June 2016) are on how Ramadan affects anti-Muslim hate crime.

One may still argue that it is still possible that non-Muslims target Muslims more

many Muslims. Still, the two police forces mentioned cover over 10% of the UK population in total.

during Ramadan, even though there is no increase in Islamophobic attacks. For example, robbers may target Muslims simply because they have less energy due to the fast. However, as mentioned, we do not see any increase in economically motivated crimes. Therefore, while it could be true that non-Muslims target Muslims more often, this effect would be balanced out by Muslims committing economically motivated crimes less often. The main result that non-economically motivated crimes increase still stands.

One may still be sceptical of this result, arguing that Muslims are unlikely to have the energy to commit a crime after as much as 18 hours of fasting. First, crime is not necessarily committed at the end of the fasting day, it could be committed in the middle of the day. (We would separately point out that studies (e.g. Schofield, 2014; Khaled and Belbraouet, 2009) find that even after one accounts for the end-of-fast meal, total caloric intake is still lower than normal. Hence it is still plausible that people are more irritable (or have less self-restraint) than normal at night. Hence even if there were to be an increase in crime at night, it would not necessarily invalidate our claim that increased irritation is driving the result. Second, the crimes for which we see an increase largely do not involve careful planning (e.g. violence and sexual offences).

To summarize, the available evidence of course does not formally rule out that the increase in crime is likely to be driven by non-Muslims. However, given that there is no rise in Islamophobia, and that economically motivated crime by non-Muslims against Muslims does not drive the result, we feel that it is not especially likely that non-Muslims are driving the result.

Possible interpretations

Since the evidence indicates that the increase in crime is likely to be driven by Muslims, we next examine possible interpretations.

While the most salient association that most people have with Ramadan is that of fasting, a number of other important events happen too, and these must be carefully considered.

One possible interpretation is that Ramadan involves large gatherings of Muslims, and that people in general tend to be emboldened to commit crime when in large groups.

Second, there are also reports that some of the crime increase is due to youths slipping out of mosques during prayer times, under the cover of anonymity[23] .

If the above interpretations are true, then we would observe that crime increases at or around the vicinity of mosques. In the two explanations given, it is unlikely that the groups of people involved would travel far.

We therefore conduct two analyses. First, we conduct the analysis at the LSOA level (recall that an LSOA is an area of around 1500 people) and analyse what happens if our sample only includes LSOAs without mosques[24] . Second, we remove all crimes that occur within a half mile radius of a mosque. In both cases, the estimated coefficient is typically higher than if mosques are included (see below table). Additionally, the area in the vicinity of mosques comprises only a small fraction of the total area of the UK. These facts suggest that the crime increases less than proportionately at or around the vicinity of mosques.

Table 2.11: Removing crime near mosques (2015 data)

	(1)	(2)	(3)	(4)
	All LSOAs	All LSOAs, except those with mosques	All LAs	All LAs, but crimes near mosques removed
$PctMuslim * R_t/10^3$	1.89***	2.48***	3.05***	3.23***
	(0.330)	(0.469)	(0.903)	(0.911)
Observations/Month	34,753	33,376	348	348

Standard errors in parentheses. All regressions include time and area fixed effects.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Therefore, it is unlikely that the increase is driven by any of the three explanations given above. At the very least, the above explanations cannot be all that is going on.

2.4 Discussion and Conclusion

Our preferred interpretation is therefore that the increased irritation caused by hunger, thirst, sleep deprivation, and lack of smoking (for smokers) drives the result. This is consistent with most of the correlational studies testing Gottfredson and Hirschi (1990)'s

[23] A report which cites this as a source of crime is available at http://www.popcenter.org/library/awards/tilley/2008/08-20(R).pdf

[24] Removing LSOAs with mosques has the effect of removing mosques and much of their vicinities from the analysis

theory that self-restraint drives crime. More importantly, it also adds to the economics literature by showing that self-restraint should not be ignored when studying crime.

While our study is the first to indicate the importance of self-restraint in a natural experimental setting, it is important to note that we only have evidence consistent with our preferred interpretation. For example, while the available evidence suggests that the crime increase is driven by Muslims rather than non-Muslims, our results do not completely rule out this possibility, as we do not have microdata showing who committed the crime. Ideally, we would also observe fasting behavior. While we do have evidence that high proportions of Muslims fast, and studies of fasting Muslims indicate that irritability increases when they are fasting (Leiper, Molla and Molla, 2003), and that non-fasting Muslims typically substitute participation in religious events for participation in civic events during Ramadans with long fasting hours (Campante and Yanagizawa-Drott, 2015), our results again do not completely rule out that a significant proportion of the crime is caused by non-fasting Muslims. We therefore remind the reader that we do not claim our results estimate causal effects resulting from quasi-random variation in fasting habits.

Most importantly, we do not claim external validity outside of the UK, or within the UK if Ramadan occurs during the winter months. In fact, we would not be surprised if our findings do not hold in these settings. For example, one might not observe the same results in Muslim-majority countries. As mentioned, they are not only located at significantly lower latitudes, resulting in shorter fasting hours, but employers are also often more flexible about working arrangements.

As such, we hope that future studies (either using administrative datasets, other natural experiments, laboratory experiments, or field experiments), can provide a more conclusive answer to this question. One promising area of future research would be to examine the effect of hunger (and more generally, malnutrition) on crime. For example, one could examine whether the provision of food stamps decreases crime more than the provision of a similar amount in cash. Other elements which could contribute to a decrease in self-restraint during Ramadan could also be examined. For example, one could in-principle use a difference-in-differences approach to analyze whether smokers' propensity to commit crime is affected differentially from non-smokers by changes in smoking bans.

Chapter 3

Rational Inattention in Valuations

The theory of rational inattention, first expounded by Sims (2003), has been offered as an explanation for a considerable number of violations of expected utility theory. Sims' original article showed that constraints on information processing capacity could explain delays in processing information. The literature has seen rapid growth in recent years: rational inattention has been used to explain sticky prices (Mackowiak and Wiederholt, 2009; Matejka, 2016), deviations from rational expectations (Carroll, 2003), stochastic choice (Matejka and McKay, 2015; Woodford, 2014), asymmetric responses to wealth shocks (Tutino, 2013), among others.

