(When) Does Ramadan Affect Daily Caloric Intake? Evidence from Rural Malawi*

Scott Weiner

Columbia University

December 16, 2021

Abstract

This paper investigates the effects of Ramadan on calorie consumption and labor supply among Muslim households in rural Malawi. Across four rounds of household survey data, I find no evidence of a decrease in calorie consumption during Ramadan on average. I do, however, find evidence that working-age people reduce their weekly work by about three hours, or nearly 20 percent, on average. This finding on calories shows substantial variation across the different rounds of data. The evidence presented calls into question the hypothesis that consumption during Ramadan should fall more dramatically when the holiday overlaps with the harvest (when baseline consumption levels are relatively high compared to the rest of the year), compared to when Ramadan falls near the annual hunger season (when baseline consumption levels tend to be much lower). I discuss potential implications of this variation for our understanding of seasonal consumption patterns.

*A version of this paper was originally included as a chapter of my Ph.D. dissertation at Columbia University. It has benefited from discussions with Supreet Kaur, Eric Verhoogen, Suresh Naidu, Kaivan Munshi, and Doug Almond, and from feedback from participants at the Development Colloquium at Columbia. Even more so, it has benefited from the discussion at my defense with my dissertation committee – Eric Verhoogen, Suresh Naidu, Rodrigo Soares, Miguel Urquiola, and Jack Willis – who helped guide me toward a more scaled back, and thus more realistic and achievable, framework for the paper. Olivier Ecker, Rachel Gilbert, and Kate Schneider generously shared time, data, programs, and other valuable information with me that were immensely helpful in the process of constructing measures of calories from the Malawi IHS data. This paper makes use of data from Malawi's IHS2 (2004–2005), IHS3 (2010–2011), IHS4 (2016--2017), and IHS5 (2019--2020), generously made available to researchers by the National Statistics Office of Malawi and the World Bank. All errors my own.

1 Introduction

Ramadan is the holiest month on in the Islamic calendar, a time for prayer, giving to charity, gathering with friends and family, and, perhaps most notably, refraining from food and drink during daylight hours. A number of papers in economics have taken advantage of (plausibly exogenous) year-to-year variation in the timing of Ramadan, as well the differences in the length of the fast due to latitude and the time of year, to study the effects of Ramadan on various outcomes, including fetal health (and related later-life consequences) (Almond and Mazumder 2011; van Ewijk 2011; Majid 2015), work output (Campante and Yanagizawa-Drott 2015; Hu and Wang 2019; Schofield 2020), and even traffic accidents (Gulek 2021). Most studies either rely on indirect evidence that the observed effects of Ramadan are attributable to decreased caloric intake, or accept that the mechanism underlying these effects is something of a black box.[1]

This paper considers the setting of rural Malawi, and looks at whether it is indeed true that caloric intake consistently falls for Muslim households during Ramadan. I find that there is not consistent evidence of such an effect in this setting. Nevertheless, I do find evidence of a decrease in labor supply, which calls into question whether the effects of Ramadan that have been observed in other settings, particularly effects on work and productivity, are necessarily attributable to a reduction in calories consumed. I also find some potentially interesting variation in this result across different rounds of data. Specifically, the variation I observe is consistent with the idea that the effect of Ramadan, at least in this context, might depend upon the timing of the holiday relative to the agricultural cycle. While it is important to note that the evidence for this interaction is far too limited to draw any firm conclusions at present, if the pattern I observe were found to apply more broadly, it would have potentially interesting implications for the broader understanding of Ramadan observance and of seasonality in consumption.

[1] To my knowledge, the only economics papers that try to directly estimate the effect on caloric intake are Schofield (2020), and its predecessor paper (Schofield 2014).

To study this question, I use data from Malawi's Integrated Household Survey (IHS), Rounds 2–5 (NSO 2005b; 2012; 2018; 2021). In each round of this survey, a module is administered asking about a large number of different food items that household members may have consumed. For any item consumed within the last week, respondents also indicate the total quantity consumed by all household members, allowing me to construct an estimate of total calories consumed by the household. Under the assumption that households interviewed during Ramadan are not systematically different from those contacted at other times of year, I use a Difference-in-Differences approach to look at whether calorie consumption per person in Muslim households is significantly different from consumption in non-Muslim households during Ramadan. Comparing this difference to households observed at other times of year then provides an estimate of the effects of Ramadan on caloric intake in Muslim households. I use an analogous methodology to study whether Ramadan affects the number of hours worked per week by Muslim working-age adults.

The primary finding in this analysis is that Ramadan has no significant effect on average caloric intake for Muslim households. Of course, the absence of evidence for such an effect does not necessarily imply that no such effect exists. However, a 95% confidence interval gives us a lower bound on this effect of a reduction by approximately 120 calories per person per day, or about a 5% reduction in calories, which would be a much smaller effect than has been identified in other research (e.g. Schofield 2020). Nonetheless, I also find a decrease of almost 20% in hours worked, a reduction of about 3 hours per week for working adults.

It is possible that the decrease in hours worked is attributable to something other than a reduction in calories. However, I cannot rule out the possibility that a decrease in caloric intake could be a contributing factor. Because my measure of calories consumed is given at the household level rather than the individual level, and because children, the elderly, and the ill would not be expected to fast during Ramadan, it is possible that I am underestimating the true effect on caloric intake among healthy, working-age adult Muslims. It is also plausible that temporary hunger due to fasting during daylight hours could make it more difficult to do work, even if the calorie deficit is negated by the post-fast meal (this would be consistent with the findings in Hu and Wang 2019, for example).

One other possibility to consider is that the effect of Ramadan could be different from year to year, depending on where it falls in the agricultural cycle, especially in a country like Malawi where there is pronounced seasonal variation in food consumption. In this light, it is interesting to note that I do find a significant decrease in consumption among Muslims during Ramadan in one round of data, in which Ramadan falls just before the onset of the annual hunger season. Interestingly, I also find some (limited) evidence of a potential increase in calorie consumption during Ramadan in two more recent rounds of survey data in which Ramadan overlapped with the main harvest season, which is a time of relative abundance for most households. This is somewhat puzzling from a theoretical perspective: if baseline consumption levels are higher during the harvest season, we would generally expect marginal utility of consumption to be lower, and thus more willingness to decrease total consumption during Ramadan when it overlaps with the harvest. Conversely, when Ramadan falls closer to the hunger season, we would expect baseline household consumption to be quite low already, such that households would be less willing to reduce consumption any further. If anything, we seem to see the opposite effect: consumption during Ramadan falls by significantly more when Ramadan occurs shortly before the hunger season compared to when it overlaps with the harvest.

We should be quite cautious in trying to explain or interpret differences between rounds – there is no evidence of a causal relationship, and with only four rounds of survey data to draw from, these differences could be purely idiosyncratic or coincidental and not indicative of any broader pattern that would be consistent over time. However, it is at least worth noting the possibility of a link between the timing of Ramadan and the agricultural season. This is particularly true in the context of Malawi, where, as in many other poor countries and regions, food insecurity is a recurring seasonal phenomenon, and this seasonal dimension has important implications for food prices, labor markets, the efficacy of informal insurance mechanisms, and the risk of famine (see Devereux et al. 2008; Chirwa et al. 2012). However, if we were able to establish evidence for such a pattern, it would be consistent with the hypothesis that when food is abundant, post-fast meals are larger, and might even outweigh the effects of the fast, while when food stocks are running low, these meals become more modest and do not cancel out the effects of the fast. It could also potentially speak to one reason we see such dramatic seasonal variation year after year in consumption in places like Malawi: if households increase consumption during Ramadan when it falls during the harvest, and only use it as an opportunity to cut back on consumption – and possibly to save a bit more food for the hunger season – when Ramadan falls shortly before the hunger season, it could mean that they are simply not taking the annual hunger season into consideration when deciding on their food consumption until it is in the very near future, and thus more salient in their minds. If this sort of salience effect plays an important role in economic decision making, it could help explain the persistence of seasonal hunger and other similar phenomena, and could also help policymakers design tools to address these sorts of concerns. Of course, more research would be needed to draw any conclusions as to this specific hypothesis.

This paper contributes to several strands of economic research. First and foremost, it contributes to the literature studying the effects of Ramadan on various social and economic outcomes. A number of papers have looked at the effects of Ramadan on overall productivity (Campante and Yanagizawa-Drott 2015; Hu and Wang 2019; Schofield 2020), subjective wellbeing (Campante and Yanagizawa-Drott 2015), fetal and long-term health outcomes (Almond and Mazumder 2011; van Ewijk 2011; Almond et al. 2014; Majid 2015; Lee et al. 2020), traffic accidents (Gulek 2021), and more (Haruvy et al. 2018; Hodler et al. 2020; Shalihin et al. 2020; Wang et al. 2020). It also contributes to a broader literature on the relationship between religion and economic and social outcomes (McCleary and Barro 2006; Hoverd and Sibley 2013; Benjamin et al. 2016; Iyer 2016; Kuran 2018; D'Haene et al. 2019). This is, to my knowledge, the first paper to directly study the impact of Ramadan on calorie intake in a poor, Sub-Saharan African country such as Malawi.

Second, this paper contributes to the vast literature using survey data, in particular from the World Bank's Living Standards Measurement Studies (LSMS), to measure caloric intake at the household level. In particular, it relates to several papers using Malawi's IHS data (which is part of the LSMS project) to study questions of food security and access to nutrition in that country (Ecker and Qaim 2011; Chirwa et al. 2012; Headey and Ecker 2013; Pauw et al. 2014; Verduzco-Gallo et al. 2014; Beck et al. 2015; Gilbert et al. 2019; Schneider 2021). One major challenge in collecting and analyzing item-by-item food consumption data is describing quantities consumed (Smith et al. 2014 discusses this challenge, along with many others). Forcing households to report quantities in standard units such as liters or kilograms is likely to result in very noisy estimates and strain on the respondent. On the other hand, allowing respondents to provide broad, non-standard units, as the Malawi IHS does, creates difficulties for the data analyst hoping to translate "heap" or "plate" to grams, in order to measure quantities of calories or other vital nutrients; several researchers have pointed out inconsistencies in the default conversion factors provided with the Malawi data, and have proposed procedures to address these concerns (Ecker and Qaim 2011; Verduzco-Gallo et al. 2014; Gilbert et al. 2019). I develop several refinements to the procedure for constructing kilogram and calorie metrics from the IHS data; this is the first paper to my knowledge to construct calorie measures for the recently-released IHS5 data, and the only one to my knowledge to pool data from all four rounds from IHS2 (2004–2005) to IHS5 (2019–2020). I will explain why I think these refinements offer a meaningful improvement to the quality of the calorie measure, and briefly discuss potential implications for the collection of similar data in the future.

Lastly, as mentioned, this paper makes a small contribution to the broad literature on seasonal hunger and food insecurity. This is a phenomenon that is widely documented in various regions throughout the world where people depend largely on rainfed agriculture (Dercon and Krishnan 2000; Devereux et al. 2008; Khandker and Mahmud 2012; Devereux et al. 2012; Bryan et al. 2014; Basu and Wong 2015). It is also poorly understood from an economic perspective, in that a standard economic model of consumption decisions would suggest that households would respond to a repeated seasonal decline in the availability of food by adjusting consumption levels in order to save in other seasons. This paper offers some preliminary suggestive evidence as to what may be the mechanisms underlying this phenomenon.

2 Background

2.1 Ramadan

Fasting during daylight hours each day of the month of Ramadan is one of the central tenets (or Five Pillars) of the Islamic faith: all adult Muslims are expected to do so, with exceptions for the ill, the elderly, etc. (Almond and Mazumder 2011). The Islamic calendar follows a lunar cycle, and each calendar year is about 11 days shorter than a year on the Gregorian calendar. This means that the timing of Ramadan changes from year to year, and cycles through each season over the course of about 33 years.

A growing body of economic literature has used the changing timing of Ramadan and length of the fast day as a source of exogenous variation to study the effects of Ramadan observance on a number of outcomes. Almond and Mazumder (2011) were to my knowledge the first in the economics literature to use Ramadan timing, and variation in the length of the fast, as a source of exogenous variation. They find that *in utero* exposure to Ramadan during the first month of pregnancy increased the likelihood of disability as an adult by around 20%. They also find that the longest observed Ramadan fast (in terms of the length of the fast day) overlapping with the first month of pregnancy decreased the male-to-female sex ratio at birth by more than 6 percentage points, and that the overlap of a longer Ramadan fast day with pregnancy at any stage decreased birth weights by about 18 grams on average. Importantly, these results are consistent with the medical literature on effects of skipping meals during pregnancy, and do not require a reduction in calories over the course of the day. van Ewijk (2011), Almond et al. (2014), and Majid (2015) build on these findings, showing that *in utero* exposure to Ramadan has persistent adverse effects throughout the life-cycle: lower reading and math scores among children; decreases in hours worked among adults; and poorer health in old age, including increases in heart problems and Type 2 diabetes.