Despite the amount of theoretical work, the empirical portion of the literature is just starting to develop. There is convincing evidence that people allocate attention in a manner consistent with rational inattention. Caplin and Dean (2013) find that rational inattention theory provides a better fit to experimental data than do standard stochastic choice models, while Cheremukhin, Popova and Tutino (2015) show that participants allocate more attention when stakes are higher. Martin (2016) finds that subjects' behavior in games are consistent with being rationally inattentive.

An unanswered question within this literature (as well as a logical next step) is whether classic biases in decision making can be explained by rational inattention. Here we experimentally test one such theoretical prediction: that of Woodford (2012).

Woodford (2012)'s theoretical paper showed that rational inattention in valuation could explain some common violations of expected utility theory, such as prospect theory (Kahneman and Tversky, 1979), focusing illusions, and decoy effects. As prospect theory (also known as reference dependence) is the most famous phenomeneon, our experiment tests this particular claim.

While Woodford's model is complicated (as are most rational inattention models), it is not difficult to explain the intuition behind it using an analogy to the visual system (which Woodford cites as the high-level principle underlying his idea). In what is called the efficient coding hypothesis (Simoncelli, 2003), the visual system rationally allocates resources to evaluate light intensities it sees the most often. As a result of this, our visual perceptions are biased. A decision maker trapped in a dark room will have difficulty distinguishing between a bright object and a very bright object; both will appear to be of the same intensity. Therefore, she on average underestimates the intensity of the very bright object, relative to the bright object. Similarly, a decision maker who has stayed in a bright environment for a long time will overestimate the brightness of a very dark object relative to a dark object. The auditory system also efficiently encodes stimuli in a manner similar to the visual system (Smith and Lewicki, 2006; Lewicki, 2002).

If our neural systems that process economic stimuli behave in a similar manner, it is not difficult to see how prospect theory preferences can arise. Suppose that in daily life, a person regularly makes decisions that involve small gains and losses (e.g. $10) but rarely on large gains and losses (e.g. $1000). Then even if a person is actually risk neutral, she will actually prefer a sure win of $500 to a 50% chance of winning $1000, because she is unable to fully appreciate the value of $1000 due to her limited information processing capacity. Likewise, she will also prefer a 50% chance of losing $1000 compared to a sure loss of $500, because losing $1000 does not seem twice as bad as losing $500. Hence, risk aversion and risk seeking behavior can co-exist. A corollary (shown in Woodford (2012)) is that exogenously lowering one's cognitive capacity should exacerbate such biases. To illustrate the intuition behind this, consider that while a risk-neutral decision maker with unlimited cognitive capacity would behave in a risk neutral manner, but cognitive constraints would cause the behavior mentioned earlier in this paragraph.

There are reasons to believe that the analogy between the economic system and other neural systems holds. From neuroscience, dopamine neurons represent (or encode) the value of an object. Roughly speaking, they record (with error) the utility an object gives. It is well known that these neurons adapt to their environment in a manner similar to neurons in other systems. If the error associated with high utility objects is large, then the error associated with low utility objects will be small, and vice-versa. Indeed, that prediction is borne out (at the neural level) by non-human primate studies (e.g. Tobler, Fiorillo and Schultz, 2005). However, we do not yet know whether such adaptive coding affects choice behavior.

In our experiment, subjects chose repeatedly between gambles. On each trial, subjects either chose between high value gambles (e.g. between winning \$40 with certainty, or a 60% chance at winning \$80), or chose between low value gambles (e.g. between winning \$4 with certainty, or a 60% chance at winning \$8).

Our primary experimental manipulation was to divide subjects into different groups, with some groups seeing high payoff gambles more often than others. If Woodford's theory holds (or if the more general analogies to the visual and auditory systems hold), then if subjects are in a group that chooses between high value gambles frequently, they should gradually reduce their underestimation of very high payoffs relative to high payoffs. As a result, they should become less risk averse. A secondary manipulation was to examine whether subjects behave as though they are more risk averse when they are distracted by a parallel task.

Our experimental results are inconsistent with our initial hypothesis that rational inattention in valuation causes prospect theory preferences. Specifically, the frequency with which a subject sees high payoff gambles does not affect their observed risk preferences over those gambles, and likewise with low payoff gambles. Although multitasking does increase observed risk aversion, such behavior can be explained by a dual-systems model, where increased cognitive load causes greater reliance on System I (Deck and Jahedi, 2015).

This paper principally contributes to the growing literature on rational inattention, much of which has been cited earlier in the Introduction. It indicates that caution is needed in reading the theoretical literature: much like the empirical adage that "correlation does not imply causation", just because rational inattention can explain a certain phenomenon does not mean that it is actually the cause. However, we feel that the main message of this paper is a positive one: that further growth in the empirical rational inattention literature (such as through the testing of theoretical models) is desirable.

Our second contribution is to suggest that some caution is needed when making direct analogies between the neural systems that process economic stimuli and other neural systems, such as the visual system. The high-level principles that many neural systems are based on may not necessarily apply to the systems that process economic stimuli. Of course, our paper does not prove that the process of adaptation cannot occur over much longer time scales (especially evolutionary time scales), and we leave such a possibility for future research.

Our third contribution is relevant to the experimental economics community in general. Despite concerns that a 100 trial experiment would be too long and repetitive, we actually found little evidence of boredom in our large (and fairly representative) sample of adult Minnesotans. In fact, behavior became more consistent as the experiment progressed, not less.

This paper continues as follows: Section 2 explains in more detail how rational inattention in valuation could lead to bias in valuation. Section 3 explains the design of the experiment. Section 4 shows the results, and Section 5 concludes.

3.1 Rational Inattention and Bias

In explaining how rational inattention can lead to bias, we rely on a simplification of Woodford's model. However, note that our experimental tests do not rely directly on Woodford's model; we are more interested in testing the more general idea that rational inattention can lead to bias in valuation. Woodford's model imposes strong parametric assumptions in order to be even numerically solvable, and a rejection of the Woodford model would not necessarily disprove the underlying idea; it could be that some parametric assumptions are wrong. Moreover, while Woodford's model is neurologically well grounded, it relies on unobservables (such as subjective representations), which would be difficult to incorporate in any empirical analysis. Finally, Woodford's paper centers on explaining observed choice anomalies, so it is only natural for us to focus on observed choice behavior too.

In the model, decision makers do not know the actual utility of an object, but have to come up with estimates[1] . In coming up with the estimates, decision makers have a prior over the objects they are choosing from. To fix ideas, suppose that the utility u of a randomly chosen object a decision maker encounters is standard normal distributed. Then, the decision maker's prior over the utility of an object has that distribution. In other words, if the decision maker were asked about the utility the object gives without any information of the object, her best guess would be that it would be drawn from a standard normal distribution. The decision maker then chooses a probabilistic mapping between u and $\hat{u}$ (denote this as $\pi(\hat{u}|u)$) to minimize expected mean squared error $E_u[E_{\hat{u}}[(\hat{u}-u)^2]]$, subject to the constraint that the probability density functions of u and $\hat{u}$ not diverge by too much for any given value of u. Note that the decision maker's choice of mapping need not be conscious, just like our eyes subconsciously adjust to their surroundings.

If the decision maker's prior is normally distributed (which is a reasonable assumption given that people frequently deal with small gains and losses, but rarely with large gains and losses[2]), and the distributions of u and $\hat{u}$ can only differ by a certain magnitude, then the values of extremely high draws from u will be underestimated. Conversely, the values of extremely low draws of u will be overestimated. Therefore, a decision maker will be risk averse with respect to gains but risk seeking with respect to losses.

3.2 Experimental Design and Results

The experiment was held at the Driven to Discover (D2D) building at the Minnesota State Fair in September 2015. The D2D building is owned by the University of Minnesota, and space is allocated for University researchers to conduct their experiments during the Fair. Recruiting subjects at the Fair has some key advantages: First, some of our experimental manipulations precluded the use of an online experiment, and it is much easier to recruit subjects at the Fair than at the lab. Second, the subject pool is much more diverse compared to a typical lab experiment, thus reducing concern about

[1] Consider again the analogy to the visual system: we do not know the actual brightness of an object, but can form estimates if asked

[2] We also verify this through our demographic questionnaire by asking people how often they make decisions on each of the following sums of money: \$10, \$30, \$50, \$70, \$90. As shown in Table 2, most subjects deal with the smaller sums much more often than the larger sums.

generalizability. Third, we wanted to monitor our subjects for visible signs of boredom, which is not possible in an online experiment.

In the experiment, each subject underwent a hundred trials on an iPad[3] ; each trial had a time limit of 10 seconds. On each trial, two gambles were displayed, and the subject chose between them. An example of a choice was a 50% chance at winning $10, or a sure win of $4. Each trial consisted of a sure win and a gamble. Figure 3.1 gives an example of such a trial.

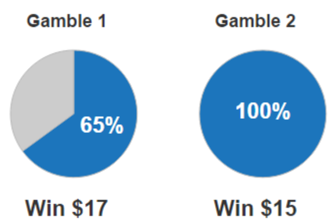

Figure 3.1: Example of a trial

Subjects were randomly assigned into six groups; their group assignment determined the type of gambles they would face. For example, subjects in Group 1 were asked to choose between high payoff gambles on all trials. Information about all groups is illustrated in Table 3.1. Note that in low payoff trials, the gamble had a winning payoff of around $12, while a high payoff gamble had winning payoff of around $85. In all cases, the payoff of the sure win was less than that of the gamble. To ensure that we estimated risk aversion as accurately as possible, the attractiveness of the sure win relative to the gamble varied a lot. Sometimes, the sure win paid off almost as much as the gamble, but there were also cases where the sure win paid less than 20% relative to the gamble. Participants were told about the types of trials they would see (and their proportions) immediately after they had filled in the demographic questionnaire, and immediately before they started choosing between gambles.

Finally, since humans often overweight low probabilities and underweight high probabilities (Kahneman and Tversky, 1979), I avoided using winning probabilities lower than 20%, or higher than 80% for the gamble. Probability distortions in this range are not significant (Prelec, 1998). Also, to ensure easy comprehension of gambles, probabilities

[3] iPads were used to make the experiment easier (compared to using a mouse and laptop)

were in multiples of 5. That is, the probabilities used were 20%, 25%, 30%, and so on, up to 80%.

Table 3.1: Experiment Groupings

Group	Gambles
1	100 high payoff trials
2	75 high payoff trials, 25 low payoff trials
3	25 high payoff trials, 75 low payoff trials
4	100 low payoff trials
5	Same as Group 1, but with multitasking
6	Same as Group 2, but with multitasking

The primary experimental manipulation was to vary the frequency with which each subject encountered high payoff trials, as well as low payoff trials. As far as we know, this experimental manipulation is novel to the literature. Notice that within each group, each subject would encounter the same trials. For example, all subjects in Group 1 chose between a 70% chance of winning $77 and a sure win of $22. However, the order of gambles was random. That is, the aforementioned choice situation was shown to different subjects at different points of the experiment. The random order of gambles controls for order effects (Harrison et al., 2005) if they exist. It also ensures that comparisons across choice situations are unconfounded by boredom, assuming boredom is a factor. That said, we do not believe that boredom was a major factor. First, only two subjects showed signs of boredom[4] , and our results are robust to removing them. Second, studies of a comparable (or longer) length are not uncommon (e.g. Levy et al., 2010; Martin, 2016; Barkley-Levenson, Leijenhorst and Galvan, 2013).

Furthermore, the trials that Group 2 and Group 3 saw were taken from Group 1 and Group 4. For example, if a subject in Group 3 had to decide between a certain high payoff gamble and a sure win, then a subject in Group 1 would also have to decide between the aforementioned gambles. Likewise, if a subject in Group 2 saw a certain low payoff gamble and a sure win, then a subject in Group 4 would also have to decide between those two choices. Finally, the trials that Group 5 saw were exactly identical to that of Group 1, and the trials that Group 6 saw were identical to that of Group 2.

[4] We considered a subject to be bored if she appeared restless to an extent that would significantly interfere with her responses (for example, mindlessly clicking).

These measures minimize any bias in results.

A second experimental manipulation was to have some subjects multitask by playing Puzzle Bobble in order to decrease their cognitive capacity. Figure 2 provides an example of a screenshot.

Figure 3.2: Puzzle Bobble

All subjects provided informed consent. A total of 450 subjects successfully completed the experiment, of which 48% were male, and 52% were female. 31% were aged 16 to 24, 20% were aged 25 to 34, and around 12% were aged 35 to 44, 13% were aged 45 to 54, and the rest were above 55. The corresponding figures for Minnesota are 14.3%, 13.0%, 14.7%, and 15.3% respectively. 62% had finished a four year college degree or higher (the corresponding figure for Minnesota is 48%).