Campante and Yanagizawa-Drott (2015) use variation in the length of the Ramadan fast based on a country's latitude and the dates of Ramadan to show that the holiday slows GDP growth in Muslim countries, but improves measures of subjective wellbeing. Hu and Wang (2019) study hourly productivity among salespeople working at a cosmetics outlet, a job requiring little physical activity but a relatively high level of cognitive effort, and find that productivity starts to wain only in the last two hours of the fast, by which point Muslim workers would have been fasting for 12 hours, and that these workers quickly recover to full productivity after sunset when they are able to eat and drink again. Schofield (2020) looks at agricultural productivity, and finds that Ramadan has no significant effect on farm work hours, but a significant decline in output of about 1 percent, consistent with a fall in productivity for Muslim workers of approximately 20–40 percent. She also finds suggestive evidence that this decline is attributable to reduced calorie intake, particularly because productivity does not recover immediately after Ramadan, consistent with the time needed for the body to recover from caloric deficits accumulated over a month of decreased food intake. Other recent papers have studied effects of Ramadan on traffic accidents Gulek (2021), prosocial behavior (measured via a Dictator Game) Haruvy et al. (2018), and public support for and incidences of terrorism Hodler et al. 2020. These latter two papers demonstrate some important aspects of Ramadan observance outside of fasting. In Haruvy et al. (2018), people were more generous during Ramadan when they were abstaining from food (but not after eating their evening meal), whereas generally people become less generous when abstaining from food, suggesting that religious fasting may serve as a reminder to be more charitable to those in need. Along similar lines, Hodler et al. 2020 find that in years and regions when the

Ramadan fast was longer due to longer daylight hours, public support for terrorism declined in polling, and an hour increase in the length of the fast day corresponded to a 2–3 percentage point drop in terrorist attacks during that year, particularly in the sorts of attack that require some level of public support or complicity (such as providing shelter or an escape route) – suggesting that the more intensive fast may perhaps have promoted greater religious introspection and rejection of violence.

A common theme throughout these papers is that it is difficult to assess the precise mechanism driving the observed effects – whether it is lack of food, dehydration, or some other social or behavioral aspect of Ramadan observance. Of the papers mentioned in the preceding discussion, the only one that attempts to directly measure the effect of Ramadan on daily calorie intake is Schofield (2020), who estimates a decrease in caloric intake among Muslims in India of approximately 600 calories per day during Ramadan.

Other papers in the medical literature (Karaağaoğlu and Yücecan 2000; Toda and Morimoto 2004; Ziaee et al. 2006) have attempted to study the direct health effects of Ramadan fasting. They each find evidence of weight loss during Ramadan, and also report some minor adverse effects such as irritability or loss of interest in work. This consistent finding of weight loss would generally imply that calorie expenditures are exceeding calorie intake, and thus that there is likely a sustained reduction in caloric intake. However, Sadeghirad et al. (2012) find in a meta-analysis of studies of health effects of Ramadan substantial heterogeneity across different settings. They do find consistent, significant evidence of weight loss, averaging to about a 1.25 kg reduction, over the course of the month; this seems to be true in most regions that they consider, and it also seems to be the case that this weight is re-gained on average by 2–6 weeks after the end of Ramadan. However, paradoxically, they do not find an alignment between weight loss and calorie reduction. In studies conducted in the Middle East and East Asia, they find calorie intake reductions of about 150–200 calories per person per day on average. On the other hand, in studies conducted in North Africa, they document an *increase* on average of over 250 calories consumed per day, despite the evidence of weight *loss* in this same set of countries at a similar scale to the weight loss observed in the Middle East and East Asia. Clearly, there is still much to be learned about the effects of Ramadan in different settings.

This paper looks at the relationship between the Ramadan fast and caloric intake, as well as labor supply, in the southern African country of Malawi. In principle, it seems intuitive that the Ramadan fast would generally cause a reduction in calorie consumption, though the evidence for this is mixed, as discussed above. In reality, people might eat more or better food during Ramadan, especially during the post-sundown *Iftar* meal, and poorer Muslims may benefit from shared meals and charity, as wealthier people in the Islamic faith are expected to give charitable contributions during Ramadan. I also consider ways in which the timing of Ramadan might interact with agricultural seasons. For example, during the yearly hunger season (see Section 2.3), it would seem logical that households who would be skipping meals even in the absence of the fast would be unlikely to reduce consumption further for the Ramadan fast. Conversely, during the peak season, we might in theory expect a larger reduction in consumption, which in turn might generate some quantity of savings to carry over for the rest of the year. On the other hand, if post-fast meals are quite large and are attended by the entire family, including those (such as children and the elderly) who are not expected to fast, we might even see an increase in calorie intake for the household overall during Ramadan. All of this suggests that it could be fruitful to study empirically how overall consumption behavior changes during the Ramadan fast. The primary question is whether the Ramadan fast leads to a reduction in food consumption overall during that month. Secondary to this question is whether/how that effect varies depending on timing of Ramadan relative to the agricultural cycle.

2.2 Challenges in measuring calorie consumption

Household calorie intake has long been recognized as a vital indicator of welfare and a necessary component in the calculation of poverty statistics. This is particularly true in developing countries, where food expenditures tend to make up the large majority of most households' expenditures Subramanian and Deaton (1996); Smith et al. (2014); Eli and Li (2020). There are a number of challenges when it comes to collecting and analyzing data on food consumption. In general, it is considered best practice to measure households' food intake by asking them to recall quantities consumed of an extensive list of food items. However, such surveys vary widely in their design and are prone to various sources of measurement error Gibson et al. (2014); Smith et al. (2014). As Beck et al. (2015) note, even when using survey data that have already been collected, methodology can make a large difference in measures of great importance to policymakers: they estimate a decrease in poverty rates of 3.4–8.4 percentage points in Malawi (depending on methodological choices) from the 2004–2005 agricultural cycle to the 2010–2011 cycle, in contrast to the disappointing official estimate of just a 1.8 percentage point decrease.

It is important to acknowledge that calories consumed are not a perfect indicator of household welfare. For one, there is the widely documented phenomenon of the Engel curve: as households get richer, the share of their budget devoted to food expenditure tends to decline. Further, wealthier households tend to substitute cheaper sources of calories to more preferred ones, which tends to result in an increase in the amount spent per calorie Subramanian and Deaton (1996). Looking exclusively at calories can also be misleading, as individuals consuming adequate levels of calories might still face significant deficiencies in vital micronutrients Ecker and Qaim (2011); Headey and Ecker (2013). [2]

Whereas caloric intake tends to be highly correlated with other measures of household welfare, particularly in developing countries, looking at the evolution of per capita calorie consumption over time can also generate misleading comparisons. Deaton and Drèze (2009) offer a stark example of this in India, in which they show that, despite rapid economic growth from the 1980s to early 2000s, per capita calorie consumption actually declined.

[2] Because of this, it is generally important to consider measures of dietary diversity, and ideally also specific nutrient intake, to get a more complete picture of household nutritional status. This is beyond the current scope of this paper.

They speculate that it might be possible to account for this at least in part through increased access to labor saving technologies and declining levels of physical activity. Eli and Li (2020) follow up on this; constructing a measure of energy expenditure, they find that totally energy expenditure decreases only minimally, largely because the proportion of children in the population falls dramatically over this period. They suggest, following evidence from Duh and Spears (2017), that improvements in sanitation and the disease environment might account for the apparent discrepancy between falling calorie consumption and improved metrics of health and nutrition. Specifically, if fewer calories are being wasted in the form of diarrhea, this might apply a reduction in total household caloric requirements. Deaton (1997) also describes complications when trying to account for different caloric needs of different household members. Of particular note, it is, for obvious reasons, impossible to use household level data to determine intra-household food allocation. This might present a challenge for measurement of the effect of Ramadan, as not every household member will generally be required to observe the fast.

Smith et al. (2014) assess surveys used to measure caloric intake across 100 different countries in order to identify best practices for data collection, and the extent to which surveys adhere to these best practices. Their results are generally underwhelming. Many of the surveys included fail to meet some of the most important criteria that they describe. Fortuitously for this paper, Malawi's IHS food intake modules stand out in this regard, largely conforming to best practices. In particular, Malawi's questionnaires collect information about a wide variety of specific food items, their provenance (own production, purchase, or gift), use a one-week recall (which is thought to be a reasonable timeframe for most households), account for seasonality in the survey design by spacing out interviews relatively evenly across the year, and make some effort to account for food consumed outside of the home.[3]

Even so, potentially important sources of measurement error remain. One of the most

[3] Perhaps even more impressive is that these practices were implemented even in the earlier rounds of the survey, well before the Smith et al. (2014) paper was released.

vexing problems in working with the Malawi data is one for which Smith et al. (2014) offer no clear-cut solution: the problem of unit measurements.[4] Malawian households are given flexibility to report quantities of consumption of each item from a wide variety of non-standard units. Converting these units to a standard metric on the back end poses a significant challenge. Various researchers (Ecker and Qaim 2011; Verduzco-Gallo et al. 2014; Beck et al. 2015; Gilbert et al. 2019) have noted that the unit conversions provided by Malawi's National Statistics Office (NSO) are incomplete and in many cases appear incorrect or internally inconsistent. Each have used different methodologies to try to address this concern. I describe the methodology that I employ in Section 3.2.

2.3 Seasonal hunger: the "father of famine"

A number of considerations demonstrate the importance of focusing attention on the seasonal nature of hunger for many of the world's poorest people. First, seasonality itself can both cause and exacerbate hunger, particularly in many of the poorest developing countries. Devereux, Vaitla, and Hauenstein Swan (2008), in a discussion of the 2002 famine in Malawi that resulted in tens of thousands of hunger-related deaths, argue that part of the reason that the famine was so severe was that households' coping resources were largely depleted during the hunger seasons of several preceding years. In this context, they refer to seasonal hunger as the "Father of Famine," arguing that it is impossible to understand and prevent famine without first tackling annual hunger seasons. With much of the developing world still heavily dependent upon rainfed subsistence agriculture, it is perhaps not surprising that seasonality plays a larger role in such places than in more developed countries. Another possible risk of ignoring seasonality is that it may lead to a severe underestimate of rates of poverty and food insecurity. Khandker and Mahmud (2012), in their extensive study of

[4] Every method presented for collecting food quantities comes with significant drawbacks, so Smith et al. (2014) do not take a firm position on which methodology is best for collecting food quantities. They do, however, recommend use of demonstration methods, in which respondents use photos or other reference points to help clarify quantities consumed. In recent rounds, the IHS has provided a photo guide for reference purposes for certain food items and units; there might be room to expand upon this practice.

seasonal hunger, suggest that "most of the world's acute hunger and undernutrition occur in the annual hunger season." Further, they argue that many households that would usually be classified as non-poor on average actually do slip below the poverty line or face food deficiencies during the hunger season. This is consistent with other literature suggesting that volatility often causes as much or more hardship for the world's poor than low average income itself (see for example, Collins et al. 2010). Seasonal hunger might also affect other aspects of economic life in developing countries. For one thing, it is well-documented in the literature that many people in developing countries depend on informal insurance arrangements to smooth economic shocks (see for example, Townsend 1994, and Munshi and Rosenzweig 2016). When there is an aggregate shock that affects the entire insurance network simultaneously, as might be the case in a hunger season, these sorts of arrangements are limited in their effectiveness. There are also several channels through which seasonal hunger could generate a poverty trap, pushing people to take actions that, while perhaps necessary for short-run survival, may be harmful in the long run. Examples of this include working on others' land for a low wage (c.f. Jayachandran 2006) at the expense of planting or harvesting on one's own land; selling off productive assets at a price far below their long-term value (c.f. Rosenzweig and Wolpin 1993); consuming crops before fully ripe, thus sacrificing much of the nutritional value; and selling crops during the peak harvest season when prices are lowest rather than saving them for later in the year when prices will usually be substantially higher (Devereux, Vaitla, and Hauenstein Swan 2008). Hunger could also have a direct effect on future earnings by decreasing labor productivity at harvest, precisely when this labor is needed to generate the next year's income. Finally, seasonal hunger might generate an intergenerational poverty trap, as childhood malnourishment, even if temporary, can have permanent effects on cognitive development, especially when coupled with other diseases to which children become more vulnerable when malnourished. Parents might also decide to take children out of school during hunger seasons due to lack of money (Devereux, Vaitla, and Hauenstein Swan 2008). Another reason to focus on seasonal hunger is that it is a widespread phenomenon, documented in many parts of the developing world. Bryan, Chowdhury, and Mobarak (2014) and Khandker and Mahmud (2012) study the "monga" season in Bangladesh, particularly in the northern region of Rangpur. Basu and Wong (2015) conduct an intervention in West Timor, Indonesia, in order to try to assist people in saving to dull the impact of the hunger season. Dercon and Krishnan (2000) visit households in Ethiopia repeatedly over a short period of time and find that many slip into and out of poverty over the course of a year. Devereux, Vaitla, and Hauenstein Swan (2008) document evidence of hunger seasons in Niger, northern Ghana, Namibia, and Malawi. In fact, though hunger seasons are primarily a feature of poor, rural areas, there is evidence of a similar phenomenon of cyclical hunger even in the US, driven by the monthly receipt of food stamps (Shapiro 2005; Hastings and Washington 2010).