Table 2 reports subject's responses to the question, "How frequently do you deal with certain sums of money?"[5] It is evident that subjects dealt with small amounts of money much more often than large amounts of money. While a substantial majority deal with amounts around $10 at least a few times weekly, half of all subjects deal with $70 and $90 less than once a month. Therefore, the distinction between high payoff trials and low payoff trials seems to be a valid one. We also note that observed risk aversion at these stakes is an anomaly: while expected utility theory implies that people

[5] A subtitle explained: "We're interested in your personal experience with certain sums of money, such as the frequency with which you buy or sell items that are around $10, and so on." The experimenters and research assistants provided any additional clarification needed.

should be approximately risk neutral over these stakes, they behave as if they are not (Rabin, 2000).

Table 3.2: Frequency of dealing with amounts of money

Amount	Option 1	Option 2	Option 3	Option 4	Option 5	Total
around \$10	7.8%	11.9%	37.3%	35.7%	7.4%	100%
around \$30	14.8%	33.2%	34.6%	13.3%	4.1%	100%
around \$50	29.9%	38.1%	23.4%	5.3%	3.3%	100%
around \$70	46.7%	35.4%	12.3%	3.1%	2.9%	100%
around \$90	53.9%	32.6%	8.0%	2.3%	3.3%	100%

Option 1: Less than once a month
Option 2: A few times monthly
Option 3: A few times weekly
Option 4: At least daily
Option 5: At least five times daily

Each subject was provided with incentives as follows:

- All participants that completed the experiment were entered into a draw for one of two \$50 Target gift cards.
- Also, subjects who multitasked were eligible for an additional (separate) draw for a \$50 Visa gift card. Each multitasking subjects score on Puzzle Bobble was equal to the number of chances he/she had at winning the additional card.
- Finally, for all participants, one trial was randomly chosen at the end of the experiment, and it's outcome was realized and recorded. At the end of the State Fair, participants were randomly chosen and mailed a cheque containing the outcome of the gamble (conditional on them winning some money), until total funds of \$900 were exhausted.

There has been a long debate over whether it is necessary to provide real incentives to participants (usually in the form of a randomly selected trial played out for real money). However, at no point of the debate was the evidence overwhelming that real incentives had to be provided. Moreover, most recent studies (and meta-studies) have found little difference between hypothetical choices (i.e. not incentivizing participants for making the correct choices) and real choices (Gneezy, Imas and List, 2015; Amir and Rand, 2012; Engel, 2011).

Therefore, we are confident that the incentives given (in part mandated by the nature of State Fair subject recruitment, which said that we had to recruit subjects for all our slots assigned and offer them the same incentives, even if we had met our recruitment target), which are somewhere in between pure hypotheticals and real incentives, should be little cause for concern.

3.2.1 Theoretical predictions

Below, we summarize what kind of behavior would be consistent with the assertion that rational inattention is a driver of prospect theory:

- Subjects that rarely choose between high payoff trials should be more risk averse with respect to high payoffs than subjects that only choose between high payoff gambles. To use the terminology of Woodford (2012), the prior of subjects that rarely deal with high payoff trials should be centered around low payoffs (since they deal with them most of the time). Therefore, they will underestimate the payoff of the gamble relative to the sure win. To argue by analogy with the visual system, subjects that rarely choose between high payoff trials are in a "dark environment", and therefore cannot correctly perceive differences between a bright and very bright object.

- Analogously, subjects that rarely choose between low payoff trials should be more risk averse with respect to low payoffs than subjects that only choose between low payoff gambles, because they overestimate the value of the sure win relative to the gamble payoff.

- Subjects that have to concurrently play Puzzle Bobble while choosing between gambles should be more risk averse than those who only choose between gambles. Recall that a person in a dark room can perceive any level of brightness accurately with unlimited cognitive capacity. However, perceived brightness becomes a concave function of actual brightness in the presence of information processing constraints. For the mathematical details, see Woodford (2012).

3.2.2 Estimation strategy

We used two methods to compare risk aversion across groups. Our first method was non-parametric. We took common trials across groups for a given payoff level (high or low), and compared the proportion of choices that were sure wins across groups. To illustrate, recall that Group 1 had 25 high payoff trials in common with Group 3. By calculating the average number of choices that were sure wins in these trials for Group 1, and the corresponding figure for Group 3, one could get an idea of whether subjects in Group 1 were more risk averse than those in Group 3.

Second, we estimated participant's utility function parametrically. Our first parametric assumption was to assume constant relative risk aversion (CRRA) in the form of power utility (see equation (1)). Because any gamble a subject faces has only one outcome that pays a positive amount, the utility u that a gamble gives can be written as

$$u = px^{\alpha} \tag{3.1}$$

where x is the payoff and p is the objective probability of winning. (As mentioned, we did not use a probability weighting function as we avoided extremely high or low probabilities.) α is the coefficient of risk aversion, which was estimated by maximum likelihood. CRRA is widely used in many areas of economics such as auction theory, and more generally, macroeconomics. Note that if a subject encountered both high and low payoff trials, it would be wrong to fit the model on all of her choices, since observed risk behavior in experiments indicates that people are more risk averse at high stakes compared to low stakes (Holt and Laury, 2002). However, people's preferences can be roughly described as CRRA over a small range of payoffs (Harrison et al., 2005). Therefore, when estimating α and γ for any subject, I only use a subject's choices in the high payoff trials, or the low payoff trials, but not both. That is, a subject who encounters both high and low payoff trials will have two estimates of both α and γ: one for low payoff trials, and another for high payoff trials.

A second parametric assumption was that the probability that a subject would choose gamble 1 is given by the logistic choice function:

$$Prob(ChooseGamble1) = \frac{1}{1 + \exp\left(\gamma(u_2 - u_1)\right)} \tag{3.2}$$

where u_i is the utility of gamble i, and γ refers to discriminability[6] , again estimated by MLE.

The use of a logistic function to model discriminability should be uncontroversial. First, the predicted discriminability functions in Figure 5 of Woodford (2012), which are reproduced here as Figure 3.3, strongly resemble logistic curves[7] . Second, the logistic function is used in a wide variety of applications, and is the solution to many Bayesian updating problems (Jordan, 1995). Last but not least, the above logistic parameterization has an intuitive property if $\gamma > 0$. Consider the case where $u_1 = u_2$ initially. If u_2 is held fixed and u_1 increases (i.e. Gamble 1 becomes more attractive relative to Gamble 2), the probability of choosing Gamble 1 increases at a decreasing rate.

3.3 Results

3.3.1 Approach 1: comparing proportions of sure wins

Recall that Group 1 had only high payoff trials, but Group 3 only faced high payoff trials 25% of the time. Comparing subject behavior on the 25 trials the groups had in common, observed behavior was the same. On average, subjects in Group 1 chose the sure win 58.2% of the time, while subjects in Group 3 chose the sure win 56.2% of the time. This difference was not statistically significant ($p = 0.57$). In fact, ignoring sampling error, the results are contrary to what we would expect.

Group 2 and Group 4's observed risk preferences over the low payoff trials they shared were essentially the same. Subjects in both groups chose the sure win 52.0% of the time on average ($p = 0.98$).

[6] Discriminability is the probability that the subject will choose the higher utility option. Here, γ measures discriminability, in the sense that if γ increases, the subject is more likely to choose the higher utility option, all else equal.

[7] The horizontal axis z refers to the normalized utility difference between two objects A and B (i.e. $u(B) - u(A)$), and the vertical axis is the probability that the decision maker will choose object B

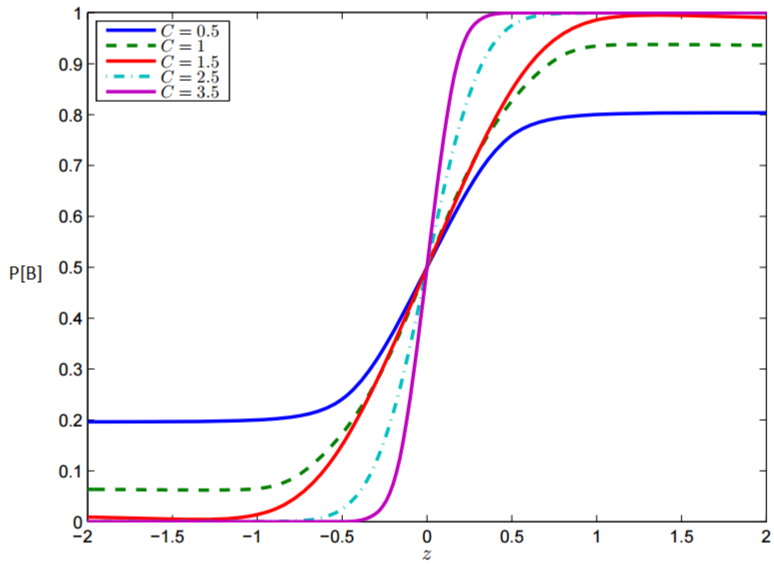

Figure 3.3: Discriminability function (Woodford, 2012)

We also could not detect any difference in behavior between Groups 1 and 2 over the 75 high payoff trials they shared. Nor could we detect any difference in behavior between Groups 3 and 4 over the 75 low payoffs they shared.

Our results are essentially unchanged if we ignore the first ten trials or only consider trials that occurred in the first half of the experiment. There is also no difference when we omit trials where subjects took less than one second to respond. Nor can we detect any change in distributions using the Kolmogorov-Smirnov test.

There are also no meaningful differences in reaction times. For the 25 trials they had in common, Group 3 subjects took and average of 3.1 seconds per trial, while Group 1 subjects took an average of 2.8 seconds per trial. Again comparing only common trials, Group 2 and Group 4 subjects both took 2.9 seconds per trial.

Comparing groups that multitasked against groups that did not (i.e. Groups 1 and 2 against Groups 5 and 6), we find that multitasking induces a large increase in risk aversion: on average, subjects who multitasked were more than 5 percentage points more likely to choose the sure win, and this difference was significant at the 5% level. This is consistent with our initial hypothesis. However, given our previous null results, the fact that such a finding can be explained by cognitive load (Deck and Jahedi, 2015), as well as fMRI evidence that the same brain regions are affected by sequential multitasking (our experimental manipulation) and cognitive load (Borst et al., 2010), we prefer to interpret our results as the effect of cognitive load. (We analyze the heterogenous effects of cognitive load in a subsequent paper.)

3.3.2 Approach 2: estimating risk aversion by MLE

Our results were similar to those we obtained in our previous approach. The mean of Group 1 subject's α was not statistically different from that of Group 3 (0.59 vs. 0.55 respectively, $p = 0.42$), and likewise for Group 2 and 4 (0.74 vs. 0.66, $p = 0.20$). Note that we still only used common trials for this approach because MLE is not necessarily unbiased, especially in small samples. If there was a bias, we wanted both groups to be affected by the bias in the same way.

Is it possible that the experiment was not long enough to induce changes in behavior? In order to explore this question further, we compared γ within groups for different halves of the experiment. From Table 3, it is evident that discriminability for each group is higher in the second half than in the first half, suggesting that some form of learning is taking place (even at low payoffs).

Table 3.3: Gamma

Group	First 50 trials	Last 50 trials	p
1	1.65	2.23	0.09
2	1.71	2.26	0.11
3	2.56	3.14	0.09
4	2.53	3.22	0.06

For group 2, only high payoff trials were used
For group 3, only low payoff trials were used

Moreover, the increase in γ suggests that boredom is not a major issue in this experiment.

3.3.3 Second experiment

In order to ensure that the null results were not Type II errors, we ran a second experiment. The key difference was that in this experiment, groups that had both high and low payoff trials had either all high payoff trials come before the low payoff trials, or vice-versa.

Table 3.4: Groupings in Second Experiment

Group	Gambles
1	90 high payoff trials
2	75 high payoff trials, and then 15 low payoff trials
3	75 low payoff trials, and then 15 high payoff trials
4	90 low payoff trials

Notice that Group 1 still saw only high payoff trials, and Group 4 only saw low payoff trials. The difference is that for Group 2, the high payoff trials all came before the low payoff trials. For Group 3, the low payoff trials came before the high payoff trials. This increased the power of the experiment to detect an effect[8] by allowing subject's priors to shift more compared to the previous experiment[9]. Note that in this experiment, participants were not told of the types of trials they would be encountering.