Lastly, the existence of seasonality in consumption is to some extent a puzzle for standard economic theory of consumption and saving. Consider the canonical Consumption Euler Equation:

$$u'(c_t) = \delta(1+r)\mathbb{E}\left[u'(c_{t+1})\right].$$

Under the common simplifying assumption[5] that $\delta = \frac{1}{1+r}$, where δ is the subjective time-discount rate and r is the interest rate, this suggests that people should be roughly trying to smooth their consumption today to match their expected consumption tomorrow.[6] Any seasonality in consumption is difficult to explain with this model; if every year, around the same time, there is a season in which income falls, this would surely not take most people by surprise. As such, this should be incorporated into expectations, and individuals should save (or borrow) to avoid seasonal hunger. If people are observed regularly failing to

[5] This rules out the idea of individuals steadily increasing or decreasing consumption over the course of their lives

[6] Technically, people are predicted to smooth expected marginal utility of consumption. If their utility function for consumption is such that $u'''(c) > 0$, they might choose to build up precautionary savings today, consuming less today than tomorrow on average, in order to insure themselves against a potentially severe negative future shock.

smooth this predictable cyclical drop in income and consumption, there are only two possible explanations: either (1) people face some form of binding constraints in both their ability to borrow and to save, such that their consumption is closely tied to their current income, or (2) something in this model is fundamentally flawed.[7]

Brune et al. (2016) conduct a randomized experiment in an attempt to investigate this question of why people in Malawi fail to save when grain prices are low in order to be able to weather or even take advantage of high prices later in the year. They randomly offer some farmers access to a standard savings account, and another group a combination of this savings account and a "commitment" savings account, for which they have to designate a date before which they will not be allowed to withdraw funds. High take up for the "commitment" account would provide suggestive evidence that people's failure to save up for the hunger season is due to a problem with self-control, rather than simply a lack of access to savings or a persistent lack of foresight. Interestingly, they find higher take up of the offer that includes the commitment option than the standard savings account, but they find that people place little money in this commitment account. They suggest that perhaps the constraint to saving is a social one: that having extra money stored up subjects one to pressure to help other family and community members with this money.

In this paper, I analyze the effects of Ramadan on calorie consumption at different points of the agricultural cycle. My results, as I discuss in Section 4, provide suggestive evidence that people do *not* seem to use the Ramadan fast as a way to shift consumption from the seasons of relative abundance (when the marginal utility of consumption should be lower) to other parts of the year. Indeed, somewhat paradoxically, I do find evidence of a decrease in consumption during Ramadan in a year when it falls shortly before the annual hunger season, despite the fact that we would expect many households to be running low on food at this point, and thus to have little leeway for a reduction in food intake from this benchmark. As this finding is based on only four rounds of cross-sectional survey data, I cannot feign any

[7] It is worth noting that these possibilities are not mutually exclusive; it could well be that both of these factors play a role in explaining seasonal hunger.

certainty that this result would replicate across other years, and I have no direct evidence for any particular explanation or mechanism underlying my results. However, I can rule out the idea that (1) savings are significantly higher when Ramadan falls closer to the peak season , and that (2) the absolute reduction in calories due to Ramadan would be weakly decreasing as the holiday falls closer to the hunger season. This could be consistent with the idea that people do not begin trying to plan and save for the hunger season until it is imminent – perhaps because of short-sightedness in planning and decision-making, or perhaps because of constraints that make it difficult or very costly to start saving earlier in the year (e.g. food spoilage, which could act as a negative and decreasing interest rate on savings). It could also be consistent with the idea that post-fast meals are simply larger affairs during the peak season, while closer to the hunger season households throughout the community lack the resources to put together a large communal meal. These are all interesting possibilities that merit further investigation, but it is beyond the scope of this paper to distinguish between different mechanisms that might account for the observed relationship between the timing of Ramadan within the agricultural cycle and its effects on consumption among Muslim households.

2.4 The setting: Malawi

Chirwa, Dorward, and Vigneri (2012)[8] document that over 80 percent of the population of Malawi is dependent on agriculture, and that this is largely rain-fed, with less than 5 percent of cultivated land being irrigated. This makes households highly susceptible to agricultural seasons. It is also a very poor country, with 56% of the rural population below the poverty line (and indeed, as has been discussed, this might underestimate the percentage who fall below the poverty line at some point each year), and 57% reporting inadequate food consumption. The primary staple crop is maize, and the pattern of seasonal hunger in the country closely tracks the crop cycle of maize. Maize is generally harvested between April and June. Prices

[8] The remainder of this description of Malawi's agricultural cycle comes largely from this same work.

are lowest during this time as crops flood the market. The proportion of households falling below the poverty line is also lowest at this time, consistent with the idea of poverty being a partly seasonal phenomenon. Most households begin to run out of their own food stocks around August or September, and maize prices jump at this point. According to survey data from 2004–2005,[9] 81% of households who grew food crops during this agricultural season had exhausted their stocks by December. Food prices continue to climb substantially from December through March, which is generally the hunger season. Many use casual labor, or *ganyu*, as a source of income during this time; large numbers of people looking for work at the same time tends to drive down wages. So the hunger season is characterized by the triple-burden of minimal food reserves, high market food prices, and low wages. On top of this, doing *ganyu* work on others' land takes away from households' ability to tend to their own land. Finally, early maize crops start to come in in March. People will boil maize in its less ripe "green" form in order to be able to eat it immediately. While this reduces the total caloric content of the crop, and thus potentially reduces the total availability of calories for the rest of the year, it enables households to put an end to the hunger season and to be reasonably well-nourished in time to harvest the bulk of their crop.

One feature of Malawi's demographics that makes it an interesting setting to study the impact of Ramadan is that Muslims are a relatively small minority population in the country – approximately 14 percent based on my calculations from the IHS5 data (NSO2021). This proportion is large enough that we should expect to encounter a reasonably large sample of Muslim households in any nationally-representative survey. But importantly, it is also small enough that it seems reasonable to think that general equilibrium effects of Ramadan or other major Muslim holidays on prices for non-Muslims should be limited. Specifically, we might be concerned that a decrease in demand for food among Muslims during Ramadan could lower food prices paid by non-Muslims, thereby increasing caloric intake among non-Muslims and exaggerating the estimated effect of Ramadan on Muslims' calorie consumption. This

[9] Specifically, the IHS2 dataset (NSO (National Statistics Office) 2005b), which will be one of my main data sources.

would be a greater threat to identification in a country with a larger proportion of Muslim households.

3 Data and Empirical Strategy

3.1 Data

The data I use for this paper come from the Malawi IHS (Integrated Household Survey), Rounds 2–5 (NSO 2005b; 2012; 2018; 2021), unless otherwise noted. I include rural households only, and for consistency across rounds also exclude households from the small island district of Likoma, which was until recently excluded from the survey. I construct the following key variables for analysis:

Interview date: In certain instances, there is conflicting or inconsistent information about the interview date in the data provided. Because the purpose of this paper was to look at how consumption changes at Ramadan, accounting for seasonal consumption patterns, knowing the precise date was fairly important. To this end, in order to verify that the interview dates were recorded correctly, I grouped households by village, knowing that teams of interviewers would generally go to a village together for a few days at a time. In cases where the interview dates were unclear or inconsistent with neighboring households, I used the interview dates of the neighboring households as well as information about the interviewers in order to impute the correct interview date.

Ramadan dates: The month of Ramadan begins at different times across the world, according to different local standards. The most common standard is the sighting of the new crescent moon, indicating that the following day will be the first day of the new month. Similarly, the end of Ramadan is also usually determined based on the sighting of the following new crescent moon, with the caveat that the month will always be no fewer than 29 days and no more than 30 days. Wherever possible, I used reports from Malawian news outlets or statements from local religious leaders to determine the start and end dates of

Ramadan. Where this information was unavailable, I consulted the Islamic Crescent Observation Project or ICOP (International Astronomical Center 2021), a crescent-sighting tracking website, which accumulates local reports of new moon sightings as well as astronomical maps of when moon sightings would have been expected in different parts of the world. Where I could not find information specific to Malawi, I relied on references from nearby countries, namely Zambia, Tanzania, and South Africa, which typically follow the same calendar; I also cross-referenced the astronomical maps on this site to make sure a crescent-sighting would have been possible on the designated date. Finally, I generated a variable indicating whether or not an interview overlapped with Ramadan. Because the interview is backward-looking, I classified it as occurring during Ramadan if it took place at least two days after the first day of Ramadan, so that at least three days of the previous week would have been fast days. If the interview took place after the final day of the fast, I did not classify it as taking place during Ramadan, even if several days of the prior week would have been fast days, because the end of Ramadan is marked with feasts and celebrations that might obscure the effects of fasting.

Muslim household: The religion of each adult household member was given separately, so designating a household as "Muslim" or "non-Muslim" was not straightforward in cases where different household members practice different religions. I identified a household as Muslim if either (1) at least half of the household members related to the household head (i.e. excluding friends, servants, etc.) were Muslim, or (2) if the household head was Muslim and at least one of his or her family members in the household was also Muslim.

Household size/adult equivalent: I included in the measure of household size all household members who were staying in the household at least one day during the week leading up to the interview.[10] Then, in order to adjust for the lower caloric requirements of children, I used information from the NSO (2005a) to construct a Malawi-specific discount factor by

[10]In cases when half or more of the listed household members would not be considered "present" by this definition, I used the full size of the household. In most cases, this seemed to produce more reasonable numbers of calories per person per day.

age group for children under 16 years of age.

Average hours worked in the last week (household): In constructing this metric, I summed hours worked for each individual across different categories (household agriculture, work in non-agricultural household business, work for a wage/salary, casual day labor, work at apprenticeship). I made a few adjustments to these measures, including truncating any hours reported above 84 hours per week total, and only including the maximum value of "working in household business" and "running household business" if both values were 10 or more hours, as these often seemed to overlap. Once this measure was constructed, I took an average of hours within each household, excluding (a) household members who were not of "working age",[11] (b) those who did not work at all, and (c) servants or lodgers included in the household roster.

In the following subsection, I explain the methodology that I use to construct my primary outcome measure of interest: calorie intake.

3.2 Measuring calories consumed

I constructed total calorie[12] counts for the household for the week using the following procedure:[13]

1. *Perform consistency checks on the default unit conversions provided with the data download.* For each food item, I compared the provided unit conversions (a) across the 3 regions; (b) across different sub-units of the same unit type (for example, comparing a "small" pail to a "medium" or "large" pail); and (c) across different units that

[11] I defined "working age" quite liberally to include anyone between 16 and 69 years old, in order to exclude as few households with working adults as possible while calculating labor supply mainly among those who would be expected to fully participate in the labor force when and if possible.

[12] Throughout this paper, I refer to "calories" under the typical colloquial definition, which in the scientific literature would generally be referred to as kilocalories or (capitalized) "Calories."

[13] Despite being widely regarded for their detailed treatment of theoretical and practical considerations for analyzing consumption data and living standards, neither Deaton (1997) nor Deaton and Zaidi (2002) give any notable guidance for the most vexing issues I encountered in working with the Malawian food survey data: converting units to standard metrics, converting standard quantities to calories, and identifying and cleaning errors, inconsistencies, or outliers in large and complex food consumption datasets.

would generally have similar conversions (for example, a "pail" and a "basin"). In all cases I corrected conversions that seemed unlikely or implausible with a number that seemed reasonable and more in line with similar observations.