The experiment was conducted on MTurk. Participants were incentivized in two ways. First, they were given a show up fee of \$1. Second, one trial was randomly selected and participants were paid 10% of the actual payoff[10].

143 subjects took part in this experiment. Again, we estimated participant's risk aversion non-parametrically and parametrically, and there was no statistically significant difference at the 10% level. For example, comparing the last 15 trials, participants in Group 1 chose the sure win 41.7% of the time while participants in Group 3 chose the sure win 41.2% of the time ($p = 0.95$). Participants in Group 2 chose the sure win 40.9% of the time while those in Group 4 chose the sure win 44.1% of the time ($p = 0.68$).

[8] We did not run our main experiment in this way, because of concerns that even if we proved a difference in risk aversion, it could be solely due to a recency effect.

[9] Notice that the analogy of a person moving from a dark environment to bright light is more suited to this experiment

[10] Participants were told that gambles were displayed in virtual currency (ECUs) which would be converted to US dollars at the rate of 10 to 1.

3.4 Discussion & Conclusion

Our results are inconsistent with the hypothesis that inattentive valuation drives prospect theory. Specifically, subjects who are exposed to an environment which generally involves low payoffs do not behave differently when asked to make decisions on high payoffs compared to subjects in a high payoff environment (and vice-versa). We do not find evidence that our results are driven by confounding factors such as boredom. On the contrary, the data indicate that the experiment is long enough for learning to take place.

To be sure, our experiment does completely not rule out the possibility that subject's priors could shift over a much longer time (especially over evolutionary time periods). That is certainly an area for future research.

We hope that our paper spurs more research in the rapidly growing and influential field of rational inattention. More generally, we hope that more tests of bounded rationality models will be done.

Chapter 4

References

Addis, M., and M. Holbroock. 2007. "Attraction, reverence, and escapism in the evaluation of films." *Psychology and Marketing*, 27: 821–845.

Almond, D., and B. Mazumder. 2011. "Health Capital and the Prenatal Environment: The Effect of Ramadan Observance during Pregnancy." *American Economic Journal: Applied Economics*, 3: 56–85.

Amir, O., and D. Rand. 2012. "Economic games on the internet: The effect of $1 stakes." *PLOS ONE*, 7: 1–13.

Barkley-Levenson, E., L. Leijenhorst, and A. Galvan. 2013. "Behavioral and neural correlates of loss aversion and risk avoidance in adolescents and adults." *Developmental Cognitive Neuroscience*, 3: 72–83.

Becker, G. 1968. "Crime and Punishment: an Economic Analysis." *Journal of Political Economy*, 78: 169–217.

Bell, B., L. Jaitman, and S. Machin. 2014. "Crime Deterrence: Evidence From the London 2011 Riots." 124: 480–506.

Bergstrom, T., L. Blume, and H. Varian. 1985. "On the private provision of public goods." *Journal of Public Economics*, 29: 25–49.

Bertrand, M., D. Karlan, S. Mullainathan, Eldar. Shafir, and J. Zinman. 2010. “What’s Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment.” *Quarterly Journal of Economics*, 968: 263–306.

Bizman, A., and Y. Yinon. 2002. “Engaging in Distancing Tactics Among Sport Fans: Effects on Self-Esteem and Emotional Responses.” *Journal of Social Psychology*, 142: 381–392.

Boen, F., N. Vanbeselaere, and J. Feys. 2002. “Behavioral Consequences of Fluctuating Group Success: An Internet Study of Soccer-Team Fans.” *Journal of Social Psychology*, 142: 769–781.

Borst, J., N. Taatgen, A. Stocco, and H. van Rijn. 2010. “The Neural Correlates of Problem States: Testing fMRI Predictions of a Computational Model of Multitasking.” *PLOS ONE*, 5: 1–17.

Butler, B. 2001. “Membership Size, Communication Activity, and Sustainability: A Resource-Based Model of Online Social Structures.” *Information Systems Research*, 12: 34662.

Cameron, A., and D. Miller. 2015. “A practitioner’s guide to cluster-robust inference.” *Journal of Human Resources*, 50: 317–372.

Campante, F., and D. Yanagizawa-Drott. 2015. “Does religion affect economic growth and happiness? Evidence from Ramadan.” *The Quarterly Journal of Economics*, 57: 615658.

Caplin, A., and M. Dean. 2013. “The Behavioral Implications of Rational Inattention with Shannon Entropy.” *Unpublished Manuscript.*

Card, D., and G. Dahl. 2011. “Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior.” *The Quarterly Journal of Economics*, 126: 103–143.

Carroll, C. 2003. “Macroeconomic Expectations of Households and Professional Forecasters.” *Quarterly Journal of Economics*, 118: 269–298.

Chaudhuri, A. 2001. "Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature." *Experimental Economics*, 14: 47–83.

Cheng, J., Danescu-Niculescu-Mizil C. Bernstein, M., and J. Leskovec. 2017. "Anyone Can Become a Troll: Causes of Trolling Behavior in Online Discussions." *Computer-Supported Cooperative Work*.

Cheremukhin, A., A. Popova, and A. Tutino. 2015. "A theory of discrete choice with information costs." *Journal of Economic Behavior and Organization*, 113: 34–50.

Crisp, R., S. Heuston, M. Farr, and R. Turner. 2007. "Seeing Red or Feeling Blue: Differentiated Intergroup Emotions and Ingroup Identification in Soccer Fans." *Group Processes and Intergroup Relations*, 10: 9–26.

Deck, C., and S. Jahedi. 2015. "The effect of cognitive load on economic decision making: A survey and new experiments." *European Economic Review*, 78: 97–119.

DeWall, C., R. Baumeister, T. Stillman, and M. Gailliot. 2007. "Violence restrained: EVects of self-regulation and its depletion on aggression." *Journal of Experimental Social Psychology*, 43: 6276.

Dinar Standard, . 2011. "2011 Productivity in Ramadan Report."

Doleac, Jennifer, and Nicholas Sanders. 2015. "Under the Cover of Darkness: How Ambient Light Influences Criminal Activity." *The Review of Economics and Statistics*, 97: 1093–1103.

Drago, F., R. Galbiati, and P. Vertova. 2009. "The Deterrent Effects of Prison: Evidence from a Natural Experiment." *Journal of Political Economy*, 117: 257–280.

Edmans, A., D. Garcia, and O. Norli. 2007. "Sports Sentiment and Stock Returns." *Journal of Finance*, 62: 1967–1998.