2. *Cross-check the provided calorie conversions* for the most recent round of data (Round 5) against those from Round 3, and fix those that seem implausible or vary excessively from other metrics. Where these metrics were in conflict or seemed implausible, I checked them against (1) the MAFOODS database (MAFOODS 2019) and (2) the USDA database (Haytowitz et al. 2019). In general, I kept the most recent provided calorie conversions unless there was a large difference and/or a compelling reason to think that one of the other sources was more accurate.

3. *Average across regions.* Although there may be good reason to think that there are systematic differences in the quantity received from a plate in one region versus another, I chose to use the average factor across regions, because differentiating by region seemed to add unnecessary noise that might have biased the calorie measure. It is also true that interviewers in all regions had the same photo guides on which they were supposed to base the unit selections; units were meant to represent a similar amount across all households.

4. *Impute missing unit types using provided units.* In many cases, I observed that certain units had a consistent relationship to others. For example, a small pail generally had a conversion factor that was quite close to the factor given for a five-liter bucket. So in cases where only one of these units was provided, I added the other unit based on this observation in order to be able to convert more observations.

5. *Match up "other" units to recognized units where possible.* In many cases, interviewers typed in custom units that exactly or nearly matched existing units. Where possible, I matched these typed-in units to a corresponding standard unit.

6. *Convert to kilograms using the exact sub-unit.* Match unit code directly to the constructed list of unit-to-kg conversions for each food item.

7. *If unmatched, try alternate sub-units and/or removing sub-unit.* For example, if the specified unit was "4B" for "medium pail", and no conversion is given for "4B", I attempt to convert using "4A" ("small pail") or just "4" ("pail").

8. *If still unmatched, assign a default unit.* For each food type, I specified a unit that seemed like a plausible default. For example, for items that usually come as pieces like mangos or cucumbers, I used "Piece" as a default. For smaller grains or flours, I used "Cup." For items that generally come in bunches, I used "Heap" or "Plate," depending on what conversions were available or were possible to plausibly construct.

9. *Complete the final conversion.* Multiply the specified quantity by the kg-per-unit conversion, then by the calorie-per-kg and edible-portion-per-kg factors.[14] This provides a sum for the total calories consumed by all household members over the last week.

10. *Divide by adjusted household size* to produce a measure of calories per adult-equivalent for the week, then divide by 7 to convert to a measure of average calories per day.

11. *Address outliers.* In all rounds of data, there are a number of households for which the raw metrics would indicate tens of thousands of calories consumed per household member per day. Because extreme outliers could significantly affect my regression analyses, I chose to censor households with calories consumed per adult-equivalent below the 2nd percentile or above the 95th percentile for each round of data. This, including the asymmetry of the cutoff points, is in line with what other researchers have done when using the Malawi IHS data.[15] An alternative to dropping these households

[14] I also implemented a few checks at this point in the process to make sure the quantity being converted seemed relatively reasonable before assigning a calorie value.

[15] Ecker and Qaim (2011), for example, in data generously shared with the author, excluded all households with calculated calories per person below 500 or above 5000, which matches roughly to my percentile cutoffs. The asymmetry in the cutoffs is also reflective of the heavily right-tailed nature of the raw distribution: the minimal level of calories I can observe is 0, whereas there is no maximum; further, very low quantities are

from the analysis of calories would be to Winsorize these outliers, for example, assigning any measure below the 2nd percentile with 2nd percentile value, and those above the 95th percentile with the 95th percentile variable. While in many circumstances, this would be preferable to excluding such cases, in this case there is good reason to think that households with very large calculated levels of caloric intake are more likely to have some food quantity recorded incorrectly, and are not necessarily more likely to fall near the top of the calorie consumption distribution. Because of this, I opted to exclude these households rather than to potentially introduce a substantial source of noise into my analysis.

This procedure is largely in line with what other researchers working with these data have done in constructing a measure of calories consumed. However, I added a number of consistency checks throughout the procedure, including cross-checks and corrections of the provided unit conversions, in an effort to try to construct an internally consistent and generally reasonable metric. I also introduced default measures, which allows me to avoid having to convert every possible unit provided without counting everything unmatched as zeros. Both of these innovations should help generate a somewhat more accurate measure of calories consumed. The primary difference between my approach and the one employed by Verduzco-Gallo et al. (2014) is that I rely on my cleaned version of NSO's provided unit conversions where possible, complemented by a set of default units, whereas Verduzco-Gallo et al. (2014) preference their own construction of "implicit" unit conversion factors, based on a comparison of the median price per unit to the price paid per kilogram, wherever feasible. While my methodology is admittedly more ad-hoc, it allows me to be somewhat more assured that the quantities imputed are based more directly on the quantities that respondents were asked to select from, and alleviates the potential concern that using consumption data to construct conversions

more likely to be accurate, especially in a country with high levels of poverty and food insecurity, whereas reports of very high levels of caloric intake (in excess of 6000 calories per day for an entire week, for example) are highly implausible and much more likely to be attributable to measurement error. I would thus argue that the measurement error itself is likely to be highly asymmetric and right-tailed, though it is difficult to prove this directly. Gilbert et al. (2019) and Verduzco-Gallo et al. (2014) drop households reporting more than 8,000 or fewer than 200 calories per person per day, which seemed perhaps excessively permissive.

for the same set of data could amplify pre-existing biases or measurement error.[16]

These concerns also speak to potential survey design issues. Allowing for a vast array of units to choose from (including an "Other" option) is theoretically appealing as a way to facilitate recording household consumption exactly as reported with minimal distortion from the interviewer. For this reason, Smith et al. (2014) neither favor nor disfavor this approach compared to potential alternatives, stating that more evidence is needed to determine what approach provides the most accurate results. However, if "pail," "heap," and "bunch" can have such vastly different interpretations from one household to the next, it is unclear how much value we truly derive from allowing for such a large set of options in terms of increased accuracy, especially if there is additional noise coming from the unit-to-kg conversions themselves, and if many units have no conversion specified at all. Given this, it might make more sense to limit the number of potential units, possibly sacrificing some ability to reflect a respondent's exact words, but perhaps increasing the likelihood of recording the actual quantity consumed with a reasonable degree of accuracy and consistency across households. Alternatively, interviewers could be trained to try to direct respondents to more readily convertible units, while still allowing them to choose from a broader menu when necessary. In addition, as Smith et al. (2014) advocate, it could be productive to expand the use of demonstration methods, such as photo examples, that help ensure that the respondent can give the interviewer a clear sense of the total quantity consumed of any particular good.

3.3 Identification strategy

The regression specification for the main analysis of this paper is as follows:

[16] Having said this, calculating implicit conversions could be a useful complement to the methodology I present; it could provide a useful consistency check, especially in cases where I resort to default units, and could also help identify systematic biases in calculated price per kg due to inaccurate unit conversions.

$$Y_{idt} = \beta^{M} Muslim_{idt} + \beta^{R} Ramadan_t + \beta^{MR} Muslim_{idt} \times Ramadan_t$$
$$+ \gamma_d \left[\times \tau_t Round_t\right] + \mu_t \times \rho_d Region_d + \varepsilon_{idt}, \quad (1)$$

where outcomes Y are observed for household i in district d at time t.

This specification sets up a difference-in-differences analysis in which I analyze differences in outcomes – namely calorie intake and labor supply – between Muslims and non-Muslims during Ramadan and in all other parts of the year. In this framework, β^{MR} represents the difference-in-differences in outcomes, and is the main coefficient of interest. Because Muslim households are disproportionately concentrated in a relatively small number of districts, I include district fixed effects (γ_d) in all specifications, and allow these to vary from round to round of the survey data by interacting them with τ_t, a set of round fixed effects.[17] In addition, in order to control for potentially large seasonal differences in consumption, I include monthly fixed effects (μ_t) interacted with region fixed effects (ρ_d) to account for differences in timing of the agricultural seasons in different parts of the country (specifically, North, Centre, or South). The identifying assumption, as with any difference-in-differences analysis, is that of parallel trends. Specifically, this assumption means that the comparison of Muslim households to non-Muslim households, outside of Ramadan, serves as a valid counterfactual for the expected difference between Muslim and non-Muslim households observed during Ramadan, if Ramadan did not take place.

On the one hand, the assumption of parallel trends seems quite reasonable, as households were selected randomly and interviews were scheduled to ensure that, within each district, interviews were spread out evenly across the months of the year. Nonetheless, there are a number of potential caveats to bear in mind. For one, surveys are cross-sectional, so each household was only interviewed once. Thus, we might see differences in various measures

[17] Of course, this interaction with round fixed effects is omitted in analyses that include only one round of data.

that are simply due to heterogeneity between households and between different villages, even within the same district. Put differently, there could be unobserved omitted variables that make households differ from one another and potentially affect our observations of Ramadan. For example, if a very wealthy household was visited during Ramadan, and much poorer households were visited at other times of the year, this could bias our estimates of the effect of Ramadan. Related to this concern, because Ramadan takes place for only one month out of the year, and because Muslims make up a relatively small proportion of the population in Malawi, the number of Muslim households observed during Ramadan for any given year will be quite small, giving us relatively low statistical power to detect differences between these households and those observed at other times of the year.

Another potential concern is that changes to consumption behavior during Ramadan affect consumption decisions for the rest of the year. If consumption decreases (increases) during Ramadan, households on a fixed budget may have more (less) money available to spend on food in the rest of the year. This could exaggerate the observed effects of Ramadan by increasing the difference from the rest of the year.[18]

A final threat to identification comes from potential measurement error. In particular, if Muslims are more likely to report consumption of certain types or quantities of foods during Ramadan, then any inaccuracies in the conversion factors used for these foods might bias our measure of the effects of Ramadan on caloric intake. Similarly, if other particular groups have a tendency to consume certain food items, and these are measured inaccurately, this could also bias our results by skewing the measures of caloric intake that we rely upon for comparison. Results might also be sensitive to different ways of addressing outliers in consumption measures, and excluding outliers could introduce selection bias. Finally, as is generally the case when there is measurement error, it could introduce attenuation bias wherein noise in the data biases our coefficient estimates towards zero and makes it more

[18]This is not necessarily a problem. It would affect the interpretation of the coefficient on $Muslim \times Ramadan$, but the coefficient would still accurately reflect the difference between Ramadan and the rest of the year.

difficult to detect the true effect that we wish to identify.

3.4 Seasonality

Appropriately controlling for seasonality is vital to any analysis of consumption data in Malawi. Considering the dramatic seasonal fluctuations in food prices, along with the large percentage of households whose consumption follows a pronounced seasonal trajectory, to not account for the time of year in which a household was interviewed would introduce significant bias into consumption measures.

In thinking about how to appropriately construct measures that will control for seasonality in consumption, it is important to note that there are regional differences in the timing of seasons in different parts of the country, and that there is some variation from one year to the next as to the actual timing of the harvest, as well as the other agricultural seasons (dry season and rainy season). To account for this, my preferred specifications all include Month$\times$Year$\times$Region fixed effects.

One potential downside of controlling for Month$\times$Year effects (e.g. March 2004, April 2004, etc.), as opposed to simple Month-of-Year effects (e.g. March, April, etc.), is that in allowing for different Month effects in each year, we may capture some meaningful inter-household variation as part of a seasonal effect, when in fact it would be correctly attributed to the specific composition of households.[19] However, given that in some rounds of survey data, I see sizable effects of Ramadan on the overall population, even after including Month$\times$Year fixed effects, these are likely seasonal effects that have not been fully accounted for with these controls, so I am not very concerned about these seasonal controls being excessively aggressive. On the other hand, we might consider controlling for a narrower timeframe than months. As an alternative, I repeated the analysis using Half-Month fixed effects (re-

[19]To illustrate this concern, consider an extreme example in which we introduced fixed effects for every interview date. While this would very specifically capture differences in the timing of seasons from year to year, it would generate very imprecise estimates of these effects and likely wash out many meaningful differences between households observed on different dates.

sults not shown).[20] This seemed to have very little effect on the overall results.

4 Results

Table 1 shows summary statistics for calories consumed per age-adjusted household member per day for each round of survey data. A few things are worth pointing out. First, note that the distribution of calorie consumption in Round 2 is substantially higher than in other 3 rounds. For a number of reasons, this is unlikely to be a reflection of truly higher levels of caloric intake in the country during this year compared to the latter years. Indeed, it directly contradicts the finding of increases in per capita calorie consumption from Round 2 to Round 3 documented in Verduzco-Gallo et al. (2014). More likely, this is due to differences in the specific implementation of the food consumption module in this round, or even perhaps in the design of the household roster and counts of household members. Fortunately, I can rule out the possibility that this anomaly is attributable to an issue in my calorie conversion methodology, as I find this same discrepancy between Round 2 and the other rounds when I construct an analogous calorie metric using the household level total calorie counts produced by other researchers (Ecker and Qaim (2011); International Food Policy Research Institute (IFPRI) (2020)).[21]

[20] These results are available from the author upon request.

[21] I remain perplexed as to why the Round 2 calorie measures are so much higher, and why this is also true of the replication data, which in my understanding comes from the *same constructed calorie counts* from which other researchers have noted an increase in calorie consumption from Round 2 to Round 3.

Table 1: Daily Calories per Adult-Equivalent by Round

	Round 2 (04-05)		Round 3 (10-11)		Round 4 (16-17)		Round 5 (19-20)	
	All HH	Incl. HH	All HH	Incl. HH	All HH	Incl. HH	All HH	Incl. HH
HH Mean Calories	2,706.42	2,687.42	2,106.49	2,049.84	2,093.05	2,001.22	2,207.89	2,148.10
	(10.47)	(11.09)	(10.64)	(10.96)	(9.69)	(10.19)	(11.03)	(12.04)
HH Median	2,539.56	2,522.56	1,979.83	1,914.88	1,970.29	1,862.61	2098.20	2028.13
HH 10th %ile	1,485.21	1,463.14	1,070.57	1,036.97	1,083.21	1,042.35	1,108.25	1,074.81
HH 90th %ile	4,219.26	4,206.27	3,499.25	3,418.68	3,370.18	3,223.91	3,713.38	3,611.14
Observations	10,491	9,143	11,412	9,368	11,577	9,431	10,633	8,687

Means and standard errors (in parentheses) for calories are calculated using sampling weights; median and percentile measures are not.
All calculations exclude households that were above the 95th percentile or below the 2nd percentile within each round in terms of calories per adult-equivalent.
"Incl. HH" excludes households in urban areas as well as in the small Likoma district, which was only visited in more recent survey rounds.

For the latter three rounds of data, average consumption levels appear to be reasonably in line with what we might expect for a relatively poor country like Malawi. Specifically, about half of the households surveyed are consuming less than 2,000 calories per adult on average. If anything, these metrics might still be somewhat overestimated, as there is a relatively large number of households reported consuming 3,000 or more calories per adult-equivalent per day.[22] The difference between the distributions among all households[23] versus those included in the analysis is also reasonably in line with expectations: the analysis includes only rural households, so it makes sense that when we add in urban households it generates a small increase in the level of calories consumed at all points along the distribution, and in all rounds of data.

Table 2 provides a sanity check as to the novel consumption metrics that I constructed, by comparing them to those used by other researchers. Across all 3 rounds for which other researchers' metrics were available, my metrics are significantly lower on average than the replication data, particularly so for Rounds 3 and 4. However, I do not necessarily take

[22]On the other hand, caloric requirements are higher for people working in agriculture and with limited access to mechanization or labor-saving household devices (see Deaton and Drèze 2009), so these figures might be a bit low even for the very poor.

[23]Note that these figures do not actually represent *all* interviewed households, as they exclude "censored" calorie counts – those below the 2nd percentile or above the 95th percentile for each survey round.

this as evidence of the measures I used being less accurate. I used similar methodologies in constructing these metrics to other researchers, but also incorporated various consistency checks in the unit and calorie conversions that were not used in constructing alternative metrics. I also specifically investigated some of the large outliers in the replication data to identify improvements I might make to the calorie conversion procedures, and compared the details of some of the largest discrepancies to look for potential errors in my own conversions. This gives me a reasonable degree of confidence that for the Round 5 data, for which there were no replication data available, my constructed calorie metrics are sensible and in line with other rounds of data.[24]

Table 2: Mean calories compared to replication data

	Constructed calories	Replication calories	Difference
Round 2 (04-05)	2,603	2,659	-55.4^{***}
	(9.36)	(9.83)	(4.78)
Round 3 (10-11)	2,120	2,569	-449.3^{***}
	(8.48)	(9.52)	(5.81)
Round 4 (16-17)	2,086	2,355	-269.0^{***}
	(7.90)	(8.17)	(4.86)
Round 5 (19-20)	2,270	–	–
	(9.65)		

$^{***}p < .01$. All means listed in this table are for censored measures of calories per adult-equivalent and calculations are unweighted. Standard errors of means (unweighted) are in parenthesis. Differences for all three rounds with replication data are signficant with a p-value $< .0001$.

Table 3 shows the results of an analysis testing the identifying assumption of the paper by looking at the difference in terms of various observable characteristics between Muslim and non-Muslim households observed during Ramadan versus at other times of year on various characteristics. In general, on most characteristics, the identifying assumption seems to hold. While there are several significant differences between Muslims and non-Muslims,

[24]I also ran all of my main analyses for Rounds 2–4 using calorie measures from these replication data, and the findings were quite similar to the ones presented here. Results from these analyses are available from the author upon request.

those differences are, for the most part, relatively similar for households observed during Ramadan compared to those observed at other times of year. Particularly encouraging is that in most cases where there does appear to be a large Muslim versus non-Muslim differential, that differential appears to be relatively stable.

Table 3: Balance in Observable Household Characteristics

	During Ramadan			Rest of Year			Diff-in-Diff					
	(1)	(2)	(3) Diff. (within	(4)	(5)	(6) Diff. (within	(7) **All**	(8) **p-value**	(9)	(10) Within Rounds:	(11)	(12)
	Muslim	Non-Mus	Distr×Round)	Muslim	Non-Mus	Distr×Round)	**Rounds**		2	3	4	5
No. HH members	4.83	4.53	0.13	4.50	4.44	0.17***	0.24**	.050	0.68**	0.01	0.16	0.22
	(0.11)	(0.05)	(0.16)	(0.04)	(0.02)	(0.05)	(0.12)		(0.33)	(0.32)	(0.30)	(0.17)
No. of children (< 16 yrs.)	2.56	2.19	0.08	2.35	2.15	0.19***	0.16	.152	0.41	−0.14	0.31	0.11
	(0.09)	(0.04)	(0.12)	(0.03)	(0.01)	(0.04)	(0.11)		(0.30)	(0.21)	(0.27)	(0.14)
No. working-age adults	2.10	2.18	−0.004	2.01	2.16	−0.023	0.08	.227	0.21	0.11	−0.02	0.07
	(0.06)	(0.04)	(0.06)	(0.02)	(0.01)	(0.025)	(0.06)		(0.15)	(0.18)	(0.11)	(0.10)
Female-headed (%)	39.8	27.4	7.65**	34.9	27.6	1.58	4.40	.126	3.1	3.5	4.0	5.5
	(3.5)	(1.2)	(3.56)	(1.1)	(0.36)	(1.07)	(2.87)		(4.6)	(4.6)	(4.0)	(6.0)
Age of HH head	43.1	44.1	1.71	43.4	43.4	0.27	−0.98	.496	1.5	1.9	−7.1***	0.68
	(1.5)	(0.48)	(1.17)	(0.32)	(0.11)	(0.38)	(1.44)		(3.1)	(1.6)	(1.7)	(2.3)
HH head married (%)	71.7	72.0	−2.52	70.4	71.4	2.92***	0.53	.903	1.2	−9.3**	14.1*	−3.7
	(3.7)	(1.0)	(3.62)	(0.80)	(0.35)	(1.12)	(4.36)		(5.7)	(4.5)	(8.3)	(5.2)
HH head always lived in village (%)	74.8	66.3	10.24**	69.8	66.7	4.79***	5.94*	.070	−5.4	−5.0	8.7**	12.5**
	(4.1)	(2.1)	(4.17)	(1.2)	(0.57)	(1.43)	(3.3)		(9.6)	(6.9)	(4.2)	(5.1)
Literate (% age ≥12)	52.2	68.1	−16.80***	51.6	66.8	−9.89***	−2.18	.413	1.1	2.1	−7.7	−2.5
	(3.1)	(1.2)	(3.40)	(1.0)	(0.41)	(1.13)	(2.67)		(4.4)	(4.6)	(5.1)	(3.4)
Ever attended school (% age ≥ 5)	72.7	82.2	−8.87***	70.2	80.9	−6.73***	0.40	.837	1.6	8.4**	−0.4	−4.3*
	(2.0)	(0.94)	(2.21)	(0.92)	(0.28)	(0.75)	(1.92)		(3.4)	(4.1)	(4.3)	(2.5)
In school (% 5≤age≤17)	88.7	91.2	−0.58	90.6	92.2	−0.79	−1.14	.612	−2.4	−0.7	−3.8	0.0
	(2.1)	(0.90)	(2.57)	(0.54)	(0.21)	(0.68)	(2.24)		(5.1)	(3.9)	(4.0)	(3.8)
Completed primary school (% age ≥14)	24.4	35.1	−14.26***	24.4	35.5	−8.23***	−1.44	.573	5.0	−4.5	−3.4	−0.6
	(2.7)	(1.5)	(3.66)	(0.76)	(0.44)	(1.11)	(2.55)		(3.8)	(4.0)	(5.3)	(4.1)
Completed secondary school (% age ≥18)	3.1	9.8	−10.41***	4.5	9.0	−5.00***	−2.89**	.035	−2.6	−0.5	−5.7*	−2.0
	(1.1)	(1.1)	(2.79)	(0.39)	(0.23)	(0.69)	(1.37)		(2.2)	(2.2)	(3.1)	(2.0)
N (households)†	355	2,141	2,496	4,465	32,242	36,707	39,203		9,840	10,038	10,015	9,310

Difference within each category (Columns 3 and 6) represents the coefficient on "Muslim" in a regression of the observable variable on "Muslim" after controlling Round × District Fixed Effects. For this reason, the Difference will not precisely equal the difference between the Muslim and non-Muslim means presented in the previous two columns. Diff-in-diff (Columns 7-12) represents the coefficient on Muslim *times* Ramadan in an analagous regression (that also controls for an "IsRamadan" dummy). Standard errors clustered by Enumeration Area within each round (the primary sampling unit) are in parentheses. Urban households excluded from analysis. All means and regressions are weighted based on sampling weights. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

† Respresents the number of households in the given category for each column. For some variables, means and regressions are based on a smaller number of observations as the variable was missing in some households.

We should, however, note two points of caution. First, some of the differences in characteristics are significant, such that it is possible that there are some underlying differences between households observed during Ramadan versus the rest of the year, even if those differences are coincidental. Second, we should be particularly cautious in reading into cross-round differences in estimated effects, as they might be purely idiosyncratic. For example, if we look at the "household head married" variable, the overall difference is very close to zero; however, in one round, the difference is negative and significant, while in another, the difference is positive and nearly significant. This could indicate some meaningful difference between the populations observed across these rounds, in violation of our identifying assumption, or it could be largely statistical noise. Thus, when it comes to our main outcome variables of interest (calorie consumption and labor supply), though we do observe some potentially interesting variation from round to round, we should be cautious in attributing differences that we observe between rounds to any particular explanation.

Table 4 shows results from the main analysis of the effects of Ramadan on consumption. Columns (1) and (2) show the results using all four rounds of data together. We can see that Muslim households overall have slightly lower levels of calorie intake on average, even after controlling for the district in which they live and seasonal consumption effects. This effect is small – about 40 calories per day, approximately a 2 percent decrease – and only marginally significant. As expected, the Ramadan effect for non-Muslims is quite close to zero on average.[25] Somewhat surprisingly, I detect no significant effect of Ramadan overall on calorie consumption among Muslims. The 95% confidence interval for the average Muslim×Ramadan effect across rounds places rough bounds on the effect of Ramadan on calorie consumption for Muslim households from a decrease of 120 calories per person per day to an increase of 267 calories (or between a 5 percent decrease and a 12 percent increase, based on the regression on log-calories).[26] This would seem to allow us to rule out a very large decrease in calories consumed on average across the four rounds of data. However, it should be emphasized that because we do not observe calories consumed at the individual level, we cannot rule out the presence of household members who would not be expected to fast (such as children and the elderly) potentially obscuring a reduction in calories among healthy, working-age adults.