Engel, C. 2011. "Dictator games: a meta study." *Experimental Economics*, 14: 583–610.

Eren, O., and N. Mocan. 2016. "Emotional Judges and Unlucky Juveniles." *NBER Working Paper*.

Fischbacher, U., S. Gachter, and E. Fehr. 2001. "Are people conditionally cooperative? Evidence from a public goods experiment." *Economics Letters*, 71: 397–404.

Fudenberg, D., and D. Levine. 2006. "A Dual-Self Model of Impulse Control." *The American Economic Review*, 96: 1449–1476.

Gneezy, U., A. Imas, and J. List. 2015. "Estimating Individual Ambiguity Aversion: A Simple Approach." *NBER Working Paper.*

Gottfredson, Michael, and Travis Hirschi. 1990. "A general theory of crime."

Guardian, The. 2011. "UK riots: every verified incident - interactive map."

Halfaker, A., A. Kittur, R. Kraut, and J. Riedl. 2009. "A jury of your peers: quality, experience and ownership in Wikipedia." *WikiSym.*

Harrison, G., E. Johnson, M. McInnes, and E. Rutstrom. 2005. "Risk aversion and incentive effects: Comment." *American Economic Review*, 95: 897–901.

Her Majesty's Inspectorate of Constabulary, . 2014. "Crime-recording: making the victim count."

Hirschman, E. 1983. "Predictors of Self-Projection, Fantasy Fulfillment, and Escapism." *Journal of Social Psychology*, 120: 63–76.

Holt, C., and S. Laury. 2002. "Risk aversion and incentive effects." *American Economic Review*, 95: 1644–1655.

Hsiang, S, M Burke, and E Miguel. 2013. "Quantifying the Influence of Climate on Human Conflict." *Science.*

Ifcher, J., and H. Zarghamee. 2011. "Happiness and Time Preference: The Effect of Positive Affect in a Random-Assignment Experiment." *American Economic Review*, 101: 3109–29.

Isen, A. 2000. "Handbook of Emotion." , ed. M. Lewis and J. Haviland-Jones, Chapter Positive affect and decision making. New York:Guilford.

Isen, A. 2007. "In Persons in Context: Building a Science of the Individual." , ed. Y. Shoda, D. Cervone and G. Downey, Chapter Positive Affect, Cognitive Flexibility, and Self-Control. New York:Guilford.

Isen, A., and N. Geva. 1987. "The Influence of Positive Affect on Acceptable Level of Risk: The Person with a Large Canoe has a Large Worry." *Organizational Behavior and Human Decision Processes*, 39: 145–54.

Jones, M., P. Coffee, D. Sheffield, M. Yanguez, and J. Barker. 2010. "Just a game? Changes in English and Spanish soccer fans emotions in the 2010 World Cup." *Psychology of Sport and Exercise*, 13: 162–169.

Jordan, M. 1995. "Why the logistic function? A tutorial discussion on probabilities and neural networks." *Computational Cognitive Science Technical Report*, 9503: 1–13.

Kadri, N, A Tilane, M El Batal, Y Taltit, SM Tahiri, and Moussaoui D. 2000. "Irritability during the month of Ramadan." *Psychosom Med*, 62: 280–285.

Kahneman, D. 2003. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *The American Economic Review*, 93: 14491475.

Kahneman, D., and A. Tversky. 1979. "Prospect theory: An analysis of decision under risk." *Econometrica*, 47: 263–292.

Khaled, B., and S. Belbraouet. 2009. "Effect of Ramadan fasting on anthropometric parameters and food consumption in 276 type 2 diabetic obese women." *International Journal of Diabetes in Developing Countries*.

Kolakowska, A. 2013. "A review of emotion recognition methods based on keystroke dynamics and mouse movements." *HSI*.

Koszegi, B., and M. Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics*, 121: 1133–1165.

Ledyard, J. 1995. "Handbook of Experimental Economics." , ed. A. Roth and J. Kagel, Chapter Public Goods: A Survey of Experimental Research. New Jersey:Princeton University Press.

Leiper, J.B., A.M. Molla, and A.M. Molla. 2003. "Effects on health of fluid restriction during fasting in Ramadan." *European Journal of Clinical Nutrition*, 57: S30S38.

Levitt, S. 1997. "Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime." *American Economic Review*, 87: 270–290.

Levitt, Steven. 2006. "White-Collar Crime Writ Small: A Case Study of Bagels, Donuts, and the Honor System." *The American Economic Review*, 96: 290–294.

Levy, I., J. Snell, A. Nelson, A. Rustichini, and P. Glimcher. 2010. "Neural representation of subjective value under risk and ambiguity." *Journal of Neurophysiology*, 103: 1036–1047.

Lewicki, M. 2002. "Efficient coding of natural sounds." *Nature Neuroscience*, 5: 356–363.

Littler, M., and M. Feldman. 2014. "Tell MAMA Reporting 201314: Anti-Muslim Overview, Analysis and 'Cumulative Extremism'."

Mackowiak, B., and M. Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99: 769–803.

Marcus, Bernd. 2004. "Self-control in the General Theory of Crime: Theoretical implications of a measurement problem." *Theoretical Criminology*, 8: 3355.

Martin, D. 2016. "Rational Inattention in Games: Experimental Evidence." *Unpublished Manuscript.*

Matejka, F. 2016. "Rationally Inattentive Seller: Sales and Discrete Pricing." *Review of Economic Studies*, 83: 1125–1155.

Matejka, F., and A. McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review*, 105: 272–98.

McCollister, K., M.T. French, and Hai. Fang. 2010. "The Cost of Crime to Society: New Crime-Specific Estimates for Policy and Program Evaluation." *Drug and Alcohol Dependence*, 108: 98109.

Mirror, The. 2011. “London riots: More than 2,000 people arrested over disorder.”

Muraven, Mark, and Roy Baumeister. 2000. “Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?” *Psychological Bulletin*, 126: 247–259.

Nunn, Nathan, and Nancy Qian. 2011. “The Potatos Contribution To Population And Urbanization: Evidence From A Historical Experiment.” *Quarterly Journal of Economics*, 126: 593650.

Oswald, A., E. Proto, and D. Sgroi. 2015. “Happiness and productivity.” *Journal of Labor Economics*, 33: 789–822.