It also seems to be the case that this observed effect varies substantially across the survey rounds; looking only at the average effect across rounds ignores these potentially interesting (or potentially spurious) differences. Looking at Columns (3) and (4), we see

[25]One benefit of studying this question in Malawi is that, because Muslims are a relatively small minority, changes in their consumption behavior should not have spillover effects into food prices or consumption levels for non-Muslims, which may be more of a concern in a country with a Muslim majority or a relatively larger Muslim minority population.

[26]Interestingly, these bounds look somewhat similar at the lower end to the estimated reduction of 150-200 calories per person that Sadeghirad et al. (2012) report as an average for people studied in the Middle East and East Asia, and also very similar at the higher end to the estimated average increase of about 260 calories per person per day in North Africa.

that in Round 2 of the survey, which took place in 2004–2005, there is a highly significant decrease in consumption among Muslim households during Ramadan. I find a decrease in calorie consumption for Muslim households of about 370 calories per adult-equivalent per day during Ramadan or about a 10% decrease when using log-calories as the outcome variable. This estimate is still somewhat less than, for example, the estimated reduction of around 600 calories per person per day among Muslim households reported by Schofield (2020). Again, this could partly be explained by the fact that some household members are not expected to fast. Additionally, Ramadan took place from mid-October to mid-November during this survey round, which in Malawi is generally just before the beginning of the hunger season. By this point in the season, most households would have run out of their own stockpile of grain, and often many will have already begun to cut back somewhat on consumption. In other words, the baseline household may already be consuming less at this time of year, such that fasting might not make as large of a difference compared to relatively more abundant times of year. Finally, it could simply be that the estimates in Schofield (2020) are context-specific and larger than the effects seen elsewhere; indeed, the meta-analysis by Sadeghirad et al. (2012) suggests an average daily individual-level calorie reduction of 150–200 calories per day in the Middle East and East Asia – by this standard, a reduction of 370 calories per day is quite large.

Table 4: Main Analysis: Calories

	All Rounds		Round 2		Round 3		Round 4		Round 5	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Cal.	ln(Cal)	Cal.	ln(Cal)	Cal.	ln(Cal)	Cal.	ln(Cal)	Cal.	ln(Cal)
Muslim	−39.85	−.021*	1.40	−.008	−38.69	−.018	−67.13	−.028	−42.38	−.024
	(24.60)	(.012)	(50.83)	(.020)	(47.34)	(.022)	(49.15)	(.025)	(47.40)	(.022)
Is Ramadan	−28.13	−.014	62.01	.008	−88.12	−.017	258.36***	.123***	−252.48***	−.118***
	(51.65)	(.022)	(80.09)	(.030)	(113.31)	(.057)	(75.22)	(.037)	(94.92)	(.039)
Muslim × Ramadan	74.01	.037	−369.75***	−.104**	−21.44	−.024	268.44	.105	180.61	.087
	(98.88)	(.043)	(109.29)	(.041)	(129.47)	(.066)	(182.46)	(.078)	(171.21)	(.077)
District × Round FEs	Y	Y								
District FEs			Y	Y	Y	Y	Y	Y	Y	Y
Month-Yr × Region FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	36,629	36,629	9,143	9,143	9,368	9,368	9,431	9,431	8,687	8,687

Standard errors clustered by Enumeration Area within each round (the primary sampling unit) are in parentheses. Urban households excluded from analysis. All regressions are weighted based on sampling weights, with weights adjusted within each district-rural area to account for observations excluded due to censoring the conusmption variable. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

Perhaps the most surprising result from this analysis is that in the remaining three rounds of survey data, despite the fact that Ramadan falls earlier in the harvest cycle, we see no decrease in caloric intake among Muslim households during Ramadan. In Round 3 (Columns 5 and 6), we see essentially no significant effect of Ramadan on calorie consumption. In Round 4 (Columns 7 and 8), we see that there is a highly significant and positive effect of Ramadan on calories consumed – but that this applies to both Muslim and non-Muslim households. If anything, Muslim households seem to experience an *even larger* increase in calorie consumption during Ramadan in this round compared to their non-Muslim counterparts. This increase is potentially quite large (the point estimate shows about 270 additional calories per person per day, on top of the 260 calorie increase in the general population), and while it is not statistically significantly different from zero, the lower bound of the 95% confidence interval is a decrease of approximately 90 calories, ruling out a decrease of similar magnitude to Round 2. In Round 5, we see a large and significant decrease in calorie consumption during Ramadan across all households, while Muslim households in this round seem to partially offset this decrease, consuming more than non-Muslims during Ramadan, though again this offset is not statistically significant. These results all remain qualitatively similar when using alternative fixed-effects measures to control for seasonality.[27]

How can we make sense of the above findings? While it is beyond the scope of this paper to provide a definitive answer to this question, I consider some possible explanations.

Understanding the timing of Ramadan in each survey round relative to the agricultural cycle provides helpful context. In Round 3, Ramadan falls between August and September, just as the harvest season is winding down in most parts of Malawi. In Round 4, Ramadan falls between June and July, which is generally the peak of the maize harvest, when prices are lowest and hunger is typically at its yearly minimum. In Round 5, Ramadan falls between May and June, which is on the early end of the main harvest season. Ideally, we would want to control for all seasonal effects so that our observed Ramadan effects would fully account for expected seasonal variations. It is clear that, at least in Round 4 and Round

[27] Specific results available from the author upon request.

5, we have not fully controlled for seasonality, as we see an increase in consumption during Ramadan in Round 4 among the general population above and beyond our seasonal controls, and a similar decrease in the general population in Round 5. That said, if the assumption of parallel trends holds, then we can still interpret the coefficient on Muslim×Ramadan as the consumption effect of Ramadan, as whatever additional seasonal effect is detected among the general population is assumed to apply to Muslims and non-Muslims analogously.

Table 5 helps to further elucidate some of this variation between rounds. First, in columns (1) -- (4), we conduct analyses similar to those shown in Table 4, but only including Rounds 4 and 5, in which Ramadan overlapped with the harvest. Here we can see more clearly that there is some suggestive evidence of a marginally significant increase in Muslim households' calorie intake during Ramadan when we pool the Muslim×Ramadan effect over these two rounds. We also see wide variation in the Ramadan effect for the population as a whole (i.e. including non-Muslims), which serves as an important reminder that I may not have completely accounted for seasonal effects, and that we should be cautious about any interpretation of round-to-round heterogeneity, as there could be important factors driving these findings that I have not fully accounted for. Columns (5) – (8) help to further illustrate this point. Columns (5) and (6) use Round 2 as a base year, and compare the effects in other rounds to this base. We see a highly significant negative coefficient on the base Muslim × Ramadan interaction, and we see a significant effect in the opposite direction in the other 3 rounds, particularly in Rounds 4 and 5. This affirms that the differences in the observed Ramadan effect on Muslims between rounds are significant. However, Columns (7) and (8) caution that the Muslim × Ramadan coefficients for Rounds 4 and 5 are not significantly different from the corresponding coefficient for Round 3, which is itself quite close to 0. The evidence for Ramadan causing an *increase* in consumption in these rounds is thus not definitive.

Table 5: Differences between Rounds: Calories

	Rounds 4 and 5 only, Round 4 as base				All Rounds, Round 2 as base		All Rounds, Round 3 as base	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Cal.	ln(Cal)	Cal.	ln(Cal)	Cal.	ln(Cal)	Cal.	ln(Cal)
Muslim [base round]	−54.4	−0.026	−63.9	−0.028	1.4	−0.008	−38.7	−0.018
	(34.1)	(0.017)	(49.0)	(0.025)	(50.8)	(0.020)	(47.3)	(0.022)
Muslim ×Round effects			17.6	0.003	Y	Y	Y	Y
			(67.2)	(0.03)				
Is Ramadan [base round]	−42.9	−0.019	266.0***	0.125***	62.0	0.008	−88.1	−0.017
	(79.5)	(0.035)	(74.5)	(0.037)	(80.0)	(0.030)	(113.3)	(0.057)
Ramadan ×Round effects			−524.9***	−0.244***	Y	Y	Y	Y
			(120.9)	(0.053)				
Muslim × Ramadan [base]	211.3	0.091	216.1*	0.094*	−369.8***	−0.104**	−21.4	−0.024
	(130.0)	(0.057)	(127.4)	(0.056)	(109.2)	(0.041)	(129.4)	(0.064)
Muslim × Ramadan×R2							−348.3**	−0.080
							(169.3)	(0.078)
Muslim × Ramadan×R3					348.3**	0.080		
					(169.3)	(0.080)		
Muslim × Ramadan×R4					638.2***	0.208**	289.9	0.128
					(212.6)	(0.088)	(223.6)	(0.103)
Muslim × Ramadan×R5					550.4***	0.190**	202.0	0.110
					(203.0)	(0.087)	(214.5)	(0.102)
District × Round FEs	Y	Y	Y	Y	Y	Y	Y	Y
Month-Yr × Region FEs	Y	Y	Y	Y	Y	Y	Y	Y
Observations	18,118	18,118	18,118	18,118	36,629	36,629	36,629	36,629

Standard errors clustered by Enumeration Area within each round (the primary sampling unit) are in parentheses. Urban households excluded from analysis. All regressions are weighted based on sampling weights, with weights adjusted within each district-rural area to account for observations excluded due to censoring the conusmption variable. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$. Muslim (Ramadan) × Round effects refers to Muslim (Ramadan) *times* Round 5 (compared to the Round 4 base) in Columns (3) and (4); for Columns (5)–(8), it indicates that interactions with all non-base round dummies are included in the regression, but not shown in the table for the sake of brevity/legibility.

One possible explanation as to why calorie intake could increase (or at least not decrease) for Muslim households during Ramadan is the importance of the post-fast *Iftar* meal. In many Muslim communities, this meal is observed as a nightly gathering in which members of the community come together to share food and celebrate. It seems plausible that the size and scope of this meal could depend on what resources the community has available. Perhaps then, in Malawi, *Iftar* could be far more extravagant during the peak harvest season, when food is relatively abundant, and much more limited as the hunger season approaches.[28] It is also the case that household members who would not be expected to fast, such as children, the elderly, and the ill would still presumably be invited to the post-fast meal, despite having eaten normally during the day. Thus, if that meal is quite substantial, increases

[28] Of course, I have no direct evidence for this, at this point it is purely speculative.

in calorie consumption among non-fasting individuals might outweigh the decreases among fasting individuals. The idea that the size of the post-fast meal might depend on the overall availability of resources could account for the observation of a decrease in calorie intake when Ramadan falls in October and a potential increase when it falls in May or June.[29]

Another intriguing possibility is that consumption behavior during Ramadan can be explained, at least in part, by myopic decision making with regard to the hunger season, which in turn could help us understand seasonal consumption behavior more broadly. Specifically, if indeed households are preparing a large feast when Ramadan falls during the peak harvest season and decreasing their overall consumption when the hunger season is about to begin, this would be consistent with the idea that households are simply failing to take measures to smooth their consumption levels over the year. We might also see a similar effect if Muslim households simply have less access to smoothing mechanisms, such as informal insurance or bank accounts, and thus are more prone to seasonal consumption fluctuations. The results from Round 5 provide some suggestive evidence against this possibility, as the Ramadan effect for Muslims, though it is not significant, appears to partially counteract the effect observed in the overall population, but certainly does not seem to exaggerate the effect.

If it is true that Muslim households are using Ramadan as a way to save resources to mitigate the effects of the hunger season, but only when Ramadan falls shortly before the hunger season, this might suggest that a similar salience effect is precisely what's driving seasonal variation in consumption more broadly. That is, the same decision-making process leading Muslims households to consume more on Ramadan when it overlaps with the peak harvest season, while starting to prepare for the hunger season only when it is imminently approaching, may be leading other households to consume at high levels around the peak season, leaving all of them with too little savings by the end of the year. In contrast, the standard economic model of intertemporal consumption decision-making would suggest that

[29]It is also plausible that charitable contributions, either from wealthier community members or from the international community, are helping the average Muslim household to consume more than they otherwise would during Ramadan.

if Ramadan is to serve as a vehicle for increased savings, it would be *most* likely to do so when it overlaps with the peak season, when baseline consumption levels are relatively high and the marginal utility of consumption is lower, while households would be least likely to increase savings when consumption is already relatively low in the lead-up to the hunger season. The pattern we observe in our data is *precisely the opposite* of what we might expect if this standard economic model of savings were correct – we can essentially rule out the possibility that this model would produce a result consistent with our empirical observations. However, we do not have sufficient evidence to conclude that this myopia-based explanation in particular is the correct one. More research would be required to establish that these results do indeed represent a consistent pattern and are not simply one-off idiosyncrasies of these specific years, and also to establish whether this explanation can account for such a pattern.[30]

It is also worth keeping in mind that measurement error and/or omitted variable bias could be affecting the presented results. As previously discussed, if certain food types are more prevalent among Muslims during Ramadan and are measured incorrectly when converting to calories, this would skew our observed effect of Ramadan. In addition, I am relying on district and month fixed effects to control for inter-household variation in (counterfactual) consumption, and those predictions surely will not provide a perfect counterfactual for the levels of consumption that we would expect from Muslim households in the absence of Ramadan. Any noise or bias in these counterfactual predictions will make it more difficult to determine the true effect of Ramadan. I cannot rule out either of these potential sources of bias.

Table 6 shows the effects of Ramadan on hours worked per week for Muslims and non-

[30]Another potential explanation that would be difficult to distinguish from the myopic planning mechanism is that households do not have access to adequate saving technology to allow them to transfer funds and/or food stocks from high-consumption times of year to the hunger season, or that such a technology is so costly that using it would result in a net welfare loss. If, for example, spoilage rates of stored food start to increase exponentially after a few months, it could help explain why Ramadan might allow some increased savings into the hunger season when it falls shortly beforehand, but not when it falls during the peak season. Direct measures of access to and use of savings mechanisms could allow us to investigate this possibility further.

Muslims. To construct this variable, I take the average hours worked across household members of working age (which I define as 16–69 years old) who reported working in the last week. I exclude lodgers, servants, and any of their family members who may be listed in the household roster but whose labor decisions are likely to be separate from those of other household members. Excluding individuals who did not work from the household average might introduce some bias into this measure, particularly if people who otherwise might have worked some number of hours choose not to do so during Ramadan. While I am more interested in studying the effects of Ramadan on intensive labor supply, under the premise that fasting might make it more difficult for people to work as much as they otherwise would have, I cannot rule out the possibility that some people would stop working altogether for parts of the month of Ramadan.[31]

Table 6: Main Analysis: Labor Supply

	All Rounds		Round 2		Round 3		Round 4		Round 5	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Hrs	ln(Hrs)	Hrs	ln(Hrs)	Hrs	ln(Hrs)	Hrs	ln(Hrs)	Hrs	ln(Hrs)
Muslim	−0.18	−0.001	−0.93	−0.040	0.12	0.009	−0.61	−0.011	0.52	0.033
	(0.44)	(0.022)	(0.99)	(0.039)	(0.76)	(0.043)	(0.99)	(0.051)	(0.72)	(0.039)
IsRamadan	0.48	0.031	0.53	0.037	1.18	0.114	0.72	0.005	0.27	0.026
	(1.09)	(0.054)	(1.20)	(0.045)	(2.94)	(0.156)	(1.63)	(0.109)	(2.45)	(0.121)
Muslim × Ramadan	-2.91^{**}	-0.177^{**}	−3.54	-0.174^{*}	−2.02	−0.265	−3.54	-0.355^{**}	−3.00	−0.070
	(1.27)	(0.079)	(2.48)	(0.099)	(3.44)	(0.214)	(2.56)	(.165)	(1.96)	(0.113)
District × Round FEs	Y	Y								
District FEs			Y	Y	Y	Y	Y	Y	Y	Y
Month-Yr × Region FEs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	33,956		9,057		8,618		7,996		8,285	
Mean hrs	22.71		28.21		21.91		19.51		21.88	

Standard errors clustered by Enumeration Area within each round (the primary sampling unit) are in parentheses. Urban households excluded from analysis. Households with no working-age members completing any work in the last week also excluded. All regressions are weighted based on sampling weights. $^{*}p < .1$, $^{**}p < .05$, $^{***}p < .01$.

As we would expect, Ramadan has no significant effect on the labor supply of non-Muslims. Consistent with other research, on the other hand, Ramadan does appear to have a significant effect on labor supply for Muslims. In particular, Table 6 shows that Muslims on average work approximately 3 fewer hours per week during Ramadan than we would otherwise expect, corresponding to approximately a 20 percent decrease in the regression on

[31] Indeed, it could be a potentially useful extension of my analysis to study Ramadan effects on the extensive margin of labor supply by using a binary variable for "any work" in the last week as the outcome variable.

log-hours.[32] Of course, the number of hours worked may not fully reflect the productivity of those hours. If hourly output decreases as well, the overall productivity effects could be larger, in line with the 20–40 percent productivity reduction that Schofield (2020) estimates. Interestingly, Schofield (2020) does not find any significant reduction in labor supply among Muslim workers.

Compared to the effects on calories, there is little variation from round to round in the measured effect of Ramadan on labor supply. It is generally encouraging that we are able to detect an effect of Ramadan on labor supply for Muslims, as it suggests that our measures of whether interviews took place during Ramadan and whether households were Muslim are reasonably accurate. However, the fact that this effect is more difficult to detect within individual rounds gives reason to be cautious about the interpretation of null effects within each round of data: we may not have sufficient power to reliably uncover true variation in the size of the effect of Ramadan from one round to the next.

The fact that we do observe an overall negative effect of Ramadan on labor supply among Muslim households, while we did not see a consistent negative effect calories on calories, merits discussion. One possible explanation is that the decrease in labor supply is not in fact due to a reduction in calories, but perhaps due to abstaining from water, tending to religious or social obligations, or sleep loss from waking up before sunrise to eat. Another possibility could simply be that my measure of calories is noisy, and that that noise masks the true effects of Ramadan on calories. It is also possible, as previously discussed, that there is indeed a reduction in calories among working-aged people, which is causing them to reduce their labor supply, but that this effect is being masked by measuring calories at the household level. Finally, it is possible (consistent with the findings in Hu and Wang 2019) that hunger during daylight hours could make it more difficult to work, even if that hunger is temporary and fully counteracted via a large post-fast meal.

[32] Compared to the mean household across the whole sample, 3 hours corresponds to about a 13% reduction in the average worker's labor supply. This slight discrepancy in estimated percentage differences is likely attributable to a baseline counterfactual predicted non-Ramadan labor supply for Muslim households that is somewhat smaller than the full sample average after including district effects, month effects, etc.

5 Conclusion

In this paper, I investigate the effects of Ramadan on calorie consumption and labor supply among Muslim households in rural Malawi. I find no evidence of a decrease in calorie consumption during Ramadan on average. I do, however, find evidence that working-age people reduce their weekly work by about three hours, or nearly 20 percent, on average. This could potentially indicate that a reduction in labor supply during Ramadan is not contingent upon an overall reduction in caloric intake. It could also be that, because I measure calories at the household level, which could include several household members who would not be expected to fast, I am underestimating the reduction in calories among working-age adults.

Another potentially interesting finding is that, in one round of survey data in which Ramadan fell just before the onset of the annual hunger season, I do see a significant decrease in caloric intake among Muslim households during Ramadan. I also find some evidence of an increase in caloric intake during Ramadan when it overlaps with the main harvest. While I cannot be sure that this finding is reflective of a larger pattern, and cannot offer evidence for any particular explanation, this evidence is, if nothing else, consistent with the possibility that the post-sunset meal that is traditionally consumed each night of Ramadan might have important implications for the overall effects of the fast on caloric intake. It is also consistent with the idea that this meal might be quite large during times of abundance and scaled back in times of scarcity.

This observation could also tentatively offer insight into the nature of seasonal food insecurity. A purely rational actor might be expected to use the opportunity of the Ramadan fast as a way to increase savings and mitigate the harm of the hunger season. We might expect this to be the case particularly when Ramadan overlaps with the harvest, and there is simply more food available to be saved. The fact that we instead see no decrease in consumption during Ramadan when it overlaps with the harvest, and some decrease in relatively lean times could indicate some potential barriers, either institutional or psychological, that make it difficult for households to smooth consumption throughout the year.

There are several reasons to be cautious about drawing any firm conclusions from these findings. For one, it is possible that there is indeed a reduction in calories, but that because calories are measured at the household level and not every household member necessarily would be expected to fast, I am underestimating the true reduction in calories among individuals who are, in fact, fasting. It is also possible that the identifying assumption underlying this analysis is violated, and that there are meaningful differences between Muslim households and/or non-Muslim households observed during Ramadan versus those observed at other times of year. While I provide some evidence for the validity of this identifying assumption, it is difficult to prove conclusively that it is upheld. Finally, while there is possible evidence of a seasonal pattern in the effects of Ramadan, we have only four rounds of survey data to draw from, limiting our ability to reliably establish any particular seasonal pattern.

One other key contribution of this paper is a novel set of calorie metrics that I constructed from the Malawi IHS survey data. In comparing my constructed metrics to replication data supplied by other researchers, I specifically looked into some of the more extreme and unlikely measures generated in the replication data in order to identify and correct for notable sources of likely inaccuracy in the existing conversions. While there are still certainly some inaccurate measures in the data I generated, and some additional steps (detailed below) that I might take to refine the measures I constructed, I made an effort to correct any major inaccuracy that I was able to identify, which could be useful to other researchers who wish to work with these data.

There are a number of robustness checks that I was not able to conduct for the present analysis, but that might be useful to bolster the credibility of the metrics I have constructed and the findings I present. First, it would be useful to check how well the calorie measures I construct align with the households' own account of whether they had enough food to eat. For example, if a household with measured calories below 1000 calories per person per day reports having more than enough food, calories were likely mismeasured for that household. Similarly, if a household reports consuming 4000 calories per person per day but also reports not having had enough food, it is likely that the true quantity of food consumed was underestimated.

There are several additional checks that I could implement in order to validate the novel calorie consumption metrics that I constructed, including both the unit-to-kg conversions and the kg-to-calorie conversions. For the unit conversions, for example, in cases in which no subunit (e.g., small, medium, or large) was provided, and a subunit was required for conversion, I typically assigned a default subunit. I could instead use the empirical distribution of subunits for each unit and each food item to approximate a conversion probabilistically when no subunit is specified. This would also help to address the fact that no subunits were collected in the earliest round of survey data. Another consistency check could be to take into account the amount spent on a given quantity of food in cases when food was purchased. Though I would expect some level of place- and time-based variation in the price per kg of any particular item, doing this could help me identify potential biases in the unit-to-kg conversions, or extreme outliers in imputed quantities or prices. Finally, in cases when I could not match the specified unit to any kg conversion, I replaced the provided unit with a default unit. While I would contend that it is a good idea to provide a set of default units, I could potentially refine this procedure by allowing for multiple defaults for a given item depending on the quantity specified and by choosing the default unit based on more rigorous and less ad-hoc criteria.

As for the kg-to-calorie conversions, when different sources gave conflicting information about the kg-to-calorie conversion for a particular food item, I attempted to choose a measure that seemed like a reasonable consensus figure. I could instead implement a more rigorous methodology for addressing such conflicts. Additionally, in cases when a food item name was typed in rather than selected out of the seemingly quite comprehensive list, I generally used the conversion provided for "other" items within the given food type. While it is unlikely that I would be able to match all typed in items, I could attempt to look for some of the more common ones that are unmatched and try to get a more precise match.

Aside from improvements to the measures of caloric intake, it might also be useful to include additional controls in the Difference-in-Differences specifications in order to give a more precise counterfactual prediction of consumption particularly for households observed during Ramadan. It also might be helpful to refine the measures I use to control for seasonality, perhaps by defining the start of an agricultural season based on observations from the survey data, such as consumption of green maize or reports of having begun harvesting crops for the most recent season.

The findings that I have presented indicate a number of potentially interesting avenues for future research. First, it could be useful to look not just at calories consumed, but also different types of food consumed, in order to develop a clearer picture of what is actually happening during Ramadan or at other times of year. For example, if particular foods are consumed as part of the post-Ramadan meal, we might be able to uncover that in the data, and then to test whether we do in fact observe variation from year to year in the amount of those particular foods consumed. This could potentially help us to understand if this post-fast meal is changing in size and scope over time.

Another potentially interesting question would be to look at the distributional effects of Ramadan. The analysis conducted in this paper relied entirely on average consumption levels. However, there could potentially be interesting differences in the effects of Ramadan on people at the lower or higher ends of the income distribution. Because Ramadan is supposed to be a time of charitable giving, we might see that the poorest households are increasing consumption at this time, while in richer households consumption is decreasing.