Pratt, Travis, and Francis Cullen. 2000. “The Empirical Status Of Gottfredson And Hirschi’s General Theory Of Crime: A Meta-Analysis.” *Criminology*, 38: 931964.

Prelec, D. 1998. “The probability weighting function.” *Econometrica*, 66: 497–527.

Rabin, M. 2000. “Risk Aversion and Expected Utility Theory: A Calibration Theorem.” *Econometrica*, 68: 1281–1292.

Schofield, H. 2014. “The Economic Costs of Low Caloric Intake: Evidence from India.” *Job Market Paper.*

Shoukry, G. 2016. “Criminals’ Response to Changing Crime Lucre.” *Economic Inquiry*, 54: 1464–1483.

Simoncelli, E. 2003. “Vision and the statistics of the visual environment.” *Current Opinion in Neurobiology*, 13: 144–149.

Sims, C. 2003. “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50: 665–690.

Smith, E., and M. Lewicki. 2006. “Efficient auditory coding.” *Nature*, 439: 978–982.

Tobler, P., C. Fiorillo, and W. Schultz. 2005. “Adaptive Coding of Reward Value by Dopamine Neurons.” *Science*, 307: 1642–1645.

Tutino, A. 2013. “Rationally inattentive consumption choices.” *Review of Economic Dynamics*, 16: 421–439.

Vazsonyi, Alexander, Lloyd Pickering, Marianne Junger, and Dick Hessing. 2001. “An Empirical Test Of A General Theory Of Crime: A Four-Nation Comparative Study Of Self-Control And The Prediction Of Deviance.” *Journal Of Research In Crime And Delinquency*, 38: 91–131.

Wann, D., J. Royalty, and A. Rochelle. 2002. “Using motivation and team identification to predict sport fans’ emotional responses to team performance.” *Journal of Sport Behavior*, 25: 207–216.

Wann, D., T. Dolan, K. McGeorge, and J. Allison. 1994. “Relationships between spectator identification and spectators’ perceptions of influence, spectators’ emotions, and competition outcome.” *Journal of Sport and Exercise Psychology*, 16: 347–364.

Woodford, M. 2012. “Inattentive Valuation and Reference-Dependent Choice.” *Unpublished Manuscript.*

Woodford, M. 2014. “Stochastic choice: An optimizing neuroeconomic model.” *American Economic Review: Papers and Proceedings*, 104: 495–500.

Appendix A

Appendix for Chapter 2

A.1 Crime categories

A.1.1 2011 and 2012

- Anti-social behavior
- Robbery
- Burglary
- Vehicle crime
- Violent crime
- Drugs
- Criminal damage and arson
- Public disorder and weapons
- Shoplifting
- Other theft
- Other crime

A.1.2 2013 onwards

- Anti-social behaviour
- Robbery
- Burglary
- Vehicle crime
- Violence and sexual offences
- Drugs
- Criminal damage and arson
- Public order
- Weapon possession
- Shoplifting
- Bicycle theft
- Theft from the person
- Other theft
- Other crime

A.2 Illustration of downward bias due to revenge attacks for Lee Rigby's murder

British soldier Lee Rigby was murdered on 22 May 2013 by Islamic extremists, and in the following weeks there were many revenge attacks against Muslims. Column 1 reproduces the results of the difference-in-difference estimation for 2013. Column 2 shows what happens if a shock $LeeRigby = PctMuslims_i \times Rigby_t$ is added to the regression. $Rigby_t$ is equal to $\frac{23}{31}$ in May 2013 and $\frac{7}{31}$ in June 2013, reflecting the fact that most revenge attacks occurred in May but there were some in June. (The weights were somewhat arbitrarily determined, but our results are not very sensitive to changes in these weights.) Notice that the estimated impact of Ramadan in specification 2 is higher than that of specification 1, consistent with classic omitted variables bias.

Table A.1: Including "Lee Rigby" shock pushes up Ramadan coefficient

	(1)	(2)
$Muslims_i * R_t$	4.62***	5.81***
	(1.09)	(1.16)
$LeeRigby$		2.66***
		(0.555)
Observations/Month	348	348

All regressions include time and area fixed effects.
Standard errors in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A.3 Some comments on crime data

Generally speaking, police record that an incident of "anti-social behavior" has occurred when:

- Citizens report that a serious offence has occurred, but there is little evidence to substantiate. For example, residents may report that drug dealing in their neighborhood, but if the alleged offenders have left by the time police arrive, the information is recorded as a case of "anti-social behavior". Only when there is sufficient information to proceed further will it be recorded as a crime.
- A misdemeanor has occurred repeatedly, and the local council is unable to resolve it. For example, reports of excessive noise from neighbors would be referred to the local council for amicable resolution. If this fails, then the police would generate a report of anti-social behavior. If police fail to resolve the issue, then the neighbors may be arrested or charged and a crime recorded accordingly.

While anti-social behavior is not taken into account when calculating aggregate crime statistics, we feel that including it into our calculations would give a more accurate picture of crime.

Indeed, British authorities realize that crime statistics (as recorded by the police) is not the only reliable indicator of crime and have commissioned the Crime Survey of England and Wales (CSEW). CSEW interviews a representative sample of people in England and Wales, asking them about whether they have experienced crimes in the past month, whether or not the crime was reported to the police.

A.4 LSOA results

These are smaller in magnitude compared to the LA estimates, likely reflecting measurement error (since people may not commit a crime in the vicinity of their residence).

Table A.2: LSOA Results

	(1) 2011	(2) 2012	(3) 2013	(4) 2014	(5) 2015
$PctMuslim_i * R_t/10^3$	1.59***	3.61***	3.74***	2.63***	1.89***
	(0.255)	(0.352)	(0.318)	(0.229)	(0.330)
Observations/Month	34,753	34,753	34,753	34,753	34,753

Standard errors in parentheses. All regressions include time and area fixed effects.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A.5 Using mosques as a proxy for the number of devout Muslims

Table A.3: Using mosque density as proxy: LA level

	(1) 2011	(2) 2012	(3) 2013	(4) 2014	(5) 2015
$TreatmentEffect$	-0.621	5.37**	7.68***	4.58***	2.98***
	(1.32)	(2.25)	(1.70)	(1.16)	(1.60)
Observations/Month	348	348	348	348	348

Standard errors in parentheses. All regressions include time and area fixed effects.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$