It could also be potentially interesting to look not just at consumption but also at income and savings to see how these are affected by Ramadan, and to see if these follow similar seasonal patterns throughout each round of survey data. Finally, it could be interesting to test whether any of the patterns and mechanisms that I have discussed in this paper, particularly surrounding the relationship between Ramadan and the agricultural cycle, would hold up in other sources of data, whether it be a future round of the Malawi IHS survey, or data coming from other countries with similar seasonal consumption patterns. Specifically, incorporating additional data of this sort could give us a way to test whether Ramadan causes an increase or decrease in calorie intake in other contexts, and whether this lines up with seasonal patterns of consumption in a similar way to the data analyzed in this paper.

References

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.

Almond, D., B. Mazumder, and R. van Ewijk (2014): "In Utero Ramadan Exposure and Children's Academic Performance," *The Economic Journal*, 125, 1501–1533.

Barro, R. J. (1999): "Ramsey Meets Laibson in the Neoclassical Growth Model," *The Quarterly Journal of Economics*, 114, 1125–1152.

Basu, K. and M. Wong (2015): "Evaluating Seasonal Food Storage and Credit Programs in East Indonesia," *Journal of Development Economics*, 115, 200–216.

Beck, U., K. Pauw, and R. Mussa (2015): "Methods Matter: The Sensitivity of Malawian Poverty Estimates to Definitions, Data, and Assumptions," WIDER Working Paper 2015/126, WIDER.

Benjamin, D. J., J. J. Choi, and G. Fisher (2016): "Religious Identity and Economic Behavior," *Review of Economics and Statistics*, 98, 617–637.

Brune, L., X. Giné, J. Goldberg, and D. Yang (2011): "Commitments to Save: A Field Experiement in Rural Malawi," World Bank Policy Research Working Paper No. 5748, World Bank.

——— (2016): "Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi," *Economic Development and Cultural Change*, 64, 187–220.

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82, 1671–1748.

Burke, M., L. F. Bergquist, and E. Miguel (2018): "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets," *The Quarterly Journal of Economics*, 134, 785–842.

Campante, F. and D. Yanagizawa-Drott (2015): "Does Religion Affect Economic Growth and Happiness? Evidence from Ramadan," *The Quarterly Journal of Economics*, 130, 615–658.

Chirwa, E. W., A. Dorward, and M. Vigneri (2012): *Seasonality, Rural Livelihoods, and Development*, Routledge, chap. Seasonality and Poverty: Evidence from Malawi, 97–116, 1st ed.

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven (2010): *Portfolios of the Poor*, Princeton University Press.

Deaton, A. (1997): *The Analysis of Household Surveys : A Microeconometric Approach to Development Policy*, Baltimore, MD: Johns Hopkins University Press.

Deaton, A. and J. Drèze (2009): "Food and Nutrition in India: Facts and Interpretations," *Economic and Political Weekly*, 44, 42–65.

Deaton, A. and S. Zaidi (2002): *Guidelines for Constructing Consumption Aggregates for Welfare Analysis*, Washington, DC: World Bank, lSMS Working Paper No. 135.

Dercon, S. and P. Krishnan (2000): "Vulnerability, Seasonality and Poverty in Ethiopia," *Journal of Development Studies*, 36, 25–53.

Devereux, S., R. Sabates-Wheeler, and R. Longhurst, eds. (2012): *Seasonality, Rural Livelihoods, and Development*, Routledge, 1st ed.

Devereux, S., B. Vaitla, and S. Hauenstein Swan (2008): *Seasons Of Hunger: Fighting Cycles of Starvation Among the World's Rural Poor*, Pluto Press.

D'Haene, E., S. Desiere, M. D'Haese, W. Verbeke, and K. Schoors (2019): "Religion, Food Choices, and Demand Seasonality: Evidence from the Ethiopian Milk Market," *Foods*, 8, 167.

Duh, J. and D. Spears (2017): "Health and Hunger: Disease, Energy Needs, and the Indian Calorie Consumption Puzzle," *The Economic Journal*, 127, 2378–2409.

Ecker, O. and M. Qaim (2011): "Analyzing Nutritional Impacts of Policies: An Empirical Study for Malawi," *World Development*, 39, 412–428.

Eli, S. and N. Li (2015): "Caloric Requirements and Food Consumption Patterns of the Poor," Working Paper No. 21697, NBER.

——— (2020): "Caloric Intake and Energy Expenditures in India," *The World Bank Economic Review*.

Gibson, J., K. Beegle, J. D. Weerdt, and J. Friedman (2014): "What does Variation in Survey Design Reveal about the Nature of Measurement Errors in Household Consumption?" *Oxford Bulletin of Economics and Statistics*, 77, 466–474.

Gilbert, R., T. Benson, and O. Ecker (2019): "Are Malawian Diets Changing? An Assessment of Nutrient Consumption and Dietary Patterns Using Household-level Evidence from 2010/11 and 2016/17," IFPRI Malawi Strategy Support Program Working Paper 30, Washington, DC.

Giné, X. and D. Yang (2009): "Insurance, Credit, and Technology Adoption: Field Experimental Evidence from Malawi," *Journal of Development Economics*, 89, 1–11.

Gulek, A. (2021): "Driving While Hungry: The Effect of Fasting on Traffic Accidents," *SSRN Electronic Journal*.

Hallak, M. H. and M. Z. A. Nomani (1988): "Body Weight Loss and Changes in

Blood Lipid Levels in Normal Men on Hypocaloric Diets during Ramadan Fasting," *The American Journal of Clinical Nutrition*, 48, 1197–1210.

Haruvy, E. E., C. A. Ioannou, and F. Golshirazi (2018): "The Religious Observance of Ramadan and Prosocial Behavior," *Economic Inquiry*, 56, 226–237.

Hastings, J. and E. Washington (2010): "The First of the Month Effect: Consumer Behavior and Store Responses," *American Economic Journal: Economic Policy*, 2, 142–162.

Haytowitz, D. B., J. K. Ahuja, X. Wu, M. Somanchi, M. Nickle, Q. A. Nguyen, J. M. Roseland, J. R. Williams, K. Y. Patterson, Y. Li, and P. R. Pehrsson (2019): "USDA National Nutrient Database for Standard Reference, Legacy Release," Available at data.nal.usda.gov/dataset/usda-national-nutrient-database-standard-reference-legacy-release.

Headey, D. and O. Ecker (2013): "Rethinking the Measurement of Food Security: From First Principles to Best Practice," *Food Security*, 5, 327–343.

Hodler, R., P. Raschky, and A. Strittmatter (2020): "Religion and Terrorism: Evidence from Ramadan Fasting," Unpublished manuscript.

Hoverd, W. J. and C. G. Sibley (2013): "Religion, Deprivation and Subjective Wellbeing: Testing a Religious Buffering Hypothesis," *International Journal of Wellbeing*, 3, 182–196.

Hu, Z. and Z. Wang (2019): "Nutrition, Labor Supply, and Productivity: Evidence from Ramadan in Indonesia," *SSRN Electronic Journal*.

International Astronomical Center (2021): "Islamic Crescents Observation Project (ICOP): Crescent Observation Results," Accessed at www.astronomycenter.net/res.html.

International Food Policy Research Institute (IFPRI) (2020): "Replication Data for: Estimates of Nutrient Availability From the Integrated Household Surveys in Malawi," Harvard Dataverse.

Iyer, S. (2016): "The New Economics of Religion," *Journal of Economic Literature*, 54, 395–441.

Jayachandran, S. (2006): "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," *Journal of Political Economy*, 114, 538–575.

Karaağaoğlu, N. and S. Yücecan (2000): "Some Behavioural Changes Observed among Fasting Subjects, Their Nutritional Habits and Energy Expenditure in Ramadan," *International Journal of Food Sciences and Nutrition*, 51, 125–134.

Kaur, S., M. Kremer, and S. Mullainathan (2015): "Self-Control at Work," *Journal of Political Economy*, 123, 1227–1277.

Khandker, S. R. and W. Mahmud (2012): *Seasonal Hunger and Public Policies: Evidence from Northwest Bangladesh*, The World Bank.

Kuran, T. (2018): "Islam and Economic Performance: Historical and Contemporary Links," *Journal of Economic Literature*, 56, 1292–1359.

Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 112, 443–478.

Lee, S., M. Nam, D. Jeong, and W. Lee (2020): "Does Ramadan Harm Infant Health? Evidence from Ethiopia," *International Economic Journal*, 34, 613–633.

MAFOODS (2019): *Malawian Food Composition Table*, Lilongwe, Malawi: South African Medical Research Council, Biostatistics Unit, 1st ed.

Majid, M. F. (2015): "The Persistent Effects of In Utero Nutrition Shocks Over the Life Cycle: Evidence from Ramadan Fasting," *Journal of Development Economics*, 117, 48–57.

McCleary, R. M. and R. J. Barro (2006): "Religion and Economy," *Journal of Economic Perspectives*, 20, 49–72.

Munshi, K. and M. Rosenzweig (2016): "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap," *American Economic Review*, 106, 46–98.

NSO (National Statistics Office) (2005a): "Note on Construction of Expenditure Aggregate and Poverty Lines for IHS2," Unpublished documentation.

——— (2005b): "Second Integrated Household Survey 2004–05 (IHS2)." Zomba, Malawi: National Statistics Office.

——— (2012): "Third Integrated Household Survey 2010–11 (IHS3)." Zomba, Malawi: National Statistics Office.

——— (2018): "Fourth Integrated Household Survey 2016–17 (IHS4)." Zomba, Malawi: National Statistics Office.

——— (2021): "Fifth Integrated Household Survey 2019–20 (IHS5)." Zomba, Malawi: National Statistics Office.

Pauw, K., U. Beck, and R. Mussa (2014): "Did Rapid Smallholder-Led Agricultural Growth Fail to Reduce Rural Poverty? Making Sense of Malawi's Poverty Puzzle," WIDER Working Paper No. 123.

Paxson, C. H. (1993): "Consumption and Income Seasonality in Thailand," *Journal of Political Economy*, 101, 39–72.

Rosenzweig, M. R. and K. I. Wolpin (1993): "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in

Low-Income Countries: Investments in Bullocks in India," *Journal of Political Economy*, 101, 223–244.

Sadeghirad, B., S. Motaghipisheh, F. Kolahdooz, M. J. Zahedi, and A. A. Haghdoost (2012): "Islamic Fasting and Weight Loss: A Systematic Review and Meta-analysis," *Public Health Nutrition*, 17, 396–406.

Schneider, K. (2021): "Nationally Representative Estimates of the Cost of Adequate Diets, Nutrient Level Drivers, and Policy Options for Households in Rural Malawi," *SSRN Electronic Journal*.

Schofield, H. (2014): "The Economic Costs of Low Caloric Intake: Evidence from India," Job Market Paper, Harvard University.

——— (2020): "Ramadan Fasting and Agricultural Output," Unpublished manuscript.

Shalihin, N., F. Firdaus, Y. Yulia, and U. Wardi (2020): "Ramadan and Strengthening of the Social Capital of Indonesian Muslim Communities," *HTS Theological Studies*, 76.

Shapiro, J. M. (2005): "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 89, 303–325.

Smith, L. C., O. Dupriez, and N. Troubat (2014): "Assessment of the Reliability and Relevance of the Food Data Collected in National Household Consumption and Expenditure Surveys," IHSN Working Paper No. 008, International Household Survey Network.

Subramanian, S. and A. Deaton (1996): "The Demand for Food and Calories," *Journal of Political Economy*, 104, 133–162.

Toda, M. and K. Morimoto (2004): "Ramadan Fasting – Effect on Healthy Muslims," *Social Behavior and Personality*, 32, 13–18.

Townsend, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, 62, 539.

van Ewijk, R. (2011): "Long-term Health Effects on the Next Generation of Ramadan Fasting during Pregnancy," *Journal of Health Economics*, 30, 1246–1260.

Verduzco-Gallo, I., O. Ecker, and K. Pauw (2014): "Changes in Food and Nutrition Security in Malawi: Analysis of Recent Survey Evidence," IFPRI Malawi Strategic Support Program Working Paper 06.

Wang, A., S. Wang, and X. Ye (2020): "Religion and Motivated Cognition: When Ramadan Meets the College and Entrance Exam," Unpublished manuscript.

Ziaee, V., M. Razaei, Z. Ahmadinejad, H. Shaikh, R. Yousefi, L. Yarmohammadi, F. Bozorgi, and M. J. Behjati (2006): "The Changes of Metabolic Profile and Weight during Ramadan Fasting," *Singapore Medical Journal*, 47, 409–414.